TK5105.7
.A7
v.2
c. 1 aa

IC

HICKLING

FEASIBILITY STUDY OF A NATIONAL HIGH SPEED COMMUNICATIONS NETWORK FOR RESEARCH, DEVELOPMENT AND EDUCATION

VOLUME A: PARTICIPANT NEEDS

FEASIBILITY STUDY OF A NATIONAL HIGH SPEED COMMUNICATIONS NETWORK FOR RESEARCH, DEVELOPMENT AND EDUCATION

VOLUME A: PARTICIPANT NEEDS

Submitted to:

INDUSTRY, SCIENCE AND TECHNOLOGY CANADA

Prepared by:

HICKLING
Science and Technology Division

and

COMGATE ENGINEERING ASSOCIATES LTD.

INDUSTRY, SCIENCE AND TECHNOLOGY CANADA LIBRARY
NOV 21 1990
BIBLIOTHÈQUE INDUSTRIE, SCIENCES ET TECHNOLOGIE CANADA

In association with:

THE ALBERTA RESEARCH COUNCIL

THE CGI GROUP

LANG MICHENER LAWRENCE & SHAW

HICKLING REF:3201
March, 1990

NOTE

This is Volume A of a study prepared by James F. Hickling Management Consultants Ltd. (HICKLING) on behalf of Industry Science and Technology Canada (ISTC), entitled "Feasibility Study of a High Speed Communications Network For Research, Development and Education". There are five volumes in this study:

1. Main Report
2. Volume A: Participant Needs
3. Volume B: Economic Analysis
4. Volume C: Technical Analysis
5. Volume D: Implementation Analysis

PREFACE

This study was commissioned by Industry, Science and Technology Canada to investigate the feasibility of establishing a National High–Speed Communications Network for the Canadian research, development and education communities. The network would have greater capacity and functionality than existing networks. While the undertaking of this study is not to be construed as a commitment by the federal government to the establishment of a network, the study will provide a solid basis for such an initiative should it be found prudent.

HICKLING is indebted to Dr. Digby Williams, Director of the Microelectronics Technology Office, and Mr. Joseph Padden, Senior Technologies Advisor for the Information Technologies Industry Branch, Industry, Science and Technology Canada, for their expert technical and managerial advice in the conduct of this study. The authors would also like to offer thanks to the more than 400 individuals who participated in expert panel sessions, in–person interviews, and surveys; the study would not have been possible without their input. Of course, any errors or omissions are the sole responsibility of HICKLING.

The report was authored by David Arthurs, Phil Kennis, and Daniel Hara of HICKLING; and Roger Choquette and Anthony Capel of COMGATE. Significant contributions were made by Dr. Saul Greenberg of the Alberta Research Council; Dr. Frederick Eshragh, Dr. Kalman Toth, and Dr. Samy Mahmoud of CGI; John Lawrence and Andree Wylie of Lang Michener Lawrence & Shaw; Elmer Hara of the University of Regina; and Fred Cassedai.

OVERVIEW

This volume, Volume A: Participant Needs, addresses the current and future needs of the various participants of a Canadian national high speed communications network for research, development, and education (referred to as the Network), and is presented in five chapters.

Chapter One contains a summary of the needs of the potential participants in this network. Also included in this section is a brief explanation of the information gathering process, the manner and rationale for the segmentation of the various participants, and a detailed list of the needs of each participant from both a technical and implementation point of view.

Chapter Two presents the results of the different instruments used to collect information over the course of the study. Details regarding who was contacted and what was said are documented in this chapter. Some interviews have been deemed to be confidential, and thus records of these sessions are not included in this package. Copies of the blank surveys used to collect information are also included here.

Chapter Three offers a detailed analysis of future networking needs and applications. Future applications of the Network must be anticipated so that the evolution of the Network can be properly planned.

Chapter Four contains a relatively brief summary of the trends in Information Equipment. In order to specify desirable features of a future national telecommunications network for R&D, a broad understanding of the anticipated developments in video, computer hardware and software, Local Area Network (LAN) and transmission technologies must be established.

The last chapter, Chapter Five, is an international review of what is happening in Canada, U.S., Europe and Japan with respect to R&D networking. This information is relevant to the theme of this volume as it is useful to understand how other countries are meeting the needs of its researchers.

TABLE OF CONTENTS

PAGE

NOTE i

PREFACE iii

OVERVIEW v

CHAPTER 1: SUMMARY OF PARTICIPANT NEEDS . . . . . 1 – 1
1. INTRODUCTION . . . . . 1 – 3
2. THE INFORMATION GATHERING PROCESS . . . . . 1 – 5
   2.1 EXPERT PANEL SESSIONS . . . . . 1 – 5
   2.2 IN–PERSON INTERVIEWS . . . . . 1 – 6
   2.3 NETWORK QUESTIONNAIRE . . . . . 1 – 6
   2.4 TELEPHONE INTERVIEWS . . . . . 1 – 7
3. NETWORK PARTICIPANTS . . . . . 1 – 9
   3.1 INTRODUCTION . . . . . 1 – 9
   3.2 CLASSIFICATION OF PARTICIPANTS . . . . . 1 – 9
4. CURRENT NEEDS OF NETWORK PARTICIPANTS . . . . . 1 – 13
   4.1 R&D USERS . . . . . 1 – 13
   4.2 IT USERS . . . . . 1 – 17
   4.3 IT SUPPLIERS . . . . . 1 – 20
   4.4 CARRIERS . . . . . 1 – 21
   4.5 EXISTING REGIONAL NETWORKS . . . . . 1 – 22
   4.6 PROVINCIAL GOVERNMENT . . . . . 1 – 24
   4.7 FEDERAL GOVERNMENT . . . . . 1 – 25
5. FUTURE NEEDS AND APPLICATIONS . . . . . 1 – 29

CHAPTER 2: INFORMATION GATHERING . . . . . 2 – 1
1. INTRODUCTION . . . . . 2 – 3
2. EXPERT PANEL SESSIONS . . . . . 2 – 5
   2.1 METHODOLOGY . . . . . 2 – 5
   2.2 RESULTS . . . . . 2 – 7
3. IN–PERSON INTERVIEWS . . . . . 2 – 9
   3.1 METHODOLOGY . . . . . 2 – 9
   3.2 RESULTS . . . . . 2 – 9
4. NETWORK QUESTIONNAIRE . . . . . 2 – 13
   4.1 METHODOLOGY . . . . . 2 – 13
   4.2 RESULTS . . . . . 2 – 13
5. TELEPHONE INTERVIEWS . . . . . 2 – 19
   5.1 METHODOLOGY . . . . . 2 – 19
   5.2 RESULTS . . . . . 2 – 19

CHAPTER 3: FUTURE APPLICATIONS . . . . . 3 – 1
1. A FRAMEWORK FOR FUTURE APPLICATIONS . . . . . 3 – 3
   1.1 INTRODUCTION . . . . . 3 – 3
   1.2 THE PROCESS OF SCIENTIFIC DISCOVERY . . . . . 3 – 3
   1.3 COLLABORATIVE FUNCTIONS AND TOOLS . . . . . 3 – 6
   1.4 OUTLINE OF THE REPORT . . . . . 3 – 7

2. SUPPORT FOR REAL–TIME REMOTE CONFERENCING . . . . . 3 – 9
2.1 INTRODUCTION . . . . . 3 – 9
2.2 TELE–PRESENCE . . . . . 3 – 10
2.3 TELE–DATA . . . . . 3 – 11
2.4 CYBERSPACE . . . . . 3 – 13
2.5 LIMITATIONS AND UNKNOWNS . . . . . 3 – 13
2.6 NETWORK REQUIREMENTS . . . . . 3 – 14
3. CASUAL REAL–TIME INTERACTION . . . . . 3 – 17
3.1 INTRODUCTION: PATTERNS AND REQUIREMENTS FOR SCIENTIFIC RESEARCH COLLABORATION . . . . . 3 – 17
3.2 VIDEO HALLWAYS . . . . . 3 – 18
3.3 SHARED ALTERNATE REALITY . . . . . 3 – 20
3.4 LIMITATIONS AND UNKNOWNS . . . . . 3 – 20
3.5 NETWORK REQUIREMENTS . . . . . 3 – 21
4. ASYNCHRONOUS MESSAGING . . . . . 3 – 23
4.1 INTRODUCTION . . . . . 3 – 23
4.2 BASIC ELECTRONIC MAIL . . . . . 3 – 24
4.3 MULTI–MEDIA MAIL . . . . . 3 – 26
4.4 SEMI–STRUCTURED AND SEMI–FORMAL MAIL . . . . . 3 – 28
4.5 NETWORK REQUIREMENTS . . . . . 3 – 31
5. BULLETIN BOARDS AND ASYNCHRONOUS CONFERENCING . . . 3 – 33
5.1 INTRODUCTION . . . . . 3 – 33
5.2 FUNCTIONS AND ISSUES . . . . . 3 – 34
5.3 LIMITATIONS AND UNKNOWNS . . . . . 3 – 37
5.4 NETWORK REQUIREMENTS . . . . . 3 – 38
6. ACCESS, MONITORING AND OPERATION OF EQUIPMENT . . . . . 3 – 39
6.1 INTRODUCTION . . . . . 3 – 39
6.2 COMPUTATIONAL ACCESS . . . . . 3 – 40
6.3 TECHNOLOGIES FOR MONITORING AND OPERATING REMOTE EQUIPMENT . . . . . 3 – 40
6.4 NETWORK REQUIREMENTS . . . . . 3 – 42
7. DIGITAL LIBRARIES . . . . . 3 – 43
7.1 INTRODUCTION . . . . . 3 – 43
7.2 CONVENTIONAL APPROACHES . . . . . 3 – 44
7.3 HYPERTEXT AND HYPERMEDIA . . . . . 3 – 45
7.4 LIMITATIONS AND UNKNOWNS . . . . . 3 – 46
7.5 NETWORK REQUIREMENTS . . . . . 3 – 47
8. APPLICATION AREAS . . . . . 3 – 49
8.1 INTRODUCTION . . . . . 3 – 49
8.2 DISTANCE EDUCATION . . . . . 3 – 49
8.3 GROUP DECISION SUPPORT SYSTEMS . . . . . 3 – 50
8.4 JOINT AUTHORING . . . . . 3 – 50
8.5 BOOTSTRAPPING COLLABORATIONS . . . . . 3 – 51
9. THE COHERENT WORKSPACE . . . . . 3 – 53
10. REFERENCES . . . . . 3 – 55

CHAPTER 4: TRENDS IN INFORMATION EQUIPMENT . . . . . 4 – 1
1. INTRODUCTION . . . . . 4 – 3
2. TECHNOLOGY BACKGROUND . . . . . 4 – 5
2.1 PERSONAL COMPUTERS . . . . . 4 – 5
2.2 COMPUTER MAINFRAMES . . . . . 4 – 5
2.3 CAD/CAM . . . . . 4 – 6
2.4 MULTIMEDIA PC TERMINALS . . . . . 4 – 6
2.5 OPTICAL DISC MEMORIES . . . . . 4 – 6
2.6 VIDEODISC SYSTEMS . . . . . 4 – 7

2.7 DIGITAL TV SETS . . . . . 4 – 8
2.8 VIDEO TELECONFERENCING . . . . . 4 – 9
2.9 LOCAL AREA NETWORKS . . . . . 4 – 10
2.10 HIGH DEFINITION TELEVISION AND HIGH RESOLUTION IMAGES . . . . . 4 – 11

CHAPTER 5: INTERNATIONAL REVIEW . . . . . 5 – 1
1. COMPARATIVE SUMMARY OF R&D NETWORKING . . . . . 5 – 3
1.1 R&D AND NETWORKING ENVIRONMENT . . . . . 5 – 3
1.2 CURRENT NETWORKING STATUS . . . . . 5 – 5
1.3 FUTURE PLANS . . . . . 5 – 7
2. R&D NETWORKS IN CANADA . . . . . 5 – 11
2.1 SUMMARY . . . . . 5 – 11
2.2 CANADIAN R&D ENVIRONMENT . . . . . 5 – 14
2.3 SCOPE OF THIS REVIEW . . . . . 5 – 15
2.4 SUMMARY OF NETWORKS IN CANADA . . . . . 5 – 16
2.5 EXISTING REGIONAL NETWORKS . . . . . 5 – 21
2.6 EXISTING NATIONAL NETWORKS . . . . . 5 – 37
2.7 FUTURE REGIONAL NETWORKS . . . . . 5 – 40
2.8 FUTURE NATIONAL NETWORKS . . . . . 5 – 41
3. R&D NETWORKS IN THE UNITED STATES . . . . . 5 – 45
3.1 SUMMARY . . . . . 5 – 45
3.2 U.S R&D ENVIRONMENT . . . . . 5 – 48
3.3 SCOPE OF THIS REVIEW . . . . . 5 – 51
3.4 SUMMARY OF NETWORKS IN UNITED STATES . . . . . 5 – 53
3.5 EXISTING NETWORKS . . . . . 5 – 61
3.6 FUTURE NETWORKS . . . . . 5 – 94
4. R&D NETWORKS IN EUROPE . . . . . 5 – 97
4.1 SUMMARY . . . . . 5 – 97
4.2 EUROPEAN R&D OVERVIEW . . . . . 5 – 97
4.3 EUROPEAN NETWORK ENVIRONMENT . . . . . 5 – 98
4.4 SCOPE OF THIS REVIEW . . . . . 5 – 101
4.5 PAN–EUROPEAN NETWORKS . . . . . 5 – 103
4.6 MAJOR NATIONAL NETWORKS . . . . . 5 – 109
4.7 NEW FUTURE NETWORKS IN EUROPE . . . . . 5 – 122
5. JAPANESE R&D NETWORKS . . . . . 5 – 125
5.1 SUMMARY . . . . . 5 – 125
5.2 JAPANESE R&D ENVIRONMENT . . . . . 5 – 126
5.3 SCOPE OF THIS REVIEW . . . . . 5 – 127
5.4 SUMMARY OF NETWORKS IN JAPAN . . . . . 5 – 128
5.5 JAPANESE NETWORK DESCRIPTIONS . . . . . 5 – 128

APPENDIX A: MINUTES OF EXPERT PANEL SESSIONS . . . . . A – 1

APPENDIX B-1: ON-SITE INTERVIEWS WITH IT RESEARCH ORGANIZATIONS . . . . . B – 1

APPENDIX B-2: ON-SITE INTERVIEWS WITH PROVINCIAL GOVERNMENTS . . . . . B – 11

APPENDIX C: NETWORK USER SURVEY . . . . . C – 1

APPENDIX D-1: CONTACTS MADE FOR PHONE INTERVIEWS . . . . . D – 1

APPENDIX E-1: LETTERS FROM PROVINCES AND UNIVERSITIES . . . . . E1 – 1

APPENDIX E-2: LETTERS FROM COMPANIES ..... E2 – 1

APPENDIX F: GLOSSARY ..... F – 1

CHAPTER 1: SUMMARY OF PARTICIPANT NEEDS

1. INTRODUCTION

The Network will be technically designed, operated and managed to satisfy, to the greatest extent possible, the needs of the Network participants. These needs will evolve over time and the Network must continue to meet these needs from 1991 into the future. The definition and description of the needs of Network participants, implies:

- A clear view of who these potential Network participants are; and
- An understanding of the objectives that these individuals or organisations may wish to achieve through their participation.

At the outset, the Study Team interpreted the term "participant" in its broadest sense, classifying potential participants as either service users, service suppliers, or sponsors. This wide interpretation proved to be critical in eventually capturing the many needs that a network as envisaged by this Study must satisfy.

The purpose of this chapter is to report on the needs which must be satisfied by the Network.

2. THE INFORMATION GATHERING PROCESS

Four main instruments were used by the study team to identify potential participants and obtain information on their needs and objectives:

- Expert Panel Sessions;
- In person interviews;
- A network questionnaire distributed via Netnorth and CDNnet; and
- Telephone Interviews.

Every effort was made to contact as many potential participants as possible. In fact, the project team received direct input from more than 400 individuals who took part in one of the four survey instruments identified above. The project team has discussed this network initiative with government officials in every province and territory in Canada, and also with Federal Government representatives (eg. NRC, DOC, EMR, GTA and CRTC).

The detailed results from each of these information gathering initiatives are documented in Chapter 2. A brief description of each instrument is provided below:

2.1 EXPERT PANEL SESSIONS

The expert panel sessions were organized as informal round–table discussions, designed to collect information from the large base of potential participants in the Network. Each of the four sessions, two in Ottawa and one in both Montreal and Vancouver, had its own regional flavour and was very useful in developing a clearer understanding of the wide ranging needs, interests and concerns of potential participants in this Network. These sessions provided excellent background information and the information gained through exchanges among panelists was something that could not have been captured any other way. However more detailed information had to be collected through on–site visits or through questionnaires.

It was decided that the expert panel sessions should be carried out in Ottawa (drawing from Toronto, Hamilton, Guelph, Waterloo, London and Washington), Montreal (drawing from Halifax) and Vancouver (drawing from Calgary and Edmonton), because the majority of researchers in Canada are located in these cities. The needs of those regions of Canada which were not involved in expert panel sessions should not be interpreted in any way as being less important or less relevant to this study. The needs of businesses and universities in all the provinces have been taken into account, through other survey instruments.

Details regarding these sessions, the minutes of the sessions, as well as a list of the panelists are included in Chapter 2. In total, 44 individuals participated, and the group can be segmented as follows:

Affiliation	Number of Participants	Percentage of Total
Industry	12	27%
Universities	12	27%
Government	10	23%
Existing Networks	6	14%
Research Organizations	4	9%
TOTAL	44	100%

At each session, representatives from the regional network, local universities, local research centres, local information technology companies and CAnet advisory board were in attendance, and participated actively in the discussion period. In addition, at some of the sessions representatives from NSERC, Centres of Excellence and the Science Council were in attendance to observe the exchange of ideas among the panelists.

2.2 IN-PERSON INTERVIEWS

It was considered essential to the information gathering process that more detailed information be collected through on–site visits to: Provincial Governments; Information Technology Companies (eg. BNR, MPR, Newbridge etc.); national and provincial carriers (Telecom Canada, CNCP, Telesat, AGT etc.); research organizations (Alberta Telecommunications Research Centre); designated Centres of Excellence (Precarn); and Universities. These visits provided an excellent opportunity to brief senior officials within government and companies on the concept of this network and to share some preliminary results. In addition, these visits provided a means to acquire a more thorough understanding of the needs of each of the players within this diverse group, and provided insight into their willingness to contribute to this venture.

Officials within all ten provinces, and both territories were contacted[1] regarding this initiative. The Provincial sessions, for the most part, were very well attended, and confirmed the high level of interest which exists across Canada in research networks. For example, in many of these meetings (Alberta, Manitoba, Saskatchewan and Ontario) representatives from industry and universities were also in attendance.

In addition to the provincial meetings, 28 companies were interviewed on–site, across Canada. Over 50% of these interviews were carried out in Ontario, as the majority of IT companies and major carriers are located in the Toronto and Ottawa areas. Approximately 20% were carried out in Quebec, and 10% in British Columbia. The remainder of these interviews took place in Alberta, Saskatchewan, Manitoba and Nova Scotia.

All information exchanged during the course of these meetings is considered to be confidential, and thus the minutes of these meetings are not available. However all companies contacted for in–person interviews are listed in the appropriate section of Chapter 2.

2.3 NETWORK QUESTIONNAIRE

A questionnaire was distributed through Netnorth and CDNnet to receive feedback from Canadian researchers who currently use the research networks available in Canada. Survey respondents provided an indication of the applications which they require (which they may or may not have access to at the present time), their requirement for a more capable network, and the general benefits of a high–speed communications network. A copy of the questionnaire is contained in Chapter 2.

Over 200 respondents completed the questionnaire. The profile of these respondents includes 60% from university professors, 18% from publicly funded research laboratories (primarily NRC), 13% from graduate students, 8% from network managers at universities, and only 1% from private sector companies engaged in R&D. A full 40% of all university responses

[1] Newfoundland, New Brunswick, PEI and both the Territories were contacted by conference call.

(includes professors, students and network managers) originated from British Columbia — which indicates significant interest in that province with respect to networks.

The main research disciplines of those who responded to the network survey are: physics (20%), communication technologies (18%), computer science (14%), earth science (8%), and mathematics (8%).

Over 70% of respondents indicated that they would like to have access to a more capable network. For a further discussion of the results please refer to Chapter 2.

2.4 TELEPHONE INTERVIEWS

As small IT companies were not polled by the expert panel sessions, the in–person interviews, or the network survey, it was judged that additional phone interviews were required to adequately cover this user group.

Over 70 organizations, all across Canada (ie. every province), were contacted by phone to answer a short questionnaire regarding this study. After briefing the respondent, a questionnaire was faxed to the participant, and the completed copy was in turn faxed back to the project team. About 32 organizations returned the survey, a response rate of about 45%; and about one half of those who returned the survey viewed this initiative as worthwhile, and something they would want to be a part of in one way or another.

For a further discussion of these results please refer to Chapter 2.

3. NETWORK PARTICIPANTS

3.1 INTRODUCTION

As noted earlier, the Study team interpreted the term "participant" in its broadest sense, classifying potential participants as either service users, service suppliers, or sponsors. Users were further classified into two categories: General R&D Users who would use the Network to carry out R&D in fields other than information technologies and Information Technology Users (IT Users) who would use the Network to carry out research and development related to information technologies, including communications.

It has been confirmed that R&D Users would use the Network as a means of accessing and/or communicating with data, fellow researchers or facilities in support of research efforts, and that IT Users would use the Network as a "Test Bed" for the development of IT products and services. These types of usages are in line with the two main objectives of the Network:

- To interconnect R&D organizations with each other, databases and special equipment
- To provide an environment to test information technologies equipment and services.

However, it became apparent as the study progressed that a finer classification of participants would be required to capture the wide range of needs associated with the broad constituency of support being sought. It also became clear that participant needs involved both the needs of individual users as well as those of organizations and institutions which were directly or indirectly supporting them. Both a "top down" view which was sensitive to the needs of all organizations, including existing R&D networks, was required to complement the "bottom up" view which focused on individual user requirements.

3.2 CLASSIFICATION OF PARTICIPANTS

Network participants have been classified into 7 groups, and have been identified as either service users, service providers or sponsors. An organization may belong to more than one group. This classification scheme is described below.

I. SERVICE USERS

1. **Research and Development Users (R&D Users).** R&D Users are defined as individual researchers doing research and development in areas other than IT such as medicine, chemistry, physics etc.. These users may be located at universities, provincial or federal research organizations or in private industry. They may be doing individual research or be affiliated with a research consortium (e.g. PRECARN). Small users are also included in this subgroup, and includes organizations who traditionally have not been able to afford access to the regional R&D networks.

2. **Information Technology Users (IT Users).** IT Users are defined as researchers in areas related to information technology. These users could be located at universities, publicly funded telecommunication research centres

(e.g. ATRC), or may be private–sector companies engaged in IT research (e.g. BNR). This subgroup also explicitly accounts for small IT companies who traditionally have not been able to afford their own test networks or access to the regional networks.

II. SERVICE PROVIDERS

3. **Information Technology Suppliers (IT Suppliers).** IT suppliers include organizations that provide IT based products (routers, user interfaces to the Network) or services (access to specialized databases, or journals etc.), excluding the Carriers.

4. **Common Carriers (Carriers).** Carriers are defined as regulated common carriers capable of providing communications service or facilities (link channel capacity) to the Network on a national scale. Carriers include the Telecom Canada member companies, CN/CP and Telesat Canada. Value added resellers (e.g. INSINC, Fonorola, Callnet) are not included in this category but are identified in the IT Supplier group above.

5. **Existing R&D Networks (R&D Networks).** Existing R&D networks (i.e. Subnetworks in the context of the Network discussion) include the academic networks currently operating in Canada. These networks are associated with organizational structures for purposes of policy, management and operation. These organizations have objectives and needs which must be addressed in their own right and which in some cases are over and above the needs of individual researchers or organizations supporting an attached device or service.

III. SPONSORS

6. **Provincial Government.** Every Province and Territory in Canada would be invited and encouraged to access this network. For the moment, regional networks are defined geographically by province boundaries, and thus the needs of the Province are closely aligned with the needs of the Regional Networks.

7. **Federal Government.** Should this network be implemented, Industry Science and Technology Canada (ISTC) would be the lead Federal Department guiding the implementation and overseeing the management structure.

In addition to the classification described above, there are additional dimensions by sector (academic, private, or government) and by region (each province and territory). Figure 1 provides a summary of all the different dimensions of the participants in this network.

Figure 1: NEEDS

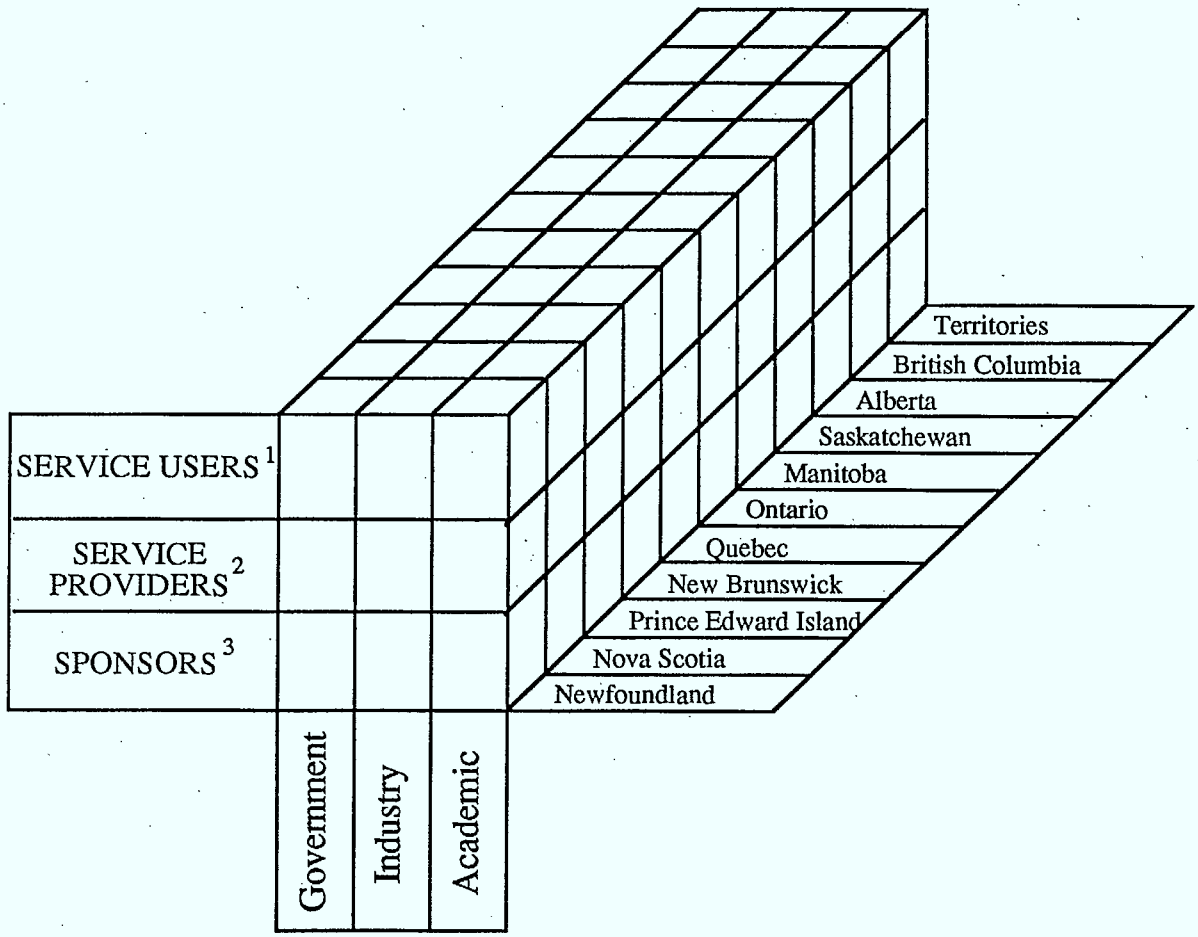

Notes:

[1] Service Users include R&D Users and IT Users
[2] Service Providers include IT Suppliers, Carriers and Regional Networks
[3] Sponsors include Federal and Provincial Governments

4. CURRENT NEEDS OF NETWORK PARTICIPANTS

This section presents a summary of the current needs of all the Network participants, and represents an amalgamation of all the information collected by the four different collection methodologies identified above.

The current needs of each user group has an impact on the technical and/or implementation design of the Network. As a result, user needs are categorized under these two broad headings (see below). The reader should be aware that this section is a summary of the opinions of each network participant, and these opinions may or may not coincide with the recommendations of the project team. The recommendations made by the project team represent a compromise between all of the opinions of all potential participants.

Table 2 provides a summary of the needs and concerns of each of the different participants.

4.1 R&D USERS

R&D Users represent the largest community of users of the Network. They may be individual researchers or be affiliated with an organization. This group encompasses researchers in all areas other than IT such as medicine, chemistry, physics, social science, humanities etc.. These users are currently located primarily at universities, but may also be from provincial or federal research organizations, or industry. In addition, the membership of the Centres of Excellence span the entire nation, and thus these centres have a demonstrated need for such a network, and represent a very important pocket of users within this group. R&D Users are typically connected via a workstation to a Local Area Network (LAN) operating in the 10 Mbps range. The Network requirement is to provide wide area connectivity to the LAN as well as user access to appropriate Network services.

The priorities of R&D Users relate to the provision of "user friendly" service with a high degree of connectivity in support of their R&D activity. The R&D User is most concerned with being able to use a service when he wants to use it. He wants reliable communications at reasonable cost and bandwidth on demand. He does not want to concern himself with the details of how communications and connectivity is achieved.

The availability of large amounts of bandwidth (i.e. T1) on demand will be required for certain applications. The Network architecture will have to address the requirement to provide users with large amounts of bandwidth for short periods of time on demand (i.e. "bursty" applications). This is particularly true for applications which support the transfer of graphical information and images in a conference mode (e.g. screen sharing) and those which are related to the transfer of very large data files (e.g. large data file dumps in real time from remote instruments).

The Network will also be required to support a large volume of electronic mail transactions.[2]

---

[2] Some carriers even suggested that electronic mail from the academic community could easily saturate any amount of available bandwidth.

TABLE 2: THE NEEDS OF NETWORK PARTICIPANTS

CONCERNS	SERVICE USERS		SERVICE PROVIDERS			SPONSORS	
	R&D USERS	IT USERS	IT SUPPLIERS	REGIONALS	CARRIERS	PROVINCIAL GOV'T	FEDERAL GOV'T (ISTC)
TECHNICAL	1. Connectivity to Regional, National and International Networks 2. Reliability 3. Access to basic services 4. Access to advanced services 5. High capacity 6. Capable of multi-media 7. User Friendly 8. Low cost access alternatives must be made available	1. Access to different levels of the network 2. High capacity (T1 as a minimum) 3. Access to real world Beta test 4. Access to all transmission media 5. Network intelligence contained on the edges of network 6. Migration to ISO 7. Smooth technology migration to T3 and higher speeds 8. Multi mode/media 9. Connectivity 10. Security 11. Low cost access	1. Open to multiple vendors 2. Beta Test site 3. State of the art equipment & services 4. Open to third party Value-Added vendors (database)	1. Inter-operability with other regional, national and international networks 2. Coexistence of TCP/IP & OSI 3. Only one sub-network for a region	1. Use existing transmission services 2. Network must consist of all types of transmission media 3. Network must interface with carrier offerings (eg. ISDN)	1. Regional Networks must continue to exist.	1. Support ISO/OSI standards 2. Trans-Canada end-to-end digital connectivity 3. Minimum technical capacity of 1.5 Mbs 4. State of the art protection and network management features 5. Offer test facilities
IMPLEMENTATION	1. Elimination of Duplication /Single initiative /Single management structure 2. T1 Tariff Rates must be reduced 3. Distance independent fees Fixed - academic Variable - industry 4. Basic services included in user fees 5. Advanced services paid in addition to basic fee 7. Representation on board 8. Access to network permitted for non-commercial purposes based on honor system 9. Consistency with International use policies	1. Clearly defined mission statement 2. Elimination of Duplication /Single initiative /Single management structure 3. Must be implemented quickly 4. OSI conformance testing through accredited organizations (eg. CIGOS) 5. Preference for usage based fees 6. Willingness to pay preference for contributions in kind 7. Full-time business-like management required 8. Propriatory interests must be minimized 9. Access to network permitted for non-commercial purposes	1. Market development is the objective 2. Discounts on Services & Equipment for Marketing purposes 3. Interested in network management 4. Tax credits for contributions/R&D	1. Must be joint participants in management of backbone network 2. Must recognize and take advantage of experience and expertise of existing regionals 3. Predatory networks must be addressed 4. Back door connections must be minimized 5. Must meet needs of current client base 6. Formal marketing and management structures must exist 7. Direct connectionsto backbone must be minimized 8. User fees fixed for a level of service	1. Market development is the objective for participation 2. Would give consideration for reduced tariff rates 3. Re-selling is a concern 4. No commercial bypass (policed by honor system) 5. Voice may cause regulatory problems 6. Interested in network management 7. Desire formal recognition of contribution (reflected in rate structure)	1. Elimination of Duplication /Single initiative /Single management structure 2. Build on CAnet, etc. 3. Accomodate the different states and needs of provinces 4. Regional Development 5. Want representation 6. Health & Education Sectors represent large beneficiaries.	1. Maximized use of Canadian Technology 2. Use of existing transmission facilities 3. Connectivity to Regional, National, International networks 4. Self supporting after 5 years 5. Work jointly with Regional Networks 6. Network will not be free (equity required from all participants)

Technical Design

The R&D users have expressed a number of needs which influence the technical design of this network. If all of these needs are to be accommodated The ISTC network must be designed to provide:

- An environment which **connects** all researchers and students across Canada independent of affiliation, location, size and research discipline. The existing federally funded Centres of Excellence provide an excellent example of a group of users that would benefit tremendously from enhanced connectivity. International connectivity to the U.S. Internet, as well as to networks in Europe and Asia are extremely important and must continue to exist. Users of existing or planned networks must be able to migrate to this network **without any loss of service.** For the moment, connectivity in Canada and the U.S., can be attained by being able to interoperate with users in the **TCP/IP domain;**

- An environment which is **highly reliable** in terms of network availability and in the services it provides. This network will not survive if it is not reliable. Therefore it may have to start out as being very simple, with proven technology. Consideration should be given to **gradually introduce upgrades** to ensure network reliability;

- An environment with **basic network services** consisting of directory and routing service, electronic mail, bulletin boards, file transfer, remote login, basic security, user support and network news must be provided by the Network. The Network should also be capable of providing **enhanced value-added services** (library services for bibliographic searches [CAN/OLE, DIALOGUE], document retrieval, and enhanced security service for those who need further assurances of the integrity of their data). At the present time, electronic mail is the predominant application available to the R&D User through existing R&D networks. The development of new applications and their use on existing networks is constrained by: the cost of commercial networking; and restricted access to user friendly workstations.

- A state–of–the–art **high capacity** environment (ie. T1 — 1.5 Mbs) to support a wide range of users and network services. Some users need to process large amounts of data (eg 100 Gigabyte) which require thousands of hours of CPU time; or require access to supercomputers where there is a common transmission requirement of 2 megabits per frame.

  With respect to capacity requirements there are three factors that must be taken into account: (1) the applications which are not possible without higher bandwidth (eg. large file transfer, transmission of high resolution, colour 3–D images, remote login to supercomputers, real–time distributed processing, and video and multimedia conferencing); and (2) many multiple users who require low/medium speeds for electronic mail and access to bulletin boards etc. results in a backbone which must be of higher speed; and (3) quick response times and adequate throughput require more bandwidth. This network must provide a level of service which is an order of magnitude above the existing situation, bearing in mind that some users have access to **dedicated** 19.2 Kbs and 56 Kbs transmission facilities;

- **A multi–media environment** flexible enough to accommodate **voice, data, text and image** (includes video). Requirements associated with real time

applications such as video and voice were mainly brought up by users with a related IT research interest. Full motion video applications were not seen as a short term requirement, although most R&D Users saw a migration to video applications such a video conferencing and education in the not too distant future;

- **Standardized, transparent and easy to use network to computer interfaces** (ie. pull down windows) which eliminate all unnecessary and intimidating procedural steps; and

- The Network should accommodate smaller users by providing **lower cost access options.**

Implementation Design:

The R&D Users have expressed a number of needs which also have an impact on the implementation and management of the ISTC network. If all of these needs are to be accommodated the ISTC network must adhere to the following:

- To the greatest degree possible, the national networking initiatives in place today must **eliminate duplication** of effort and funding. Only one national research network should physically exist, sponsored by one leading agency. The ISTC network must build on and extend the CAnet, CDNnet and Netnorth initiatives to ensure that a network is put into place that provides the highest level of service to the highest number of users, at reasonable cost. Users want to pay **one connection fee**, and have access to a wide assortment of services;

- User fees should be **distance independent.** Users are willing to pay more for higher connection speeds but are not willing to pay more to send data further. It was also felt that user fees should be **fixed** for a given level of service (ie. scaled fee based on speed of connection). For example the fee for dial–up service would be less than high speed access. In addition the fixed fee is attractive from a budgeting perspective. The rate structure should also **not change dramatically** from year to year. For example if the topology changes or the Network gets upgraded to T3 etc., government subsidy must continue to be provided. Predicted reductions in T1 rates (up to 50%) suggest that even though the ISTC network would provide a higher level of service, user fees for this network would be approximately the same as existing user fees (ie. prior to any changes to the T1 rate);

- The existing tariff rates for T1 service pose a barrier for organizations to connect to the Network. Unless rates come down significantly, ongoing **government subsidy** will be required in order to provide affordable access to all participants;

- The costs to provide the **basic services**, outlined above, should be included in the user fees. **Higher level value–added services** should not be embedded in fees (ie. CAN/OLE, DIALOGUE, library access, security service for those who want enhanced data integrity, access to specialized databases etc). The board of directors for this network must decide what gets put on the Network, and what user services are embedded in the basic fees, but users want to make their own choices when it comes to higher level services;

- Users, within a region or within an institution, want **representation** on the Network board of directors. This elected representative would protect the interests of his constituency and also keep the users informed of future directions and services provided by the Network; and

- Traffic on the Network should be **non-commercial**, and the definition of acceptable usage has to be clearly conveyed. Access and use of the Network should be based on the honour system, due to the difficulty in policing the transmission of packets. This network must interoperate with other international networks (ie. Internet in the U.S.), therefore some consideration must be given to **ensure consistency** in the acceptable use definitions, between these networks.

4.2 IT USERS

IT Users are defined as researchers and developers in areas related to information technology. For the most part these organizations are private–sector based, but users could be located at universities or publicly funded telecommunication research centres (e.g. ATRC). This subgroup also explicitly accounts for small IT companies who traditionally have not had access to the regional networks.

In general, the requirements of the IT User are quite different from those of the R&D User. The R&D user is application oriented while the IT user is utilizing the Network itself for research and development. The first group desires user friendly access, as previously discussed, while the second group wishes to be fully involved with the communications mechanism. The Network architecture must support this second group as well.

The IT Users objectives vis a vis the Network may be any one of the following:

- Stimulate the demand for a new service or product within the R&D community;

- Use the Network as a vehicle for Beta testing; and

- Use the Network as a vehicle to attain a better understanding of user requirements, as part of the process of product development.

In order for these objectives to be attained, the Network must provide an architecture which gives IT users access to the communication mechanisms of the upper and lower layers of the Network. The architecture must provide this capability to IT Users while simultaneously supporting the provision of reliable user friendly applications to the R&D User community. IT Users can be classified into two subgroups: those wishing to supply or experiment with upper layer protocols (above the Network Layer) and those interested in the lower layers.[3]

The upper layer subgroup is interested in using the Network to do research and development in a number of areas, including network management services and application services related to file transfer, job control, electronic mail, voice, video and others. The lower layer subgroup is interested in research and development activities associated with the physical communications medium. For both subgroups it is important that the development and research activity should be based on international standards.

---

[3] See the Network Architecture and Migration Plan issue paper for a detailed discussion of the protocol layers.

Technical Design

The IT users have expressed a number of needs which influence the technical design of this network. If all of these needs are to be accommodated, the Network must be designed to provide:

- An environment which provides access to the various levels of the Network itself **facilitating** the development and testing of new products (ie. OSI conformance, high speed protocol development, error correction and avoidance protocols) and new services (ie. value–added database service, OSI conformance centre). Some developers wish to supply or experiment with upper layer protocols (above the network layer) and others are interested in the lower layers;

- A state–of–the–art **high capacity** environment (ie. T1 — 1.5 Mbs initially moving to higher speeds when available) to support a wide range of users and network services. This network must always be at least one step ahead of commercial offerings and support the development of fibre–optic transmission and high speed switching technologies for **future broadband services**;

- An environment which can supply the traffic levels required to **simulate a realistic beta test site**, which goes beyond the testing that can be done in a laboratory experiment;

- An environment which utilizes **all transmission media** (ie. fibre–optic, microwave and satellite facilities);

- An environment where the **network intelligence** resides on the periphery of the network (ie. end systems) The intelligence will not be imbedded in the switches and routers on the network. However the Network must be intelligent enough to manage itself;

- An environment which provides a **migration path** to ISO/OSI;

- An environment which will ensure **a smooth technology migration** to higher speeds. The first migration will be T3.

- An environment which can accommodate multi–media formats, including **voice, data, text and image**, and explicitly allows for **video** applications;

- The Network should accommodate smaller users by providing **low speed access dial–up.**

- An environment which **connects** all IT users inside and outside the main urban cores of Canada. Many IT developers are located in industrial park complexes, 20–30 Km outside the central core — and require access to this network. The replacement of the existing infrastructure with fibre is an expensive and long term proposition. International connectivity is also very important; and

- An environment which will make available state–of–the–art **security and access control** and network management services.[4]

[4] Refer to the "Network Security and Access Control" section of Volume C.

Implementation Design

The implementation and management of the ISTC network must adhere to the following:

- A **mission statement** which clearly expresses the objectives for the R&D user and the IT user;[5]

- To the greatest degree possible, the existing national networking initiatives must **eliminate duplication** of effort and funding. Only one national research network should physically exist, sponsored by one leading agency. The ISTC network must build on the CAnet[6], CDNnet and Netnorth initiatives to ensure that a network is put into place that provides the highest level of service to the highest number of users, at a reasonable cost. Users want to pay **one connection fee**, and have access to a wide assortment of services;

- For this network to achieve its goals it must be **implemented quickly** (by early 1991). There is a window of opportunity to sell OSI–based products to the U.S. and Europe. If the implementation of this Networks drags on, this window may be lost;

- The creation of a **certified OSI conformance test centre** (possibly spearheaded by CIGOS) on the Network may provide the incentive and the knowledge base required for companies to develop OSI compatible hardware and software;

- A billing system based on **usage-based fees** (or a combination of fixed and usage) must be put into place;

- Some organizations have expressed an interest in helping to **finance** this network. Their preference is to make a contribution in kind by providing the equipment and/or expertise required to implement and manage the Network. The specific details of their involvement cannot be ascertained until a business plan is in place, and participant roles clearly defined;

- The Network must be **managed** by a full–time, professional organization or association. The technical centre must be open 24 hours a day, seven days a week. The network manager must be selected so as to **minimize proprietary interests** which might conflict with, or bias participation by the fullest spectrum of Canada's information technology industry;

- **Access** to the Network should be for R&D purposes, and commercial and administrative traffic should not be allowed. However commercial traffic should be permitted when testing a particular product;

- Consideration must be given to **intellectual property rights**. Academics will want to publish immediately; whereas the industrial partners will want exclusive patents or licenses, and the option of delaying publication until their proprietary plans are enacted.
- **R&D tax credit recognition** must be guaranteed in exchange for helping to

---

5 For further discussion of the mission and goals of this Network refer to the appropriate section of Volume D.

6 The relationship between CAnet and the ISTC Network is discussed in Volume D.

fund the Network;

4.3 IT SUPPLIERS

Excluding carriers, IT Suppliers encompass a broad range of organizations which provide IT based products or services. These suppliers may include private sector, government as well as academic organizations. IT Suppliers include suppliers of communications products to the Network (e.g. Northern Telecom etc.), suppliers of database related services (e.g. CISTI, the National Library etc.), suppliers of communications related services (e.g. U of T etc.), suppliers of computational services (e.g. SuperNet) as well as suppliers of other services which may be developed by IT users in the future. In fact, many organizations have been identified as both IT users and Suppliers to this network, but are treated separately because their needs and objectives are quite different.

The IT supplier is principally interested in:

- Generating immediate revenue or profit as a Network provider. For example, at CISTI, there is an increasing interest in revenue generation and reaching out to more than traditional clients; or
- Using the Network as a showcase for new products and services.

Technical Design

IT suppliers have expressed a number of needs which influence the technical design of this network. From their perspective, the Network must be designed such that:

- The network architecture is **open** to multiple vendors;
- The Network would utilize **state-of-the-art technologies**, and the initial network design would use network products which are currently under development; and
- The Network could be used as a **beta test site** for new products.

Implementation Design

IT suppliers have also expressed a number of needs which have a bearing on the implementation design of this network. From their perspective, the Network must be implemented such that:

- The formation of a **broad-based user community** must happen quickly in order to attract value–added services to be developed; and
- The **fee structure** which is put into place must be consistent with the marketing and developmental needs of suppliers.
- Any arrangements made with suppliers to attain **discounts** on products/services must be accompanied by formal recognition of R&D expenditure.
- Consideration should be given to a consortium of IT suppliers to manage the

Network, as a number of companies have expressed such an interest.

4.4 CARRIERS

The Carriers are defined as regulated common carriers capable of providing communications service or facilities (link channel capacity) to the Network on a national scale. Carriers include the Telecom Canada member companies, CN/CP and Telesat Canada.[7]

The Carriers represent the largest group in Canada involved with the provision of communication services and the installation of wide area network facilities. Their participation is key to the success of the Network for a number of reasons. The Carriers have a mandate to provide commercial communications services in Canada, and it is through the Carriers that the Network will obtain wide area network services or channel capacity, as appropriate. In the ideal scenario, Network requirements would be satisfied through commercial Carrier offerings. However this may not be possible within the next few years. In addition, should the Network attempt to obtain channel capacity or services at below tariff rates, Carrier support will be required[8]. The provision of link channel capacity will be one of the significant costs associated with the Network's implementation[9].

The Carriers objectives and needs centre on obtaining technical, operational and marketing information and developing new markets. Their interest in the Network is directly related to the leading edge broadband data networking aspects of the Network. The carriers are not interested in the Network if it is not in line with the long term strategic direction for their own networks. The carriers for the moment are open to suggestions as to how they might play a role in **funding** this network, in terms of equipment, expertise and direct funding.

Technical Design

The Carriers have expressed a number of needs which influence the technical design of this network. If all of these needs are to be accommodated the Network must be designed to:

- Lease **existing transmission facilities**, owned and operated by carriers;
- Fibre, satellite and microwave **subnetworks must coexist**, and interoperate with each other, with no implied hierarchy; and
- It is important that the Network be **capable of interfacing** with the carrier offerings, both current and future, and be seen by the carriers as significantly advanced and leading edge in terms of the technologies being used and the Network architecture. The Network must interface with public data and voice services as well as the **evolving ISDN service** offerings. The carriers would not be interested in this network if not in line with the strategic direction of their own networks.

Implementation Design

---

7 Value added resellers (e.g. INSINC, Fonorola, Callnet) are not included in this category but are identified in the IT Supplier group below.

8 See the Regulatory section of Volume D for a discussion of this point.

9 See the section on Network Costs contained within Volume B.

Carriers have also expressed a number of needs which have a bearing on the implementation design of the Network. The implementation and management of the ISTC network should adhere to the following:

- Should the Network attempt to obtain channel capacity or services at below tariff rates, carrier support will be required[10];
- Although the acceleration of this network is outside the **timing** of the plans of carriers to provide these services, they view the concept positively as it may bring them differentiation with respect to commercial offerings (ie. video conferencing) and an opportunity to **develop new markets**;
- There is a concern that **excess capacity on** the Network may be resold;
- The Network should not carry **commercial bypass** traffic. The carriers are comfortable with the honour system being used to limit the commercial traffic on this network;
- If voice is brought onto this network this may open the door to a whole host of other **regulatory issues**;
- Some consideration should be given to the Carriers **managing this network**;[11]
- **Formal recognition**, from the CRTC, of any contribution (ie. contributions must be reflected in tariffs); and
- The issues related to "**intellectual property**" must be treated cautiously. If a carrier spends money on developing a product jointly with an R&D partner, measures must exist to insure a competitor does not acquire access to this information and develop a similar product.

4.5 EXISTING REGIONAL NETWORKS

Existing R&D networks, in addition to having specific technical characteristics, are associated with organizational structures for purposes of policy, management and operation. These organizations have objectives and needs which must be addressed in their own right, and which in some cases are over and above the needs of individual researchers or the needs of end user organizations supporting an attached device or service.

The existing R&D networks in Canada support the majority of R&D Users in the academic sector. The constituency represented by these networks will be the Network's initial user base, a critical element of this Network's successful implementation.

With respect to the backbone network, the existing R&D networks have two fundamental

---

10 Refer to the regulatory chapter in Volume D for a discussion of the options available for obtaining channel capacity or services below tariff rates.

11 Refer to the section on network management in Volume D, for a further discussion of the role of the network manager.

needs:

- Connection with each other and the International Internet; and
- Joint participation in the management and development of the Network.

The interoperability requirement centres on the coexistence of both the TCP\IP and OSI domains on the Network. The joint participation requirement centres on a recognized position on the Network's Board of Directors.[12]

Technical design

The regional network mangers have expressed a number of needs which influence the technical design of this network. If these needs are to be accommodated the ISTC network must be designed such that:

- The regional networks must continue to operate and exist, and be able to **interoperate** with each other and with other **international** networks;
- The interoperability discussed above centres on the **coexistence of both the TCP/IP and OSI domains** on the Network[13];
- Only **one subnetwork** for each geographic region should be maintained[14]. A "region" may be defined as either a province, an area within a province, or a group of provinces. The concern here is that "**predatory networks**", may emerge in the more highly populated areas, and undercut the existing regional networks and attract traffic away from the existing networks.[15]
- The backbone network must be able to **accommodate** the emergence of new regional networks.

Implementation Design:

From the perspective of regional network managers the implementation and management of the ISTC network must be designed such that:

- The needs of the **current client base**, predominately R&D users in the academic sector, must be met from the first day of operation.
- The management structures in place at the regional level should be **joint**

---

12 Refer to the recommended management structure described in Volume D.

13 For a further discussion of protocols and use of standards refer to the appropriate technical paper in Volume C.

14 For further discussion, refer to the section on the role of the regional network, contained in Volume D.

15 For further discussion, refer to the section on the role of the regional network, contained in Volume D.

participants in the management of the backbone[16].

- The current body of technical and management experience attained by regional network managers must be **incorporated** into an implementation plan for this network. Universities have been involved in network management over the last five or six years, and have attained vast amounts of expertise with respect to network management. To the maximum extent possible this expertise must be utilized, if this Network is to be successful.

- **Full-time** network manager(s)[17], and **formalized marketing and management structures** must be put into place at the regional and at the backbone network level. Some regional networks already have these formal structures in place, whereas other networks are more informal in their management approach. There is a requirement then to make sure that adequate marketing and management structures are in place for each regional network to ensure that the objectives and goals of the backbone are attained.

- Some consideration should be given to preventing users from **directly connecting to the backbone**, instead of connecting to the backbone via the regional network. Direct connections to the backbone would deny regional networks from collecting revenue from large users, which could jeopardize their operation.[18]

- User fees should be **fixed** for a given level of service (ie. scaled fee based on speed of connection), as a usage–based fee would lead to prohibitively high overhead costs. In general, it is felt that if the fee was usage based, billing costs would double. The fixed fee is also attractive from a budgeting perspective. In addition some regional networks require industrial users to pay a higher fee than university users.

- **Back-door connections** should be minimized, and access to the Network must be handled formally.

4.6 PROVINCIAL GOVERNMENT

Users and service providers in every province and territory in Canada would be invited and encouraged to access and provide services to this network. For the moment, regional networks are defined geographically by provincial boundaries, and thus the needs of the regional network are closely aligned with the needs of the Province.

In the context of this section the provincial government is seen as a potential sponsor of this network.

Technical Design

---

16 This is described in more detail in the section on Management Structure, contained in Volume D.

17 Refer to the issue paper on "Network management and the Technical Centre", in Volume C, for more information regarding the role of the Network Manager.

18 For further discussion of this point, refer to the section on " the role of the regional network" in Volume D.

Provincial representatives have expressed a number of needs which may influence the technical design of this network. If these needs are to be accommodated the ISTC network must be designed such that:

- Data collected by provincial governments should eventually be put on the Network. Each province could use this network as an opportunity to build databases in mining, pulp and paper, fisheries, agriculture etc..

- Access to video lectures on a number of topic areas should eventually be made available to all students (high–school, college and universities) as well as to business and other organizations who require retraining, or access to specialized information.

Implementation Design

As listed below, provincial representatives have also indicated a number of issues which may influence how this network is implemented:

- The Federal Government has to provide **a unified direction** for networking in Canada. ISTC, NRC and DOC must cooperate and work together in bringing about one national R&D network. To the greatest degree possible, the existing national networking initiatives must eliminate duplication of effort and funding. Therefore only one national research network should physically exist, sponsored by one leading agency. The ISTC network must build on and extend the CAnet, CDNnet and Netnorth initiatives to ensure that a network is put into place that provides the highest level of service to the highest number of users, at the lowest possible cost. Users want to pay **one connection fee**, and have access to a wide assortment of services.

- Not all provinces are in the same position with regards to networking. Some provinces have a regional network in place others don't. Some provinces are looking to next–generation technology whereas other provinces are looking just to catch–up. This Network must accommodate the different states of the provinces.

- Some provinces view the Network as a vehicle for regional development.

- The potential benefits to the health and educational sectors (remote education) are very large. In many industries, employees require retraining. This network could be used for that purpose.

4.7 FEDERAL GOVERNMENT

If it is decided that this network should be implemented, Industry Science and Technology Canada (ISTC) would be the lead Federal Department guiding its implementation and overseeing its management.

The Core Mandate of ISTC, as set out in Part 1 of Bill C–3, is to develop national policies and programs to strengthen the economy; increase the international competitiveness of Canadian industry; encourage the fullest and most efficient use of science and technology in Canada; foster entrepreneurship and the start–up growth and expansion of small business.

This mandate is consistent with the mission and goals of the Network.[19]

Technical Design

ISTC has suggested minimum requirements for the technical design of this network. These requirements are as follows:

- The Network must support the use of standards, and must be consistent with developing ISO/OSI standards.
- The Network must have Trans–Canada, end–to–end digital connectivity;
- The Network must offer transparent service for multi–media (voice, data and image) communications;
- The Network capacity must be 1.5 Mbs minimum on fibre–optic, microwave and satellite media;
- The Network must offer state–of–the–art protection features and network management;
- The Network must offer test facilities for new product services without interference with operation traffic;
- The implementation of the Network must be based on Canadian technology available in Canada, to the maximum extent possible;
- The Network must use existing transmission facilities where possible; and
- The architecture must enable access to federal, regional and international R&D networks.

The requirements listed above are consistent with the needs of all the potential participants of the Network ie. R&D users, IT users, IT suppliers, carriers, regional networks and provincial governments.

Implementation Design

As listed below, ISTC has also indicated a number of issues which may influence how this network is implemented and funded:

- This network must be seen as an **industrial development tool**, making Canadian industry more internationally competitive;
- This network must be used to **disseminate information** collected by government departments;
- The Network must promote the **technology transfer** between universities and industry;

[19] Refer to the Mission and Goals section of Volume D.

- The Network may be used as a vehicle to promote **increased collaboration** among industry, and may help facilitate the development of strategic alliances;
- The operating portion of this network must be **self-supporting** after five years. The Federal Government should support this long enough into the future to allow the user base to build;
- The Federal government may have to have to **work jointly** with regional networks and the provinces to ensure that the goals of this network are attained; and
- All Network participants must be willing to share in the costs, risks and the benefits of this network (ie. some form of **equity** is required from all participants). This network will not be free.

5. FUTURE NEEDS AND APPLICATIONS

The proposed network will be designed to serve the needs of researchers throughout the 1990's and beyond. It is speculated that the Network could begin operations at the T1 rate early in 1991. However, the Network must be positioned for a smooth migration to higher speeds (ie T3 first migration). New standards and technologies will be incorporated into the Network as they evolve. In fact, the Network would serve as a vehicle for the development and testing of future information technology products and services. Future applications of the Network must be anticipated so that the evolution of the Network can be properly planned for. The information required about these future applications is the subject of Chapter 3 of this volume.

The proposed network will be designed to serve the needs of researchers throughout the 1990's and beyond. It is speculated that the Network could begin operations at the T1 rate early in 1991. However, the Network must be positioned for a smooth migration to higher speeds (ie T3 first migration). New standards and technologies will be incorporated into the Network as they evolve. In fact, the Network would serve as a vehicle for the development and testing of future information technology products and services. Future applications of the Network must be anticipated so that the evolution of the Network can be properly planned for. These future applications are the subject of Chapter 3 of this volume.

The needs of distributed scientists can be structured into three general approaches: supporting communications between researchers; providing researchers access and control over remote resources and equipment; and tools that let researchers store and manipulate information. Central to each section in Chapter 3 is a description of current visions, with emphasis placed upon research systems that are near commercialization or likely to be commercialized in some form over the next decade. Limitations and unknowns of the technologies are raised, and network requirements estimated.

Supporting Communications Between Researchers

Real-time remote conferencing permits formal or semi formal meetings among geographically distributed people. Techniques for transmitting a sense of participatory presence and for sharing data over distances are necessary.

A 1.5 Mbps network can address some of the requirements of remote conferencing. In particular, the tele–data requirements that have already been developed and used on LAN networks can migrate to the 1.5 Mbps network with little performance penalty. These include screen sharing of existing single user application packages, and multi–user workspaces that allow participants to simultaneously draw and type in the view.

The network may also suffice for some limited types of tele–presence. Textual real–time communication systems such as Cantata require modest bandwidth, while the shared gesturing ability found in workspaces probably requires little more than the transmission of a person's pointer location. As mentioned above, low bandwidth versions of Talking Heads does not require video transmission (although there is still a question about how well a voice–animated image provides tele–presence). Similarly, slow–frame video can capture at least some aspects of tele–presence.

Yet the dynamics between attendees in a meeting can be so diverse that it is unlikely that any limited channel will capture all the subtle yet important cues. It stands to reason that real tele–presence would require at least a full video channel and a corresponding increase in bandwidth over the network. At the very least, one video connection is needed for room to room contact. Bandwidth demands multiply with the connections made, as seen when the number of rooms increase, or when video screens represent each individual in the meeting.

A similar high bandwidth requirement is needed when tele–data requires video, as may occur when the data is not available in digitized form, is not easily scanned in, or is animated.

The actual network requirements for cyberspace is unclear. Three–dimensional real–time animations must be modelled and exchanged; interactions between participants with each other and with the simulated environment transmitted and synchronized; and sound (including its coordinates of origin) transmitted. Although simple cyberspace interactions using distributed models may run over the 1.5 Mbps network, a true realization of an alternate reality would likely require considerable bandwidth.

Casual real–time interaction is a crucial yet often overlooked component of scientific collaboration. While formal meetings take care of key events in the scientific process, casual interaction at the interpersonal level is the glue that keeps the collaboration working. Several innovative systems are described that bring people together through frequent, unplanned, and high–quality interactions.

A prerequisite of casual interaction is that sessions must be of high–quality and real–time. By definition, then, the network requirements will be high. At the very least, there must be enough bandwidth on the network to support: a method for allowing potential collaborators to know who is around; a voice channel; a shared workspace (tele–data). These demands can be met within the constraints of a 1.5 Mbps network.

Truly effective interaction likely requires a video link. As with remote conferencing, slow–scan video could work at the lower bandwidths, while compressions schemes may allow full bandwidth video across a 1.5 Mbps network. This allows, however, for only a single connection. While this is reasonable for a system like Video Wall, both Cruiser and the TeleCollaboration project would require multiple video links. The alternative is to allow only one person to monopolize the network while wandering the virtual hallways and while establishing a link. A similar problem occurs within SharedArk. While initial contact is made through the multi–user "virtual world," the ensuing interaction establishes a video link.

Asynchronous messaging covers the familiar electronic mail systems we are now used to. The domain also includes more sophisticated versions of electronic mail, including multi–media, semi–structured, and semi–formal mail systems.

Most basic electronic mail requirements are well–handled by a fairly low–bandwidth network. This is usually due to the small size of the mail packet, and because the time to transmit is not critical due to the asynchronous aspect of the communication. As mail becomes more sophisticated, as with the addition of structure and protocol into the message, the average message size will rise with the increase in transmitted information.

More bandwidth is required as mail contents become increasingly sophisticated. Taking document preparation systems as a model, we have seen a progression from simple ASCII text to quite sophisticated typeset documents including complex figures and bit–mapped images. There is every reason to believe that users will demand electronic mail to follow the

same course.

The real increase in bandwidth requirements will occur as true multi–media mail becomes available. Animated graphics, digitized speech and full video all require high bandwidth. Since these messages have high storage needs, only a pointer to the multi–media message may be transmitted for notification purposes. The actual contents may be sent on demand, and in real time when the reader actually refers to it.

Bulletin boards and asynchronous conferencing are systems that allow people to interact with a large community by posting and viewing messages, and joining and actively participating in on–going discussions of some topic of interest.

The network requirements for conferencing are range from low to extreme. As with electronic mail, low–bandwidth systems can handle even large numbers of postings of limited size as long as immediate delivery is not necessary.

Multi–media postings require significant bandwidth. Consider a centralized architecture, where the user must log on to the remotely–located system. Since postings are transmitted only when they are selected by the user from the subject heading, they must be sent in real–time. The network must be able to transmit the posting with little delay, for excessive waiting would deter the reader from scanning all but the most interesting articles. With a distributed architecture, articles must be broadcast to all subscribed sites as they are posted. Given a large number of sites and a high volume of postings, considerable demands may be made of the network.

Access and Control Over Remote Resources and Equipment

The access, monitoring and operation of *remote equipment* has two aspects: access to remote computational power, and real–time systems for remote operation. Technologies such as Scada systems, master–slave tele–operation, and tele–robotics are included.

Computational access requires modest network speeds. Normal telephone speeds are acceptable for ASCII–terminal emulation; and a 1.5 Mbps network is adequate for window–based interfaces. Network requirements for parallel systems can range from low to high, depending upon the size of the computational unit and the amount of communication between these units.

Shifting the emphasis towards monitoring and operation, existing SCADA systems require relatively low–bandwidths. They are well served by current telephone lines and would certainly be well covered by the 1.5 Mbps Network. True real–time systems, especially those controlling robots, have much stricter bandwidth demands and time tolerances. Speed is critical. In master–server tele–operations, the operator may need a high–resolution video image of the remote robot to see what is going on. Additionally, the feedback cycle between action, consequence, and notification cycle must be short: the maximum tolerable feedback delay for effective real–time control of the robot by the human is one tenth of a second (Scott, 1984).

Computer controlled tele–robotics usually require high–bandwidth reflecting the data rates for control and feedback of servers, and perhaps information transmission by any attached vision systems. The preliminary NASA/NBS Standard Reference Model for Tele–robotic Control System Architecture (NASREM) describes timing tolerances of a few milli–seconds for synchronizing activity, especially at the servo level (Albus, McCain and Lumia, 1989). Even semi–autonomous robots may require bandwidth for monitoring (perhaps through

video), data collection, transmission of high–level control information, and synchronization with other devices. Other criteria are, of course, system dependant.

Robotics have a good migration path to a nation–wide network, simply because they already rely on software supervision and often use high–speed local area networks for communication.

Storing and Manipulating Information

The *digital library* includes conventional approaches such as the digital repository and bibliographic databases, as well as the attractive alternative of hypertext. The digital library is the backbone behind the communal knowledge base.

Conventional bibliographic databases have low network requirements due to the small amount of information generated; they now run adequately over modems. In marked contrast, file transfer from digital repositories can involve many megabytes of information. Some repositories have software releases available of 50 megabytes or more! Fortunately file transfer is rarely time critical, and are still well–served by low–bandwidth lines.

In contrast, hypertext can require significantly more bandwidth. If a hypertext system is maintained at a centralized location, then users must access it through the network. Information retrieval is time–critical. If a reader chases a link, the node contents should appear as quickly as possible. Delays will deter readers from chasing paths, contradicting the intention of hypertext. While moderate speed lines may suffice for pure text systems, hypermedia nodes would require the bandwidth needed for rapid transmission of audio, graphics, video, and so on. Additional bandwidth is also needed to drive the graphical user interface employed by most hypertext systems.

Conclusion

Many applications can support high–level collaborative functions. Application areas include distance education, group decision support systems, and joint authoring.

The networked nation is more than a means of overcoming distance barriers, for it has the ability to augment researchers' and developers' knowledge and activities beyond those now found traditionally. Each tool and function mentioned in this report can, by itself, provide added value. Communication vehicles will fit tasks, digital libraries will become knowledge bases; tele–robotics will extend our physical limitations. Yet the real power will arrive when tools and people are integrated together into the *Coherent Workspace.*

The coherent workspace is an idealized integrated platform that supports distributed researchers. It occurs when people and their machines have a common knowledge and sense of purpose, with individual and group activities being well–coordinated.

CHAPTER 2: INFORMATION GATHERING

1. INTRODUCTION

This chapter contains a summary of the methodology used to gather information for this report, and a synopsis of the results.

Four main instruments were used by the study team to identify potential participants and obtain information on their needs and objectives:

- Expert Panel Sessions;
- In–person interviews;
- A questionnaire distributed via the Netnorth and CDNnet networks; and
- Telephone interviews.

Every effort was made to contact as many potential participants as possible. In fact the project team received direct input from more than 400 individuals who took part in one of the four survey instruments identified above. The project team has discussed this network initiative with government officials in every province and territory in Canada, and also with Federal Government representatives (eg. NRC, DOC, EMR, GTA and CRTC).

The detailed results from each of the information gathering instruments identified above are contained in the following sections.

2. EXPERT PANEL SESSIONS

The expert panel sessions were organized as informal round–table discussions, designed to collect information from potential participants. Each of the four sessions had its own regional flavour and was very useful in developing an understanding of the wide ranging needs, interests and concerns of all key players. The information gained through these discussions could not have been captured in any other way, and provided the project team with backgrond for the on–site interviews used to collect more detailed information.

At each session, representatives from the regional networks, local universities, local research centres, local information technology companies and the CAnet advisory board were in attendance and participated actively in the discussion period. In addition, at some of the sessions representatives from NSERC, Centres of Excellence, and the Science Council were in attendance to observe the exchange of ideas among the panelists.

2.1 METHODOLOGY

The project team hosted four expert panel sessions; two in Ottawa, one in Montreal, and one in Vancouver (logistical details of each session are summarized in Table 1). It was decided that two sessions were required in Ottawa in order to cover the large number of IT companies in the area, as well as the different government departments that might play a role in the implementation of the Network. The guiding principles used for selecting the participants for these sessions are as follows:

- representation from CAnet and the regional network (ie. Onet, RISQ, NSTN, ARnet and BCnet);
- representation from different geographical areas[1];
- representation from the designated Centres of Excellence;[2]
- representation from researchers with high computational requirements;
- balanced representation among universities, business and government[3];
- representation from a range of research areas (ie. oceanography, physics, medicine, genetics etc.), within universities, business and government; and

---

1 The sessions were attended by individuals located in Ottawa, Washington, Toronto, Guelph, Hamilton, Waterloo, London, Montreal, Halifax, Vancouver, Victoria, Edmonton and Calgary.

2 Many of the university researchers selected for the expert panel sessions also represented one of the designated Centres of Excellence.

3 In total the expert panel sessions involved 44 participants; 12 from industry; 12 from universities; 10 from government; 6 from existing networks and 4 from other research organizations.

TABLE 1: LOGISTICS OF EXPERT PANEL SESSIONS

SESSION	LOCATION	DATE	THEME
1	Ottawa; Delta Hotel	18/12/89 9:00 am – 3:00 pm	User needs from a wide variety of participants in industry, universities, government and existing regional networks
2A	Ottawa; Delta Hotel	18/01/90 9:00 am – 12:00 pm	Priorities of various government departments and agencies
2B	Ottawa; Delta Hotel	18/01/90 1:00 pm – 4:00 pm	User needs from a general R & D perspective and issues with respect to regional networks
3	Montreal; Auberge du Govenour	29/01/90 1:00 pm – 4:00 pm	User needs from a general R & D perspective and issues with respect to regional networks
4A	Vancouver; Hotel Vancouver	08/02/90 9:00 am – 12:00 pm	User needs from an IT R & D perspective
4B	Vancouver; Hotel Vancouver	08/02/90 1:00 pm – 4:00 pm	User needs from a general R & D perspective and with respect to regional networks

- representation from a wide range of IT companies[4] with different specializations and expertise;

2.2 RESULTS

The first expert panel session in Ottawa (a day long event), included a wide mix of participants representing four different IT companies, two government departments, three universities, the regional network (Onet) and CAnet. The sessionwas lively, however the diversity within the group made it very difficult to focus the discussion, and to reach consensus on specific issues. As a result it was decided that subsequent sessions be split into separate forums; with each forum focusing on the needs of a specific participant group.

The second expert panel session, held in Ottawa, was split into morning and afternoon sessions. Michael Roberts, a representative from EDUCOM (based in Washington) was invited to participate in this session and to brief the panel regarding the current research networks in the U.S.. The morning forum focused on the needs of the Canadian Government and included representation from NRC (and CAnet), CISTI, DOC, GTA, CRTC and Treasury Board. The afternoon forum addressed the needs of general R&D researchers and was attended by researchers within government (Canadian Space Agency), industry (Northern Telecom, Connaught Laboratories) and Universities (Guelph and McMaster).

The third session took place in Montreal. The morning was spent speaking with IT researchers at Memotech, Spar, and McGill University (which is spearheading the Centre of Excellence for Telecommunications). The afternoon was spent discussing the needs of regional networks and general R&D users. It was attended by representatives from RISQ, Nova Scotia Technical Network (NSTN), CAnet and the Oceanography Centre of Excellence.

The morning session in Vancouver was attended by five different IT organizations and centred on their needs and requirements. The afternoon forum concentrated on the experiences and status of BCnet, ARnet, CDNnet and CAnet, and the needs of the two Centres of Excellence in attendance.

The Minutes of these sessions are contained in Appendix A.

[4] Large IT companies were chosen for the expert panel sessions. The smaller IT companies were surveyed by phone.

3. IN-PERSON INTERVIEWS

It was considered essential to the information gathering process that more detailed information be collected through on site visits to: Provincial Governments; Information Technology Companies (eg. BNR, MPR, Newbridge etc.); national and provincial carriers (Telecom Canada, CNCP, Telesat, AGT etc.); research organizations (Alberta Telecommunications Research Centre); designated Centres of Excellence (Precarn); and Universities. These visits provided an excellent opportunity to brief senior officials within the organizations on the concept of this network and to share with them some preliminary results. In addition, these visits provided a means to acquire a more thorough understanding of the needs of each of the players and provided insight into their willingness to contribute to this venture.

3.1 METHODOLOGY

Officials within all ten provinces, and both territories, in Canada, were contacted[5] regarding this initiative. The Provincial sessions, for the most part, were very well attended and confirmed the high level of interest which exists across Canada in research networks. For example, at many of these meetings, representatives from industry and universities were also in attendance (ie. Alberta, Manitoba, Saskatchewan and Ontario).

In addition to the provincial meetings, 28 companies were interviewed on–site, across Canada. Over 50% of these interviews were carried out in Ontario, as the majority of IT companies and major carriers are located in the Toronto and Ottawa areas. Approximately 20% were carried out in Quebec, and 10% in British Columbia. The remainder of these interviews took place in Alberta, Saskatchewan, Manitoba and Nova Scotia.

To better gauge the amount of support for this network, each person that was interviewed, was invited to write a letter outlining the viewpoint of the organization/province they represent. These letters are contained in Appendix E. Appendix E–1 contains the letters from the provinces and Appendix E–2 includes letters from research organizations.

Although the content of these interviews are considered confidential, all the names of the organizations and provincial representatives that were contacted are listed in Appendix B.

3.2 RESULTS

The meetings with the provinces proved to be extremely valuable in raising the exposure of this initiative. A large proportion of these meetings were spent explaining the differences and similarities between the ISTC initiative and CAnet[6]. The provincial representatives have made it clear that any future national network must build on the CAnet, CDNnet and

---

5 Newfoundland, New Brunswick, PEI and both the Territories were contacted by conference call.

6 For a further discussion of the relationship between the ISTC initiative, CAnet, and other national initiatives (like Vision 2000) refer to Volume D, section 9.

Netnorth initiatives, take into account the needs of regional networks, and ensure that networking services are provided to the highest number of users at reasonable cost. To the greatest degree possible, this initiative must eliminate duplication of effort and funding, as Canada does not have a large enough population to support disjointed efforts in networking. However, not all provinces are at the same level with respect to networking. Where some provinces are trying to move forward with next generation networking technology, other provinces do not yet have regional networks in place, and are trying to catch–up with current technology. Therefore, there is general agreement among all provinces that this network is essential to meet the needs of researchers and students across Canada. To better understand the positions of the provinces refer to Appendix E–1, which contains copies of the letters received by the project team.

Many of the companies which were visited had participated in a previous expert panel session. The on–site interview provided an opportunity to discuss, more openly, their involvement and interest in the network.

Although every IT organization contacted (28 in total) saw merit in this initiative, they had many concerns and reservations regarding its implementation. It was generally agreed that the network, as presented, was technically feasible. It was also quite evident from the vast majority of interviews that the network could be made to be feasible from both an economic and implementation point of view, if certain course of action were adopted. Therefore, in many cases the letters received from IT organizations offer qualified support for the network (See Appendix E–2). The common interests, comments and priorities of these companies are as follows:

- For the most part these companies perceive the concept of this network as a technology push rather than a market pull. Although if researchers within a Centre of Excellence are to communicate effectively, this type of infrastructure must exist across Canada. Therefore a market pull argument could be made for the Centres of Excellence.

- The need for multi–media conferencing will increase dramatically over the next several years;

- If the Network were to assist in the development and testing of new products and services, significant productivity advantages would be attained;

- The management aspects of this network are crucial, in order to plan the migration from TCP/IP (and other protocols) to ISO/OSI;

- Companies are currently focussing, and will continue to focus, research dollars on broadband applications and technology, and thus the concept of this network is in line with the strategic direction of these companies;

- For this network to achieve its goals it must be implemented quickly;

- A T1 network is the minimum that should be put into place, and the transition plan of this network is very important; and

- The relationship between the ISTC network and CAnet must be addressed, and in the end, this relationship must be synergistic;

The majority of companies contacted indicated that they would be open to alternatives as to how they could contribute to financing this network. Contributions in kind by providing discounts on equipment or providing expertise appear to be the most popular options.

However, it has been made clear that these contributions must be seen as R&D expenditures, and the associated tax credits must be assured. A great interest exists among the companies interviewed in managing the network.

4. NETWORK QUESTIONNAIRE

4.1 METHODOLOGY

A questionnaire was distributed across Netnorth and CDNnet to obtain feedback from Canadian researchers who currently use the research networks available in Canada. Survey respondents provided an indication of how they currently use network facilities, whether or not they needed access to a more capable network, the applications which they require but currently are unable to obtain on existing networks, and the associated benefits. A copy of the questionnaire is contained in Appendix C.

4.2 RESULTS

Over 200 respondents completed the questionnaire. The profile of these respondents includes 60% from university professors, 18% from publicly funded research laboratories (primarily NRC), 13% from graduate students, 8% from network managers at universities, and only 1% from private sector companies engaged in R&D.

In total 81% of the respondents were located at universities (includes professors, students and network managers). Of this total a full 40% originated from British Columbia — which indicates significant interest in this province with respect to networking. Just under 25% were located in Ontario, 12% from Nova Scotia, 8% from both Quebec and Alberta, and the remaining 7% divided evenly between New Brunswick, Manitoba and Saskatchewan.

The main research disciplines of those who responded to the network survey are: physics (20%), communication technologies (18%), computer science (14%), earth science (8%), expert systems and artificial intelligence (5%), mathematics (8%), biology (5%), Oceanography (3%), health and medical (2%), social (3%) and other (14%).

Over 70% of respondents indicated that they would like to have access to a more capable network.

The following is a question–by–question summary of the answers provided by the respondents to this questionnaire.

Q1 **If your general area of research is telecommunications, would you consider using the network to develop and test new products and new services?**

18% of survey respondents indicated that their general area of research is telecommunications. Of this total a full 70% said they would use this Network to develop and test new products and services.

Q2 **If "Yes" to the above question, please elaborate on the products and services you would develop and test.**

A number of products and services were identified as follows:

- OSI Protocol development and testing
- OSI presentation and application level development and testing
- Shared Hypermedia spaces

- Standards for File and Data Transfers
- Transmission of Images
- Tele–operation
- Development of Standardized User Interfaces
- Distributed Directory Services
- Distributed CASE Tools
- Hypertext
- Cooperative Work Environments
- Electronic Conferencing Software
- Electronically Published Journals
- User Directory Services
- Multi–format Mail Systems
- Networking Software
- Distributed Applications and Systems for Group Collaboration
- Network Distributed Computation
- Large Scale Image Analysis Computation
- News Facilities
- Large Scale Distributed Database Lookup and retrieval
- Transferring LAN models for Interprocess Communication to WANs.
- VAX to Fax products
- ISDN
- Collaborative Drawing Tools
- Multi–user Editors
- Multi–user Network Talk Programs

Q3 Briefly describe the role of data communications in your research.

The majority of respondents indicated that they use existing data communication facilities to send and receive information (data, computer programs and software, analysis tools, manuscripts, precise position information etc.) to and from research colleagues in Canada and professional associations around the world, through electronic mail and file transfer;

Respondents also indicated that data communications is used for the following:

- Collaborative research.
- Processing large amounts of data (eg. 100 Gigabyte). These calculations require thousands of CPU hours.
- Review work of others
- Remote login to supercomputers and remote use of telescopes
- Access to bulletin boards
- Access to databases
- Quick and efficient way to update source code and databases
- Prototyping and Performance Measurement of distributed applications
- Useful in establishing contact with researchers working in a particular area
- Exchanging ideas, references and questions.
- Collaborative writing of papers
- Message Handling Systems
- Distributed Directory Services
- Scanning Internet mailing lists and news groups, and briefings of international scientific meetings
- Ask and answer questions
- Design communication protocols
- Distributed software development

Q4. Please give examples of the services/applications you require.

A number of services/applications were identified:

- Software to facilitate the use of network services
- E–Mail
- Conferencing with live digital video
- Directory Services for E–mail addresses
- Cheaper leasing costs of channels (more in line with U.S.)
- High–speed transmission of large data files
- Remote Login (the existing Onet – Internet line is saturated which makes remote logins impossible)
- Access to source code of basic protocol software at low cost
- Terrestrial radio
- To work in real–time conferencing areas
- Convenient way to communicate between UNIX and DOS systems.
- Network News
- Distributed hypertext and multi–media libraries
- Telnet, FTP, Mail, Usenet news
- Access to national bulletin boards
- Transmittal of satellite derived data
- Fast communication to sites that can put quality images on tape
- Send and receive files in both ASCII and binary formats
- Use network for peripheral sharing
- Access to databases and directories
- Scanning and digitizing images remotely
- Sharing costly high quality printers
- Gateway services
- Ability to selectively permit/deny access to any branch/node/user
- Network management and monitoring, file archive access
- X–windows
- Bibliographic searching facilities
- Access to catalogues of other Canadian and world–wide libraries — arrange inter–library loan
- Byte stream access to session layer

Q5 Indicate the manner in which you communicate with your research colleagues. Please specify what percentage of the total time you spend communicating with your research colleagues.

Some of the highlights of this question include:

- Approximately 66% of the respondents suggested that they spend under 5 hours a week communicating with their colleagues; 25% spend between 5 and 10 hours per week; and only 9% spend over 10 hours;

- Almost 60% of all respondents said that over 60% of all verbal communication is in person; and 67% indicated that less than 40% of all verbal communication is by telephone;

- Over 60% of all respondents said that 60% of all written communication is handled by electronic mail; and 78% indicated that less than 20% of all written communication is by fax; and 73% said less than 20% of all written communication is sent by mail.

Q6. Which of the following aspects of a communications network are most important to you? Please rank in order of importance, using a scale of 0-10, with 10 being the highest.

(a) Connectivity to whom you wish to communicate.

Over 50% of all respondents gave this feature of the network a weighting of 10; and a further 15% gave it a weighting of 9. This feature was considered to be the most important aspect of this Network.

(b) Availability of the network when you wish to use it.

15% gave this a weighting of 10, and over 37% gave it a weighting of 9. This feature was considered to be the second most important feature to provide.

(c) Capability of the network in terms of capacity to support high speed applications.

Only 10% of all respondents gave this a weighting over 8. Of all the features listed this was identified as being least important.

(d) Cost of the network to use.

11% of all respondents gave this aspect of the network a weighting of 10, but almost 40% gave it a weighting of 8 or 9.

(e) Ease of use.

Relative to some of the other aspects of the network this was not considered to be as important. Only 10% gave it a weighting of 10 and a further 20% a weighting of 8 or 9.

(f) Communication services supported by the Network

This feature was also not considered to be as important as some of the other conditions. Only 16% of all respondents gave it a weighting of over 8.

Q7. Would you like to have access to a more capable network?

Over 70% of all respondents said they would like to have access to a more capable network.

Q8. If "YES" to the above question please specify the applications that you would run on this network, that you cannot currently run using existing networks.

Respondents suggested the following applications:

- Video Conferencing;
- Large File Transfer at high speed;
- Need higher speed to remote locations;
- Transmission of graphics quickly (better resolution);
- Transmission of all kinds of files –text, binary and graphics;
- Quicker response time (provide interactive talking through network);
- Single point access to all networks;
- Reach more colleagues;

- Higher aggregate throughput;
- Ability to include charts, graphs in each manuscript;
- Transmission of non–english text complete with accents; and mathematical formulae;
- Better FTP to Europe;
- Access to distributed resources;
- Link to supercomputers;
- Exchange of computer graphics information;
- Better access to library facilities;
- Access to distributed directories and databases;
- Access to more sites;
- Access to existing and forthcoming government reports;
- Centralized database of scientific information;
- Standardized image transmission;
- Real–time communications;
- High resolution and colour imaging — 3D graphics;
- Multimedia electronic mail;
- Standardized addressing;
- A gateway to FAX addresses;
- A distributed name server service;
- Exchanging geographic data (maps);
- Talk on the phone and transmit sketches (drawn by mouse); to colleague;
- Transfer of medical images.

Q9. Identify the benefits that you would gain from the existence of such a network.

A number of benefits were identified with this Network:

- Increased interactive, real–time collaboration;
- Elimination of complicated procedural steps to use the Network;
- Ability to test candidate solutions to research problems;
- Wider usage, therefore wider connectivity;
- Do not have to operate and install own private networks;
- Speed up decision–making and results;
- Enhanced graphics;
- Reduced reliance on magnetic tape for file transfer;
- Improved access to up–to–date software;
- Higher productivity and efficient use of computer facilities;
- Ability to experiment in multimedia electronic mail save shipping and travel costs;
- Collaborative R&D would be enhanced at cheaper cost;
- Opens new business and research possibilities;
- Reduction in travel costs;
- More reliable electronic mail
- Environmental benefits (reduction of paper)

Q10. Please provide any other comments you feel are relevant to this study.

Some of the respondents had these additional comments to make:

- Do not create multiple national networks. Supporting users on multiple networks is impractical.

- It is difficult to predict with certainty all the applications and benefits of this type of network. Many of the applications will evolve, once users get used to having access to the network. We do not yet understand the full impact of the use of networks. More applications will be discovered as the networks become established.

- It is time Canada regains its lead in networking. This network should be done right the first time;

- Canada has the large land area and long distances to make networking a very important priority;

- Canadian researchers are at a distinct disadvantage when compared to American researchers who have much better access to network services.

- The advent of high speed networks in the U.S. is only now beginning to allow the development of the kinds of software that will truly empower science. The departmental minicomputer and workstation were the most recent revolutions in scientific computing; networking is undoubtedly the next.

- It is a shame that Canada is about 5 years behind the States in terms of communications technology. Canada has been slow to develop the level of networking needed today to support their research community. The creation of a network like the one proposed is the most important contribution that the Federal Government could make to promote Canadian R&D.

- It is access to global networks, not connectivity between Canadian networks that will provide maximum leverage for Canadian Science. For example, Canadian scientists tend to collaborate with non–Canadians because we tend to develop small groups with high specialization, generally not duplicated in the country.

- Networking in Canada is significantly hampered by the costs of leased lines (which are up to 5–10 times an equivalent line in the States).

- The network should provide universal access to all researchers in Canada.

- Main problem with existing networks like Bitnet is that about 20% of all inter–institutional mail gets lost, without any notification to the sender.

- Network reliability goes down considerably when going offshore.

- It would be of great advantage if the Canadian research community could lead the world in electronically supported distributed research — it is a logical scenario given the distances that divide us.

- The biggest problem in networking today are not technical but organizational and political. It is quite easy to set up a relatively high speed network either locally or on a larger scale. However it is not so easy to get people to agree on the system to use, or to make the system easy to use. What happens is networking efforts are not coordinated and a number of incompatible networks are set up.

5. TELEPHONE INTERVIEWS

5.1 METHODOLOGY

Small IT companies (seen as potential users or suppliers) were not polled by the expert panel sessions, the in–person interviews, nor the network survey. Therefore, it was judged that additional phone interviews were required to adequately cover this user group.

Over 70 organizations across Canada (ie. every province), were contacted by phone to answer a short questionnaire regarding this study.[7] About 32 organizations returned the survey, a response rate of about 45%. Appendix D–1 contains a list of all the companies contacted and Appendix D–2 includes the questionnaire itself.

5.2 RESULTS

This section summarizes the information collected from these telephone interviews. The results are presented question by question.

Q1. Briefly identify the types of communication products and/or services that your organization provides:

85% of the companies responding to this survey produce some kind of "communications product or device" whereas only 15% provide some type of communication service. Some of the products that these companies produce are as follows: communication software, multiplexers, fibre optic systems, communication cable, fax networking equipment, network access/routing products, LAN servers and modems. The services provided by these companies include fibre optic cable splicing and testing, systems engineering, systems integration, network management, and packet switching.

Q2. Briefly describe the nature of your research, current and future:

The research areas being pursued by these companies are as follows:

- Specification and design of ISDN switches and value–added features;
- Designing new transaction terminals/software for banking and retail;
- Development of telecommunication equipment such as fibre optic and microwave transmission systems and digital switching equipment;
- Application software;
- Fast Packet Switching;
- Transmission of image;
- Communication Software (Network Management);
- Algorithms and concepts to meet call accounting market requirements
- Data Communication Protocols and Technology

[7] These companies were selected from the CANTECH database which is owned and maintained by Hutchison Research. This database of canadian high technology companies is the most comprehensive and reliable index available.

- Compression, compaction and other performance issues
- Faster network transport protocols;
- Different access techniques;
- Distribution of digital and audio video;
- Improved integrated LAN workstation networked with mainframes, providing data, voice, and video capability on a broadband backbone

Q3. Which Network(s) are accessed by your organization/department?

Only two of the companies contacted have access to their own network, and only one company makes use of Netnorth and the Internet in the States. Datapac and Infoswitch are used extensively to accommodate the data communications requirement within these companies.

Q4. Would you like to have access to a more capable network?

50% of all respondents indicated they would in fact like to have access to a more capable network.

If "YES" to Q4 please indicate the features you would like to see.

These features include the following:

- Capability to transmit text and graphics electronically rather than fax;
- Connection to the U.S;
- Access to specialized databases;
- High speed access;
- A "reasonable" connection fee;
- OSI Protocol testbed;
- Testbed for ISDN Services;
- Certification facilities recognized internationally;
- Test suite for protocol verification
- Links to library services; and
- E–mail

Q5. If a national high speed R&D network were available, would you consider using it as a testbed for the investigation and initial development of communication products?

Approximately 50% of respondents suggested they would use the network as a testbed. This testbed would be valuable for the following reasons:

- Currently limited in the volume of data that can be tested when using own test facilities;
- Testing chip designs;
- Do not currently have access to a network and believe it would be useful for the type of products offered;
- Accelerate conformance to international standards;
- Would facilitate the development of communication software;
- Provides an environment that cannot be simulated in a laboratory;
- Would provide access to multiple hosts using multiple disciplines; and
- Development of broadband ISDN products requires higher speeds.

Q6. What percentage of your work is undertaken in collaboration with colleagues/branch offices in remote locations?

70% of all respondents said that less than 10% of their work is undertaken in collaboration with colleagues; 20% said 10% –20% and only 10% stated that over 20% of their time is spent collaborating with colleagues.

Q7. How would this percentage change, if at all, with a convenient means of contact through *The Network*?

Just over 60% of all respondents believed that access to the network would lead to a 10% increase in the time spent collaborating with colleagues/branch offices.

Q8. What would be the benefits of *The Network* be to your organization/department? (productivity improvement, marketing and increased competitiveness, new products and services opportunities)?

The benefits identified by the survey respondents are listed below:

- Time would be saved if did not have to physically load fax machine;
- New product opportunities;
- Development of specialized databases will significantly reduce the time required to access information;
- Improved productivity leading to improved competitiveness;
- Improved contact with market of peers;
- Improves time to market;
- Facilitates new product ideas; and
- Will increase market awareness.

Q9. What is your organization's total annual budget for R&D?

50% of the companies contacted spend less than $2M per year on R&D; 42% spend between $2M – $4M; and only 8% spend more than $8M.

Q10. What is the average annual value of sales of your organization?

25% of respondents to this question indicated that sales were less than $10M per year; 30% stated that sales were between $10M –$20M; 25% between $20M – $30M; and only 20% above $30M.

By combining Q9 and Q10 it was found that 60% of respondents spend less than 10% of sales on R&D; 20% between 10% – 20%; and 20% between 20% – 30%.

Q11. Would you expect access to *the network* to have an impact on your gross sales in the long run?

50% of respondents to this question stated that they believed the network would have a positive impact on gross sales in the long run. The expected value of this increase is in the range of about 5%.

Q12. In light of the benefits of *The Network*, and assuming it is non-profit, how would you prefer to support the development and operation of the network? -- by user fees, or would you prefer to support the network through other means such as donating staff time, computing facility time, and/or equipment?

20% of respondents felt they would be willing to support this network through user fees, as long as the fees were less than $10,000 per year. However several users indicated that they would be willing to pay as much as $25,000 per year, for access to this network. More than 90% of the companies who support user fees also favour a user fee based on usage.

Less than 10% of respondents indicated their willingness to support this initiative through the donation (or discounts) in expertise, computing facilities, hardware or software.

CHAPTER 3: FUTURE APPLICATIONS

1. A FRAMEWORK FOR FUTURE APPLICATIONS

1.1 INTRODUCTION

The proposed network will be designed to serve the needs of researchers throughout the 1990's and beyond. It is speculated that the network could begin operations at the T1 rate early in 1991. However, the network must be positioned for a smooth migration to higher speeds (ie T3 first migration). New standards and technologies will be incorporated into the network as they evolve. In fact, the network would serve as a vehicle for the development and testing of future information technology products and services. Future applications of the network must be anticipated so that the evolution of the network can be properly planned for. The information required about these future applications is the subject of this report.

This section provides a conceptual framework for future applications of the network. The framework borrows heavily from a report summarizing a National Science Foundation (USA) workshop on the *National Collaboratory*, a vision of a nation–wide structure that provides tools enabling researchers to perform their work and collaborations without regard to geographical location (Lederberg and Uncapher, 1989). The framework is a multi–tiered hierarchy (see Figure 1, adapted from Lederberg and Uncapher, 1989). At the top are the research functions in the scientific process that researchers and developers cycle through. Scientists achieve these functions through a variety of collaboration functions, which in turn are implemented via collaboration tools and underlying technologies.

1.2 THE PROCESS OF SCIENTIFIC DISCOVERY

Many of the research and development functions required by scientists map onto the steps taken in the process of scientific discovery. (Lederberg, 1989)

1. Gain the education, motivation, and social perception of an area.
2. Build a theoretical framework and formulate questions.
3. Design an experimental approach to answer the question.
4. Assemble resources, including space, people, equipment and money.
5. Conduct experiment, possibly using project management techniques.
6. Analyze the data.
7. Form and test conclusions and theories.
8. Present preliminary results to colleagues through draft manuscripts and informal presentations.
9. Formally publish and present the results in journals, conferences, etc.
10. Obtain community reaction.
11. Establish technical applications and/or commercialize the results.
12. Attract new participants; establish a new organization; extend the project.

Research and development functions	Educate, socialize, motivate investigator	Build theoretical framework, formulate questions	Design approach	Assemble resources (space, people, equipment, money)		Project management & organization
Collaboration functions		Idea generation	Supporting community knowledge base (acquisition, assimilation & dissemination)	Collaborative resource finding Finding a collaborator	Project management & organization Capture, display, editing state of collaboration Negotiating & allocating scarce resources Decision support	Remote & collaborative experiments Telepresence Rapid remote prototyping and transport of artifacts
Collaboration tools	User training, education tools	Digital library Standards for Access	Information retrieval systems	Structured discussion	Group decision support systems	Remote experiment & resource scheduler
Enabling technologies	Software engineering Engineering technology for group interfaces Software engineering tools	Meta technologies for building required technologies & tools Instrumentation of tools to understand their impact	Standards (graphics, window systems, information...)	Graphics, human-computer interface & multimedia Windows & virtual graphics across different machines Shared windows and workspaces	Speech/sound Video in a network & in a window	Display hardware Large public displays Portable displays & input systems

The collaborative framework.
Adapted from Lederberg and Uncapher, 1989
(continued on next page)

Research and development functions *(continued)*	Analyze data	Form & test conclusions, theories	Presentation to colleagues	Publish	Obtain community reaction	Commercialize	Attract new participants/form new organizations
Collaboration functions *(continued)*	Managing physical environment Service via email		Ongoing/open scientific review & commentary	Prepare documents Publish documents Review documents	Meetings Distributed seminars & meetings Video teleconferencing Synchronous & asynchronous communication Shared gestures space		Informal collaboration Social Browsing Real time meeting management Stuctured group history (note taking)
Collaboration tools *(continued)*	Bulletin boards Asynchronous conferencing		Remote meetings and editing Co-authoring tools Shared drawing surfaces & sketching tools	Meeting room to meeting room protocols Packet video for video meetings	Video telephones & conference rooms Virtual hallways		Conference coordination Cursor-passing protocols & controls for group editing
Enabling technologies *(continued)*	Robotics Tele-operation		Communication High bandwidth networks (gigabits?) Portable communication Link services	Video transmission Tools for ntetwork performance, debugging & management	Information databases, indexing & retrieval Document representation, editing, access Fast linguistic search Automated image segmentation & feature extraction		Digital image bank Hypertext & Hypermedia Sharable, persistent databases

The collaborative framework.
Adapted from Lederberg and Uncapher, 1989
(continued from previous page)

Regardless of discipline, these stages of research are a reasonable approximation of the high–level tasks performed by scientists. As the process applies to distance–separated scientists just as it does to co–located teams and individual investigators, the stages provide a meaningful motivation driving the applications the network should support.

1.3 COLLABORATIVE FUNCTIONS AND TOOLS

To accomplish these steps listed above, distance–separated scientists draw upon a variety of collaboration functions. They must interact casually and formally with remote colleagues; share data and resources; monitor distant instrumentation and control equipment; share information in a knowledge base; manage the project; make decisions, and so on (Figure 1, second row). Several of these functions are briefly described here, along with the collaboration tools that support these functions (Figure 1, third row); later sections of this report will expand on these.

Community knowledge base: Researchers need to acquire information from a communal knowledge base to assist their idea generation and project development. They also need to disseminate their own works into it. The community knowledge base is still mostly resident in the paper medium, leading to problems of publication delays and difficulties of finding the required information. A digital library that can overcome these problems by offering good information retrieval tools, and by allowing new information to be well–assimilated in a timely manner. Digital library should support multi–media and set standards for access.

Finding collaborators through casual interaction: When probing for collaborators and initiating a project, researchers usually identify shared interests, argue assumptions, generate ideas, probe for suitable partnerships and extend commitments. The process often requires frequent, brief and spontaneous communications between potential collaborators. Since interaction between people is well–proven to have an exponential decay with distance (Kraut, Egido and Galegher, 1988a), technical support for casual interaction is aimed at overcoming the distance barrier.

Project management and organization: Managing a research project–especially a large ones–is a daunting task for distributed researchers. Participants require technology to assist project management. This includes not only the generation of the project plan, but also the ability to capture, display and edit the state of the collaboration. In addition, group decision support systems can assist the members to come to consensus. Technologies included project management tools, structured electronic mail tools that track researcher's commitments to each other (Winograd, 1988a), and structured argumentation tools to assist the decision–making process (Conklin, 1987).

Remote and collaborative experimentation: Remote experimentation occurs for several reasons. Because resources are scarce or expensive, they may be held by outside organizations and agencies. In this case, researchers are required to negotiate, allocate, schedule and use tools that are situated elsewhere. Alternatively, the experiment may be performed at one scientist's site, with distant collaborators being involved through remote control. Technical support of these activities falls into several areas. Robotics and tele–operation are required for remote control of the experiment. Resource schedulers assist the booking of scarce tools. Tele–data–the ability to see data at a variety of sites–is required if collaborators are to see and react to the generated experimental data. Remote prototyping is necessary if collaborators are forming the experiment interactively.

Ongoing scientific review and commentary: An important part of the scientific process is to obtain reaction from one's peers during the complete life–cycle of the experiment. For this to succeed, a researcher must be able to disseminate informally the on–going aspects of the

work, and to collect the commentary back from the community. Technology can support this process in several ways. Through electronic mail, the scientist can send the information to specific people, while information can be broadcast to the community at large through a bulletin board or asynchronous conferencing system. Replies can be returned to the bulletin board, or personally to the researcher. Real–time support for distant casual interactions is another way for researchers to disseminate and gain feedback on the fly.

Collaborative document preparation: Preparing formal documents depends heavily on the styles and responsibilities of the researchers involved. Minimally, collaborators must be able to exchange drafts and comments, possibly through electronic mail. However, the future will see tools that let people interactively create and revise documents together. The technology includes real–time interactions with a document displayed simultaneously on all participants' workstations (Greenberg, 1990), or through co–authoring tools that provide structured access and annotation to a document (Leland, Fish and Kraut, 1988). The final document can be disseminated through the community knowledge base.

Meetings: The real–time meeting is perhaps the cornerstone behind the collaborative process. It is vital for brainstorming, rapid evolution of ideas and arguments, and for formalization of goals and methodology. Technology can support a variety of real–time meetings in several ways (Greenberg and Chang, 1989). Video tele–conferencing gives a sense of "tele–presence" by allowing distributed participants to see one another. Shared workspaces, along with annotation tools and a gesture space, permits participants to work together through a shared media. Meeting schedulers assist the daunting task of trying to collect all participants together at the same time. Other technologies can help a group capture and save the important points as a meeting progresses.

1.4 OUTLINE OF THE REPORT

This report structures the needs of distributed scientists into three general approaches: supporting communications between researchers (Sections 2 through 5); providing researchers access and control over remote resources and equipment (Section 6); and tools that let researchers store and manipulate information (Section 7). Although the sections do not map directly onto the framework described above, they do cover most of the collaboration functions and tools listed in Figure 1. These sections begin with an introduction to the collaborative approach provided, relating it to the specific needs of researchers and developers. Central to each section is a description of current visions, with emphasis placed upon research systems that are near commercialization or likely to be commercialized in some form over the next decade. Limitations and unknowns of the technologies are raised, and network requirements estimated. In contrast, Section 8 highlights three application areas as examples collaborative functions. Section 9 closes the report with the vision of the coherent workspace.

Computer support for real–time remote conferencing is discussed in Section 2. These are the formal or semi formal meetings that geographically–distributed people attend. Techniques for transmitting a sense of participatory presence and for sharing data over distances are covered.

Casual real–time interaction, covered in Section 3, is a crucial yet often overlooked component of scientific collaboration. While formal meetings take care of key events in the scientific process, casual interaction at the interpersonal level is the glue that keeps the collaboration working. Several innovative systems are described that bring people together through frequent, unplanned, and high–quality interactions.
Asynchronous messaging is explained in Section 4. This covers the familiar electronic mail systems we are now used to. The domain is extended to more sophisticated versions of

electronic mail, including multi–media, semi–structured, and semi–formal mail systems. Bulletin boards and asynchronous conferencing is described in Section 5. These are systems that allow people to interact with a large community by posting and viewing messages, and joining and actively participating in on–going discussions of some topic of interest. Technical variations and issues behind these systems are provided.

Access, monitoring and operation of distributed equipment is covered in Section 6. Two aspects are distinguished: access to remote computational power, and real–time systems for remote operation. Technologies such as Scada systems, master–slave tele–operation, and tele–robotics are included.

The digital library is detailed in Section 7, and lists conventional approaches such as the digital repository and bibliographic databases, as well as the attractive alternative of hypertext. The digital library is the backbone behind the communal knowledge base.

Several applications are provided in Section 8 to illustrate how the approaches of the previous sections can support high–level collaborative functions. Application areas include distance education, group decision support systems, and joint authoring. A brief note describing the need for the network for bootstrap collaborations is provided.

A vision of the *coherent workspace* closes the report in Section 9. The coherent workspace is an idealized integrated platform that supports distributed researchers. It occurs when people and their machines have a common knowledge and sense of purpose, with individual and group activities being well–coordinated.

2. SUPPORT FOR REAL-TIME REMOTE CONFERENCING

2.1 INTRODUCTION

Real–time remote conferencing brings participants of formal or semi–formal scheduled meetings together at the same time, even when some or all are physically distributed over different locations. This purpose is reflected in the several other names that remote conferencing goes by: multi–site tele–conferencing; distributed meetings; and same–time, different place meetings.

Remote conferencing is crucial for geographically distributed researchers and developers, with values and functions similar to face–to–face meetings. The list below indicates a few benefits.

1. Remote conferencing is an education tool where people can take interactive and participatory classes from distant experts (who can also be considered a "scarce resource" in a way similar to expensive and rare equipment).

2. Meetings are crucial for preliminary and on–going design sessions that involve brainstorming, sketching out ideas, creating outlines, and so on.

3. Meetings are needed by collaborators for on–going project development. Examples include setting goals, scheduling, formalizing the various stages of the project (design, methodology and analysis), and discussing and solving problems.

4. Draft seminars are presented to other collaborators for feedback and further development before going public.

5. Meetings are the backbone behind informal and formal research presentations to the community at large. Not only can one researcher present preliminary, on–going and final results to distributed colleagues, but the audience can provide immediate feedback, ask questions, and so on.

This section will describe computer support for real–time remote conferencing.
Two aspects of conferencing are distinguished in the following two subsections: *tele–presence,* a way of giving distributed participants in a meeting a feeling that they are in the same meeting room and the means to orchestrate the conversation, and *tele–data,* a way of having participants bring into the meeting the materials and on–going work they wish to share with others. Illustrations of several research systems furnish an image of what we can expect from tele–presence and tele–data over the next few years. In the next sub–section, an innovative and futuristic approach called *cyberspace* is described, a potentially powerful means for conferencing, data sharing and visualization that we may see in limited use by the end of the decade. The subsequent sub–section raises several limitations and unknowns that, in spite of the apparent needs for remote conferencing, hinder its use. The section closes by matching the network capabilities to the requirements of real–time remote conferencing.

2.2 TELE–PRESENCE

Perhaps the most interesting aspect of remote conferencing has been in the field of *tele–presence*–a way of giving distributed participants a feeling that they are in the same meeting room. Tele–presence concentrates on transmitting both the explicit and subtle dynamics that occur between participants. These include body language, gestures, eye contact, meta–level communication cues, knowing who is speaking and who is listening, voice cues, and so on. Tele–presence facilitates effective management and orchestration of remote meetings by the natural and practised techniques used in face–to–face meetings.

One simple experiment in tele–presence is found in *Cantata*, a text–based remote conferencing tool that includes a multi–window broadcast environment (Chang, 1986; Chang, 1989). Participants converse with each other through text, and each participant's text is displayed through a window on the screen. An attendee has the option of selectively viewing any other participant's input to the conversation by raising that participant's window. When text is typed by one or more "speakers," it is broadcast to all participants and displayed in the window representing that speaker.

Although simple, Cantata supports tele–presence through several devices.

1. *Knowing who is speaking*. Listeners know who is "speaking" because they can see text appearing as a sender is composing it. This is especially important because, unlike voice, many people can simultaneously broadcast text.

2. *Focus of attention*. The "listener" has the option of paying less attention to specific speakers by adjusting a *focus of attention* gauge, resulting in the text being filtered to show only occasional words. The less the attention, the less actual words displayed. Listeners still get a "background hum" by seeing dots printed as the words are composed. The speaker has a corresponding gauge that indicates how much overall attention the group is paying to him.

3. *Interruptions*. A person can force everyone else to pay attention to him via a text–based equivalent of interrupting a meeting through shouting. Cantata allows any participant to compose and broadcast an interruptive message to other attendees. The sent message appears in its own window popped up on top of all other windows on the receiver's screen. Unlike the normal broadcast environment, participants cannot disallow, hide, or filter the "shouted" text.

Text–based communication, although applicable in some situations, is likely ineffective for the majority of remote real–time meetings. Most research and commercial efforts have investigated visual and audio tele–presence through *video conferencing* across dedicated meeting rooms. One or more people meet in a room; other participants in the conference meet in equivalent rooms at the distant sites. Video images of the attendees are then transmitted between these rooms across a high–bandwidth communication channel. In the simplest case, a camera will just transmit an image of all participants in the room, perhaps with the camera focusing on the active speaker. A more complex scenario would see a single screen for every participant, where monitors and speakers are all located in the same relative position across all rooms so that eye contact and directional sound cues are maintained.

Perhaps the best effort in tele–presence is MIT Media Lab's *Talking Heads* (MIT, 1983a). A remote participant is represented by a translucent mask (cast at a previous time) of his face. The video image is projected into the mask, giving the effect of a 3–d "hologram." The innovative aspect is that the mask rotates to reflect the actual head movement of the person, as picked up by motion sensors. The feeling of presence is striking. Participants can naturally and effectively orchestrate the conversation through normal eye contact and head movement.

Low bandwidth versions of Talking Heads do away with the video signal by transmitting only the head movements and the audio signal, and then selecting for display one of several pre–stored images or caricatures of the speaker that best match the incoming signala speech recognizer is used to match lip movements (MIT, 1983b). Although this facade is not nearly as rich as the full bandwidth version, participants still see who is speaking and are able to direct attention to each other.

2.3 TELE-DATA

Most real meetings require not only the people, but also the materials and on–going work participants wish to share with others. These include notes, documents, plans and drawings, as well as some common work surface that allows each person to annotate, draw, brainstorm, record, and convey ideas during the meeting's progress. Given that an individual's work is commonly centred around a workstation, the networked computer can become a valuable medium for people to share on–line work with each other. *Tele–data* allows participants at a meeting to access, present and interact with materials that would normally be inaccessible to the distributed group.

For example, one multi–site tele–conferencing setup that uses several types of tele–data is the Multipoint Interactive Audiovisual Communication (MIAC) audiographic conferencing system (Clark, 1989). MIAC supports remote communication through transmission of high quality audio, facsimiles, still picture TV frames, real–time *tele–writing*, and control of interactions over a 64 kilobit/second communication channel. Its salient features follow.

> *Audio.* Each participant has his own microphone. Listeners receive an indication on their display of who is currently speaking. In a non–conducted meeting, anyone can speak at any time. In a conducted meeting, the chairperson can speak whenever he wishes, while other participants must explicitly request the floor from the chair.
>
> *Video.* Still picture TV is used to transmit a single video frame between meeting rooms.
>
> *Writing / Sketching.* Using a data tablet, participants can exchange handwritten information in real–time via the *tele–writer*. Three tele–writing scenarios are possible: exchanging the tele–writer image only; superimposing the image on the still–picture TV; and moving the cursor over the display.
>
> *Facsimiles.* A facsimile can be loaded and sent from one site to another through a facsimile machine. MIAC mediates the point to multipoint communication.
>
> *Messaging.* Short text messages can be sent between participants.

Another approach to tele–data stems from taking a standard computer application and sharing it between participants of a remote conference through a "shared screen" or "shared window." Each participant sees the same image of the running application on his own screen, and has opportunity to interact with it by taking turns. Special "*view–sharing*" software would allow *any* unaltered single–user application to be brought into a meeting; the application itself would have no awareness that more than one person was using it. The view–sharing software's responsibilities include registering participants, maintaining consistent shared views, managing floor control for serial input to the application, and allowing attendees to gesture and annotate around the view (Greenberg, 1990). Although simpler in idea than true multi–user applications that are aware of and cater to all participants, the capability of sharing views and interactions with the many single–user applications now available can augment significantly people's ability to work together, both in face–to–face and remote encounters.

Shared view systems are far from new. Over twenty years ago, the visionary Doug Engelbart held what was probably the first shared screen conference through his NLS system (Engelbart and English, 1968), where six displays were arranged on a table so that a group of twenty participants could see the screens. While only one participant could control the screen, other participants could control a large arrow (the first tele–pointer). Since then, shared screen systems have evolved to match current interface capabilities. *MBLINK*, for example, not only allows multiple workstations to share a screen bitmap, but also displays each participant's distinctive cursor on the view (Sarin and Greif, 1985). Several research systems now permit people to share and arrange individual windows rather than the complete screen, achieving greater flexibility by allowing one to arrange his personal display to include both private work and shared windows (Ensor, 1989; Ensor, Ahuja, Horn et al., 1988; Gust, 1989; Lantz, 1986). At the Alberta Research Council, Greenberg has decoupled the view–sharing kernel from the interface required for explicit floor control, resulting in a system that can be readily specialized to the needs of the participants and to the hardware requirements (Greenberg, 1990). On the commercial front, Farralon Software sells a simple, inexpensive but surprisingly effective shared–screen facility for the Macintosh called *Timbuktu* (Farallon, 1988). A detailed description of shared view systems is available in Greenberg (1990).

Other quite sophisticated applications of tele–data are based upon a communal multi–user workspace that allows all participants to share, view and simultaneously interact with a work surface. Perhaps the best and most well–know example is the software suite available in *CoLab* (Stefik, Foster, Bobrow et al, 1987b). Although designed by Xerox PARC to support face–to–face meetings by small groups, its software is relevant to remote conferencing. The CoLab software suite consists of three tools: *Boardnoter*, a shared chalkboard; *Cognoter*, a tool for brainstorming and idea organization; and *Argnoter*, a tool to organize and evaluate arguments.

Boardnoter supports and extends the notion of a chalkboard, and is a good example of how computer–based technology can mimic and augment the capabilities of existing media devices. Like a chalkboard, the Boardnoter supports informal freestyle sketching and erasing on a communal work area by participants. Unlike a chalkboard, the Boardnoter extends the drawing power of the group:

- text and figures to be easily movable;
- images to be re–organized on the display and stored in databases;
- active involvement of participants from their seats;
- multi–person gesturing via a tele–pointer (Stefik, Bobrow, Foster et al., 1987a).

Meeting tools need not be as generic as a chalkboard. Using existing equipment, a meeting can bring in a particular flavour of meeting tool designed to facilitate a special meeting task. Cognoter is an example of such a tool, for it is explicitly designed to separate the planning of a presentation into three separate stages (Foster and Stefik, 1986; Stefik et al., 1987b). In the first stage, ideas generated by each participant are entered simultaneously as independent "catch–words" in the free space on a public window. Catchwords can be further annotated with supporting text in sub–windows. In the second, ideas are organized by allowing participants to order, link and cluster the displayed catchwords. Physical and conceptual clutter is reduced by allowing people to chunk grouped items into a hierarchy. The third stage involves idea evaluation. The overall structure is reviewed, details added, and irrelevant ideas removed.

2.4 CYBERSPACE

The most innovative and futuristic approach to remote conferencing may lie in *Cyberspace*. Cyberspace immerses a person's senses into a three–dimensional (3–d) simulated virtual world. Seeing the world in a stereoscopic head–mounted display that has a screen for each eye, one moves through it by head and body gestures. Motion sensors pick up and translate real movements to virtual ones, and the view is adjusted accordingly. Users interact with the simulated world through a data–glove or data–suit that allows them to grasp and manipulate the virtual objects they see. They hear sounds through a 3–d audio display. The effect, although still primitive, is to exist and interact within a virtual reality–cyberspace.

The relevance of cyberspace to remote conferencing becomes apparent when two or more people interact within the virtual space. Imagine a conference held in virtual room, with attendees milling about, holding private conversations, and viewing and manipulating some of the available 3–d data entities. Science fiction? Not quite. The first demonstration of VPL Research Inc's *shared virtual reality* system occurred on June 7, 1989 in San Francisco[1].

The potential of cyberspace is far beyond anything offered by other remote conferencing systems. Not only can participants interact with each other (eg by moving around a virtual room), but the meeting place can simulate any environment. For example, the cyberspace model may be constructed and animated from incoming data points of an instrumented experiment (or the simulation of an experiment), with scientists viewing the data and adjusting the control parameters interactively. Alternatively, the model may be one of a conventional meeting room with a sophisticated whiteboard, so that participants can act within a cyberspace meeting as they would in a real meeting.

2.5 LIMITATIONS AND UNKNOWNS

Although both tele–presence and tele–data are clearly important, their effective implementation and the role they play in particular types of meetings are not well known. Requirements of meetings vary greatly; group dynamics are volatile; people are not used to computers in meetings; software for meeting support is (at best) at the prototype level; technology is obtrusive.

Video conferencing, for example, has fallen far short of its promise for several reasons.

1. Vendors gave video conferencing an ill–conceived image as a means of reducing the need for travel to face–to–face meetings, which it does not do (Egido, 1988). Travel actually increases, for the need for direct meetings grows with the frequency of the interpersonal contacts made over video.

2. Although video conferencing has proven suitable for passive meetings emphasising presentations, it appears to be a poor medium for supporting the more common highly interactive style of meeting where there is much inter–personal interaction (Egido, 1988).

3. Video presence of participants seems to add little to communication (Chapanis, 1975; Johansen and Bullen, 1984), and therefore may not warrant the technological and physical restrictions it places on the meeting.

---

[1] Off the shelf equipment to realize 3–d cyberspace is available through VPL Research Inc, Redwood City, California.

Anecdotal evidence supports these views. One account mentions that after a six–month novelty period had worn off, the day–to–day use of the video conferencing facility shifted from display of participants to simply pointing the camera to the data. Participants were content to talk anonymously over a speaker phone (Guttman, 1989). It is the data that was considered more important than the simple form of tele–presence offered by video.

Some of the failures above may not be due to the notion of tele–presence, but rather to limitations in technology. Special video conferencing rooms, for example, means that participants must schedule and limit their meetings to these rooms rather than use their own offices. (Although cyberspace conferencing may remove these restrictions, the technology is costly and the effects still primitive.) The high cost of bandwidth means that video transmission often uses a compression scheme that severely impacts on the quality of the displayed image. Furthermore, keyboards, monitors, microphones, wires, and cameras may be a significant intrusion to the meeting, particularly to those participants who do not feel comfortable with the technology.

Given that face–to–face meetings are considered more effective than remote ones, there is still much room for improvement. Some progress is being made. Xerox PARC, for example, uses multi–disciplinary teams of sociologists, anthropologists and computer scientists to study how people communicate through a video channel (Harrison, 1989; Stults, 1988). The knowledge acquired through this process will help determine the design requirements behind truly useful remote conferencing systems.

These limitations should not discourage use of remote conferencing. Technical problems will be overcome and the human factors issues solved. The real potential exists because remote conferencing is vital to overcoming the distance barrier faced by distributed researchers and developers.

2.6 NETWORK REQUIREMENTS

A 1.5 Mbps network can address some of the requirements of remote conferencing. In particular, the tele–data requirements that have already been developed and used on LAN networks can migrate to the 1.5 Mbps network with little performance penalty. These include screen sharing of existing single user application packages, and multi–user workspaces that allow participants to simultaneously draw and type in the view.

The network may also suffice for some limited types of tele–presence. Textual real–time communication systems such as Cantata require modest bandwidth, while the shared gesturing ability found in workspaces probably requires little more than the transmission of a person's pointer location. As mentioned above, low bandwidth versions of Talking Heads does not require video transmission (although there is still a question about how well a voice–animated image provides tele–presence). Similarly, slow–frame video can capture at least some aspects of tele–presence.

Yet the dynamics between attendees in a meeting can be so diverse that it is unlikely that any limited channel will capture all the subtle yet important cues. It stands to reason that real tele–presence would require at least a full video channel and a corresponding increase in bandwidth over the network. At the very least, one video connection is needed for room to room contact. Bandwidth demands multiply with the connections made, as seen when the number of rooms increase, or when video screens represent each individual in the meeting.

A similar high bandwidth requirement is needed when tele–data requires video, as may occur when the data is not available in digitized form, is not easily scanned in, or is animated.

40

The actual network requirements for cyberspace is unclear. Three–dimensional real–time animations must be modelled and exchanged; interactions between participants with each other and with the simulated environment transmitted and synchronized; and sound (including its coordinates of origin) transmitted. Although simple cyberspace interactions using distributed models may run over the 1.5 Mbps network, a true realization of an alternate reality would likely require considerable bandwidth.

3. CASUAL REAL-TIME INTERACTION

3.1 INTRODUCTION: PATTERNS AND REQUIREMENTS FOR SCIENTIFIC RESEARCH COLLABORATION

Real–time conferencing addresses scheduled formal and semi–formal meetings. Yet it is not necessarily pre–planned, purposeful meetings that are best supported through computer mediation, but casual unplanned meetings as well. This section argues that real–time support for casual interaction is vital to scientific research collaborations and must be supported by the network. Several innovative systems indicate the potential shape of these future applications.

Research is fundamentally a social process, with many interactions required for people to initiate and execute collaborative research. Yet only a small part of this process relies on the formal communications supplied through scheduled meetings. Consider the model of research collaborations below, developed from interviews with scientists in various areas of research (Kraut, Galegher and Egido, 1988b).

Relationship level	Finding a partner	Supervising and sustaining progress	Establishing division of credit
	Sharing background assumptions	Establishing division of labour	
	Identify shared interests	Establishing trust	
Task level	Generating ideas and	Sharing Information	Writing the manuscript
		Coordinating activities	
		Doing the work	
	Initiation	**Execution**	**Public presentation**

The process of many of these stages clearly rely on frequent casual contacts to maintain collaboration on both the task and the interpersonal level. In initiation of a research project, for example, researchers usually identify shared interests, generate ideas, share assumptions, probe for suitable partnerships and extend commitments. During execution of the project, collaborators must share information and coordinate activities, reshape project goals, establish sub–goals and divisions of labour, comment on the other's work, and so on. Although scheduled meetings may be held during this process, most communication is informal. Through brief communications occurring in quite casual settings over coffee or lunch, in hallways, or perhaps during an otherwise uninteresting meeting) collaborators exchange ideas and information, feedback, encouragement, status changes and minor (yet important) points. The technologies used are simple; a scrap of paper to make a point or jot down ideas, or perhaps a whiteboard for brainstorming. Kraut notes that collaborators rarely used formal project management techniques to assist their task (Kraut et al., 1988b) .

Kraut concludes that the establishment and maintenance of a personal relationship is just as vital to the collaborative research effort as the content of the work itself. Three criteria are usually necessary for informal communication and the resulting collaborations to succeed (Root, 1988):

1. *High frequency of communication,* where potential collaborators have ample opportunity to communicate to each other;

2. *High–quality real–time interactions* allowing intense face–to–face sessions and the sharing of materials;

3. *Low cost interactions* so that potential collaborators can have spontaneous or one–person initiated meetings with little effort.

Root especially emphasises the importance of brief unplanned encounters where bits of technical and personal information are exchanged "on the fly."

Given these criteria, it should be no surprise that the bottleneck to rich spontaneous interactions is distance. Frequency of communication decreases, communication media is often low bandwidth and of low quality, and interactions come at a high cost (eg telephone tag or travel time). Interaction between people is well–proven to have an exponential decay with distance, as shown not only by the reduction in the frequency of face–to–face interactions (Allen, 1977; Kraut et al., 1988a), but also by the lower exchange of telephone calls (Mayer, 1977) and electronic mail messages (Eveland and Bikson, 1986). For example, the number of collaborations drops off sharply when one contrasts people working on the same floor, on different floors, and in different buildings. Further study indicated that the effects of physical proximity and interaction are not merely an artifact of people with similar interests usually being co–located (Kraut et al., 1988a).

The conclusion formed is that distributed researchers are less likely to go through the critical initiation process if they cannot get connected, leading to a decrease in potential collaborations. Similarly, the communication problems encountered during the execution and presentation stage may be significant enough to hinder or abort the project.

Technology has potential to bring distance–separated people into contact through frequent, unplanned, high–quality, and real–time interactions that come at low personal cost. Two visions are described in the sub–sections below: *video hallways* and *shared alternate reality*. Both require high–bandwidth communication networks to succeed. Limitations and bandwidths are raised in later sub–sections.

3.2 VIDEO HALLWAYS

Several research laboratories are exploring the possibility of "video hallways" for casual interaction between remote sites. The first case was Xerox's *Video Wall*, which placed a slow–scan video connection between two research laboratories located in California and Oregon (Goodman and Abel, 1987; Stults, 1988), also summarized by (Root, 1988). Spontaneous "drop–in" interactions between people at the two sites were encouraged by placing large video screens in common areas. Point to point connections between individual offices were also allowed to a limited extent.

Video Wall worked. Goodman and Abel reported that 70% of all Video Wall interactions were spontaneous, and the other 30% planned (Goodman and Abel, 1987). A different breakdown indicated that one third of all communication was social in nature and two thirds technical. Users reported that these interactions would probably not have taken place without

the link. On the other hand, users reported dissatisfaction with the poor image quality of slow scan video.

A second prototype video hallway is *Cruiser* (Fish, 1989; Root, 1988). While Xerox's Video Wall directly connects two physical locations, Cruiser attempts to create a virtual community where everyone has instant access to everyone else. Cruiser is designed on two premises: 1) users can browse a virtual world seeking social encounters, and 2) users can construct, organize and populate the virtual world independent (within reason) of the physical world. There are three methods for browsing.

1. A *jump* supports a direct planned movement to a physical location. A user can select a specific location, and the image projected from the camera at that location appears on his screen.

2. A *path* extends the jump idea by listing a sequence of locations and the order in which to visit them. This, in effect, becomes a "virtual hallway" through which the user can walk through.

3. A *random walk* is similar to a path, except that the Cruiser system generates the locations in the sequence. The selected locations can be purely random, or they can be chosen as a function satisfying some user desire.

What do people actually see when using Cruiser? A visitor peeking into a person's office (via a video camera) will see and hear whatever image and sound the camera projects. The occupant, on the other hand, sees a virtual hallway on his screen and the image of the visitor as he is passing through.

While "peeking" into offices raises the spectre of George Orwell's "Big Brother," the Cruiser design recognizes an individual's desire to control privacy. As in real–world offices, people have the option of metaphorically keeping the door open (seen by the visitor as full video); semi–opened (seen by visitor as partially–drawn blinds across the image); or completely closed (no image projected). Root identifies several variables that can be controlled for setting privacy levels: access of visitors to the video and audio channel; ability of visitors to interrupt; ability of selected visitors to over–ride other settings; and a privacy value (Root, 1988). For example, a closed door policy is implemented by setting video and audio transmission to none but allowing interruptions by people with high priority levels.

Many interactions are not a result of someone "cruising for action." Rather, they arise from people bumping into one another while performing their everyday work, and in joining in on conversations already in progress. Cruiser supports this style of "situated interaction" by allowing people to attach images to several work activities. For example, sending a document to the printer will automatically invoke a random walk that encounters the other people using the printer.

Root warns in his paper that Cruiser is very much a prototype, and much remains to be built and tested. In contrast, the *US West* TeleCollaboration project supports a high–speed full video link between two research laboratories in Colorado separated by a distance of 100 kilometres (Corey, Abel, Bulick et al., 1989) . Their system allows users to "video walk" through the physical hallways and offices of distant sites. They can even scan larger rooms by tele–operation of remote–controlled video cameras. On their own screen, they can search for people, see who is around, start informal conversations, engage others in coffee room chit–chat, and so on.

3.3 SHARED ALTERNATE REALITY

Video hallways use computer support only to help establish personal encounters. Its users have to decide to go cruising, or they have to leave themselves open to intentional or accidental encounters. Randall Smith's vision of the *Shared Alternate Reality Kit* (SharedArk) takes another approach (Smith, 1988). Unlike virtual hallways, SharedArk is more than a medium for meeting people.

SharedArk is based on a model of a shared virtual yet physical world (a two–dimensional "flatland"used for teaching students physics. Students can wander through flatland, and manipulate physical objects with a mouse–operated hand. Unlike most virtual worlds, flatland is populated by all the people travelling in it. Students may accidentally encounter each other (ie one will see another person's hand). They then have opportunity to open an auxiliary video and audio connection for more direct communication. Within SharedArk, students can form collaborations on simulated physics experiments and jointly edit text and graphics.

SharedArk has several other features. First, a person can see who and what is around him through a "radar view" that provides a miniature of the surrounding space. Second, people (and their objects) can quickly move from one virtual site to another by stepping into a "tele–porter." Third, people can set up private regions within the virtual world that excludes other people from travelling or looking within it.

A natural extension to SharedArk is, of course, cyberspace. Pushing Smith's vision a bit further, imagine a multi–world version of cyberspace, where each "world" represents a topic of interest. People travelling through and exploring a particular world of their choice have opportunity to accidentally encounter and make acquaintance with other people with similar interests, just as travellers do in real life.

3.4 LIMITATIONS AND UNKNOWNS

While video hallways are exciting, they have several serious limitations. First, the cost alone is prohibitive. High–bandwidth video between distant sites incurs expensive communication tariffs, and low–bandwidth video is only marginally acceptable to participants. When this is combined with the cost of cabling and adding cameras and monitors to all offices, it is likely that most institutions would consider the simpler video wall setup before video hallways, at least in the short term. A second limitation is social. Although many people are comfortable with working in common work areas or within open offices, video "peeking" raises the concern of invasion of privacy and abuse by management. A third limitation is one of effectiveness. While informal interaction is important and does seem to work for small groups, the possible influx of visitors to an office may lead to people (especially popular people) to adapt a closed door policy. Distance, after all, does provide a buffer against excessive drop–in visitors.

Video hallways have explored personal encounters only. If these meetings could be augmented by importing work tools as well (eg drawing surfaces), we may see the start of the "back of the envelope" computer–supported spontaneous meeting that allows people to develop, jot down and share ideas in a non–formal situation. SharedArk is, of course, closer to this data–oriented model. Smith, however, points out severe performance problems in his system, mostly due to the inadequacies of present–day network and computational technology for handling real–time distributed interaction and animation. In practice, SharedArk will only work well with three or four users and a few dozen objects (Smith, 1988).

Another aspect that must be considered in casual interaction is the "startup problem." Random encounters, while potentially rewarding, can be wasteful of effort, particularly when one

person actively seeks but cannot find other specific people to meet with. The task of establishing a casual encounter with a specific group is even more frustrating than scheduling a meeting. Not only must each person be contacted individually, but it is virtually certain that not everyone is available then and there. Even collaborating on a casual basis with those who are available electronically (ie actively working on their terminals) is difficult without mechanisms for easily identifying that availability. For example, experience with Cantata has shown that people tend not to use it because they have to phone or walk down the hall to get others to activate the program; the cost for a casual remote meeting is just too great.

The problem is that people connected on networks through workstations have few ways of knowing (without video hallways) who is present and available for conferencing, even if they are signed on. Two mechanisms are being implemented by Chang and Copping at the Alberta Research Council. The first, *Messenger*, is a Macintosh desk accessory that shows a person all others who have recently moved a mouse or touched a key within the last few moments (currently set at 60 seconds). A message can then be sent to a subset or all those identified as recently active. Those receiving the message hear a bell and see a flashing icon and can then view the message and reply. The second, *Golf Ball*, is a common knowledge mechanism based on the *Messenger* kernel. A message containing a proposed action (eg calling a meeting) is sent to a group of recently active persons, and the message acceptance and acknowledgments and responses to the proposal are seen by all who accept the message. Thus the entire group is aware of the responses of the members as they accumulate. The success of these mechanisms in promoting casual group formation in networks will be seen with experience.

3.5 NETWORK REQUIREMENTS

A prerequisite of casual interaction is that sessions must be of high–quality and real–time. By definition, then, the network requirements will be high. At the very least, there must be enough bandwidth on the network to support: a method for allowing potential collaborators to know who is around; a voice channel; a shared workspace (tele–data). These demands can be met within the constraints of a 1.5 Mbps network.

Truly effective interaction likely requires a video link. As with remote conferencing, slow–scan video could work at the lower bandwidths, while compressions schemes may allow full bandwidth video across a 1.5 Mbps network. This allows, however, for only a single connection. While this is reasonable for a system like Video Wall, both Cruiser and the TeleCollaboration project would require multiple video links. The alternative is to allow only one person to monopolize the network while wandering the virtual hallways and while establishing a link. A similar problem occurs within SharedArk. While initial contact is made through the multi–user "virtual world," the ensuing interaction establishes a video link.

4. ASYNCHRONOUS MESSAGING

4.1 INTRODUCTION

Perhaps the greatest success story in bringing people together over a network is in computer support for asynchronous communication–non real–time discourse and information exchange between people. This section will concentrate on one aspect of asynchronous communication: electronic mail and its variations.

Electronic mail (email) is an important support tool for geographically–distributed and co–located researchers and developers. It serves a role far beyond standard surface mail, extending into roles traditionally held by inter–office notes and memos, the telephone, and real–time meetings. Although benefits overlap, email affords different capabilities.

- It addresses communication where the sender does not require immediate response. Similarly, it is well suited for those cases where the recipient needs time to formulate an answer.
- It is usually one–person initiated.
- It overcomes "telephone–tag."
- Delivery time is significantly faster than surface mail, with message transmission in the order of minutes and even seconds being common. In this way, email should be considered more of a messaging system rather than a mail system.
- It comes at low personal cost. Electronic mail is usually quite informal in content and format, and fairly easy to compose and read.
- It is a means for "principals" to communicate with each other. Intermediaries (such as secretaries) do not usually compose electronic mail as they would with written mail.
- It crosses the boundaries between organizational and political hierarchies. Partially due to the lack of formality in message contents, the absence of a person's organizational title in a message address and the low personal cost of sending a message, junior people are more aggressive at communicating with senior people. This is in marked contrast to the way other communication devices such as telephones, face–to–face contacts, and surface mail are used.

The points below indicate several benefits of email to researchers and developers in a distributed community. This list can be extended to almost all aspects of communication.

1. *Establishing contact with other researchers.* One person may initiate contact with another person, even though no previous direct communication had taken place. The contact may be motivate by several reasons, such as when a reader of an article wishes to pass comments on to an author, or when one is aware of another's common interests.

2. *Supporting communication in all stages of a research project.* Virtually all stages of research demands communication between collaborators, from project initiation, design, management, analysis and dissemination.

3. *Acquiring information.* Researchers may request information directly from the most appropriate person. For example, acquisition for a rare resource may be preceded by a request asking about its availability, cost, and so on.

4. *Dissemination of research.* Researchers may post early draft versions of papers to colleagues and collect back comments. Similarly, developers may send out product information, press releases, and announcements to specific parties.

5. *Dissemination of information to a known group of people.* Since mail also supports one–to–many communication, people may post items of interest to members of a distribution list.

6. *Document transfer.* Email is not restricted to personal messages. It may also be employed as a simple way to transfer documents or other digital resources.

7. *Social contacts.* As mentioned in the previous section, continual social contact is the glue that helps collaborations persist. Email is a reasonable way for distributed people to exchange social notes, keep people updated, and so on.

The section begins with a background of the basic electronic mail service. Users compose and read textual mail through any one of several user interfaces, while the mail itself is conveyed through a separate sub–system that routes the message over various networks until the desired destination is reached. I detail why such a simple service has proven effective, and what general hurdles are still in place. The section continues with a description of several variations on email that go beyond simple unstructured text messages. Through *multi–media mail*, people can exchange almost any digital information over the wire; graphics, animations, voice, and so on. Through *semi–structured messages*, people can compose typed messages through form filling, and have the system filter incoming mail through its knowledge of the message structure. Finally, *semi–formal messages* restrict the inter–personal exchange of messages to a pre–defined protocol.

4.2 BASIC ELECTRONIC MAIL

4.2.1 Overview

Through a basic electronic mail service, people can compose a text message with a text editor and send that message to one or more other people. This is the service that most of us are familiar with. The software (and the user interface) usually distinguishes a message into a header and a body. The header normally contains routing information such as who sent it, who will receive it, the address, time of mailing, whether the note has been carbon copied, forwarded, replied to, and so on. A subject line summarizing the note is usually included in the header. The body, on the other hand, contains the actual text message being mailed and is usually just passed through as is to the receiver.

Electronic mail has proven a surprisingly effective means for asynchronous communication. Its advantage over surface mail is its speed. Despite even large distances, turn–around time can be seconds or minutes. Because of this swiftness, email is closer to a "messaging" system, and often augments the roles more conventionally assumed by inter–office memos and even telephone calls. Email is especially good at getting around the time–wasting "telephone tag" that haunts attempts at real–time conversation. Sproull and Kiesler summarize this view.

> Most organizational analyses of electronic mail view it simply as an information accelerator, a tool that reduces the amount of time it takes for people to get information they otherwise would have received more slowly
>
> (Sproull and Kiesler, 1986).

There are several other points worth making.

- Email correspondence is typically informal, especially when compared to surface mail. Short, misspelled and grammatically incorrect email messages are common.

- The availability of distribution lists means that a user can trivially send the same message to many people.

Studies of electronic mail within organizations have verified its value. For example, consider the following case study where two slightly different groups were observed in a natural office setting (Eveland and Bikson, 1988). Both groups comprised two types of people: normal employees who worked in the office, and ex–employees (retirees) who were usually at home. While members within each group could communicate between each other in conventional ways, one group also had basic email facilities. They found that retirees using email had a much higher rate of communication with other members of the group when contrasted with those who did not have email (the control group). Also, communication did not cluster as much when email was used; the boundary between cliques was not as sharp. In this case, email was successful in keeping retirees and their valuable expertise involved with the office. In contrast, the office eventually lost contact with retirees in the control group. The overall interpretation is that email significantly and directly affects the outcomes and the process of cooperative work.

4.2.2 Limitations and unknowns

Electronic mail use is not without its limitations. Several technical problems are described below (see (Pliskin, 1989) for added detail).

- There is no good general directory service for querying a person's address, and addresses are often unstable.

- Addresses are rarely independent of a network or particular mail routing service. Even when an address is found, a person may not have access to the transport network specified by the address.

- Email is often only as reliable as the nodes in the network. When (say) a recipient's computer has been down for even a moderate amount of time, the mail may be returned to the sender.

- Interfaces to email systems range from good to terrible. In some interfaces, sending a quick note can be a tedious process, and sorting and filtering incoming mail a difficult chore.

Socio–emotive problems and issues mostly stem from an individual's acceptance of email, having a critical mass of people using email, and relying on people to log onto the system on a regular basis.

- The use or non–use of email by key people in the organization sends a clear signal to subordinates on how they should use it.

- Some people do not like electronic mail. They may have prejudices against working on a computer in general, or they consider the overhead of learning an operating system, a mail sub–system, and an editor not worth the bother.

- If people do not have ready access to a terminal, the time to receive a message will be longer. As a corollary, email is used most heavily by people with a terminal on their desk (Sproull and Kiesler, 1986).

- As electronic mail is not delegated as easily as paper mail (ie to a secretary), senior managers may be deterred from using it.

- Junk mail (even well–informed junk mail) is beginning to inundate email systems. Unlike physical junk mail, the recipient often has to read part of the note before recognizing it as junk; there are no packaging cues.

- Email usually lacks social context cues. Little is sent about an individual, such as one's position within the organization, job title, age, and appearance. Email also lacks cues that indicates the situation in which the mail was created and how it should be received (eg corporate memos, announcements) (Sproull and Kiesler, 1986).

Sproull and Kiesler (1986) noticed several interesting social effects from their study of email use.

- People preferred to send email messages to superiors than to subordinates.

- People preferred email for sending bad news.

- People behaved irresponsibly more often with email than with face–to–face conversations.

Electronic mail is more than just a communication system, for it also demands a variety of time and task management activities from its users. Mackay studied email users and rated them in several categories, each with quite different habits and objectives (Mackay, 1988). Prioritizers concentrate on the problem of managing incoming messages. Archivers concentrate on archiving information for subsequent use, and delegators delegate mail by passing it on to others. Mackay's study indicates that mail use is strikingly diverse, and that designers of email interfaces should recognize this diversity by creating systems that provide flexibility over a wide range of users.

4.3 MULTI–MEDIA MAIL

4.3.1 Overview

Basic email achieves a high degree of success because its messages presume a simple yet ubiquitous standard; ASCII text. Since most computers are well equipped to handle and manipulate ASCII, the person has a rich and usually familiar environment in which to compose, edit, and save mail. While email does not demand this standard (any binary stream can be sent in the body of a message) the interface is rarely set up to interpret the stream as anything but text. It is up to the recipient to recognize unusual message bodies and to view it through the appropriate external software.

Of course, other non–textual standards exist. Facsimile machines, for example, scan and transmit documents as images. The advantage is that any black and white document can be

sent as a facsimile. The disadvantage is that the units of the message (characters, drawing components) are no longer available for manipulation by the receiver.

Multi–media mail, on the other hand, is built on the premise that mail is more than just text. Mail systems should allow many different types of media objects (or combination thereof) to be composed, sent and received as messages (Borenstein, Everhart, Rosenberg et al., 1988). A multi–media message, for example, may comprise a mixture of plain text, a drawing, a voice annotation, a video. The receiver not only sees the message in its proper form, but also has the ability to manipulate the different media objects through integrated on–line tools.

Perhaps the best example of multi–media mail is the Carnegie–Mellon University *Andrew Messaging System* (Borenstein et al., 1988; Borenstein and Thyberg, 1988). Its important points are that it is a combined mail/bulletin board facility, and that it is multi–media. One can, for example:

- transmit line drawings, rasters, animations, and spreadsheets;
- ask for responses to a message via mail that asks its reader to select from a list of choices;

Interestingly enough, the multi–media aspects of the Andrew Messaging System did not catch on immediately with its user population. Instead, Borenstein and Thyberg (1988) describe a gradual "raising of consciousness" over several months. The major breakthrough came when a user mailed a bug report by including a screen snapshot of the clearly visible bug instead of a verbal description. Through the examples of how Andrew was actually used over a long period of time at Carnegie Mellon University, Borenstein and Thyberg (1988) leave a positive impression of what advanced multi–media technology can offer.

Diamond is another example of multi–media messaging system that allows users to create, edit, transmit, and manage multimedia documents (Thomas, Forsdick, Crowley et al., 1985). The emphasis on documents instead of messages reflects the designers' belief that messaging is only one part of the document handling process when using electronic mail. As with an Andrew message, a single *Diamond* document might include various media forms: structured text and object–based graphics, bitmaps, voice, and even active spreadsheets. *Diamond* also maintains its documents in a distributed database. As a consequence, users do not store the actual documents, but rather a *citation* (ie a pointer) to it. This approach means that multiple copies of documents need not be stored, a considerable savings in storage requirements given the potential size of multimedia documents.

Digitized speech and voice–messaging systems are also becoming readily available. IBM has had the speech filing system (commercialized as the Audio Distribution System) available since 1981 (Gould and Boies, 1983; Gould and Boies, 1984). Technical workstations, such as Sun Microsystems and the NeXT machine, have or are beginning to include voice mail within the standard mail system.

4.3.2 Limitations and unknowns

Multi–media messaging is fraught with severe difficulties and unknowns. Several problems are listed below.

- Aside from ASCII text and facsimiles, there is no agreed–upon standard for sending multi–media.

- Multi–media often requires high bandwidth, which is expensive in terms of both transmission and storage.

- Few workstations can handle "esoteric" media forms, such as sound, video, and color. Media requiring raster capabilities are similarly affected. Although bitmapped displays are common, they are still heavily outnumbered by character–based displays.

- Multi–media will require senders (and receivers) to know how to use potentially complex multi–media editors.

4.4 SEMI-STRUCTURED AND SEMI-FORMAL MAIL

Basic electronic mail has very little explicit structure beyond separating the address field from the body of the message. Yet many messages (and sequences of messages) implicitly contain some type of structure. Personal notes, memos, standard forms, meeting announcements, and equipment requests are some examples of different semi–structured messages, each containing a different structure. Proposals/rebuttals, requests/confirmations, and questions/answers are examples of semi–formal messages that follow different between–message protocols. If the structure and protocol is made explicit and knowable to the computer, there may be potential for supplying better email systems.

4.4.1 Semi-structured messages

Malone et al defines a semi–structured message as

> "messages of identifiable types, with each type containing a known set of fields, but with some of the fields containing unstructured text or other information
> (Malone, Grant, Lai et al., 1987).

He gives as an example a seminar message, where structured fields would include the seminar time, place, speaker, and topic. Some of these fields may be typed to allow only permissible values. For example, the place slot may only allow a room number from some enumerated set of possible rooms. The seminar abstract would be an example of an unstructured text field. Another example of a semi–structured message is a "request for action" that includes a deadline field.

Several advantages occur when structure is made explicit (Lai and Malone, 1988; Malone et al., 1987; Malone, Grant and Turbak, 1986).

- Computers can process structured fields more easily than free text fields. For example, unread notices about seminars that have already taken place may be removed automatically, and messages passing some specified criteria may be placed in different folders.

- The system may fill in default values or offer a list of choices to the composer of the message. For example, the "place" slot in the seminar message may default to the most commonly used meeting room, or may have a pop–up menu attached to the field allowing the person to select from a choice of rooms.

- Semi–structured mail reflects how people naturally generate, process and categorize routine messages.

- Templates encourage a consistent (and perhaps organization–wide) approach to sending particular message types. For example, if a request for action template includes a deadline slot, the composer will be more likely to use it than if he was composing a completely unstructured message.

One example of a semi–structured messaging system is Object Lens (Lai and Malone, 1988), a second–generation version of Malone's fairly well–known Information Lens (Malone et al., 1986). Object Lens contains two fundamental ideas. First, passive information can be represented as *semi–structured objects*, where each object is defined as part of an inheritance hierarchy. Consider the added structure as one specializes through the following hierarchy branch:

"Thing – Message – Action Request – Meeting Proposal"

Whereas a completely unstructured message may be a primitive mail form, the meeting proposal includes specific information about time, place, decision requests, and so on. By defining and modifying templates for these objects, users can represent and interact with many different kinds of information.

In the second idea, Object Lens profits from the added structure through *semi–autonomous agents* and *active rules* for processing information. When creating these agents, users specify rules for automatically processing information in different situations. A rule triggered by incoming within–organization mail may, for example, sort messages from superiors into an "Urgent" folder, and discard those messages arriving from an uninteresting distribution list. With these two ideas, Object Lens integrates object–oriented databases, hypertext, and electronic messaging with intelligent routing.

4.4.2 Semi–formal messages

Structure not only exists within a message, but also between messages. A semi–formal message is:

> a message exchange following a given protocol, where the message type and between–message sequences are restricted by the protocol.

A well–known and well–documented semi–formal messaging system is Winograd's Coordinator (Winograd, 1988a; Winograd, 1988b; Winograd and Flores, 1986). The protocol used by the Coordinator is based upon speech act theory that asserts

> "every utterance falls into a small number of categories and an act within one of these categories does not occur at random. If a person requests another person to perform a task then there is only a limited set of valid responses to this request.
> (Bair and Gale, 1988)

There are five things you can do with an utterance (Winograd, 1988a).

Assertive. Commit the speaker to the truth of what is being said.

Directive. Get the hearer to do something.

Commissive. Commit the speaker to some course of action.

Declaration. Declare some correspondence between the propositional contents of the speech act and reality.

Expressive. Express a psychological state about some state of affairs.

Winograd goes on to describe a structure showing how speech acts relate to each other in a larger "conversation for action" scheme that represents how one person requests another person to do something. Drawn as a state transition network, the structure defines the possible interactions between utterances, where utterances are typed by request, promise, counter–offer, declination, cancellation, or report.

Winograd's Coordinator implements the "conversation for action" protocol mentioned above on an IBM PC–based electronic mail system. Because the system is aware of the protocol, it can help a user keep track of where things stand (Winograd, 1988b). The summary screen displays the new conversations that have started, the on–going conversations one is involved in, what commitments have been made, what responses are expected by other people, and matters that have been completed. Depending on where one is within the protocol, a list of choices of allowed message types is offered to the person who wishes to respond to a message. Winograd summarizes the appeal of the Coordinator.

The Coordinator has no magic to coerce people to come through with what they promise, but it provides a straight–forward structure in which they can review the status of their commitments, alter commitments they are no longer in a position to fulfil, anticipate coming breakdowns, make new commitments to take care of breakdowns and opportunities appearing in their conversations, and generally be clear (with themselves and others) about the state of their work.

(Winograd, 1988b)

4.4.3 Limitations and unknowns

Although semi–formal messaging has not been around for very long, early usage studies are encouraging.

Semi–formal communication as exemplified by the Coordinator was studied by Bair and Gale (1988). The field study includes six companies that had installed the Coordinator on seven to 900 nodes. Their anecdotal findings are summarized below.

- Some (but not all) people, especially managers, found the Coordinator's ability to track commitments to be a powerful management tool, particularly when subordinates were geographically widespread.
- Structured conversation reduced junk mail.
- The Coordinator encouraged reflection and thoughtfulness in management and thinking.
- The Coordinator was an aid to planning through its listing of outstanding commitments.
- The Coordinator enabled formal negotiations to proceed smoothly, even when participants had no face–to–face relationship.
- A minority of companies/people reported commitment tracking to be socially unacceptable.

- Negative reactions to the Coordinator usually came from people used to a different email system and interface or to those who did not understand the theory and paradigm behind the system.

- A critical mass of users was necessary before the Coordinator became beneficial.

- Since not all "speech acts" occurred over the Coordinator (eg face–to–face and telephone meetings are other communication channels used), the tracking process was sometimes incomplete.

- The strongest predictor of success was the use of the Coordinator by upper management (Bair and Gale, 1988)

Although the comments above seem mostly positive, the Coordinator also has the dubious honor of being known as an example of "fascist software" that dictates that users do things a certain way. Since the Coordinator fundamentally alters the way people work, Bair and Gale suggest that its introduction to an organization must be accompanied by workshops that allows employees to learn the ideas behind conversation for action and to buy into the approach.

4.5 NETWORK REQUIREMENTS

Most basic electronic mail requirements are well–handled by a fairly low–bandwidth network. This is usually due to the small size of the mail packet, and because the time to transmit is not critical due to the asynchronous aspect of the communication. As mail becomes more sophisticated, as with the addition of structure and protocol into the message, the average message size will rise with the increase in transmitted information.

More bandwidth is required as mail contents become increasingly sophisticated. Taking document preparation systems as a model, we have seen a progression from simple ASCII text to quite sophisticated typeset documents including complex figures and bit–mapped images. There is every reason to believe that users will demand electronic mail to follow the same course.

The real increase in bandwidth requirements will occur as true multi–media mail becomes available. Animated graphics, digitized speech and full video all require high bandwidth. Since these messages have high storage needs, only a pointer to the multi–media message may be transmitted for notification purposes. The actual contents may be sent on demand, and in real time when the reader actually refers to it.

5. BULLETIN BOARDS AND ASYNCHRONOUS CONFERENCING

5.1 INTRODUCTION

Bulletin boards emphasise allowing people to post messages to the on–line community, while asynchronous conferencing enables one to create, join, and actively participate in an on–going discussion of some topic of interest.

Electronic mail is used most often for one–to–one communication, and to a limited extent for one–to–many communication. Mail is sent directly to a recipient through their address. Bulletin boards, on the other hand, act as repositories for one–to–many messages, usually organized by subject matter. A person posts a message to a subject folder that is forwarded to the on–line community at large. Unlike electronic mail, the recipients are not addressed directly. Instead, anyone interested in that subject may choose to read it.

Asynchronous computer conferencing differs in emphasis from bulletin boards by promoting many–to–many communications[2]. Conferencing systems emphasise allowing people to create, join, and actively participate in on–going asynchronous discussions in topics of common interest (Baecker and Buxton, 1987). Messages and commentaries are organized by topic, with individual postings usually sorted in chronological order.

Existing systems often support both asynchronous conferencing and bulletin boards, with little differentiation between the two. Consequently, this discussion will not distinguish between the two any further . Electronic mail is normally available as well (usually as a sub–system), so community members can communicate directly as a follow–on of a posting.

Asynchronous conferencing/bulletin boards can provide considerable benefit to a research and development community. A few salient points are listed below.

1. It is a considerably faster method of distributing information than paper.
2. It fosters discussion and creativity between participants.
3. It keeps everyone well informed about current activities and topics of interest.
4. People can join and actively participate with discussions even though they may not know anyone in the group personally.
5. People may respond directly to a person's posting via integrated electronic mail rather than to the group.
6. People may broadcast information requests over the bulletin board. For example "Does anyone know if there are any published articles on the topic of..." or "Has anyone seen this problem and found out how to solve it"?

---

2 – Computer conferencing – is the popular term for asynchronous conferencing. However, it is a poor term, for the name itself does not differentiate between the ways that computers can support conferencing (eg face to face, real–time, and so on).

7. People may post abstracts of work in progress, internal reports or publications in press, allowing interested parties to request the full paper.

This section will describe the standard functions that are usually provided by asynchronous conferencing systems, raising several issues in current and future designs. Limitations and unknowns are noted, and network requirements discussed.

5.2 FUNCTIONS AND ISSUES

Current systems offer a wide range of functionality to the community. Some standard features are listed below.

Organization. Articles are organized into groups representing topics of interests or discussions.

Search facilities. People can find items of interest by searching for items in article headers (dates, people and topic names, words, topics, keyword lists), and for strings appearing within a posting.

Sorting capabilities. People can sort the message headings within a particular conference in several ways. Some examples are chronological order, alphabetic order of authors, alphabetic order of subject header.

Connections through electronic mail. People can respond directly to the author of a posting without leaving the conferencing system.

Access control. Access control may be provided so that partail arerated by one person, or have a restricted membership, or control who the active participants are (ie the ones who can post to the group) separately from who is allowed to read the postings.

Information management. To remove the clutter found in large conferencing systems, people can un–subscribe to particular groups and delete unwanted postings, removing them from view. The system automatically archives those postings that are read but not deleted.

An important architectural issue with conferencing systems is where the information actually resides and how people access it. With a *centralized* architecture, postings and the group structure is maintained on one computer. People access the system through remote dial–up. With a *distributed* architecture, postings are broadcast to all subscribers. The consequences are two–fold. First, centralized architectures are often complete services that tightly couple the availability of the postings with the user interface employed to present the messages to the user. Distributed architectures, on the other hand, often have several interfaces that take advantage of the postings being distributed. Second, it is more difficult to capture the relationships of postings within a distributed architecture due to synchronization problems. For example, one site may receive a reply to a message before receiving the original message. This issue is particularly relevant when the inter–relationship between messages is important. An example is a group that wishes to separate a discussion's central theme from side–arguments.

The main advantages of asynchronous conferencing is its ability to bring together a large number of geographically separated people with common interest over an extended period of time (Jackson, 1989). Because it is asynchronous, people can contribute to the "conversation" at their leisure. Because it is text based, old material can be reviewed as a structured record of messages (Jackson, 1989).
The icular groups are moderated by one person, or have a restricted membership, or control

who the active participants are (ie the ones who can post to the group) separately from who is allowed to read the postings.

> *Information management*. To remove the clutter found in large conferencing systems, people can un–subscribe to particular groups and delete unwanted postings, removing them from view. The system automatically archives those postings that are read but not deleted.

An important architectural issue with conferencing systems is where the information actually resides and how people access it. With a *centralized* architecture, postings and the group structure is maintained on one computer. People access the system through remote dial–up. With a *distributed* architecture, postings are broadcast to all subscribers. The consequences are two–fold. First, centralized architectures are often complete services that tightly couple the availability of the postings with the user interface employed to present the messages to the user. Distributed architectures, on the other hand, often have several interfaces that take advantage of the postings being distributed. Second, it is more difficult to capture the relationships of postings within a distributed architecture due to synchronization problems. For example, one site may receive a reply to a message before receiving the original message. This issue is particularly relevant when the inter–relationship between messages is important. An example is a group that wishes to separate a discussion's central theme from side–arguments.

The main advantages of asynchronous conferencing is its ability to bring together a large number of geographically separated people with common interest over an extended period of time (Jackson, 1989). Because it is asynchronous, people can contribute to the "conversation" at their leisure. Because it is text based, old material can be reviewed as a structured record of messages (Jackson, 1989).

The asynchronous aspect also provides for possibilities of non–sequentiality within the conference. For example, consider the temporal sequence, shown in the figure below, where person A initiates some topic, B and C reply directly to A and D alters the topic. A new person E now reading this chain of messages may reply to A's original posting, comment on C's reply, and then discuss D's topic change.

The figure also shows a logical re–construction of this simple conversation as a set of topics and responses . As can be seen, the resulting conversational pattern is quite a bit different from the temporal one.

The difference between the conversational and temporal organization of computer conferences highlights several needs of the people using the system. As Jackson writes,

> "conversations can grow and diverge into complex non–linear structures, so participants need ways to follow branching discussion paths, and to ascertain the context of any particular item in a conversation."
>
> (Jackson, 1989)

If the system structures messages by their temporal arrival (as most commercial systems do), it is increasingly difficult for a reader to:

- follow conversational paths;
- understand the context of replies;
- recognize, pursue and/or prune conversations on divergent conversations;
- get a feel for the gestalt of the conversation.

As a result, many conversations have a fragmented feel about them. Some commercially available conferencing systems pay some service to the conversational ordering of messages. For example, *Participate* is a text–based service that allows each note in a conference to become the root of another conference (Stevens, 1986). More interesting is *Banyan*, a research prototype of an asynchronous conferencing system (Jackson, 1989). Banyan represents conversations and messages as nodes within a graphical hypertext system. Links within the nodes explicitly indicate the logical ordering of the conversation. Not only can one see the context of a reply (backward chaining), but one can also inspect the original message and follow replies from that point on (forward chaining). Banyan also allows people to create views on temporary collections of messages resulting from searches or queries. People may hide uninteresting topics. Banyon is, in fact, similar to Object Lens, for they are both based upon an underlying hypertext model of message storing (Lai and Malone, 1988).

The previously discussed *Andrew Messaging System* is worth revisiting for the novel and interesting features in its bulletin board system (Borenstein and Thyberg, 1988).

- As with email, all *Andrew* messages may comprise one or more multi–media objects.
- A person is allowed to create a privately editable, publicly readable electronic "magazine". Borenstein and Thyborg give an example of a user with a strong interest in music compiling key articles from several music–related newsgroups into a single magazine (Borenstein and Thyberg, 1988).
- Protection mechanisms permit creation of bulletin boards with a varicontai Examples include:
  - group oriented: readable and writable by only a designated group of people;
  - administrative and advisory: writable by all, readable by few;
  - secretary support of email: postable by a few email recipients, readable by one.

- Some messages can, when read, invoke a message–specific interaction with the user. These have been used for system–facilitated voting, and for return–receipt confirmations from the reader.

Andrew also shares several common features with Object Lens.

- Users may specify how incoming mail is to be filtered (eg deleted, re–directed, placed in special folders, and so on).
- Message templates are available for replies.

The widespread use and the extremely large content size of even primitive asynchronous conferencing systems indicates how important they are. Unlike all other forms of computer–supported cooperative work mentioned so for, conferencing systems travel well beyond the small group by bringing together an extended community of people with common interests.

Semi–formal messaging, described in the last section, can be extended as semi–formal conferencing. The best example is gIBIS, a system that captures early design deliberations on large complex problems (Conklin, 1988). It is based upon the Issue Based Information Design (IBIS) methodology that views design as a rhetorical process, with a set of issues that can be generalized, specialized, responded to, questioned, argued and so on. As with Object Lens, gIBIS is based to a large part on semi–structured messages. Through its well–designed interface, participants propose and respond to issues in structured ways that eliminate unconstructive moves such as name–calling and argument by repetition. A key aspect of gIBIS is that all messages are captured and structured within a group–accessible hypertext database. As with most hypertext systems, the relationship between text fragments–the particular arguments made–are illustrated and manipulated graphically by the user.

5.3 LIMITATIONS AND UNKNOWNS

There are several issues that plague conferencing systems. The first is the user interface. A conferencing system will only work if many people can access it. Because of the wide variety of display hardware and the lack of standards, this has restricted most systems to glass–teletype ascii–based interfaces, a poor medium for representing the rich inter–connections between postings. Fortunately, the advent of X window standard will likely see new conferencing interfaces window based upon graphical systems.

The second issue is the difficulty of handling and filtering the tremendous amount of postings that become available in some conferencing system. In the Unix–based Usenet, for example, not only are hundreds of news groups available, but each group may have tens to hundreds of postings made to it daily or weekly. To make matters worse, un–moderated groups are often subject to "junk" postings (often called flames) that add little to the group's discussion. New subscribers are often over–whelmed, while even faithful readers are often hard–pressed to keep up to date. The result is that most users subscribe to just a few selected groups. When they fall behind in their reading, they often prefer to skip over old material so that they can keep up with the new. Relevant postings may be lost as a consequence. The tentative solution is to incorporate intelligent filtering into the system, so that un–interesting articles are pruned and interesting ones made prominent. While simple filters exist in some systems, intelligent filters must wait until ideas similar to those in Object Lens (Lai and Malone, 1988) are commercialized.

The third issue arises from semi–formal conferencing. Conklin and Begeman evaluated and criticized the gIBIS system work upon preliminary observations of its use (Conklin, 1987). On the positive side, users found the IBIS formalism (and the gIBIS tool) to be a powerful

method for research thinking and design deliberation, especially for detecting incompleteness and inconsistency in thinking and rhetoric that added little to the argumentation. Users also found the semi–structured aspects of gIBIS aided message composition, search, and comprehension.

Several negative aspects arose with the IBIS formalism as implemented by gIBIS (Conklin, 1987). IBIS sometimes proved inadequate for capturing all aspects of the problem being considered, and it did not deal well with the vague, contradictory and incomplete thoughts usually expressed in the early stages of a problem. Meta–level communications is another issue. While gIBIS supports the substantive part of the work, it does not support annotative comments about the work, nor procedural comments about the communications itself. In spite of these criticisms, we should remember that gIBIS is an early prototype and has much room for improvement.

5.4 NETWORK REQUIREMENTS

The network requirements for conferencing are range from low to extreme. As with electronic mail, low–bandwidth systems can handle even large numbers of postings of limited size as long as immediate delivery is not necessary.

Multi–media postings require significant bandwidth. Consider a centralized architecture, where the user must log on to the remotely–located system. Since postings are transmitted only when they are selected by the user from the subject heading, they must be sent in real–time. The network must be able to transmit the posting with little delay, for excessive waiting would deter the reader from scanning all but the most interesting articles. With a distributed architecture, articles must be broadcast to all subscribed sites as they are posted. Given a large number of sites and a high volume of postings, considerable demands may be made of the network.

6. ACCESS, MONITORING AND OPERATION OF EQUIPMENT

6.1 INTRODUCTION

Real–time monitoring and operation of distributed equipment is now commonplace within industrial settings, usually through real–time systems. Systems include supervisory control and data acquisition, tele–operation, and tele–robotics. Availability remote computational power is discussed here as well. All are highly relevant to researchers and developers who require access to remote equipment.

1. Equipment is a shared resource available to non–collaborators. This often occurs when high equipment costs must be amortized over a large number of people (for example, a super–computer or a radio telescope) and when the equipment is rare (perhaps due to it being a novel technology).

2. Intensive computation may be distributed in parallel to idle computers in the network.

3. Some instrumentation is inaccessible due to its location (underground, space, deep sea) and hazardous environments (nuclear power plants).

4. Members of a distributed team monitor and control equipment located at one collaborator's site. One example is simultaneous engineering.

5. A scientific community observes an experiment in progress (one example recently occurred when astronomers from around the world gained limited access to the Voyageur II's Uranus data as it was transmitted to Earth). This is valuable to obtain early community reaction to methodology and analysis, or when the data is so complex that good analysis requires study of data by multiple independent experts.

6. Remote observation and operation is an extension of "tele–data" by including observation and control of physical devices as well as software. Not only can people see what the other is referring too, but have greater opportunity to interact through these devices.

7. Hands–on distance education and support becomes possible. This is especially important for keeping the community up to date on the latest technology, and for the on–going support requirements that occur when one consults an expert.

8. Maintenance of specialized equipment may occur through specialists remotely monitoring and servicing the system. This is important when the users of the equipment are not expected to have the expertise to maintain it.

9. Programmable robotic systems allow flexible re–tooling and re–configuration of distant systems so that it can handle a variety of situations.

6.2 COMPUTATIONAL ACCESS

People often need to access and work on remote computers. They may be working through a terminal with no compute power, their own machine may be too slow, or necessary hardware and information is only available on the remote system. This sub–section reviews the several ways that people access remote computers.

ASCII–based interfaces. The simplest way to access a remote computer is through a modem or network. Since most personal computers and workstations can emulate ASCII terminals, the person simply logs in as normal once the connection has been made. Typical connection speeds through modems are 1200 to 2400 baud, with modem speeds up to 9600 baud now becoming available. The later is fairly close to the line speeds of terminals hooked directly into the computer. Networked workstations can, of course, offer greater speeds for remote logins.

Networked window systems. ASCII–based systems are reaching the end of their lifetime. Window systems running graphical applications are now becoming the standard platform for human/computer interfaces. Until quite recently, these window systems required a person to be logged in directly to a bit–mapped workstation; they would not run over a network. However modern networked–based window systems, such as Sun Microsystem's NeWS and the Open Software Foundations X Windows, have removed this limitation. A person may have several windows on display, with each perhaps running an application being executed on a completely different computer. In essence, the communication protocol for the user interface has shifted away from low–bandwidth ASCII to a higher–bandwidth graphics standard. Consequently, the connect speeds once adequate for ASCII terminal emulation are painfully slow for emulating window systems.

Distributed computation. While workstations are becoming increasingly powerful, the demands being made of them are also rising in complexity. Distributed computing is one method of overcoming this problem. At its simplest, a user connects to a powerful computer (say a supercomputer), using a less–powerful computer (such as a workstation) to handle the more mundane computational tasks. But increasing attention is now being paid to *parallel computation,* where many computers are linked together to solve a problem that has been functionally decomposed into portions that can be executed in parallel. A low–cost approach to parallelism is to have one's machine search for idle or under–utilized computers on the network and to distribute the decomposed problem accordingly. Another approach is to use the dedicated parallel computers now coming onto the market. In spite of these advances, the bottleneck to parallel computation is the network speed. The smaller the parallel units to be distributed, the greater the communication demands.

6.3 TECHNOLOGIES FOR MONITORING AND OPERATING REMOTE EQUIPMENT

There are a variety of technologies relevant to monitoring and controlling remote equipment. Many are now in use in specialized settings (ie heavy industry, research laboratories, space), but their use should become more available to the community as costs and availability decrease.

At the heart of control and operation of remote equipment is the area of *real–time systems.* As its name indicates, these systems demand response in real time. Significant time delays for relaying information to and from a controller and the equipment is not acceptable. Several different flavours of real–time systems exists.

SCADA–Supervisory Control and Data Acquisition systems: are the slow end of real–time systems. They commonly operate on low speed networks (such as telephone lines with

modems). SCADA systems are responsible for monitoring and occasionally over–riding the (usually) automatic operation of remote sites–the latter is called supervisory control. In typical use, *remote terminal units* send data points from tens or thousands of widely–distributed and heavily instrumented sites to a centralized database, where they are checked against pre–defined constraints. Alarm conditions are raised on an operator's console when allowable bounds are passed. Operators then view the collected data points through mimic diagrams displayed on a screen, apply their knowledge to control the remote system from their consoles, and see the resulting changes almost immediately. Examples of SCADA system use range from operation of elevators within a building, city–wide electrical utilities, and international gas pipeline system control.

As applications become less tolerant of time delays and the need for control becomes more rigorous, we move closer to the heart of real–time systems where SCADA techniques are no longer acceptable. We will differentiate "true" real–time systems from SCADA by its need for continual rather than supervisory control. Time frames also differ. Network transmission of seconds or even minutes is often acceptable for SCADA; real–time systems often require tolerance levels of milli–seconds. Real–time systems are used for a variety of purposes, from close monitoring and control of critical operations within a plant (such as a nuclear power plant) to tele–operations of robots.

Master–slave tele–operation: covers any manipulation system that is remotely controlled by a human operator in response to sensory information transmitted from the workplace (Scott, 1984). One example is direct human control of a robot's actions in real–time, usually by moving a device that mimics the motion of the actual robot. These systems are required when the need for human involvement is high, the working environment is remotely located or hazardous, and when the size/strength/manipulation differential between human and machine is great. Because the robot is not autonomous, a high degree of operator feedback is needed.

Tele–robotics: covers computer–controlled tele–operation of robots, and is one of the most promising application areas to support distributed research and development. Because robots are generic in operation (such as robot arms), they can be programmed to perform a variety of specific tasks related to specific needs. Tele–robotics are already being seen in *flexible manufacturing*, a robot–based system for the manufacture of diverse products.

There are two basic methods of tele–robotic operation.

1. A computer sends high–level instructions or downloads programs to a remote robot, so that many low–level actions are handled by a controller local to the robot. In this situation, the robot is considered semi–autonomous. It can react to immediate situations (such as dealing with obstacles) without referring back to the distant computer.

2. The centralized computer directly controls the remote robot. This may be necessary when there is not enough compute power available to the robot, or when part of the computer's task is to synchronize the robot's actions with other machines it may be controlling (such as other robots or an assembly line).

The semi–autonomous robot is the goal of much robotics research, for the greater the on–board intelligence of the robot, the less bandwidth and rapid reaction time required by the distant controller. This is especially critical in space robotics, where speed of light constraints introduce unacceptable response for distant control. Even so, semi–autonomous systems are usually linked to a central system so that humans can monitor and intervene in the operation, and for data collection purposes.

6.4 NETWORK REQUIREMENTS

Computational access requires modest network speeds. Normal telephone speeds are acceptable for ASCII–terminal emulation; and a 1.5 Mbps network is adequate for window–based interfaces. Network requirements for parallel systems can range from low to high, depending upon the size of the computational unit and the amount of communication between these units.

Shifting the emphasis towards monitoring and operation, existing SCADA systems require relatively low–bandwidths. They are well served by current telephone lines and would certainly be well covered by the 1.5 Mbps Network. True real–time systems, especially those controlling robots, have much stricter bandwidth demands and time tolerances. Speed is critical. In master–server tele–operations, the operator may need a high–resolution video image of the remote robot to see what is going on. Additionally, the feedback cycle between action, consequence, and notification cycle must be short: the maximum tolerable feedback delay for effective real–time control of the robot by the human is one tenth of a second (Scott, 1984).

Computer controlled tele–robotics usually require high–bandwidth reflecting the data rates for control and feedback of servers, and perhaps information transmission by any attached vision systems. The preliminary NASA/NBS Standard Reference Model for Tele–robotic Control System Architecture (NASREM) describes timing tolerances of a few milli–seconds for synchronizing activity, especially at the servo level (Albus, McCain and Lumia, 1989). Even semi–autonomous robots may require bandwidth for monitoring (perhaps through video), data collection, transmission of high–level control information, and synchronization with other devices. Other criteria are, of course, system dependant.

Robotics have a good migration path to a nation–wide network, simply because they already rely on software supervision and often use high–speed local area networks for communication.

7. DIGITAL LIBRARIES

7.1 INTRODUCTION

Discoveries in research, technology, and development are exploding. Not only is the amount of information produced growing, but it is becoming increasingly more specialized.

As far back as 1945, Vannevar Bush described the dilemma we are now surmounted with:

> "there is increased evidence that we are being bogged down today as specialization extends. The investigator is staggered by the findings and conclusions of thousands of other workers–conclusions which he cannot find time to grasp, much less to remember, as they appear. Yet specialization becomes increasingly necessary for progress, and the effort to bridge between disciplines is correspondingly superficial.
> (Bush, 1945)

How will information be handled in the next decade? It is clear that traditional technologies, such as paper books and journals, are unable to accommodate well a researcher's need for timely and relevant information. Review and publication processes for conference proceedings and journals normally span the better part of a year, while book production is even longer. Information searching and retrieval is difficult even with one's area of specialty. As an expert moves outside of his specialty, it becomes almost impossible to find, let alone be informed of the key related material.

We require new ways of manipulating information. At the heart of these methods will be the *digital library*–a computational way of storing, relating and retrieving information. In many respects, digital libraries are related to bulletin boards/asynchronous conferencing system in that both supply similar benefits. They are two ends of a spectrum, with bulletin boards addressing the *process* of creating information (drafts, discussions), and digital libraries addressing the *archiving, relating* and *access* of information. A digital library has potential to provide many benefits to research and development.

1. Researchers can see their works published immediately.

2. Published works become available to the community immediately.

3. People may discover and access "grey" publications–technical reports published internally in an institution–through the library. Grey reports are currently difficult to find out about, and even more difficult to acquire.

4. People may review information presented in alternate and more relevant forms. The technology of publication is extended to include multi–media documents, such as video, sound, animation, active programs, and so on.

5. A specialist can search, filter and retrieve relevant information from large databases. Search and retrieval facilities need not be restricted to specialized archival sources (as they are now).

6. Since there is no notion of being "out of print," researchers may obtain historical articles as easily as contemporary ones.

7. Readers can follow cross–references easily (such as bibliographic citations) by retrieving the reference directly.

8. Similarly, readers may review attached commentaries, rebuttals or extensions to articles.

9. A researcher may personalize the information by adding notes and links to relevant material. Additionally, one may gain a new perspective to the information by viewing another person's notes and links.

The list can grow almost indefinitely. The impact of a digital library to society is potentially as beneficial as the introduction of the printing press.

This section will briefly describe the conventional approaches related to digital libraries, and then discuss the hypertext approach. Limitations and unknowns are raised, and the network requirements provided.

7.2 CONVENTIONAL APPROACHES

There are two standard ways of storing information on computers–as hierarchical file systems and within a highly–structured database. As both methods are likely well–understood by the reader of this document, they will not be described. Rather, their use, suitability and influence to the digital library will be discussed through their roles as a digital repository and a bibliographic database.

Digital repositories are file systems maintained by an organization such that files can be accessed and copied by people at another organization. They are currently used by institutions that make information–usually software and documentation–available in the public domain. Access is normally through a network such as the Internet. Consider the *Free Software Foundation*, an organization that maintains a digital repository of Unix–based software. Although people can buy a tape of the library at a nominal cost, one may use the file transfer protocol (ftp) across the Internet to list and transfer part or all of the files from the foundation's machine to one's local machine. Another digital repository is available through the White Sands Missile Range, who maintain ADA–related software and documents. Yet another is a software repository held by Sun Microsystems, who allow people to request and transfer software through an automated electronic mail server. Many universities also have similar repositories to encourage information exchange and cross–fertilization–in essence, a distributed library system.

Of course, digital repositories are extremely limited. They do not have any search capability, and people only find out about a repository's contents by mundane means, such as an announcement on a bulletin board or by remotely accessing and searching around the file hierarchy. Still, they provide a simple mechanism for distributing information at relatively low cost.

Bibliographic databases are databases containing citations of publications. Most are limited in scope. They either reflect a special topic of interest or the holdings of a physical library. Although abstracts of publications are sometimes included in a citation, the full text is rarely available.

A person would use a bibliographic database for several reasons. First, one may be interested in *searching* for all titles related to a particular, quite specific topic. If abstracts are available, one could then scan them for direct relevancy. Second, one may wish to *locate* a publication at a physical library, possibly arranging to receive a hardcopy through an inter–library loan.

As a highly specialized form of a digital library, bibliographic databases have several deficiencies. First, the skills required at choosing appropriate search terms usually requires a librarian to act as an intermediary. Second, they rarely provide the actual publication, and they do not list relevant cross–references and commentaries. Third, grey publications are rarely included. Fourth, they are centralized. Users must usually log onto them remotely over a network or modem via a terminal emulator, with little direct capability for transferring the found information to one's local machine.

7.3 HYPERTEXT AND HYPERMEDIA

Neither files systems nor traditional databases are an adequate backbone for future digital libraries. While files lack the structure needed, databases are too rigid in the structure they provide. An alternative lies somewhere between these extremes–hypertext and hypermedia.

Hypertext extends the notion of a document beyond sequential text by allowing complex interwoven structures to be created and manipulated by linking text fragments. The fundamental idea is simple: links can be added anywhere in the text database which, when followed, will transport the reader to another location. Associating types with links extends the power for enhancing semi–structured access to a document's contents, with instant availability of:

- related information;
- chasing references;
- visualization of the information structure;
- examples illustrating the text;
- rich searching and indexing facilities;
- selective and personal in–depth explorations;
- annotations comprising definitions, footnotes and asides;
- marked trails of the path through the information taken by the reader;
- guided annotated tours defining a "tour leader's" perspective and comments (Trigg, 1988);
- convenient opportunities for activities such as adding personal annotations and place–marking.

Hypertext becomes *hypermedia* when any media form can be used and linked into the document. An author's point may be annotated with an instantly accessible image, sound track, or video clip. Sometimes active sections (ie a running program) may be incorporated into otherwise passive documents to permit user interaction. When this ability is added, hypermedia becomes a rich new metaphor for interacting with computers and file stores.

Fundamental to hypertext and hypermedia is the user interface used to navigate information and traverse links. Most interfaces are based upon *browsers*, where node contents are displayed in a window and links to other nodes are marked through in–text icons. Browsers usually include some type of overview facility that provides a map through the information, usually by displaying the network graphically. Readers typically search for information by

following the links embedded in the text, by searching the network, and by selecting an item directly from the overview map (Conklin, 1987).

There are now a large number of research and commercial hypertext systems. Early systems were envisioned as macro literary systems that focus on the integration of colossal volumes of information (Conklin, 1987). A few are listed below.

1. The first is Bush's un–implemented vision of the *Memex* system (Bush, 1945). His idea was not based upon digital technology; rather, large stores of scientific literature would be linked through microfilm and photocells.

2. In the 1960's and 1970's, Doug Engelbart proposed and built NLS/Augment, a system that emphasizes an environment for "knowledge workers" (Engelbart and English, 1968). Not only was NLS/Augment important as an early implementation, but it still stands out as an environment that integrates hypertext, communication, document processing, idea exchange and information management.

3. Perhaps the most well–known macro–literary system is Project Xanadu, developed and now being implemented by Ted Nelson (Nelson, 1987). Nelson's goal is ambitious, for he sees placing the entire world's literary corpus on line (Conklin, 1987). Rather than concentrate on the user interface, Xanadu investigates the database requirements needed to maintain such a large, heavily linked information store, as well as providing mechanisms for accounting and royalty distribution.

While the early visions of macro–literary systems are suitable to a large, centralized digital library, most current commercial systems are not, for they are designed to accommodate small information structures equivalent to (say) a conventional book. But hypertext will continue to grow over the next decade, and will catch up to the early visions of Bush, Engelbart and Nelson.

7.4 LIMITATIONS AND UNKNOWNS

As with any new form of literary presentation, there are many problems yet to be solved in hypertext and hypermedia. While some are due to current shortcomings of hypertext interfaces, two critical issues arrive from human behaviour: disorientation and cognitive overhead (Conklin, 1987).

Disorientation is the "lost in hyperspace" problem that occurs when people lose track of where they are in the network, or are unclear of how to get to some other place in the network. The difficulty arises because, unlike sequential text, readers are free to wander off in any direction. Technical solutions may overcome some difficulties. Graphical browsers may trace and visually show a person's location in the hypertext document, perhaps distinguishing between mainline progress and side–excursions. Sophisticated query facilities similar to those available in databases can help searchers find and navigate to desired nodes. Filters may remove unnecessary clutter from display, showing only links of interest to the reader. "Meta–structure" may be added to the hypertext system as "guided tours" through selected paths (Trigg, 1988).

Cognitive overhead arises from the additional effort and concentration necessary to decide if links are worth following, and when one maintains several tasks or "trails" through a document at one time (Conklin, 1987). Is a side path worth taking? How does one switch conceptual context back to the main path after chasing one or more side paths? Again,

technical solution may ease the problem. Brief explanations of node contents may be made available from the link itself to help a reader decide if a link is worth pursuing. Graphical browsers can help people track excursions, helping them maintain and return to previous contexts.

7.5 NETWORK REQUIREMENTS

Conventional bibliographic databases have low network requirements due to the small amount of information generated; they now run adequately over modems. In marked contrast, file transfer from digital repositories can involve many megabytes of information. Some repositories have software releases available of 50 megabytes or more! Fortunately file transfer is rarely time critical, and are still well–served by low–bandwidth lines.

In contrast, hypertext can require significantly more bandwidth. If a hypertext system is maintained at a centralized location, then users must access it through the network. Information retrieval is time–critical. If a reader chases a link, the node contents should appear as quickly as possible. Delays will deter readers from chasing paths, contradicting the intention of hypertext. While moderate speed lines may suffice for pure text systems, hypermedia nodes would require the bandwidth needed for rapid transmission of audio, graphics, video, and so on. Additional bandwidth is also needed to drive the graphical user interface employed by most hypertext systems.

8. APPLICATION AREAS

8.1 INTRODUCTION

This report began by separating the primary research and collaboration functions from the collaboration tools required to implement those functions. The previous sections have concentrated on these collaboration tools, mentioning several of the functions as examples that can benefit a distributed research and development community. This section will highlight three application areas to indicate several ways that the collaboration tools can be applied: distance education, group decision support systems, and joint authoring. It closes with a note describing the need for the network to bootstrap collaborations.

8.2 DISTANCE EDUCATION

Certainly one of the most critical issues facing a research and development community is in training and continuing education. Excellence of work demands that researchers continue to gain familiarity with new equipment, learn new trends in their own area of expertise, and to extend their horizon towards other disciplines. Many of the systems already described in this report can provide the backbone of a distance training and education system running over a network.

Current computer–aided instruction (CAI) usually revolves around learner interaction with information in a database. Courses are prepared instructional modules where learners can go through the material at their own discretion and pace. Both distributed libraries and hypermedia support CAI. Packaged courses can be distributed to remote learners through a digital library. Using some of the ideas available in hypermedia, directed courses can be augmented by the relevant scientific readings, allowing learners to pursue a topic as lightly or deeply as desired.

However, good education demands learner interaction with the instructor and other learners. By augmenting CAI with electronic mail and bulletin boards, learners may contact the course organizer for questions and clarification, or even each other for on–going "class" discussions of the material. When courseware is built as a simulation system such as SharedArk and cyberspace, the richness of the visualization and interaction and the possibility of real–time multi–student collaboration extend beyond anything now available.

Equally as exciting is the ability to hold real–time distributed lectures and seminars, where the lecture halls are the sites connected through real–time tele–conferencing. This is particularly advantageous to organizations that cannot afford to bring in prestigious speakers. By linking in to larger organizations or by cost–sharing, researchers at even small institutions can gain access to presentations made by prominent figures. Multi–way interactions between the presenter and audience, and the audience members with each other create a rich environment encouraging questions, comments and discussions. The tools of the classroom are provided through tele–data–the whiteboard, the shared notepad. Even real–time "hands–on" training could be available through tele–operation and tele–robotics.

Distance education must satisfy many criteria to work. The educational resources and technologies must be physically accessible; the material must be topical to the learner; people involved should be collaborative–accessible (Quigley, 1989). The proposed Network is critical for achieving the first and third points.

8.3 GROUP DECISION SUPPORT SYSTEMS

Group decision support systems (GDSS) combines computer–supported communication and decision support technologies to facilitate formulation and solution of unstructured problems by a group of people (DeSanctis and Gallupe, 1987). The goals of GDSS are:

> "to reduce the "process loss" associated with disorganized activity, member domination, social pressure, inhibition of expression, and other difficulties commonly encountered in groups and, at the same time, to increase the efficiency and quality of the resulting group decision.
>
> (Watson, DeSanctis and Poole, 1988)

DeSanctis offers three levels of GDSS systems (DeSanctis and Gallupe, 1987).

> *Level 1:* Remove common communication barriers to facilitate information exchange. These are the areas usually handled by the generic collaboration tools mentioned in the previous sections: tele–conferencing, email, shared workspaces, and so on.
>
> *Level 2:* Provide decision modelling and group decision techniques aimed at reducing uncertainty and noise that occur in the group's decision process. These include planning tools, agenda setting, decision modelling methods (such as decision trees, risk analysis, and forecasting methods).
>
> *Level 3:* Structure the group communication pattern (the rules of discourse) to one appropriate to achieving the goals of the meeting. Structured group methods include the Nominal Group and Delphi techniques, as well as the IBIS and speech act approach described earlier.

A GDSS clearly has an important role in distributed research and development, especially when inefficiencies and uncertainties related to the normal decision process are exaggerated due to distance. The collaboration tools mentioned in previous sections only fulfil the first level of GDSS; they must be extended to include the specialized techniques for structured decision analysis, and at directing and structuring the discussion. Unlike most generic communication systems, GDSS systems intervene in the group's natural decision process to help them come to a higher quality decision.

8.4 JOINT AUTHORING

As with GDSS, distributed joint authoring is a specialization and extension of the previously defined collaboration tools. Joint authoring can take many flavours. At its simplest, document drafts and commentary may be exchanged by electronic mail, or perhaps worked on in real time through tele–data. More sophisticated authoring tools would understand and support directly the nature of the interaction between authors, reviewers, and readers.

Consider Quilt, a tool for collaborative document production (Leland et al., 1988). Unlike other collaborative document systems which support only direct authoring aspects, Quilt emphasizes and supports the communication vital to good collaboration. For example, structured hypermedia links allow people to attach text and voice annotations to the document, specialized as revision suggestions, public comments, and directed messages. The necessary coordination between collaborators is enhanced via activity logging, notification and triggering mechanisms. Access permissions can be set by the author to reflect the varying roles of collaborators (as writers, commenters, reviewers), while user–customizable definitions for such things as document and annotation types make the system both flexible and extensible.

Specialized joint authoring systems can make use of many of the technologies mentioned in previous sections: real–time and asynchronous communication; bulletin boards and asynchronous conferencing for communal reaction; digital libraries acting as a centralized store of the evolving paper that is accessible by many.

8.5 BOOTSTRAPPING COLLABORATIONS

Unlike applications intended for a single users or sites, collaboration tools operating over distance have a bootstrapping problem. Researchers and developers will not use systems until they are available and ubiquitous, while vendors will not build the applications until these people will commit to buying them. In essence, a critical mass of people and equipment is required to bootstrap the process.

High speed networks are currently too expensive and too difficult to obtain for the average researcher. Their use is restrained to organizations with exceptional needs. As long as the status quo remains, collaboration tools are unlikely to make any impact.

The delivery of the Network can overcome this problem in several ways. First, it will act as a testbed for research and development for the network itself, and of future collaboration tools and information technologies. Second, because of its availability to a diverse set of people, members of the community with real needs can be used in pilot studies for further testing and tuning of the collaboration applications. As the applications become closer to commercialization, their introduction to the community becomes a matter of extending the audience already addressed over the network, instead of starting afresh. The act of having the application available before–hand will allow part of the community, perhaps a critical mass, to experience, buy into, and advertise the application.

9. THE COHERENT WORKSPACE

The networked nation is more than a means of overcoming distance barriers, for it has the ability to augment researchers' and developers' knowledge and activities beyond those now found traditionally. Each tool and function mentioned in this report can, by itself, provide added value. Communication vehicles will fit tasks, digital libraries will become knowledge bases; tele–robotics will extend our physical limitations. Yet the real power will arrive when tools and people are integrated together into the *Coherent Workspace*.

The Coherent Workspace occurs when people and their machines have a common knowledge and sense of purpose, with individual and group activities being well–coordinated. Coherent systems have their roots in several early visions. Vanevar Bush saw *Memex* as a way to tie information and people together through a dynamic richly–connected knowledge base (Bush, 1945). Doug Engelbart extended the Memex theme through the *Community Handbook*, a coordinated handling of a very large, complex and continually evolving knowledge base that comprehensively represents the current state of a collaborative endeavour (Engelbart and Lehtman, 1988). These include:

- *formative information* for a collaboration such as purpose, working hypotheses, goals and expectations;
- *project management* such as commitment tracking, status schedules, project and meeting management tools, and standards employed;
- *project content*, such as observations, data, and results;
- *training and education*, such as how–to methods;
- *support and incorporation of communications* to support and track the on–going dialogs about the project and contents of the database.

A coherent workspace includes these notions through four principles: persistence, inter–operability, integration, and coordination.

The *principle of persistence* covers the integrity of an activity across time and distance, across group size, and across changing group membership. Asynchronous bulletin boards are a simple but good example. The topic and on–going discussions of a "meeting" are captured, so that old and new group members can review what has happened and how the thread of the dialog has altered over time. The discussion persists even when discussants drop out, new ones come in, or existing ones change roles from, say, observers to active speakers. But persistent systems could do much more. First, they could capture all of a group's activity, independent of time and place. Second, the system could explicitly maintain the goals of the group, perhaps moderating the activity to keep it on track.

Consider an anecdote of a group of engineers designing brakes for an automobile[3]. The design and final part produced included a burr of metal attached to the assembly to help

[3] Example provided by Dr. Marilyn Mantei, University of Toronto.

dissipate heat build–up. Several years later, a new team re–designed the brake. Although they had the final blueprints of the earlier brake, they did not have the reasons for the design decisions made. They removed the seemingly non–functional burr for cosmetic reasons. The result was a recall of all cars with the new, now deficient brake. A persistent system, on the other hand, would have captured the on–going discussion, the decisions made, and the reasons behind the decisions, and extended those across the "null meeting" period until the new group met.

The *principle of inter–operability* applies when independent systems operate well together. This implies standards for communication, common representations of problems, standard interfaces, machine independence, and so on. Commercial applications are now beginning to follow inter–operability, at least on a low level. In the user interface domain, look and feel standards are now available, such as Unix International's *Open Look* and the Open Software Foundation's *Motif*. Machine independence is provided by networked–based window and graphics protocols such as X windows, which allows a program running on one vendor's workstation to be viewed on another vendor's machine. Inter–program communication is also making some headway. For example, many Apple Macintosh programs now follow an interchange format allowing items copied from one program's view to be pasted into another program's view without either being aware of the other. In artificial intelligence, blackboard systems allow programs to make use of and to update a common knowledge representation.

Inter–operability becomes the *principle of integration* when sub–systems are tightly coupled. Updating one component would directly and immediately be reflected in the update of another. Perhaps all components would be built upon a single underlying representation of the on–going activity, with the components merely being different views and perspectives of the knowledge base and group.

Finally, the *principle of coordination* integrates all activity in the sense of group awareness, common knowledge, purpose, relevant training, and coordination of activities. It maintains the goals of the group, schedules tasks, commitments and dependencies, the progress to date. It keeps members up to date on the relevant progress of the group and of the external community.

Although elusive, the coherent workspace is under active development. One example is a program now being developed at McDonnell Douglas, a complete information system that supports very large aerospace programs (thousands of people), over the whole life–cycle of the project (e.g., three decades) of extremely complex products[4].

[4] Example provided by Doug Engelbart.

10. REFERENCES

Albus, J. S., McCain, H. G. and Lumia, R. (1989) "NASA/NBS Standard reference model for telerobot control system architecture (NASREM)." NIST Technical Note 1235, Robot Systems Division, Center for Manufacturing Engineering, National Institute of Standards and Technology, Gaithersburg, MD,

Allen, T. J. (1977) *Managing the flow of technology*. Cambridge, Mass, MIT Press,

Baecker, R. M. and Buxton, W. A. S. (1987) *Readings in human-computer interaction: A multidisciplinary approach*. Los Altos, California, Morgan Kaufmann,

Bair, J. H. and Gale, S. (1988) "An investigation of the Coordinator as an example of computer supported cooperative work." Hewlett Packard Laboratories, California, Unpublished.

Borenstein, N., Everhart, C., Rosenberg, J. and Stoller, A. (1988) "A multi–media message system for Andrew." In *Proceedings of the USENIX Winter Conference*, 37–42, Dallas, Texas, February 9–12,

Borenstein, N. and Thyberg, A. (1988) "Cooperative work in the Andrew message system." In *Proceedings of the Conference on Computer-Supported Cooperative Work*, 306–323, Portland, Oregon, September 26–28, ACM Press.

Bush, V. (1945) "As we may think." *Atlantic Monthly*, 176(1), 101–108, June. Reprinted in Greif, 1988.

Chang, E. (1986) "Participant systems." *Future Computing Systems*, 1(3), 253–270,

Chang, E. (1989) "Protocols for group coordination in participant systems." In *The Structure of Multimodal Dialogue*, edition). 229–247, North–Holland, Taylor, M. M., Neel, F. and Bouwhuis, D. G., Elsevier Science Publishers B.V.

Chapanis, A. (1975) "Interactive human communication." *Scientific American*, **232**(3), 36–42, Reprinted in Greif, 1988.

Clark, W. J. (1989) "The MIAC audiographic conferencing system: A practical implementation of the audiovisual service infrastructure." In *Multimedia '89: Proceedings of the 2nd International Workshop on Multimedia Communications*, Montebello, Quebec, April 20–23,

Conklin, J. (1987) "Hypertext: An introduction and survey." *IEEE Computer*, **20**(9), 17–41, Reprinted in Greif, 1988.

Conklin, J. (1988) "gIBIS: A hypertext tool for exploratory policy discussion." In *Proceedings of the Conference on Computer-Supported Cooperative Work*, 140–152, Portland, Oregon, September 26–28, ACM Press.

Corey, D., Abel, M., Bulick, S. and Schmidt, J. (1989) "Multi–media communication: The US WEST advanced technologies prototype system." In *Fifth IEEE Workshop on Telematics*, Denver, Colorado, September 17–21, Submitted paper.

DeSanctis, G. and Gallupe, R. B. (1987) "A foundation for the study of group decision support systems." *Management Science*, 33(5), 589–609, May.

Egido, C. (1988) "Video conferencing as a technology to support group work: A review of its failures." In *Proceedings of the Conference on Computer-Supported Cooperative Work*, 13–24, Portland, Oregon, September 26–28, ACM Press.

Engelbart, D. and English, W. (1968) "A research center for augmenting human intellect." In *Proceedings of the Fall Joint Computing Conference*, edition). 33, 395–410, Montvale, NY, Fall, AFIPS Press. Reprinted in Greif, 1988.

Engelbart, D. and Lehtman, H. (1988) "Working together." *Byte Magazine*, December.

Ensor, J. R. (1989) "Rapport: A multimedia conferencing system." *The ACM SIGGRAPH Video Review Supplement to Computer Graphics*, **45**(5), ACM Press, Baltimore, MD. Videotape.

Ensor, J. R., Ahuja, S. R., Horn, D. N. and Lucco, S. E. (1988) "The Rapport multimedia conferencing system—a software overview." In *Proceedings of the 2nd IEEE Conference of Computer Workstations*, 52–58, Santa Clara, March 7–10,

Eveland, J. D. and Bikson, T. (1986) "Evolving electronic communication networks: An empirical assessment." In *Proceedings of the Conference on Computer-Supported Cooperative Work*, 91–101, Austin, Texas, December 3–5, ACM Press.

Eveland, J. D. and Bikson, T. K. (1988) "Work group structures and computer support: A field experiment." In *Proceedings of the Conference on Computer-Supported Cooperative Work*, 324–343, Portland, Oregon, September 26–28, ACM Press.

Farallon (1988) "Timbuktu user's guide." edition). Farallon Computing Inc., Berkely, California,

Fish, R. S. (1989) "Cruiser: A multi–media system for social browsing." *The ACM SIGGRAPH Video Review Supplement to Computer Graphics*, **45**(6), ACM Press, Baltimore, MD. Videotape.

Foster, G. and Stefik, M. (1986) "Cognoter: Theory and practice of a Colab–orative tool." In *Proceedings of the Conference on Computer-Supported Cooperative Work*, 7–15, Austin, Texas, December 3–5, ACM Press.

Goodman, G. and Abel, M. (1987) "Communication and collaboration: Facilitating cooperative work through communication." *Office: Technology and People*, **3**(2), 129–146,

Gould, J. D. and Boies, S. J. (1983) "Human factors challenges in creating a principal support office system – The speech filing system approach." *ACM Transactions on Office Information Systems*, **1**(4), 273–298, Also published in Baecker & Buxton, 1987.

Gould, J. D. and Boies, S. J. (1984) "Speech filing – An office system for principles." *IBM Systems Journal*, **23**(1), 65–81, Also published in Baecker & Buxton, 1987.

Greenberg, S. (1990) "Sharing views and interactions with single–user applications." In *Proceedings of the Conference on Office Information Systems*, Boston, April, Also available as Research Report, Advanced Technologies, Alberta Research Council, Calgary, Alberta, Canada.

Greenberg, S. and Chang, E. (1989) "Computer support for real time collaborative work." In *Proceedings of the Conference on Numerical Mathematics and Computing*, Winnipeg, Manitoba, September 28–30, Available in Congressus Numerantium vol 74 and 75.

Greif, I. (1988) *Computer-supported cooperative work: A book of readings*. San Mateo, California, Morgan Kaufmann Publishers Inc,

Gust, P. (1989) "Multi–user interfaces for extended group collaboration." In *Proceedings of the Groupware Technology Workshop*, Xerox PARC, Palo Alto, California, August,

Guttman, S. (1989) "Farallon: Building the ultimate network." *MicroTimes*, 6(7), July.

Harrison, S. (1989) "The office design project." *The ACM SIGGRAPH Video Review Supplement to Computer Graphics*, 45(4), Videotape.

Jackson, S. L. (1989) "Hypertex

CHAPTER 4: TRENDS IN INFORMATION EQUIPMENT

1. INTRODUCTION

Performance specifications of a telecommunication network are necessarily dictated by the needs expressed by the user. The needs, of course, arise from the use of various office equipment for OA (Office Automation), computers, data storage equipment, measurement instruments and other electronic equipment. The network technology, therefore, cannot be considered in isolation.

Hence, in order to specify desirable features of a future national telecommunication network for R & D, a broad understanding of the anticipated developments in video, computer hardware and software, LAN (Local Area Network) and transmission technologies must be established.

2. TECHNOLOGY BACKGROUND

More often than not, network performance specifications expressed by users are limited by their knowledge of the available technology and ignorance of the technology that is forthcoming.

A market survey for the sale of ice cream based on interviews of a population that has never tasted ice cream will certainly lead to incorrect conclusions. A more logical approach would be to offer taste samples before an opinion is solicited.

The telecommunication network under consideration suffers from a problem similar to that of the ice cream mentioned above. Since it is impossible to offer a taste sample of the new technologies to the user, projections based on a solid knowledge of the state of technologies must be used to arrive at what, in the long term, will be the correct conclusions.

2.1 PERSONAL COMPUTERS

The computing speed of PCs (Personal Computers) are faster than 30 MHz today and it can be expected to become faster in the future. The memory capacities are becoming larger and in particular, the CD (Compact Disc) storage system allows a PC to store and manipulate giga–bytes of information. The CD storage system is already available in a WORM (Write Once Read Many) format and erasable CDs are just now being commercialized.

It should be noted that WORM storage systems which do not allow erasure are no less attractive than the erasable type in view of the fact that the majority of PC users seldom purge their hard disc memories unless they are forced to do so by lack of memory space. Also, almost all PC users have a large collection of old floppy discs covering many years of work.

The availability of large storage capacities and high computing speeds point to the need for high speed transmission links such as T1 = 1.544 Mb/s because users will not readily tolerate waiting times of minutes for the completion of a file transfer of a long document.

The higher computing speeds will also require data transmission rates such as T1 = 1.544 Mb/s in order to avoid a bottleneck when transferring data between PCs.

2.2 COMPUTER MAINFRAMES

Computing speeds of mainframes are such that many of them can make good use of data transfer rates in excess of 100 Mb/s between mainframes. Optical disc storage systems have also increased the memory capacities more than 10 fold.

The increase in computing speeds can be expected to continue in the future as the technologies for fifth generation computers are developed further. The need to provide a T2 = 6.312 Mb/s interconnection might be anticipated.

2.3 CAD/CAM

Unless the high resolution screen of a CAD/CAM (Computer Aided Design/Computer Aided Manufacturing) is filled within a second or so, valuable professional time is wasted. With the improvement in computing speeds, the bottleneck could well be the transmission bit rate of the interconnection between the computer and the CAD/CAM display terminal.

In order to provide near instantaneous creation of displays on CAD/CAM terminals, provision of at least T1 = 1.544 Mb/s, and in some cases, T2 = 6.312 Mb/s interconnection rates should be considered.

2.4 MULTIMEDIA PC TERMINALS

Communication interface cards for PCs now allow them to be used as facsimile terminals where the received facsimile document is displayed on the CRT (Cathode Ray Tube). Of course, a text composed by word processor software can also be transmitted in the opposite direction as a facsimile signal through such commu–nication interface cards.

Video image interface cards are also available today. They allow photographs and pictures to be scanned by an image scanner and stored in the PC memory. Reproduction on the CRT display is, of course, possible. Software to manipulate such images is also available today.

In addition, graphic software packages for PCs allow creation and manipulation of various charts, graphs and graphic designs. Some have reached a level where they might be classed as a new medium for artistic expression.

Aside from the use as a facsimile terminal and an image and graphics work station, today the PC is used frequently as a word processor. Many software packages are available.

In order to allow the PC to become a multimedia terminal for the different signal formats mentioned above, the Japanese computer terminal manufacturers are cooperating to establish a common standard for a multimedia interface box to which a PC can be connected.

Emergence of such a box will most likely accelerate the move towards one PC per staff–member in the office environment. Such is already the case for most R & D establishments today.

The multimedia terminal will require a network that supports standard analogue telephone lines as well as video image, graphic and text signals. In the future, signal formats that are expected to be supplied through ISDN (Integrated Digital Service Network), such as 2B + 2D = 160 kb/s (B = 64 kb/s, D = 16 kb/s) and T1 = 1.544 Mb/s will need to be accommodated by the interface box. Development in this direction will undoubtedly take place.

2.5 OPTICAL DISC MEMORIES

Document filing systems using WORM optical discs are available in Japan now. They will be offered to the North American market in the next few years at a cost of about $50,000 per system. At present they are designed to deal with text format primarily, but eventually signals from multi–media PC terminals will undoubtedly be supported as well.

Such document filing systems will realize the paperless office in the future. Incoming letters can be read into the system through a facsimile reader and stored on the optical disc as a facsimile signal. Electronic mail which is supported by the system can be used to transmit the letter to the proper person.

The recipient of the letter will view it on a multimedia PC terminal and make annotations in the margins using the graphics program software that is installed in the PC. The letter can then be filed anew in the optical disc and transmitted to the next person through the electronic mail system.

Eventually, a letter in response will be prepared in the word processor text mode and printed out for a signature. The signed letter can then be read into the optical disc storage by a file reader before delivery through the postal service.

Only two paper documents are involved in the scenario that is presented above; the original letter and the final letter in response.

A telecommunication link between such document filing systems will eliminate the need to use the postal service and thereby reduce the need of paper because a letter can be transmitted electronically. The problem of affixing a signature to the letter can be solved by using a graphic tablet.

A large scale optical disc storage system can conceivably serve Canada's R & D community as an information data bank on a demand basis through a national network.

The signal formats used to interconnect multimedia optical document filing systems will most likely conform to ISDN standard bit rates such as B = 64 kb/s and T1 = 1.544 Mb/s.

2.6 VIDEODISC SYSTEMS

Over the past several years, sales of laserdisc players for videodiscs have increased significantly in Japan primarily because of the extra picture quality that is provided by the full 4 MHz bandwidth of the luminosity signal. A better than average quality signal for NTSC (National Television Standards Committee) is provided by most laserdisc players.

The erasable videodisc system is being readied for commercialization in Japan today. Also, there are laboratory versions of HDTV (High Definition Television) videodisc systems. More than 45 minutes of a compressed HDTV signal can be recorded on a CD sized videodisc.

It should be noted that videodiscs not only record moving pictures but they can also store many still–pictures as well. A standard 30 cm diameter videodisc has about 27,000 pictures stored on one side in pulsed–analogue form.

When text and video still–pictures are combined, documents and books such as an encyclopedia can be stored on a videodisc. In fact, Collier's Encyclopedia has been recorded on a single CD size videodisc and offered commercially.

A videodisc system can, therefore, be used as a data bank that provides visual, text and sound information. The system can be in the form of a jukebox where multiple discs are mounted on a single spindle. In the future, the videodisc jukebox can become a valuable information source as well as an interactive educational system for R & D.

In order to facilitate access to a videodisc data bank through a telecommunications network, the most convenient signal format is in digital form because digital compression and

processing techniques can be applied to reduce the bandwidth (i.e. bit rate) requirement. Standard analogue transmission of NTSC TV signals require a bandwidth of 4 MHz. A network that supports such a bandwidth is prohibitively expensive.

When video information is recorded in digital form the number of frames that can be stored on a disc will be reduced by more than 1/5 in comparison to the pulsed–analogue format for standard videodiscs. However, the capacity of a video disc system can be increased readily by using larger size discs and by placing more discs on a single spindle. Also, to avoid blocking (i.e. busy signals), each disc can have multiple playback heads.

Probably the most convenient digital signal formats for videodisc data bank will be those standardized in ISDN. For text material, 64 kb/s should be more than sufficient. For images and large text data transfers, T1 = 1.544 Mb/s should be adequate.

2.7 DIGITAL TV SETS

Internally digitized TV sets are ready to be marketed by manufacturers in Japan. Already a top–of–the–line product is being test marketed in U.S.A. by TOSHIBA. The reason for using digital circuitry inside a standard TV set is the fact that the component count can be reduced by as much as 30 %. This means lower manufacturing costs and an increase in reliability.

Since the digital circuitry also contains a frame storage memory, digital image processing can be performed. Ghosts can be removed and picture–in–picture displays can be provided. The latter means that 3 or 4 TV channels can be followed at the same time on a split–screen basis.

Signal processing can be taken one step further by additional digital circuitry. The NTSC standard interlaced–frame can be converted to a non–interlaced format and the frames presented at a 60 Hz rate. Image definition becomes sharper and colour presentation is also improved.

Also, the number of scan lines can be doubled to 1,050 lines per frame and the resolution of HDTV matched. Spatial interpolation between the lines provides the missing pictorial information. Temporal interpolation can also be performed if extra memory is available to store several frames at once.

Thus, by digital techniques, the standard NTSC TV image can be converted to a pseudo–HDTV image which can match that of a 35 mm film slide. Such a TV set using a projection system to obtain a large screen size has been demonstrated in a laboratory at TOSHIBA of Japan.

When a full–motion NTSC TV signal is digitized, the resulting bit rate is over 100 Mb/s. Such a bit rate is too high for any network to support economically but digital processing can reduce the bit rate significantly depending on the degree of movement that is present in the image. Of course, the complexity of the digital circuitry increases as the bit rate is lowered.

A more economical approach is the use of stop–motion (i.e. freeze–frame) for the transmission of moving images. A trans–mission bit rate of T1 = 1.544 Mb/s should be sufficient to support a video teleconferencing system. In fact, the standard for a T1 rate videophone signal is about to be established internationally.

2.8 VIDEO TELECONFERENCING

Video teleconferencing services are offered commercially in a number of countries today, although they are not as popular as anticipated initially. One reason for their unpopularity is the high transmission costs. The conference setting may also be acting as a deterrent by being very formal.

The screen size in some cases is too small and the lack of eye contact arising from an improperly placed video camera in some instances conveys an incorrect impression between the participating parties.

Many of the difficulties encountered by video teleconferencing systems can be overcome today by utilizing in full the technologies that are available.

A large projection screen for the pseudo–HDTV display mentioned above can be built into a wall and a cafeteria table placed in front of it. The concept is to provide an informal setting where the participants sit and have a discussion over coffee.

The lens of the camera is chosen so as to provide a life–size image on the screen. Eye contact is made possible by placing the camera directly behind the projection screen that has a half–silvered mirror canted at a 45 °C angle.

The system should also be designed to allow on occasion teleconferencing among more than two locations. A split screen display can accommodate the multi–parties. Some thought will need to be given to the system design on how to preserve eye contact between the participants.

Aside from being able to talk face–to–face, another important factor in discussions associated with R & D is the desire to share a common writing space such as that represented by a piece of 8.5" X 11" paper. A facsimile machine can provide a partial solution but a more satisfactory method would be to share a flat white CRT display on a teleconference basis. Light–pens and software for such computer teleconferencing is available today. The CRT display might be imbedded horizontally into the cafeteria table to simulate a piece of paper.

An electronic blackboard that reproduces the drawings remotely might be used as the common writing space but such a product is not on the market yet although the technology has been available for some time now.

The audio associated with the video teleconferencing service should have very low distortion but the bandwidth probably can be limited to the standard 4 kHz of the telephone circuit because of cost concerns. Preferably, the press–to–talk type voice actuated controls should be avoided to maintain a natural sound setting.

The positive acoustic feedback (i.e. howling) that occurs when fully bi–directional audio is implemented, can be suppressed by AGC (Automatic Gain Control) or an appropriate arrangement of comb–filters.

In order to keep the cost of the transmission link to a minimum, stop–motion transmission techniques can be used for the video signal. A bit rate of T1 = 1.544 Mb/s should be more than adequate. The telecommunication link for the shared writing space can be a standard RS–232C link at 4800 b/s.

2.9 LOCAL AREA NETWORKS

The predicted integration of voice and data has not been realized and the LAN (Local Area Network) companies have ousted the PBX (Private Branch Exchange) from the data network field.

One of the major factors against the PBX is its high line cost which averages about $ 1,000 per telephone. Equipment to connect a terminal to the telephone line is an extra–cost item as well. In contrast, many LANs can be installed for less than $ 800 per terminal.

Also, an impression was created that a terminal–to–terminal guaranteed throughput of 64 kb/s through the PBX was inadequate whereas in reality most applications could be more than satisfied with such a bit rate.

Most LANs have the advantage in terms of non–blocking interconnection where 2 or more terminals can communicate with a single terminal on a TDM (Time Division Multiplex) basis. This means that sharing of a single resource such as a data file is more efficient because of the TDM feature of interaction which eliminates the waiting period (i.e. busy signal) that can be encountered in a PBX telephone–type interconnection.

Of course, in most LANs the time required to complete a transaction between the single resource and a terminal becomes longer in proportion to the number of terminals that are trying to use the single resource at the same time.

Since most LANs are based on TDM techniques, if a high terminal–to–terminal throughput such as 1.544 Mb/s is to be guaranteed, the operational bit rate of the LAN must be high. For example, if 100 terminal pairs are communicating with each other the operational bit rate must be higher than 150 Mb/s. Not many LANs have such high operational bit rates yet.

When operational bit rates exceed 100 Mb/s, expensive high speed components such as ECL (Emitter Coupled Logic) ICs (Integrated Circuits) which consume much more electrical power than standard TTL (Transistor–Transistor–Logic) ICs, or high cost GaAs ICs must be used. If fiber optic transmission links are employed, expensive SM (Single Mode) technology must be used.

An alternate approach to designing LANs that can guarantee a terminal–to–terminal throughput of 1.544 Mb/s or higher is the centrally switched star network which is exactly what the PBX employs. The terminal–to–terminal guaranteed throughput is then determined by the maximum bit rate that can be supported through the centrally located switch and the transmission lines used in the arms of the star network.

If circuit (i.e. spatial) switching methods such as that used in analogue PBXs are employed, switching of 50 Mb/s signals is not out of the question. In fact a 32 X 32 digital matrix switch capable of circuit switching 50 Mb/s signals is available now (e.g. Sierra Semiconductor Corporation: SC11320CU). Even the 4 X 4 analogue matrix switch used in the MITEL SX–200 analogue PBX can circuit switch 5 Mb/s with no problems in high cross talk levels and instances of low on/off isolation.

Since the centrally switched star network interconnects terminals on a one–on–one basis, the transmission lines of the network need only support the maximum operational bit rate of the terminals. Copper–wire pairs are sufficient for T1 = 1.544 Mb/s and if T3 = 44.736 Mb/s is to be supported, multimode GI (Graded Index) fiber optic cables can be used.

Thus, the technology exists to manufacture a PBX/LAN that can offer interconnection to hundreds of terminals requiring high bit rate interconnections such as T1 = 1.544 Mb/s. It

is a matter of time before such a product is offered commercially. In fact, Net–Five Telecom Corporation of Regina, Saskatchewan is developing a 1.2 Gb/s fiber optic star network LAN that has a guaranteed non–blocking throughput of 10 Mb/s. It is based on a 128 X 128 digital matrix switch that can circuit switch 10 Mb/s signals.

The emergence of a PBX/LAN that can support T1 = 1.544 Mb/s and higher bit rate interconnections is almost certain. Also, that future fiber optic ring–type LANs will also support such bit rates per terminal is apparent from the direction of developments in fiber optics technologies although they may be more expensive than circuit switched LANs in a star configuration.

2.10 HIGH DEFINITION TELEVISION AND HIGH RESOLUTION IMAGES

A number of HDTV (High Definition Television) systems are being promoted today. Their displays can match the quality of 35 mm film slides even when the screen size is 6 ft diagonal. A set of digital ICs have also been produced for processing the signals digitally inside a HDTV receiver.

When the HDTV signal is digitized, the resulting bit rate is about 800 Mb/s. This can be digitally processed to 300 Mb/s or less and recorded onto a laserdisc.

The application of HDTV technology to R & D can be seen in medicine and areas where high resolution imagery is important. Film images from X–Rays, CAT (Computer Aided Tomography) and MRI (Magnetic Resonance Imaging) systems as well as an ERTS (Earth Resources Satellite) can be scanned by a high resolution scanner and stored digitally, either on a magnetic medium or an optical disc. In many cases the intermediate step of photographic film can be bypassed and the image information stored directly as digital information. This is possible with CAT and MRI systems as well as many X–Ray systems.

For research, diagnostic or training purposes, these digitized images can be transmitted over a telecommunication network to locations separated by long distances. Because of the large number of bits that must be transmitted, bit rates of T1 = 1.544 Mb/s and higher need to be considered.

It should be noted that long distance transmission of high resolution CAT, MRI and X–Ray images not only has an importance in R & D but can also play a vital role in offering expert medical care universally across Canada. The cost of providing health care could also be reduced through the expert consultation made available through these high resolution images. In some cases observation of high resolution stop–motion real time images of a patient may lead to life saving treatment in emergencies.

CHAPTER 5: INTERNATIONAL REVIEW

1. COMPARATIVE SUMMARY OF R&D NETWORKING

This chapter summarizes Research, Development and Education (R&D) networking in Canada and compares this situation with what is happening in the U.S., Europe, and Japan. The comparison is made under three broad sections.

In Section One, entitled R&D and Networking Environment, the political, economic, and technological infrastructures that support R&D, and more specifically R&D networks, are discussed.

Under Section Two, Current Network Status, the existing Canadian networks are compared to the existing networks in the U.S., Europe, and Japan. These networks are compared in terms of performance, service, use, administration, and funding.

The last section, Future Network Developments, compares Canadian proposals, approved additions and expansions, with future networks being planned in other industrialized countries.

More detail is provided in the specific chapters dedicated to each country review.

1.1 R&D AND NETWORKING ENVIRONMENT

This section focuses on the following four areas:

- General R&D environment in Canada compared to other advanced nations;
- Telecommunications Regulatory Environment in Canada compared to other advanced nations;
- Network environment in Canada compared to other advanced nations; and
- The effect of R&D networks on international competitiveness.

R&D Environment

By whatever measure of performance one chooses, Canada's R&D record is poor when compared to other industrialized countries. For example in 1988 Canada spent 1.32% of GDP on R&D (GERD ratio). In Japan, the U.S. and many European countries, like West Germany and Sweden, GERD ratios are at least double the Canadian figure.[1] In addition industrial R&D is substantially lower in Canada (0.75% of GDP) then in the U.S., Japan and Europe, where industrial R&D comprises over 2% of GDP. However it is important to note that the "information technology" industry is Canada's best performer of R&D.

The relatively small population of Canadian scientists, engineers, and researchers, the relative lack of Canadian industry to fund R&D, and the shrinking ability of the Federal Government to finance investment, means that Canadian research efforts must be coordinated among institutions both in and outside Canada. In order to have access to all the necessary

---

[1] Ottawa Citizen; Research Spending Comes Up Short; Thursday March 8, 1990.

equipment and expertise, Canadian researchers must colloborate among themselves and also with researchers around the world.

The Centres of Excellence Program recently sponsored by Industry Science and Technology Canada (ISTC) puts additional pressure on the need for researchers across Canada to communicate. Researchers within the same centre of excellence are geographically separated by long distances (which is true for all researchers in Canada), which creates the need for a communications network. This $240 million, five year program provides funding to 14 centres, involving scientists from a score of Canadian universities, government laboratories and private corporations. A national high–speed communications network would contribute substantially to the infrastructure required to make this program a success.

Telecommunications Environment

The telecommunications environment in Canada suffers in comparison to those in other countries. The telecommunications systems in Canada are currently run by a number of geographical and political monopolies, each with its own technologies, policies, and costs. The result is that design and operation of a national electronic network must be integrated into a widely varying set of constraints, at much higher cost than elsewhere. Long distance communications costs in Canada are an order of magnitude higher than in the U.S., and significantly more than in Europe. In both the U.S. and Japan, the telecom monopolies have recently been opened to competition; some European countries have followed suit, and there is significant pressure from the EC to eliminate government telecom monopolies and standardize the European telecom industry. R&D networking is, of course, only a small part of Canadian activities that suffer from the fragmented telecom monopolies in Canada; there is strong pressure from business to follow the lead of the rest of the world in deregulating and standardizing the telecommunications environment, and eliminating monopolies in Canada.

The fragmented monopoly situation of Canada's telecommunication industry shows signs of cracking under intense pressure from business, commerce, and the research community. Most recently, Telecom Canada, the association of regional telecom monopolies which sets long–distance rates for links crossing monopoly borders, has announced substantial reductions in long–distance rates for digital service of 70 % for 56 Kb/s service and 50 % for T1 1.544 Mb/s service. While the cuts do not bring rates down to U.S. levels, Telecom officials suggest further cuts may be in the offing as useage grows. The telecom companies, fearful of losing their monopolies, have simultaneously:

- weakened the argument that competition is needed to bring rates down;
- weakened the existing competition, whose rates are tied to Telecom rates by regulatory rulings, but which cannot offer certain services; and
- undermined criticism that high pricing makes resale and sharing of line capacity uneconomical in many cases.

Only time can tell whether this move allows the current operational and regulatory tangle to remain in place or whether Canada will be forced to join the rest of the world in deregulating and demonopolizing its telecommunications industry.

Rapidly Changing Network Environment

The current situation of R&D networking in the world is one of rapid development. The history of networking extends less than two decades during which the developments in computers, switching, and transmission technology have been phenomenal. Networks are

moving away from the original, informal, non–standardized, and often experimental connections between individual researchers, to a highly organized, standardized tool of national and international R&D.

Canada has a long history of leadership in telecommunications technology research, starting with the telephone, through satellite communications and digital switching technology, to the development of the first ISO/OSI X.400 mail software. However, the impetus for network development in America, Japan, and Europe has switched to government support and to networks as an R&D production tool. As a result, Canada is now losing ground rapidly in terms of speed, services, and connectivity.

The computer and electronic switching hardware available for networking in Canada is comparable to what is used in other countries. However, with the exception of BCnet, Canada does not have examples of high speed fibre–optic research networks. In the U.S., the smallest of the three long–distance competitors has three operating national fibre networks. Most Canadian network connections use the Datapac service or equivalent at 9.6 Kb/s, or the MegaStream equivalent at 56 – 64 Kb/s. The situation is similar in European and Japanese networks. In terms of speed, NSFNet in the U.S. is the only major network to provide T1 speed between a significant number of network nodes, and JANET in the U.K. provides 512 Kb/s service between its eight backbone nodes.

Canadian networks are able to take advantage of Anikom satellite relay service from Telesat Canada for long–distance high speed links. These can be significantly cheaper than terrestrial links using Telecom Canada, a situation caused by the artificially high rates. While satellite links are used for some inter–network gateways, they are not currently used as links within R&D networks elsewhere in the world. Currently, only the AECLnet and SUPERNET in Canada use Anikom satellite links. Neither of these are T1 connections, which Anikom is capable of providing.

Research Networks and National Competitiveness

First–world nations today exhibit an increasing attitude of "techno–nationalism", in which scientific and technological advances are guarded as a key component of national competitiveness. For example the prevalent attitude in the U.S. appears to be that high–speed, high–capacity R&D communications networks are essential tools to support Government policies and mandates with respect to international competitiveness. As a result the U.S Government treats the investment in these networks as a prerequisite cost required to maintain the competitiveness of the nation. Detailed cost benefit analysis are not necessary to build an appropriate case for the investment.

The situation in Canada is somewhat different then the U.S. situation described above. In Canada, the impact of research networks on international competitiveness is a link that has not been forged nearly as strong as in the U.S. Japan Europe. In fact historically, Canada has had difficulty justifying investment in the expensive research tools which form key parts of advanced research programs in aerospace, biotechnology, telecommunications, energy production, etc..

1.2 CURRENT NETWORKING STATUS

This section draws some general conclusions regarding how Canadian research networks compare to research networks in other countries. These networks are compared in terms of protocol development, performance, network applications, access and use, and funding.

Protocol Development

Canadian networks are generally based on American computer equipment and thus use communications protocols developed for use on these computers, such as DECNET, NJE, and TCP/IP. However, Canada has been in the forefront of software development such as the EAN implementation of X.400 mail specification (ISO/OSI), which is also used in the European EAN network. Major efforts are underway in Europe, mostly with government support, to complete and test a full set of ISO/OSI specification software covering all network services (mail, news, directory, remote job entry, remote login, etc.) in preparation for an international migration of network protocol to the OSI standard. Much effort in Japan has been devoted to developing conversions and protocols for transmission, translation, and display of Japanese characters. Although the U.S. has formally recognized its commitment to moving towards the OSI standard, the TCP/IP standard is so deeply entrenched that it is generally accepted that the U.S. will migrate to OSI at a much slower rate then the rest of the world. As a result, there may be a window of opportunity to export OSI products to the U.S.

Network Performance

The existing national networks in Canada (Netnorth and CDNnet) are considerabley slower than existing networks in the U.S. (NSFNET), U.K (JANET), Germany (DFN) and Japan (NACSIS). In all of these countries, with the exception of Canada, T1 backbone links exist or are being implemented. The U.S. is also considering whether to upgrade the existing NSFNET backbone to T3 speeds.

Network Applications

Use of network services does not differ greatly in North America and Europe. The most popular services are mail and news, followed by file transfer, remote login, and remote job entry where these services are offered. Mail and news are low–volume, low speed applications which can adequately be serviced by existing telecom links at reasonable cost. Network users can be expected to pay the cost of the link, or university budgets can cover the cost under general computing budgets. Higher performance services requires that additional funding be obtained directly for the network because of the increased costs of higher speed links. The costs imposed on users, as a result of enhanced network performance (ie. T1 links), are much more onerous in Canada then in any of the other countries.

There are a number of networks around the world which are quite different than the traditional operational R&D networks, and thus deserve special recognition. For instance the major purpose of ARISTOTE, a French national R&D network is to act as a testbed for French research institutions to develop networking technology. In addition the Sigma Network in Japan is a research and development testbed network designed for use in the Sigma Project, which is intended to promote increases in software productivity by producing a standard work station environment for use in Japan.

In addition to the two exceptions mentioned above, an important component of NSFNET (and ARPANET to a certain extent) is research into network design and performance. An independent research network is used to develop and test new technology for the backbone and provide a facility for the anticipated migration of OSI protocols in the coming years and for the development of gigabit technology. The testing of products and services does not in any way affect the performance or reliability of the operational network — the two are kept separate. The test and engineering network consists of a loop from Merit in Ann Arbour, Michigan to IBM Technical Computing Systems in Milford, Conneticut to IBM research in Yorktown, New York to MCI in Reston, Virginia back to Merit. Packet video is being experimented with on this network (with potential application to video conferencing).

The Japanese networks are quite different from European and North American networks as they are used only peripherally for mail. The major purpose of the NACSIS system is fast, local access to one of the largest databases of scientific and technical information in the world; electronic mail and other networking services are provided but assume less importance in the system design. In order to offer an online information retrieval service and catalog information service to researchers at universities, a high performance computer system is installed to meet the needs of large–scale database services. An optical disc storage system of 83 gigabytes of memory has been implemented for this purpose.
Japan's R&D networks are more oriented towards public access to information databases and news than networks in other countries. The JUNET system is intended as a testbed for networking development rather than an R&D production network; its traffic is estimated at 85% news and only 15% mail. This may be a result of the relatively high number of FAX machines in Japan and the difficulty arising from typing Japanese characters on a standard keyboard.

Access and Use

The participation of industrial companies in R&D networks in all countries appears to be minimal. Many networks specifically exclude corporate or profit–oriented organizations from membership or connection. In the U.S., commercial, for–profit use of research networks is forbidden; in Canada, use of research networks by commercial for–profit enterprises is not generally forbidden, but they are usually assessed a higher membership fee (which acts as a barrier for small companies). As such, the participation rate by industry in Canadian networks is also quite low. Some European networks have commercial enterprises as members; especially those that are funded by membership fees. Most networks around the world which are funded by government science and/or education departments are not open to commercial enterprises. In Japan, the NACSIS/N1net network currently does not have direct connections to commercial organizations. However, there are 176 companies participating in HEPnet, and a number of companies work closely with researchers to develop Japanese character conversion applications and ISO/OSI protocol implementations, which may be tested on JUNET. In general the participation of industrial companies in R&D networks in all countries is quite low.

Funding

The importance placed on networking by governments in each country is reflected by the way national networks are funded. Here, Canada is an exception: for the most part, Canadian networks have been funded by member organizations. While most of these member organizations are universities and therefore government–funded, the money they spend on network membership usually comes from general research or computer centre budgets. In the U.S., the government has funded capital expenditures for most national networks, and provides initial capital and partial funding to mid–level networks (equivalent to Canadian provincial networks). The majority of national networks in Europe are directly funded by the national government concerned, and continental network research is funded by the European Community. The Japanese government has completely financed the NACSIS network and database development.

1.3 FUTURE PLANS

In virtually all the countries reviewed, there are active plans for enhancing, expanding, and/or unifying their R&D networks. Governments are becoming the driving force behind network development. The political organization of networks along national lines will become more pronounced as national governments take on more of the funding, operating, and policy responsibility.

Future Canadian Networks

In Canada there are plans to develop or enhance regional networks in Nova Scotia, New Brunswick/PEI, Alberta and Manitoba. Although there is great interest in developing regional networks in Saskatechewan, Newfoundland and the Territorities, formal plans have not yet been put in place. In terms of a national backbone network, the CANET proposal is modest compared to plans for NREN in the United States. The CANET plan calls for an initial backbone speed of 56 Kb/s connecting regional networks. While there is significant support for T1 speeds at the outset, which is comparable to NSFNet's current backbone speed in the U.S., the cost of the cheapest Canadian T1 long distance service, Telesat Anikom 1000, is more than 8 times that of compareable service in the U.S..

Future U.S. Networks

There is general consensus in the U.S for the establishment of a National Research and Educational Network (NREN), with NSF as its lead agency. This stems from the existing U.S. Internet being far from uniform in the type and quality of service. It also does not yet reach the entire research community. Agencies like DoD, DoE, NASA and NSF have created their own networks (or in some cases several networks) to support individual missions. The results are that interconnectability of these networks is poor, protocols are different and a network user may have 5 or 6 identification numbers. The traffic on NSFNET and other research networks in the Internet is primarily electronic mail. Real time connections to advanced computers are frustrating because of delays and unexpected disconnects. Breakthroughs in enhancing service requirements depend on cross–disciplinary research and, therefore, the need for one cohesive national network.

A three phase plan exists to build the NREN:

> **Phase 1:** Is complete and involved the NSFNET backbone being upgraded to full T1 service.
>
> **Phase 2:** Is currently underway and involves upgrading the service provided to about 200 research institutions using a shared backbone with T3 (45 Mb/s) capacity.
>
> **Phase 3:** Will provide one to three gigabit per second networking services to selected research facilities and 45 Mbs networking to 1000 sites. Deployment is not expected until middle to late 1990s. Research in gigabit technology is currently underway. Several key development areas include (1) switching systems for ultra–high speed communications links to move information at billions of bits per second; (2) network to computer interfaces making it possible to exchange data in multiple formats and media at high speeds; (3) software programs to manage the network; (4) telecommunication links to support the network.

As NREN evolves from NSFNET, a three level hierarchy will be maintained:

- Level 1: Campus or research organization;
- Level 2: Mid–level network or regional network (defined on a regional basis or discipline basis) linking individual sites; and
- Level 3: The backbone itself which would physically link the different regional networks.

There is a general commitment in the U.S. to move to ISO–OSI as soon as possible. DEC and IBM have announced to their international clients and the U.S. government that they fully intend to implement OSI by this fall. A transport layer of OSI running along side the

same machines with TCP/IP already exists. Most of the experts in this field have indicated that interoperability between these two domains is not an area of concern .

Universality of access is also a feature of NREN. It is anticipated that the NREN will give researchers, students and colleges of all sizes, as well as large and small companies in every state, access to the same high performance computing tools, data banks, supercomputers, libraries, specialized research facilities and educational technologies.

Future European Networks

A report submitted in Spring 1989 to the European Economic Commission by RARE, the European body for coordinating national academic networks, has proposed a European High–Speed Networking Initiative, to consist of a central backbone called VENUS, regional networks, and supporting networks for specific disciplines. Plans call for 20 to 40 nodes across Europe, serving from 200 to 500 sites. A phased improvement in network speeds is considered, with 2 Mbs immediately, 100 Mbs in a few years time, and gigabit speeds by the end of the century. Initial costs were stated as $6 Million annually, rising to $24 – $30.

In West Germany, a national fibre optic network is being installed by the Deutsche Bundespost, the national PTT. This "Forerunner" network, with capacities up to 140 Mbs, is intended for video teleconferencing and other high–bandwidth requirements, and will be available for scientific networking including graphics transmission and access to supercomputers. At the Regional Computing Center of Stuttgart University, applications are being developed and installed to make use of this capacity. Two projects are in progress to connect research centres located 1000 km and 150 km from Stuttgart with the Cray 2 supercomputer there at speeds of 100 Mb/s. Several years will be required to reach this speed, and costs of the high–speed link must still be negotiated with the PTT, which is in the midst of deregulation.

The campus of Cambridge University is spread over relatively long distances within the city of Cambridge, U.K. A new network using about 200 kilometres of optical cable has recently been completed to connect existing local networks. Both the local networks and the connecting network use a design called the Cambridge Fast Ring, which permits relatively cheap hardware and low complexity. The new network, built with support from Olivetti Corp., has potential speeds in the Gigabit range and has been tested at 600 Mb/s.

An organization has been set up by the Italian Ministry of Scientific Research and Technology to develop a high–speed national backbone connecting all research institutes and universities in Italy. This organization is called "Gruppo Armonizziazione delle Reti per la Ricercha (GARR)". The system will use 2 Mb/s leased lines and offer four protocols: X.25, TCP/IP, DECNET, and SNA. Migration to full ISO/OSI protocols will be undertaken as soon as feasible. The backbone will also have a high–speed link to CERN, the European high–energy physics centre located near Geneva.

The network will be financed by the Ministry of Scientific Research and Technology at a estimated rate of 9.7 Million ECU/year (about $14 Million), with half to pay for leased lines and the remainder for control units and software.

Future Japanese Networks

The Science Information Network (also referred to as NACSIS) appears to be the unifying force to realize a single network that serves the R&D interests of Japan.

2. R&D NETWORKS IN CANADA

2.1 SUMMARY

Canada is lagging behind the U.S. and other countries in the Western World in the development of national high speed communications networks for research and development. This is very regrettable since only ten years ago Canada was at the forefront of network technology, Although we are still a technical leader in some aspects of network technology, we have lost ground to the U.S., Great Britain, Germany, France and Japan. In many of these countries communication lines capable of transmission up to 40Mb/s are now being installed.

Despite the fact that a full–function Canada–wide network does not yet exist, there has been significant activity, over the last several years, in the design and implementation of regional networks in Canada. Regional networks are firmly established in British Columbia, Ontario and Quebec and are in the the process of being developed or extended in Alberta, Nova Scotia, New Brunswick, and PEI. Although there is great interest in developing regional networks in Saskatechewan, Newfoundland and the Territorities, formal plans have not yet been put in place. The Federal Government's interest in a national backbone network provides an incentive for regional networks to be implemented. A regional structure must be in place for a national network to of benefit.

In general, the networks used for research, including research administration and messaging, are almost exclusively based, funded, and operated by university computer/communications departments, though some networks do include private–industrial companies. In some cases, funds in support of networks are identified and provided directly by provincial governments or related agencies. For the most part, however, many of the direct costs are covered by contributions from general university operating budgets. Many, if not most, of the administrative costs (salaries, travel etc.) are not identified and accounted for explicitly, but rather are buried in general departmental budgets. In fact, the benefit of the regional networks as they exist today is that universities pour significant amounts of "free" labour into designing and managing these networks. Universities have been involved in network management for five or six years in Canada.

Canadian Telecommunications Industry

One of the relative bright spots in Canadian high technology development has been in the telecommunications field, in which Canadian companies are among the world leaders. This is due, in no small part, to a large R&D effort funded by the telecommunications monopolies in various regions of the country. This success seems to be in spite of our R&D environment, however, rather than a result of it. In a recent report submitted by Dr. McPhail at the University of Calgary[2] he states that:

> "Canada's regrettable R&D track record can be attributed, in part, to a major structural problem. There is a lack of cooperation among the major players in the Canadian telecom environment, a deficiency that is reinforced by the patchwork regulatory system. There are no meaningful, large–scale joint efforts by unions, universities, researchers, governments (municipal, provincial, or federal), professional

---

2 McPhail Research; Telecom 2001: A Strategic Forecast; October 1989.

> associations or the telecom industry itself, that seek to collectively support either major research and development activities or significant off–shore marketing efforts, both of which are vital to the success of Canadian telecommunication in the future."

There has also been a lot of debate over Canadian telecommunications policy. In a discussion paper commissioned by the Royal Bank of Canada as a contribution to the public debate on telecommunications policy changes, Janisch and Schultz[3] report on voice and data communications in Canada as growth industries, and provide evidence of the overall data communications market in Canada growing at an annual rate of almost 30 per cent. Specifically, public data networks have grown at an annual rate of 25 per cent from 1985 to 1989. Even with this growth, they note however, in a quote from the Economic Council of Canada, that:

> "[We] are lagging in comparison with other industrial nations, particularly the United States and Western Europe. Our message, therefore, is that Canadians must improve their performance with respect to the introduction and use of computer–based technologies. Failure to do so will lead to a loss of prosperity and jobs as Canada's competitive position deteriorates."

Janisch and Schultz go on to say:

> "Canada's persistent lag in information technology diffusion should be seen as a serious problem of national proportions. This concern has been raised repeatedly by the Economic Council of Canada and by the federal Department of Communications, and by such private sector organizations as the Information Technology Association of Canada and the Canadian Advanced Technology Association."

and

> "The failure to inspire the national imagination as to the critical importance of information technologies to Canada's future, combined with structural impediments, have been particularly noticeable with respect to the federal government's inability to develop a positive national telecommunications policy. This failure to "romance the microchip" has meant that initiatives undertaken by the federal government have always been seen by provincial authorities as threatening to erode their jurisdiction rather than creating a national market in which all could benefit from the successful transition to an information economy."

Janisch and Schultz then describe, and take serious issue with, the jurisdictional division of responsibilities in Canada and the differences in provincial policies which lead to quite different choices available to businesses across the country:

> "The result has been to fragment national markets, to balkanize national carriers and to discourage the development and introduction of new service offerings and new service providers."

In conclusion, they strongly recommend a move toward a national policy which

> "unleashes the full creative and regulatory potential of competition in telecommunications in order to permit user–driven advanced systems and services".

---

3 Royal Bank; Exploiting The Information Revolution: Telecommunication Issues and Options for Canada; October 1989.

In a recent speech to the Canadian Club of Toronto, Allan Taylor, Chairman and Chief Executive Officer of the Royal Bank of Canada[4], describes telecommunications regulation in Canada:

> "Ottawa regulates the largest regional telephone companies. This is true in Quebec, Ontario, British Columbia, and the North.
>
> "Except that in Ontario and Quebec, the provinces regulate some small telephone companies. Except that in some provinces, the larger telephone companies are regulated or run by the provincial governments. Except that in Edmonton, Thunder Bay and Prince Rupert, municipal governments are in the telephone business.
>
> "In Ontario, once technical standards are met, any kind of phone can be hooked into the telephone system. But, in Manitoba, you have to rent a telephone from the government–owned company before you install another one bought from someone else. Businesses may not own an internal switchboard. They must go through the provincial phone company.
>
> "In Saskatchewan, until recently, you couldn't even advertise telephones for sale. Newfoundland and Prince Edward Island still do not have cellular phones. The list of different services, prices, and equipment available – or not available – goes on and on."

Many observers of the telecommunications scene have commented on the impact of regulation and the absence of real competition in the Canadian market, and blame this situation for the significant differences in rates for various telecom services between Canada and the United States. Hoey compares the cost of an interstate AT&T T1 service to a similar Telecom Canada T1 service, using B.C. Tel's tariffs[5]. Depending on the mileage involved, the Telecom Canada T1 rates are between five and eight times the AT&T rates. Hoey points out that the AT&T rates are not even the lowest in the USA, and that other carriers, such as MCI and Sprint, offer T1 services at rates 8% to 40% below AT&T's, and that some specialized carriers offer additional reductions for long–term contracts. Hoey goes on to blame this disparity in rates on Canadian regulatory policy, and on the lack of effective competition, and recommends a revamping of the regulatory approach not only to T1 but to other competitive services.

The problems and barriers faced by Canadian R&D networks stand in stark contrast to the activities and plans of the NSF in the United States, specifically with regard to NSFNET. In this regard, Quarterman notes[6]:

> "Although originally intended specifically to connect supercomputer centers, NSFNET now provides other services as well, and has in fact become the national U.S. Research network, especially since the demise of ARPANET."

---

4 Royal Bank; The New National Dream: A Telecommunications Policy for the 21st Century; November 1989.

5 Eamon Hoey; The Angus report on Communications Systems, Services, and Strategies; Number 69; October 1989.

6 John S. Quarterman; The Matrix: Computer Networks and Conferencing Systems Worldwide; January 1990.

It should be noted that NSF was heavily involved in the original, and continued, funding for the five NSF–sponsored supercomputer centers, in addition to their above–noted support for NSFNET. In a brief presented to NSERC by the OCLSC Supercomputer Users' Group at the University of Toronto and the Inter–University Supercomputer Advisory Board, the authors note that, in the U.S., Federal (NSF) support for these supercomputer centres was about 60% of their total funding. The authors then go on to compare this level of U.S. Federal funding for supercomputers with Canadian Federal funding for supercomputers, and estimate that the Canadian Federal contribution was, at best, one two hundredth of the NSF supercomputer budget, and that a more realistic fraction is closer to one five hundredth. Given that the Ontario Centre for Large Scale Computing (OCLSC) supercomputer is now used by academic researchers literally from Vancouver Island to Newfoundland, with approximately one–third of the full capacity of the machine being consumed by researchers from outside Ontario, this question of support for the Centre, including remote access to it, may not be entirely inappropriate in this National Overview of the "science support climate" in Canada.

Further to this discussion on supercomputing support, NSERC sometime ago appointed Dr. Alan George, of the University of Waterloo, to undertake an evaluation of Research Computing in Canada. It is understood that Dr. George, in a report already delivered, has acknowledged the need to make plans for research networks with T3 capability.

The funding problem does not rest with NSERC, however. Arthur May, President of NSERC, over a year ago pointed out a basic problem when he noted that the average operating grant which NSERC could provide to researchers in Canada was only of the order of $25,000 or so, because of the very limited funds made available to NSERC for redistribution. The funding problem, it would seem, is much bigger than NSERC.

The project team has contacted many of the key players in the business of providing data network services in Canada, both in the university research environment and in the private sector. Throughout the country, and indeed, the world, there are many initiatives taking place in the form of development and marketing of new services; initiation of joint agreements between university groups, government, and the private sector; and detailed planning for major changes and upgrades in existing network arrangements.

2.2 CANADIAN R&D ENVIRONMENT

There has been much public debate over the state of science in Canada, the appropriate roles of government and the private sector in supporting research and development in Canada, and over the more general state of Canadian telecommunications policy.

On the issue of support for research and development in Canada, there has been much criticism of each other by the federal government and the private sector, each of which has regularly accused the other of failing to provide adequate support and funding. McPhail makes reference to Canada's "poor record in research and development":

> "Canada ranks eleventh among the world's nations in terms of R&D expenditures as a percentage of GDP. Less than 1.5 per cent of the Canadian GDP is directed toward research and development as compared with more than 2.5 per cent in the United States, Japan, West Germany, and Sweden. A comparison of Canada's R&D performance with that of Japan, the United States, West Germany, France, Sweden, the United Kingdom, and the Netherlands is striking (McMillan, p18, CMA Magazine, July/August, 1988). Canada ranks lowest in terms of gross R&D expenditures as a percentage of GDP, industry–funded R&D as a percentage of GDP, international patents granted per 100,000 inhabitants, number of scientists and

engineers in the labour force per 100,000 inhabitants, and number of technology intensive industries with a positive trade balance. The country is second lowest in terms of government–funded R&D as a percentage of GDP, higher education R&D as a percentage of GDP, and domestic patents granted per 100,000 inhabitants. Only with respect to government–performed R&D and the number of advanced university degrees awarded per 100,000 inhabitants does Canada approach the middle of the scale."

The failure on the part of Canadian industry to invest amounts comparable to other countries in R&D can be attributed to several uniquely Canadian factors such as:

- a preponderance of mature industries in resource extraction and primary processing which make up a large proportion of GDP but in which technological advances are not of primary importance;
- a high level of foreign ownership; research for these companies is generally done in the home country;
- a high level of concentration of ownership in industry, which reduces the competitive advantage of technological leadership; and
- a business, investment, and taxation climate relatively unfavourable to R&D.

This situation is likely to be further exacerbated in the future by increasing government funding constraints. Government funding for the National Research Council has recently been reduced, with an emphasis put on NRC carrying out industry–driven research funded by industrial beneficiaries (a significant change of direction for one of the largest R&D institutions in the country). Governments increasingly hobbled by expanding social programs and interest payments on massive debt are less able to justify expenditures on the long–term and ill–defined benefits of R&D. Similarily, funding for post–secondary education has been stagnant or dropping for many years, and many universities are reaching a crisis point in terms of being able to attract and provide facilities for promising researchers.

2.3 SCOPE OF THIS REVIEW

This review presents details on existing and proposed regional and national research networks in Canada. The following networks are discussed in this review:

I. Existing Regional Research/University Networks

1. AECLNET
2. AHENET
3. ARNET
4. BCNET
5. DRENET
6. HALIFAXNET
7. MANITOBANET
8. NBPEIECNET
9. ONET
10. RISQ
11. SUPERNET

II Existing National Reserach Networks

1. CDNNET
2. Netnorth

III. Future Regional Initiatives

1. MANITOBANET (Extended)
2. MEMORIALNET
3. NSTN

IV Future National Research Networks

1. CANET (NRCNET)

2.4 SUMMARY OF NETWORKS IN CANADA

Over the past several years, there have been many significant cooperative initiatives undertaken by various organizations across the country with regard to research networking. Several powerful networks and/or sub–nets have come into being (for example, BCNET in British Columbia, ONET in Ontario, RISQ in Quebec, and SUPERNET in Ontario, Quebec and Nova Scotia) and more are under development in Nova Scotia, Manitoba, Alberta and New Brunswick/PEI. In addition, key university computer/communications individuals from across the country have been meeting to organize their efforts in response to the CANET initiative (previously known as NRCNET), and to proceed with implementation once Treasury Board approval is secured.

Tables 2.1, 2.2 and 2.3 compare and contrast the characteristics of existing and future Canadian networks, both regional and national. Size, purpose, protocol used, capacity/speed, media used, acceptable use, management and administration costs and benefits as well as a number of other discriminators are summarized by these tables. The main conclusions are described below:

Purpose: The predominance of regional R&D networks in Canada suggests that their purpose and function is as varied as their number. Generally, the networks have evolved, in part, to meet specific requirements/capabilities for a group or organization whose needs were not being serviced by existing networks. The fundamental motives for the formation of new networks include such desires as access to supercomputers, data bases and libraries that existing networks do not or can not access, additional capacity and speed, and increased connectivity to researchers, remote facilities and to other R&D networks.

Several Canadian R&D networks appear to have evolved from the need to provide simple mail and messaging capability to its users, such as that provided by RISQ, CDNNET and NETNORTH. Connections to supercomputers such as the CRAY, the CYBER 900 or the CYBER 205 through NETNORTH, SUPERNET, ONET, and AECLNET, permit a multitude of analyses such as image analysis, remote sensing, climate research and expert systems analysis. Other networks, such as NBPEIECNET, are connected to local university computers. Service–oriented networks such as ARNET link research organizations in Alberta as well as provide gateways to other research and educational networks. AHENET consists of a series of data links connecting colleges, hospitals and oil companies (and their supercomputers), providing an inexpensive means of networking in the province. As the number of regional and national networks grow, the desirability to link to

TABLE 2.1
COMPARISON OF CANADIAN RESEARCH NETWORKS

EXISTING REGIONAL NETWORKS

NETWORK NAME	SIZE OF NETWORK	LINE SPEED Kb/s	PROTOCOLS	FIRST YEAR OF OPERATION	ACCESS	FUNDING	MANAGEMENT	DISCRIMINATING CHARACTERISTICS
AECLNET	4 Nodes	19.2 - 56	DECNET CDCNET	1987	Government - AECL, EMR	Atomic Energy Research $350,000/year	Atomic Energy Research	Star network that provides "wide-area LAN" for multiple locations. Vitalink Translans links local Ethernets.
AHENET	40 Hosts	9.6	DECNET	Not obtained	Alberta colleges, hospitals, oil exploration companies	Self-supporting through user fees.	Not obtained	Links VAXes. Inexpensive but limited capability. Plan to expand and improve.
ARNET	40 Hosts 30 - 40 Sites	56 T1 upgrades when feasible	TCP/IP ISO-OSI	1990	Alberta Research Council, University of Calgary, University of Alberta, UBC, Athabasca University, hospitals libraries and various private companies.	Provincial - Minsitry of Technology Research and Telecommunications.	Professional management as independent not-for-profit organization.	Replaces AHEN. Sevice oriented, not used to advance new network products. Interconnection with NSFNET is anticipated.
BCNET	8 Nodes	19.2 - 1,544	TCP/IP DECNET ISO (x.25)	1988	Open to any organization in B.C. engaged in education and research.	Start-up funds from Educational Research and Development Agency, with matching funds from federal and provincial sources; user fees - self-sufficiency expected soon.	BCNET	Star network on UBC with links to seven other sites. One of most advanced networks in Canada. Connected to NSFNET. BCNET private-dial SLIP service. Has full-time manager.
DRENET	over 45 Hosts 9 Sites	9.6 - 19.2; (56 Kbs link to NYSERNET)	TCP/IP (x.25 PDN)	1983	Defence Research Establishments, Military colleges	Federal - DRE	Under Contract by Software Kinetics Ltd.	Link to Internet, future links to CDNNET and NETNORTH
HALIFAXNET	over 80 Hosts	9.6 - 19.2	DECNET	In place for several years	Not Obtained	User fees	Not Obtained	Eventual replacement/absorption into NSTNET.
MANITOBANET	Not obtained	9.6	RJE; TCP/IP and Ethernet to be added	Not obtained	University of Manitoba Campuses and University of Winnipeg	Provincial	Not Obtained	Mail and remote job entry are supported on Manitoba Telephone System services. Plans to upgrade and expand the network.
NBPEIECNET	7 Nodes 6 Sites	19.2	RJE; TCP/IP links planned	1972	UNB and facilities at remote sites in New Brunswick, P.E.I., and Nova Scotia.	Provincial - Higher Education Commission	University of Moncton, College St. Louis Maillet Sippegon, Mount Allison, St. John Campus and St. Thomas.	Star network, connected to Nova Scotia Network via NETNORTH. Datapac access.
ONET	8 Nodes	56 - 112	TCP/IP some DECNET	1988	Currently 15 members: universities, centres of excellence, industry, High Energy Physics Network (HEPNET)	Start-up partly funded by the Ontario Centre for Large Scale Computation; user fees; space and power provided by members.	Management Committee composed of voting members paying full fee	Access to CRAY Supercomputer, and to Ontario Centres of Excellence. Connected to NSFNET; plans to merge with NETNORTH.
RISQ	9 Nodes	56	DECNET; TCP/IP with special techniques	1989	University community, Le Centre de Reserche Informatique de Montreal (CRIM), Industrial users.	User fees	McGill University and CRIM	Star network with hub located at CRIM. 56 Kb interconnections to CSNET and to NYSERNET. Possibility of sharing resources with SUPERNET.
SUPERNET	About 10 Nodes	9.6 - 56	Not obtained	1987	Federal government departments requiring use of CRAY super computer facilities (EMR, AES, DND)	User fees	Operating board comprised of representatives from each member organization	Star network, with links ranging from Datapak to Telesat Anikom satelite. May be absorbed into proposed CANET if T1 capable.

Source: [1] The Matrix; Computer Networks and Conferencing Systems Worldwide; John S. Quarterman; 1990.
[2] HICKLING Research and Personal Interviews.

TABLE 2.2
COMPARISON OF CANADIAN RESEARCH NETWORKS

FUTURE REGIONAL NETWORKS

NETWORK NAME	SIZE OF NETWORK	LINE SPEED Kb/s	PROTOCOLS	FIRST YEAR OF OPERATION	ACCESS	FUNDING	MANAGEMENT	DISCRIMINATING CHARACTERISTICS
NSTN	Over 100 industrial users interested	56	TCP/IP, OSI	Future	Research and development community; Nova Scotia universities and businesses in technology sector.	Technology Transfer and Industrial Innovation; ultimately self-sustaining.	Software Kinetics and Dalhousie University as a for-profit organization.	Technology the same as CANET; no plans to restrict the types of traffic.
MANITOBANET (Extended)	Not Obtained	56	IP	Future	See MANITOABNET under Existing Regional Networks New links to CITT, Fort Garry campus to UOM, Medical capuses, AECL in Pinawa and Agriculture Canada	See MANITOBANET under Existing Regional Networks	See MANITOBANET under Existing Regional Networks	Motivated in part to provide medical researchers access to Magnetic Resonance device proposed for CIIT location.
MEMORIALNET	Not Obtained	Not Obtained	TCP/IP DECNET	Future	Academic institutions, government agencies and Newfoundland and Labrador Computer Services crown corporation.	Discussions underway with Atlantic Canada Opportunities Agency.	Not Obtained	High-speed metropolitan network in St. John's, Nfld. Campus network at Memorial University which will run over fibre optic links.

Source: [1] The Matrix; Computer Networks and Conferencing Systems Worldwide; John S. Quarterman; 1990.
[2] HICKLING Research and Personal Interviews.

TABLE 2.3
COMPARISON OF CANADIAN RESEARCH NETWORKS

EXISTING AND FUTURE NATIONAL NETWORKS

NETWORK NAME	SIZE OF NETWORK	LINE SPEED Kb/s	PROTOCOLS	FIRST YEAR OF OPERATION	ACCESS	FUNDING	MANAGEMENT	DISCRIMINATING CHARACTERISTICS
CDNNET	175 Hosts 5,250 Users	2.4	EAN x.400; TCP/IP and TTXP support in some links.	1981	open	User-fees; NSERC funding ended in 1988	Administered by CDNNET HQ at UBC; policy set by committee of one rep from each member organization.	Star network providing reliable mail service as primary function; first network to use x.400 protocol; electronic mail bridge to OSI networks in Europe and Asia; gateway to non-ISO research networks.
NETNORTH	167 Nodes 57 Institutions	2.4 - 19.2 Plan to 56	NJE	1982	Primarily Universities	User-fees; start-up funding provided by IBM	Committee of Executive Directors elected by member organizations.	Canadian component of BITNET/EARN international network; main gateway to European Networks; plan upgrade to 56Kb between Cornell and University of Toronto with direct connections to NSFNET; likely to evolve into proposed CANET.
CANET (formerly NRCNET)	1,000 Users initially 30,000 Users by 1991; 5 - 9 Nodes	56 - TI	TCP/IP, OSI	Future	Regional Networks - institutions that have a research or eductation interest.	Funding from Regional Networks that are connected to CANET	University of Toronto and IBM consortium controlled by a management board made up of representatives from each of the regions.	A national high-speed backbone would link various existing regional networks; similar to NSFNET in terms of protocols and implementation; viewed as as extension of INTERNET; NETNORTH or CDNNET would join or be replaced by CANET.

Source: [1] The Matrix; Computer Networks and Conferencing Systems Worldwide; John S. Quarterman; 1990.
[2] HICKLING Research and Personal Interviews.

other networks increases. In Canada, connections to regional and international networks are supported by NETNORTH, CDNNET and CANET (formerly NRCNET).

Protocol: The most common protocol in Canada is TCP/IP, with DECNET a close second. Eventual migration to OSI/ISO will be considered for several existing and future networks (NSTN and CANET). BCNET, one of the most advanced regional networks in Canada, uses all three protocols.

Capacity/Speed: Typically, regional R&D networks cannot support speeds greater than 56 Kb/s. The common speeds range from 9.6 to 19.2 Kb/s, with a few offering speeds up to 56 Kb/s (AECLNET, RISQ, SUPERNET, and NETNORTH planning to go to 56). BCNET offers speeds up to T1, and ONET has a connection to NSFNET through Cornel at 112 Kb/s. All future networks plan speeds of at least 56 Kb/s.

Media: Most transmission media for R&D networks in Canada are leased lines from telephone companies, many having links to Datapac. Only the more advanced networks, such as BCNET and SUPERNET use fibre optic lines (on relatively short links), or satellite media (SUPERNET has a 56 Kb/s line to a Telesat Anikom satellite link from Valcartier to Dorval). Some future networks, such as MEMORIALNET, are proposing use of fibre optic links. If T1 speeds are deemed economically feasible for CANET, then a national fibre optic line would become necessary.

Acceptable Users: Membership to existing national and regional research networks tend to concentrate on universities, research councils and centres of excellence. Some networks, such as AECLNET and DRENET and SUPERNET are primarily government networks used to satisfy specific departmental requirements. Several networks serve specific research areas, such as energy (oil exploration) and medicine (hospitals). Networks with available capacity, and those who are planning to expand capacity, are generally interested in extending membership to the private sector, if they already have not done so. In such cases, the private sector usually must be conducting work in education or research, and be interested in technology transfer. Unlike many U.S. R&D networks, access to several Canadian networks is also provided to commercial enterprises satisfying certain criteria.

Fee Structure: The fee structure is about as diverse as the number of networks. BCNET charges according to transmission speed used, and does not vary by frequency of transmission or connect time. RISQ has a similar fee structure to this, but differentiates between university and non–university members. Another fee structure, typified by ONET, is one of a one time "start–up" fee plus a variable tariff based on utilization. CDNNET charges are based on the type of user (university, commercial) and the number of computers accessed. Other networks are funded through other means.

Management and Administration: The management structure varies, but the most common arrangement consists of a management committee (or operating board, steering committee or directorship) composed of representation from each full–fee paying member organization. Networks managed in such a manner include ONET, SUPERNET, CDNNET and NETNORTH. NBPEIECNET is managed by a group of directors from the various member universities and colleges. DRENET is managed under contract with a private firm, as will be the proposed NSTN.

Funding: Other than through user fees, networks used for research are often funded by university computer/communications departments. In such cases the direct costs may be buried in existing operating budgets, or other resources are provided to the network in the form of free labour to design and/or manage the network. Funding is also often provided by provincial government or a related agency. NSERC provides some assistance (on average $25,000 per grant), but is limited by the funds made available to it. Occasionally, private sector will contribute to the start–up, operation or management of a network (IBM provided some funding for NETNORTH, which ended in 1987).

Costs: Costs are much higher in Canada than in the United States, primarily due to the much higher communication rates here. This is, in part, a consequence of regulation and Canada's system of telecommunication monopolies. Typically, costs are approximately five to eight times higher than in the U.S. When compared to fees paid by large users, the rates may be more than 10 times higher in Canada.

Beyond the start–up period, annual operating costs tend to lie in the $350,000 to $500,000 range, with some as low as $200,000 (ONET and BCNET) and others as high as $1,000,000 (NETNORTH). The new CANET is expected to require an annual budget of about $1.2 million.

2.5 EXISTING REGIONAL NETWORKS

This section describes the characteristics of the existing regional networks in Canada.

2.5.1 AECLNET

Background

- Started in 1987 to link Atomic Energy Research (AECL) offices
- Star network with hub in Ottawa and connections to Chalk River, Sheridan Park (Mississauga), and Pinawa, Manitoba
- Provides "wide–area LAN" for the multiple locations and access to CYBER 900 computer in Chalk River

Technical Design

- 56 Kb satellite circuits link Ottawa with Chalk River and Pinawa using Telesat Anikom 500 service, and backed up by 19.2 Kb landline circuits, plus a 19.2 Kb landline from Ottawa to Sheridan Park.
- Supports DECNET, CDCNET, and are ready to convert to TCP/IP when SUPERNET is converted
- Vitalink Translans links local Ethernets

Implementation Design

- Network administered and funded by Atomic Energy Research, a subsidiary of Atomic Energy of Canada Ltd., a crown corporation.

- Close links to government include SUPERNET link between Energy Mines and Resources office and AECL office in Ottawa.
- Current budget for data communications is about $350,000/year.

2.5.2 AHENET: Alberta Higher Education Network

(Also known as AHEN, ADHOCNET and various other names)

- Developed as a series of data links between colleges, hospitals, oil exploration companies etc. in south and central Alberta over the past several years.
- Network is self–supporting, and consists of 9600 baud lines linking VAXes under DECNET.
- Access to supercomputers at Calgary oil companies available.
- Network is considered to be an inexpensive, but necessarily limited, solution to some of the province's networking needs.
- Major plan to expand and improve is described under ARNET

2.5.3 ARNET

Background

- Network recently implemented January 90 (Phase 1) linking the Alberta Research Council (ARC) in Edmonton, ARC in Calgary, the University of Calgary, the University of Alberta in Edmonton and UBC via 56 Kb lines. Phase 2 involves a larger provincial wide effort providing an opportunity for the private sector to be involved;
- ARNET will essentially replace AHEN, the existing regional network.
- The five–year business plan called for total annual expenses growing from $540K to $960K over the first five years, and seeks a provincial contribution of $1.5 M over the five years (Ministry of Technology, Research and Telecommunications).
- The goal of the network is to carry throughout the province between research organizations and to provide gateways to other provincial, national and international research and educational networks. Becoming an integral part of a national network and providing industrial researchers access to educational institutions are also important goals;
- The network has a service–orientation and will not be used to advance the development of new state–of–the–art network products.

Technical Design

- The network uses TCP/IP and will move to ISO–OSI when appropriate.
- Initially Two 56 Kbps installed between Edmonton and Calgary, one connected through Red Deer to form a redundant loop;
- T1 upgrades will be implemented when available and feasible

ARNet Phase I Topology

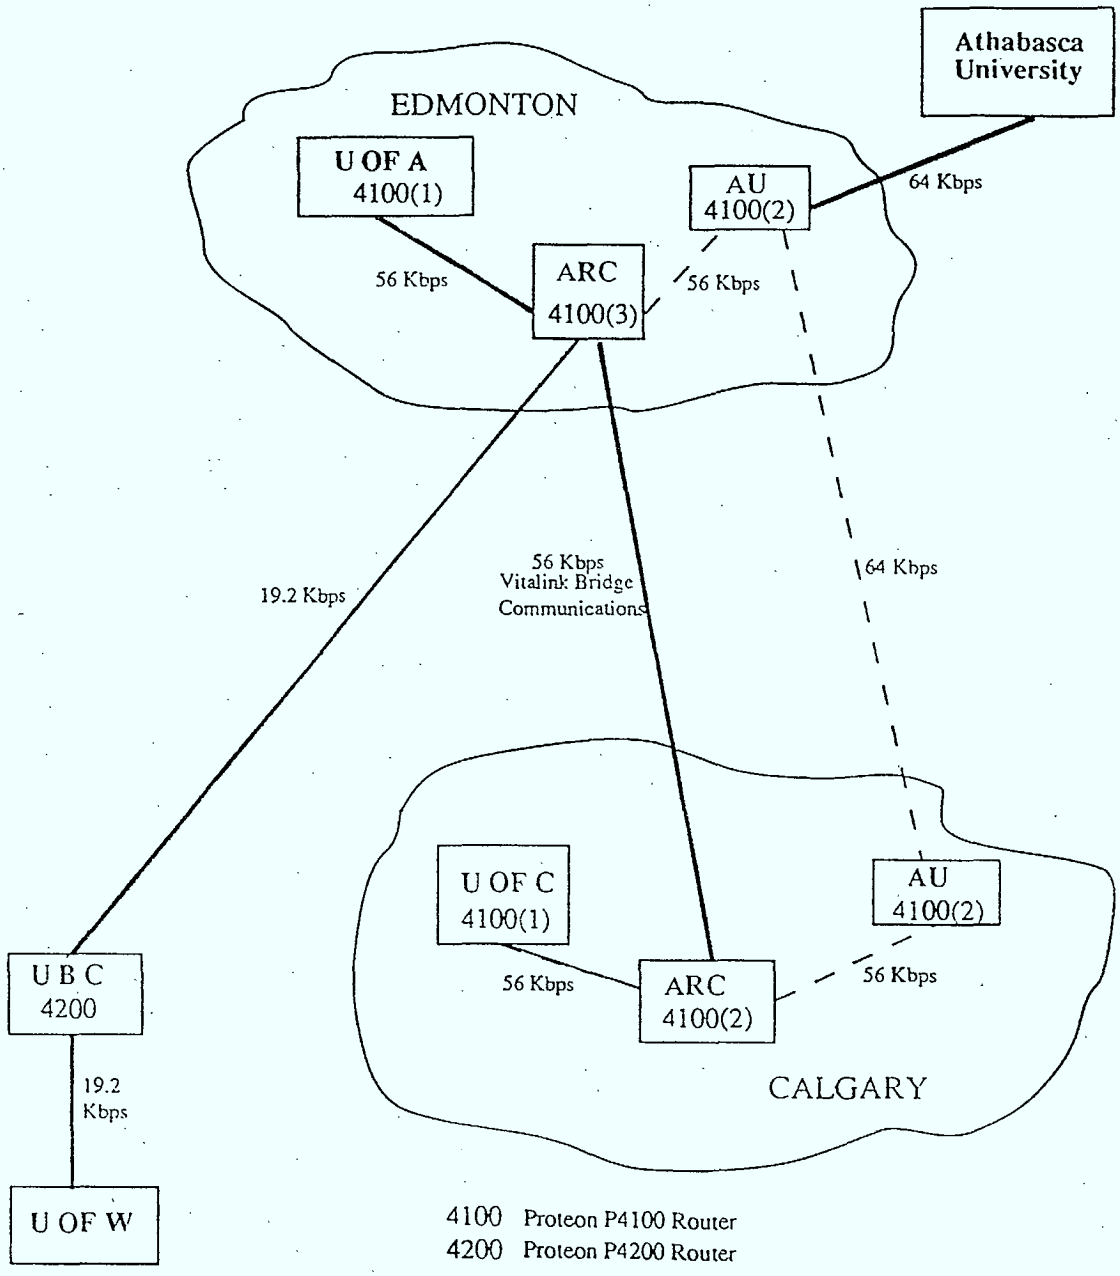

- IMPLEMENTED FEBRUARY 1990

Source: The Creation of a Regional Netowrk for Alberta; Prepared by The Alberta Research Council; November 23

Proposed Alberta Research Network

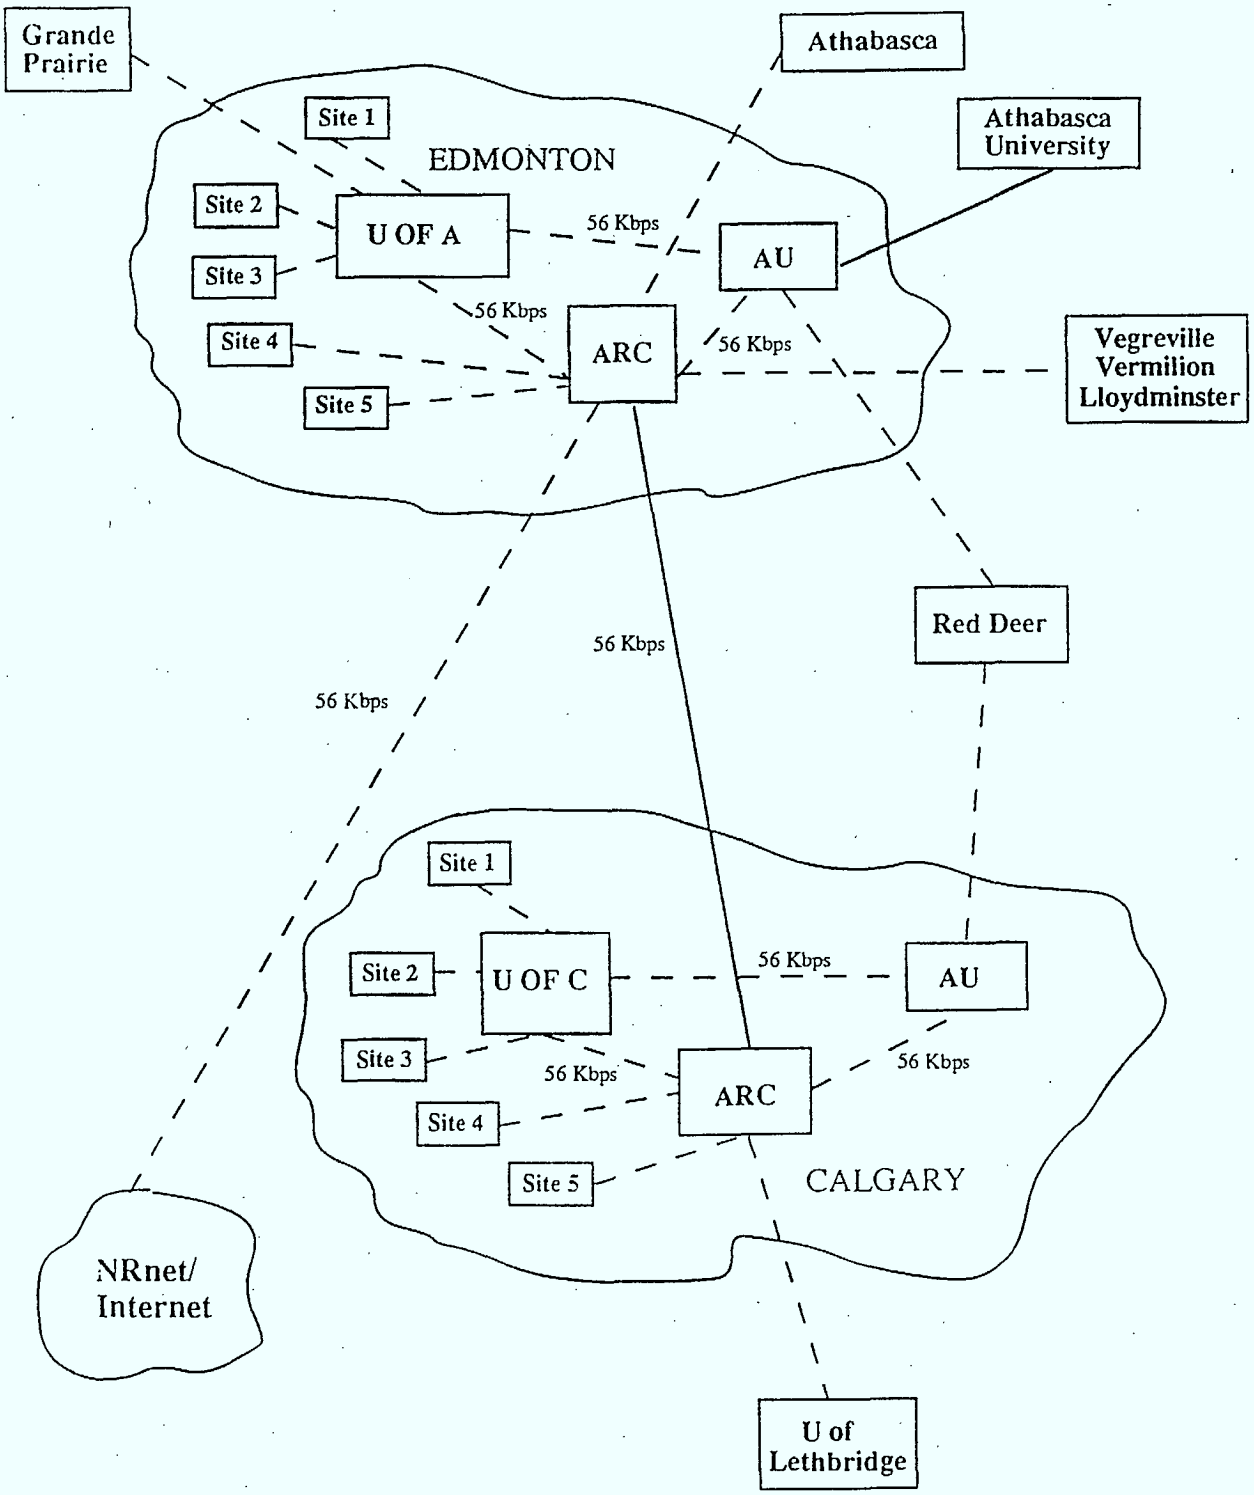

Source: The Creation of a Regional Netowrk for Alberta; Prepared by The Alberta Research Council; November 23

Implementation Design

- Membership will be extended to Athabasca University and possibly to Myrius (super computer manufacturer) and other interested industrial partners. An interconnection to NSFNET is also anticipated through BCnet. The network will be extended to certain hospitals (endorsed by Alberta Cancer Board), libraries, and various private organizations involved in research, reaching perhaps 30–40 sites in total;
- Strong interest has been shown by industry (Convex Computers Canada Ltd., Nova Husky Research Corporation, Byrius Research Corporation and DMR);
- Professional management is expected set up as an independent not–for–profit organization;

2.5.4 BCNET

Background

- One of the most advanced regional networks in Canada;
- Officially opened June 1988;
- Links UBC, University of Victoria, Simon Fraser University, Tri University Meson facility, Microtel Pacific Research, Dominion Astrophysical Observatory and the Advanced Systems Institute. (Charter members include UBC, Uinversity of Victoria, Simon Fraser University and TRIUMF)
- Started with funding through Educational Research and Development Agency agreement providing matching federal and provincial funds

Technical Design

- Star network centred on University of British Columbia with links to seven other sites in Lower Mainland and Vancouver Island
- Links range from 19.2 Kb/s line to T1 fibre optic link and 1.544 Mb/s T1 broadband link
- Connected to NSFNET via link to University of Washington;
- Also connected to CDNnet.
- TCP/IP, DECNET, and ISO (X.25) are all used; Vitalink bridges and Proteon routers also are used.
- BCnet private–dial SLIP service is an inexpensive alternative to access the Internet. It is designed for users who do not need the capacity and the response time of a high speed link.

Implementation Design

- Policy set by BCnet steering committee. Membership is open to any organization in B.C engaged in education and research;

BCnet Configuration - May 1/89

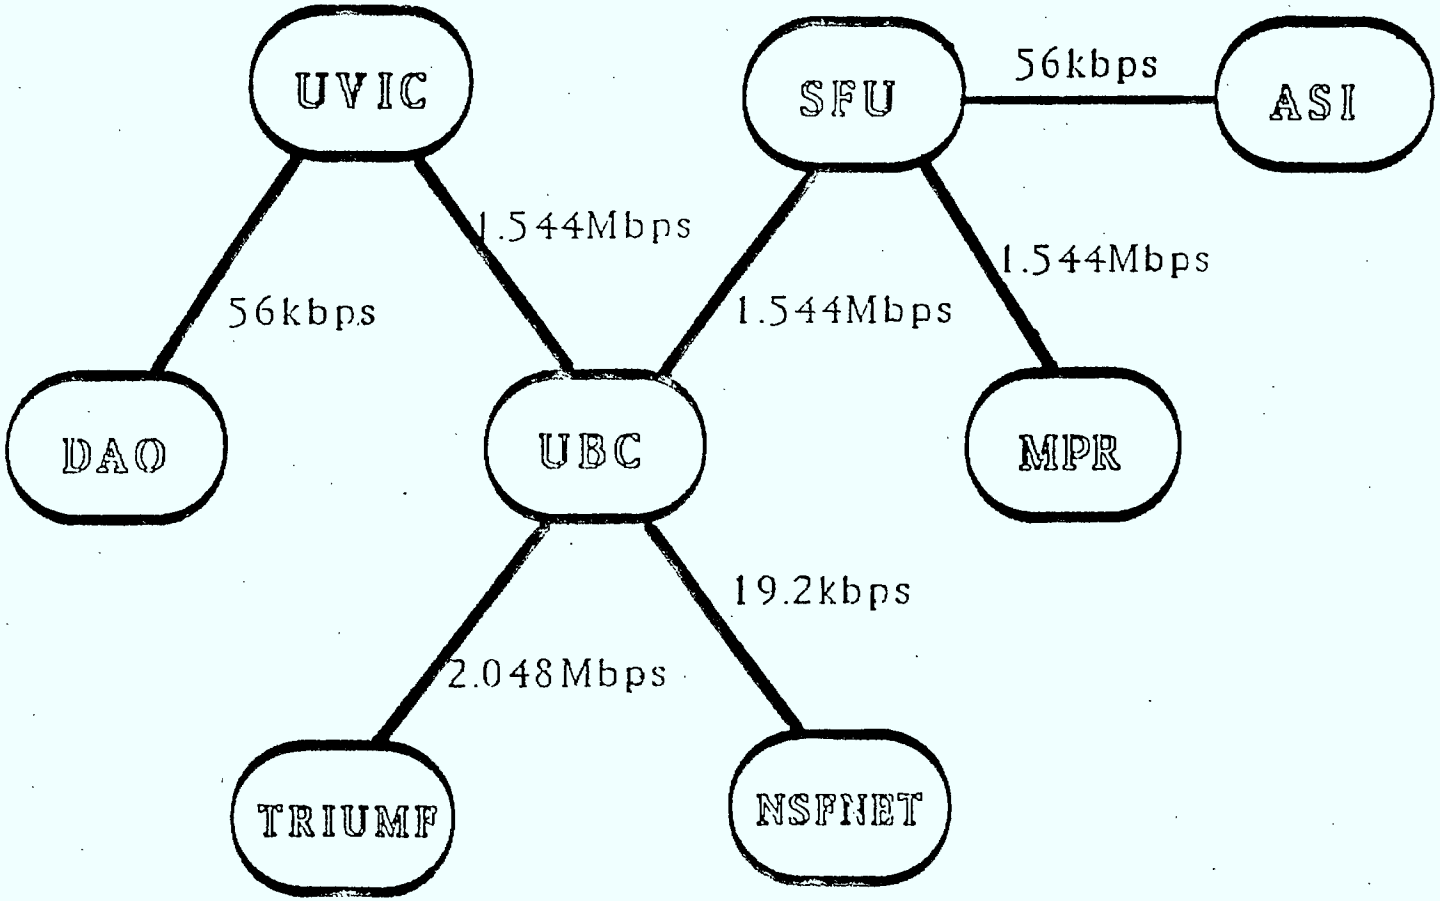

Source: The Creation of a Regional Network for Alberta; Prepared by the Alberta Research Council; November 23.

BCnet Pie Chart

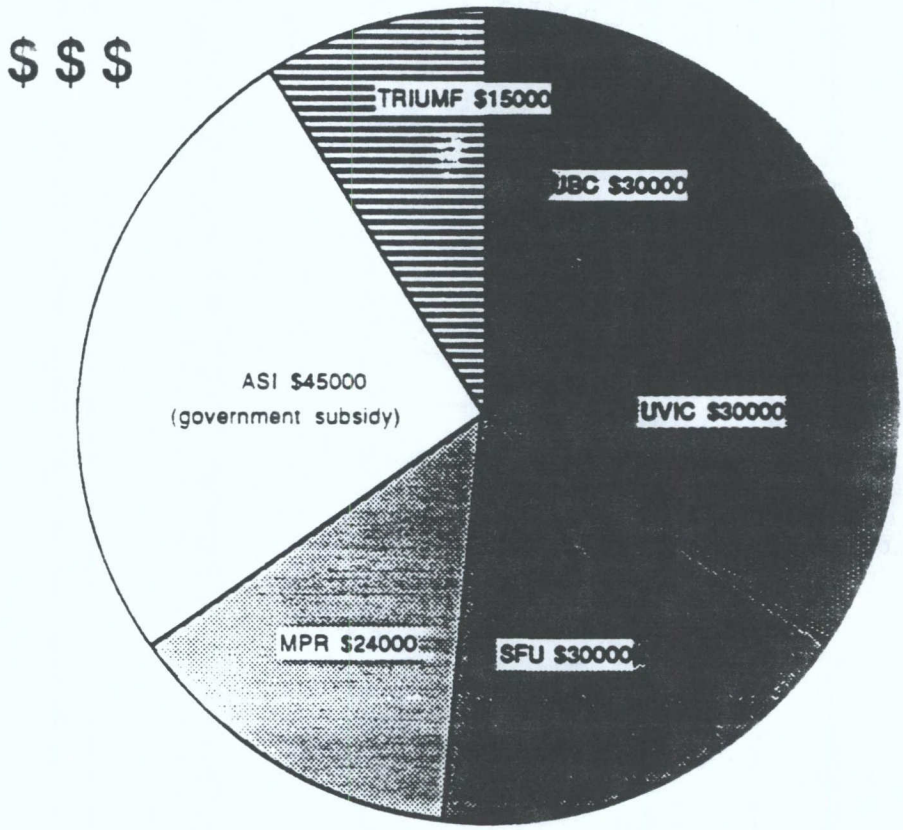

Division of 1988/89
BCnet Operating Costs
(estimated at $177,000)

Source: The Creation of a Regional Netowrk for Alberta; Prepared by The Alberta Research Council; November 23

- BCnet headquarters responsible for overall management and day–to–day operation of the network and the central BCnet Hub.

- Rates are set by BCnet steering Committee and are reviewed annually. The costs associated with BCnet HQ operation, leasing communication lines, maintenance of central and secondary hubs, and external gateways are shared by all members.

- Self–sufficiency expected soon; members to be charged about $30,000/year to support annual budget of $200,000

- Connection costs vary between $15K – $45K with major universities paying $30K; and there are no per–transmission charges. Connection fees depend solely on connection speed. The fee structure is contained below:

Speed (bps)	Annual Charge
4,800	$ 640
9,600	$ 1,280
19,200	$ 2,570
38,400	$ 5,140
56,000	$ 7,500
112,000	$ 15,000
224,000	$ 30,000
1,544,000	$ 30,000

- Several new members are expected soon, including community colleges, teaching hospital, and industries;

- Unique among existing networks in Canada in that it has had a full–time manager since July, 1989.

2.5.5 DREnet

Background

- Begun in 1983 to link Defence Research Establishments in Ottawa and Dartmouth;

- Internet connecting sites, systems, and LANs at DREs in nine sites with over 45 hosts;

Technical Design

- Long–haul network is called XDRENET; uses TCP/IP over X.25 PDN with a mix of 9.6 kB and 19.2 kB links;

- Link to Internet through a 56 kB link to NYSERNET;

Implementation Design

- Network managed under contract by Software Kinetics Ltd.;

- Future plans include addition of two more DRE sites plus the three Canadian military colleges, and direct connections to CDNnet and NETNORTH;

2.5.6 HALIFAXNET

- A DECNET network in the Halifax area that has been in place for several years;
- There are more than 80 hosts;
- Links the five universities and two industrial sites in the Halifax area, through a mix of 9600 bps and 19.2 Kb lines;
- The network is scheduled for replacement/absorption into NSTN (Nova Scotia Technical Network) under a contract which has just been awarded;

2.5.7 MANITOBANET

- A multi–line 9600 baud network which links the University of Manitoba campuses at Winnipeg (Fort Garry) and Brandon, and the University of Winnipeg.
- Mail and remote job entry are supported using X.25 and RJE on Manitoba Telephone System services.
- Annual operating costs are about $200,000 for communications lines and $200,000 for computing operations
- Network is funded by grants from the Provincial government
- TCP/IP and Ethernet support will soon be added to the University of Manitoba mainframe
- Other organizations in the province are participating in plans to expand and upgrade the network

2.5.8 NBPEIECNET

Background

- Started in 1972 to provide remote access to computer facilities at University of New Brunswick. This objective of accessing facilities at remote sites is still very important;
- Low speed network connecting six universities;

Technical Design

- This star network provides 19.2 Kb/s interactive access (Bisync, SNA) from six sites in New Brunswick and PEI to the IBM 3090 Vector Facility "Supercomputer" at UNB;
- The 19.2 Kb lines are multiplexed providing RJE and asynchronous services via modems;
- Connected to Nova Scotia Network via Netnorth;
- Datapac access is provided;

Implementation Design

- Funding is provided by the Higher Education Commission with a budget of $500K per year for computer charges, telecommunication costs, maintenance and hardware upgrades;
- Managed by a group of directors from University of Moncton, College St. Louis maillet Sippegon, Mount Allison, St John Campus and St Thomas;
- Circuits are leased from the telephone companies, at a cost of approximately $1200/month each;
- Costs are paid for directly by the provinces;
- Plans are being developed for a TCP/IP link between the two provinces to take the interactive load within the next six to eight months;
- It is also expected that UNB will be a node on CANET;

2.5.9 ONET

Background

- Organized as a cooperative solution to meet the needs among universities and other organizations in Ontario; began in sping of 1988.
- Grew from the need of researchers belonging to the Centres of Excellence Program (primarily ITRC and ISTS)
- Sponsoured by the Ontario Centre for Large Scale Computing — operators of the CRAY supercomputer at the University of Toronto;
- One major function is to provide access to the CRAY supercomputer at the University of Toronto; another is access to resources available at the Ontario Centres of Excellence;

Technical Design

- Network recently updated to 56Kbps with 112 kbps to the NSFNET node at Cornell;
- A 56 kbps link to RISQ in Quebec;
- Main protocol is TCP/IP although some DECNET functionality is supported at the University of Toronto;

Implementation Design

- Currently 15 memebers which includes six university members (U of T, York, Ottawa, Waterloo, McMaster and Western); two Centres of Excellence which have industrial affiliates like Motorola and Hewlett Packard; as well as industry (Gandalf, Royal Trust Technology Evaluation Centre, and BNR); and the High Energy Physics network (HEPNET) has also joined;

ONet

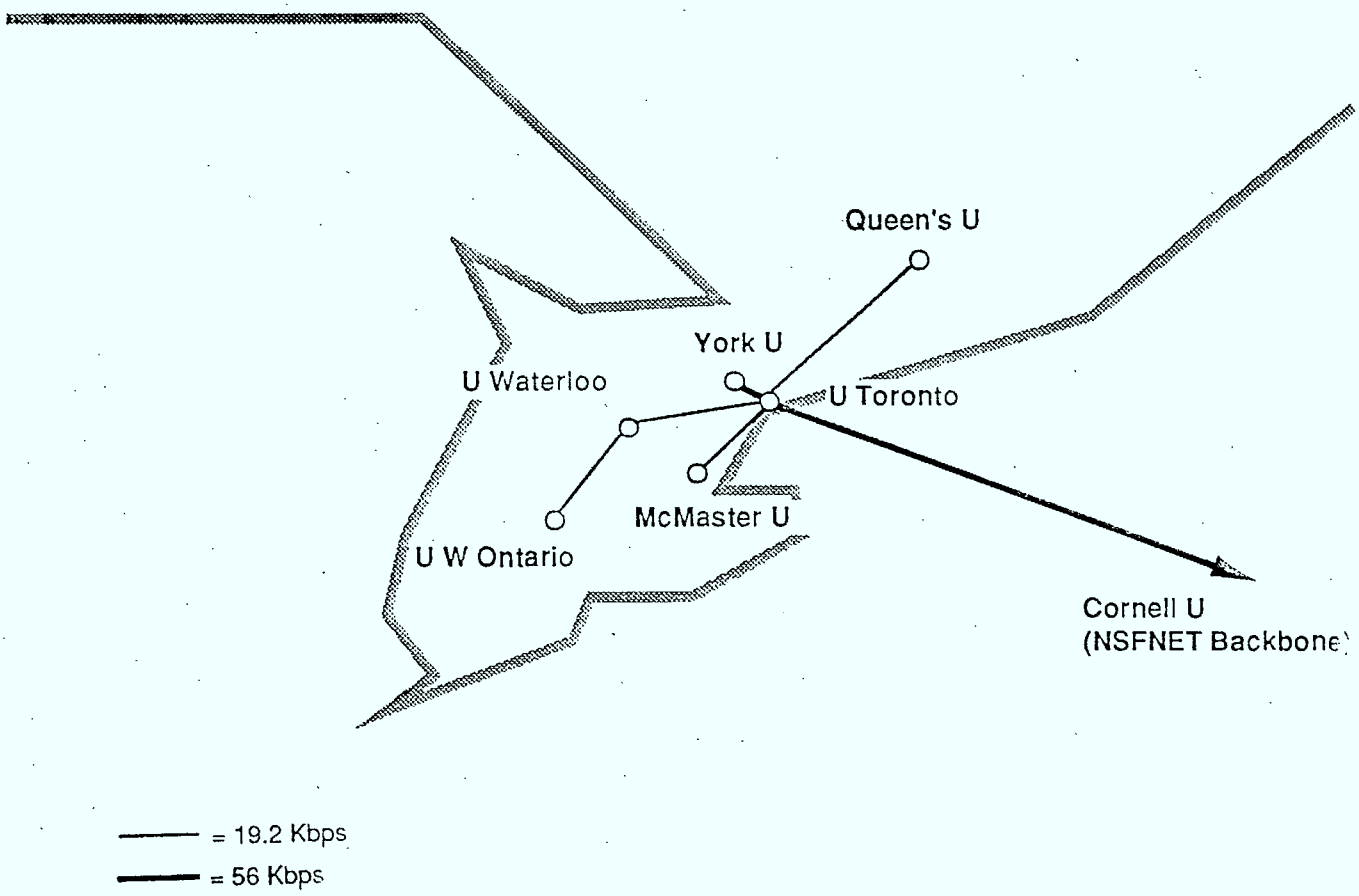

Source: Users' Directory of Computer Networks; Compiled by Tracy Lynn LaQuey, University of Texas; July 1989.

ONet: Summer 1989 (projected)

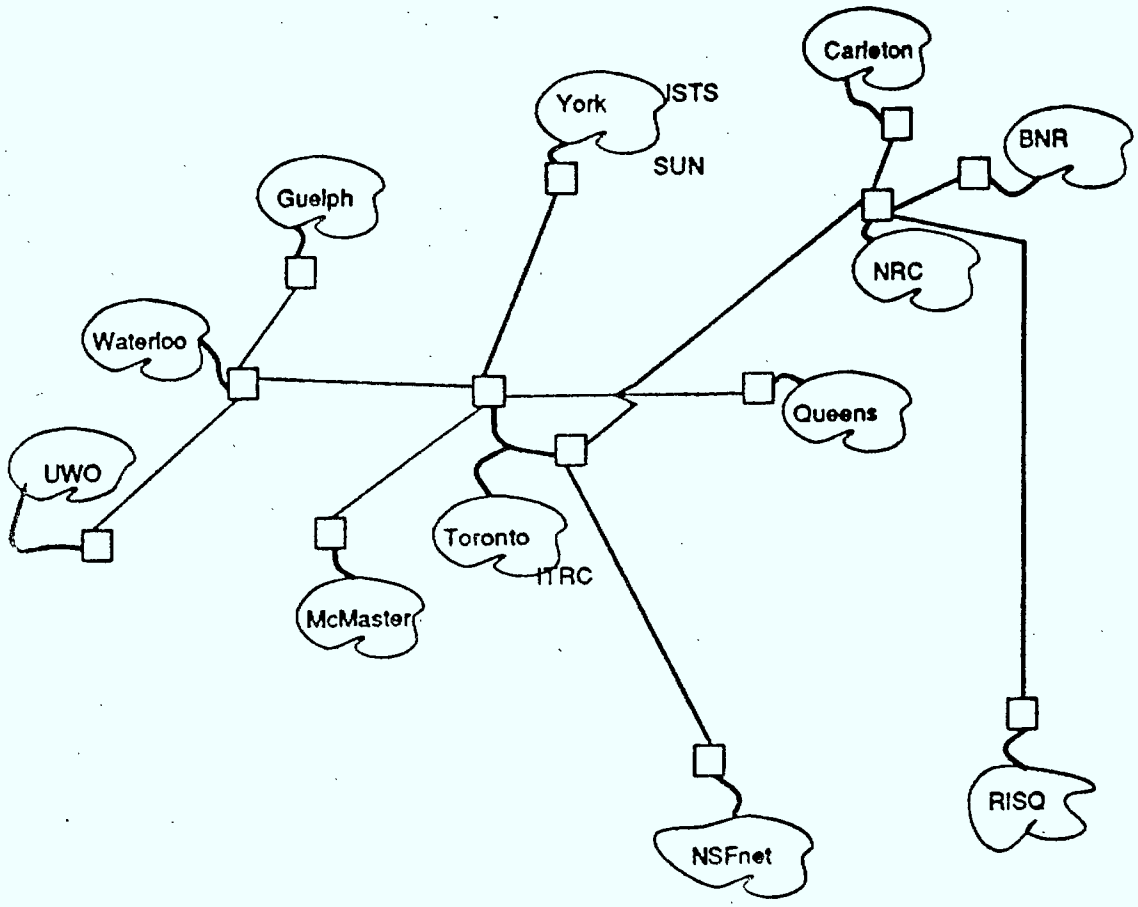

Source: Focus; Spring/Summer 1989; Volume 4 No. 2.

- Funded by contributions from each member of approximately $18,000/year — ie. flat fee. This cost includes the startup cost of buying a proteon router (about $8,000 each). Industrial partners are charged the same fee as universities;
- Some users who have high utilization (ie. special case) pay an increment above the $18,000 fee (ie. the super computer facility pays $60,000 and HEPNET pays $45,000).
- Members provide space and power for required routers; indirect costs are buried in departmental operating budgets;
- Startup costs were approximately $100,000, part of which was funded by the Ontario Centre for Large Scale Computation, the operating authority for the CRAY X–MP;
- Management Committee (composed of voting members paying full fee) handle all operational concerns. A separate committe has been set up to handle technical issues;
- Annual costs are estimated at $210,000 for lines, maintenance etc.;
- Plans are underway to merge with NetNorth; also several other institutions and private companies are considering membership;

2.5.10 RISQ

Reseau Inter–ordinateur Scientifique Quebecois

Background

- 56 kbps network operational since July 1989 between university community (University of Montreal, McGill, Universite du Quebec, Concordia, Laval, Sherbrooke and UQAM), Le Centre de Recherche Informatique de Montreal (CRIM) and Environment Canada's supercomputer in Doral;
- McGill and University of Montreal initiated the creation of the network out of the need to attain higher bandwidth and as a result of pressure from researchers.
- A feasibility study was not carried out because the committee felt that the need for such a network was so large that the benefits were obvious (ie. a matter of catching up);

Technical Design

- A star network with the hub located at CRIM, Centre de Recherche Informatique de Montreal;
- Services include mail, file transfer, and remote login;
- Nine nodes are connected in the Montreal area;
- 56 Kb leased lines are used;
- The basic protocol is DECNET (original network DECNET connecting computer science departments at universities), but TCP/IP is also handled using special encapsulation techniques and Digital Ultrix systems which run both protocols;

RISQ - Réseau interordinateurs scientifique québécois

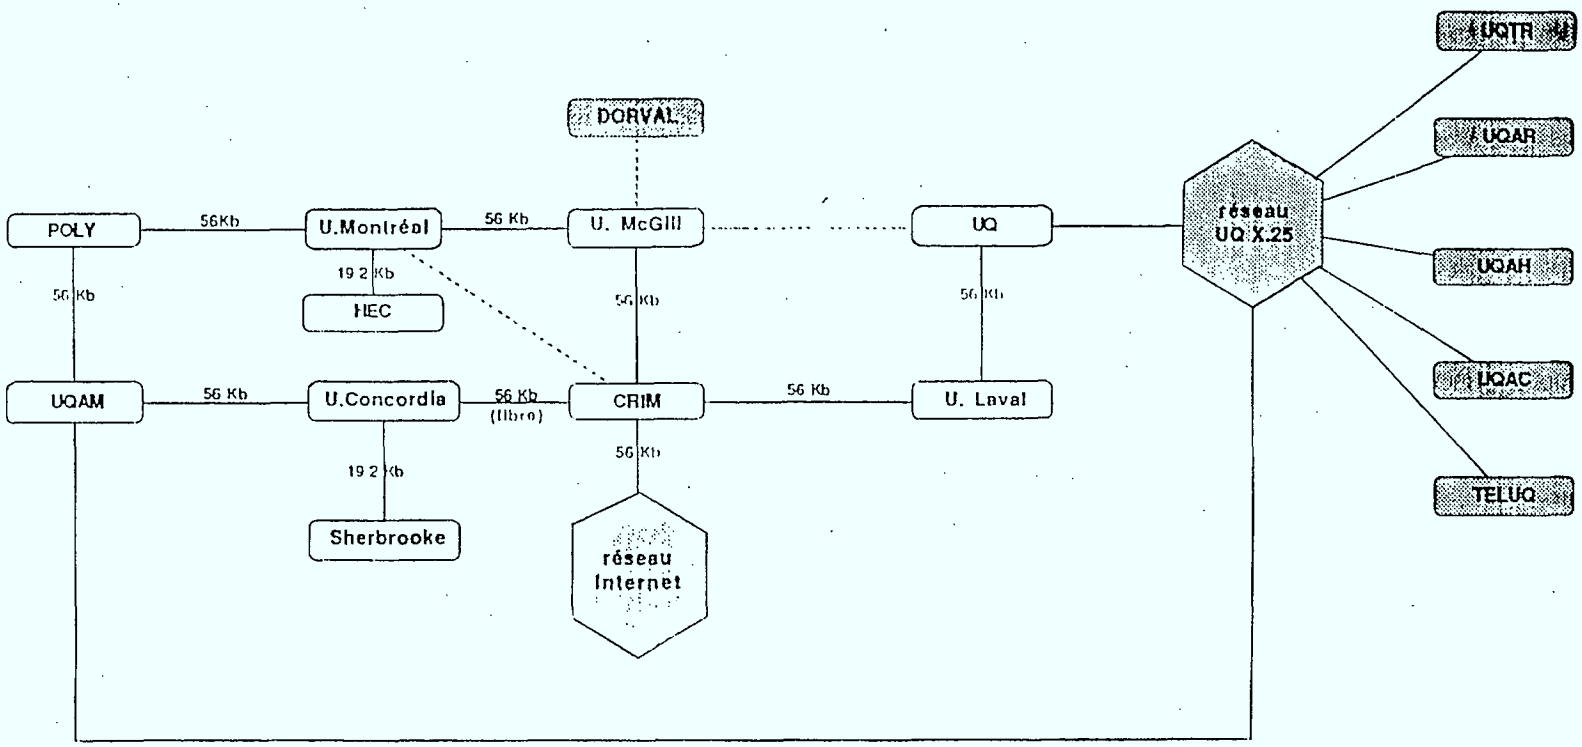

Source: The Creation of a Regional Network for Alberta; Prepared by the Alberta Research Council; November 23.

- There are 56 Kb interconnections to CSNET at Cambridge, Mass., to NYSERNET, and to Onet in Toronto

Implementation Design

- RISQ is a joint venture of CRIM and the universities. There is a board which is the final decision maker with respect to policy;
- McGill University contributes the Network Operations Centre; CRIM contributes the Network Information Centre, together they share management responsibilities;
- Funding comes from membership fees, set as follows:

  $15,000/year for CRIM members
  $ 6,250/year for 9600 bps line, non–university member
  $12,500/year for 19.2 Kb line, non–university member
  $25,000/year for 56 Kb line, non–university member

- The network owns all the equipment;
- Total operating costs are estimated at $2 Million over three years;
- RISQ and the SUPERNET consortium are working out arrangements to share resources of the two networks;
- Industrial users have expressed an interest in joining. The Hydro Quebec Research Department is currently the only industrial user;

2.5.11 SUPERNET

Background

- Formed in 1987 primarily to link Federal government users with a CRAY computer run by Atmospheric Environment Canada which had excess capacity
- Responsible for providing supercomputer services at wholesale rates and the provision of accounting and security reporting services as well as support information and technical assistance;
- Original members were Energy, Mines, and Resources Canada, Environment Canada, and Department of National Defence. Represents a consortium of federal government participants who have a common need to gain access to super computer facilities.
- Supernet being used for image analysis, analysis of remote sensing data, climate research, water flow modelling, expert systems etc. which requires high bandwidth communications technology

Technical Design

- Star network based on the AES facility at Dorval;
- About 10 nodes connected in Ontario, Quebec, and the Maritimes;

SUPERNET INFRASTRUCTURE

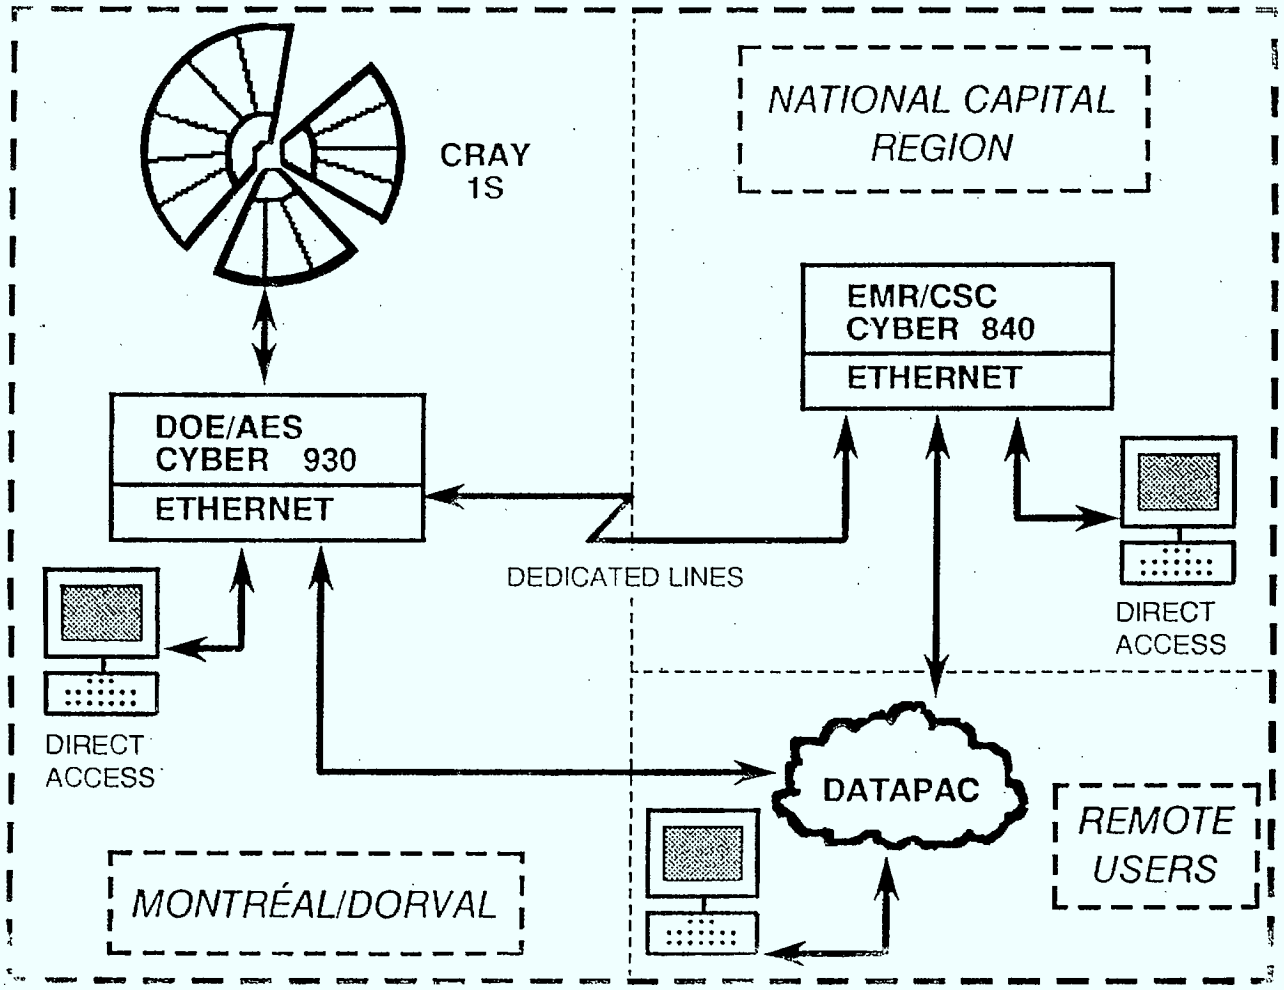

Source: Government of Canada; Supernet: Linking you to SuperComputers.

- Links range from Datapak connections through 56 Kb lines to a Telesat Anikom satellite link from Valcartier, Quebec to Dorval;
- Supernet Operating Board comprised of representatives from each member organization;

Implementation Design

- Supernet provides end user support through cooperative arrangements between computer organizations;
- Departments become Supernet participants by implementing a connection with the network and negotiating an agreement with a supplier;
- It is assumed that the network will be absorbed into the proposed CANET if it provides T1 speeds;

2.6 EXISTING NATIONAL NETWORKS

This section describes the characteristics of the existing national networks in Canada.

2.6.1 CDNNET

Background

- Started in 1981 as an NSERC–supported research project into an electronic messaging system, for educators and researchers, conforming to international standards which have become known as X.400;
- Initial version of the EAN software system released to selected test sites to support communication with each other in 1983; first network in world to use X.400 mail protocols;
- Software used to handle messages and file transfer meets international standards as promoted by CCITT and ISO.
- Main service is mail (4,500 messages per day); main users 60 % research, 30 % student, 10 % admin;
- Coverage is national, number of hosts estimated as 175 in August 1988

Technical Design

- Has implemented sophisticated X.400 electronic mail services (includes delivery notification, non–delivery notification and deferred delivery) to universities and research centres;
- Directory service for locating other CDNnet users also available;
- Electronic mail bridge to OSI networks in Europe and Asia;
- Gateway to non–ISO research networks (CSNET, UUCP, Bitnet and the Internet);
- Format is star with extra connections; the hub is at UBC (supported by UBC);

- Most long–distance links are X.25 at 2.4 Kb using Datapac.
- Besides EAN, TCP/IP and/or TTXP are supported in some links
- Reliability very high within network, mail delivery usually within minutes

Implementation Design

- Administered by CDNnet Headquarters at UBC;
- Has over 30 memebres from coast to coast;
- Policy set by committee of one rep from each member organization; executive elected by reps;
- Annual budget for operation of the network is approximately $250,000, primarily to cover staff expenses;
- Annual dues range from $1,320 (University/College – one computer) to $19,800 (commercial – multiple computers). These dues include software (EAN) licence fees for universities and colleges, but not for commercial purposes;
- CDNnet considers the reduction of dues for small organizations;
- NSERC funding ended in 1988

2.6.2 NETNORTH

Background

- The NETNORTH consortium and network started in 1982 as a part of the plan for an Ontario Universities Network (OUnet); however membership expanded outside Ontario;
- It is the Canadian component of the BITNET/EARN international network and shares the same technology and objectives;
- Originally intended as a store–and–forward message/mail network for academics, it has also successfully been used for mailing jobs to supercomputers such as the CRAY X–MP at the Ontario Centre for Large–Scale Computation and the CYBER 205 at the University of Calgary;
- Membership stands at 57 institutions with 167 NJE nodes at August 1988;
- For Canadian universities Netnorth is the main gateway to European Networks;

Technical Design

- Store and forward network — which means that bottlenecks may occur;
- Netnorth and Bitnet are joined by a link between the University of Toronto and Cornell University;
- Link speeds range from 2400 bps to 19.2 Kb; arranged in a tree structure

NETNORTH TOPOLOGY
As of February 6, 1989

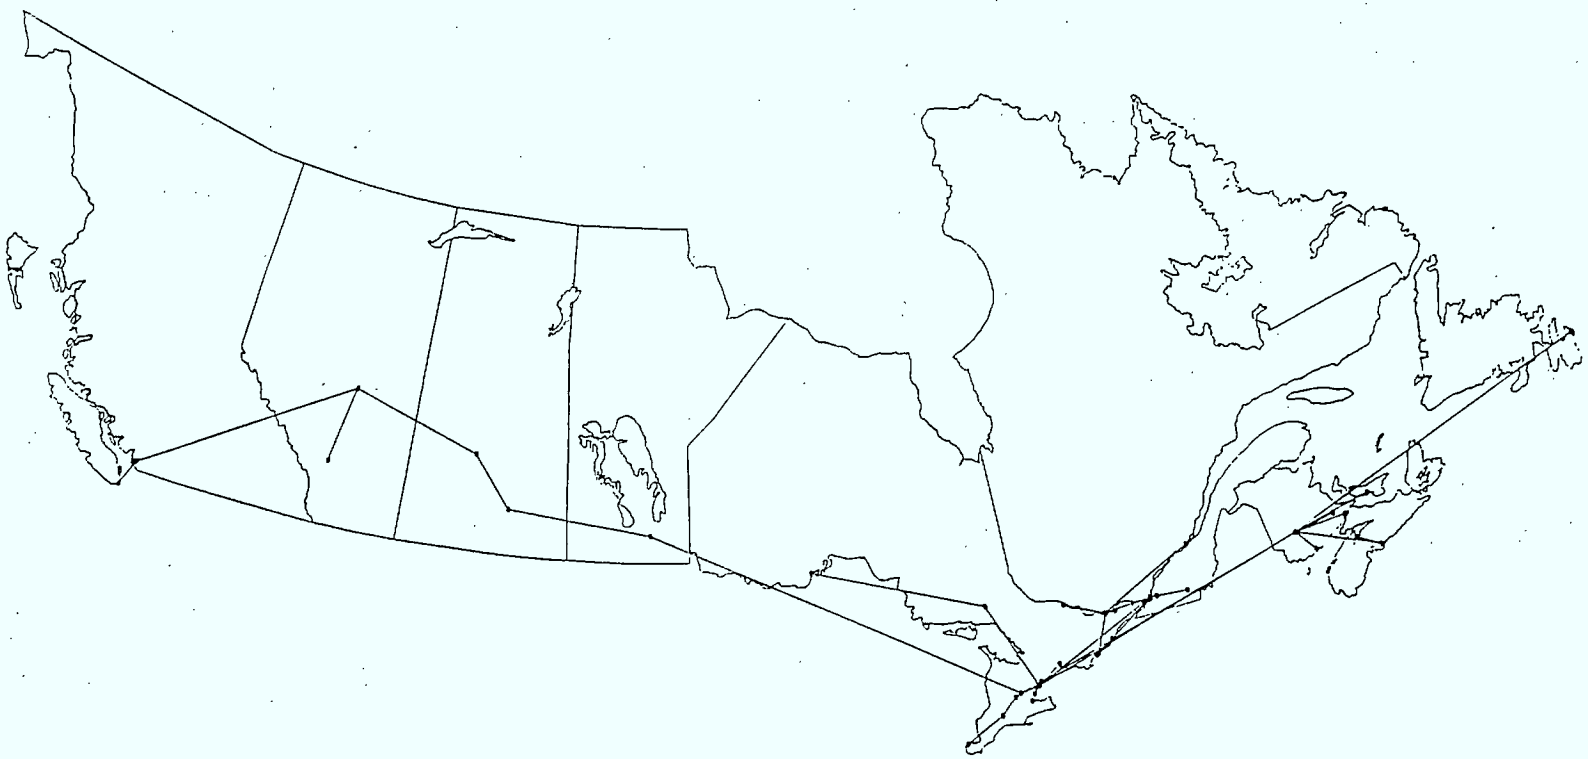

Source: Users' Directory of Computer Networks; Compiled by Tracy Lynn LaQuey, University of Texas; July 1989.

- Connection to BITNET is by 9600 bps leased line between University of Guelph and Cornell University;
- Protocol is NJE; many hosts are gateways to campus systems using TCP/IP or DECNET;

Implementation Design

- Member organizations designate reps and site directors; directors elect Committee of Executive Directors who carry out management and planning;
- User costs are approximately $6,000;
- Operating costs are funded through membership fees; they are about $250,000 annually for national links and $750,000 for the regional links;
- Much of the capital cost and early operating costs were covered by a three year grant from IBM which ended in 1987
- A plan to upgrade the NetNorth connection to 56 kB between Cornell and the University of Toronto would also provide direct connections to NSFNET and the Internet, and possibly participation in BITNET II using NJE over IP.
- Likely to evolve into the proposed CAnet;

2.7 FUTURE REGIONAL NETWORKS

This section describes the characteristics of the future (ie. proposed) regional networks in Canada.

2.7.1 NSTN

Nova Scotia Technical Network

Background

- Project funded under the Technology Transfer and Industrial Innovation. $2 million over two years has been provided to implement NSTN. After this it is anticipated that the network would become self sustaining;
- Caters to the business community in general but is specifically oriented towards the research and development community;
- This network emphasises involvement from the private sector (Systemhouse study concluded that over 100 industrial users were interested in joining the network);
- The network is targeted at all Nova Scotia universities, as well as small and large businesses in the technology sector. One of the expected benefits is technology transfer between these groups;

Technical Design

- The network will start out at 56 Kb and TCP/IP and wil go to OSI as soon as possible;

- NSTN consists of a backbone to which customer sites are attached. A redundant mesh topology is being planned to minimize the impact of communication line and node failure.

- The technology is the same as CANET;

Implementation Design

- Software Kinetics will establish and manage the business operations of the network through its wholly owned subsidiary (NSTN Inc.), which is a joint venture of Dalhousie University and Software Kinetics — and have been set up as a for profit organization;

- Software Kinetics will run the business side of the network and Dalhousie University will install and manage the technical operations of the network;

- The fee structure is currently under discussion. However there is general agreement that not all users will be equal. Different types of access will be available. Bandwidth and level of service will be variable depending on the needs of the users. Fixed fees for fixed service is the notion — fees will not be metered based on frequency of use;

- Unlike many of the other existing networks, there are no plans to restrict the types of traffic to be sent over the network;

2.7.2 MANITOBANET (Extended)

- Plans are being developed to implement an expanded 56 Kb IP network in the province of Manitoba, with a view to linking the local NRC office (CIIT in Winnipeg), the Fort Garry campus of the University of Manitoba, the Medical campus of the University of Manitoba, the University of Winnipeg, and several other candidate locations, including AECL in Pinawa, and an Agriculture Canada office;

- Part of the motivation is to provide medical researchers and other staff access to a Magnetic Resonance device which has also been proposed for the CIIT location;

2.7.3 MEMORIALNET

- There is great interest in establishing a high-speed metropolitan network in St. John's, Newfoundland, to link academic institutions, government agencies, and the Newfoundland and Labrador Computer Services crown corporation. Discussions are underway with the Atlantic Canada Opportunities Agency with regard to funding etc.;

- A campus network is being installed at Memorial University which will run TCP/IP and DECNET over fibre optic links.

2.8 FUTURE NATIONAL NETWORKS

This section describes the characteristics of the future (ie. proposed) national networks in Canada.

2.8.1 CANET (Formerly NRCNET)[7]

Purpose The purpose of CAnet is to interconnect the existing and emerging Canadian regional networks, and thereby to support data communications related to the research, academic and technology transfer needs of Canada.

Participants The University of Toronto (U of T) will be the prime contractor for the implementation and operation of CAnet. Other participants in the U of T group will be IBM Canada and INSINC, a telecommunications services reseller. U of T will manage all phases of the implementation project, including the installation and testing of both nodal facilities and communications lines, and will act as facility manager and operatior.

The staff required for implementation and operation of the network are a combination of new and exisiting staff. U of T will contribute the time and expertise of a number of existing staff. This contribution will be phased out over the three years, such that full costs will be recovered starting in the fourth year. The total value of U of T's contribution amounts to about $292,000 over the three years.

The nodal processors required at each site for connection to regional networks and the appropriate software will be supplied to U of T by IBM at substantial discounts. IBM's contribution is valued at more than $230,000. IBM will also provide development assistance and full access to the software and networking developments being done for NSFNET.

INSINC will supply the communications links between the nodal locations in each backbone city. INSINC's contribution is to provide these connections at cost plus a small administration charge in areas where it already buys bulk services. In areas where INSINC does not yet buy in bulk, services will be provided at cost.

Network Management Board The client regional networks will have a very strong role in directing the operation and funding of CAnet. A Network Management Board (NMB) has been established, with one representative from each province. The representatives have been selected by the regional networks in provinces where these already exist, and by mutual agreement of the research networking communities in provinces where regional networks are developing.

The NBM will oversee the operation of CAnet; it will make descisions regarding changes to topology and levels of service, determine how the costs are to be shared among the regional networks, and develop policies and procedures, including those relating to the connection of new members.

The NMB has chosen the intial network topology and level of service, has agreed on the budget for the first four years of operation, and has agreed on a funding formula to share the costs of operation among the regional networks. The NMB is considering incorportating CAnet as a means of taking full ownership and responsibility for the national network.

Relationship to NetNorth The NetNorth network was established in 1984 to provide mail and file–transfer services between Canadian academic, research and related organizations. It is an order of magnitude slower than CAnet, and uses protocols that are not as rich or non–propretary as those of CAnet. NetNorth is very supportive of CAnet plans and is expecting to merge with CAnet. NetNorth members are expected to join the existing or emerging regional networks in their provinces.

---

[7] This is copied from the CAnet Business Plan prepared by the Computing Services department of U of T, as approved by the Network Management Board.

When the merger has been completed, the current regional and national expenditures on NetNorth should be available to support the enhancement of regional networks and CAnet.

Topology The intial topology of CAnet will be that of a backbone network of ten nodes, interconnecting regional networks in every province. All links will support 56Kbps.

Connections to the NSFnet are presently provided from regional networks in Vancouver, Toronto and Montreal. NRC is currently negotiating with NSF with the aim of obtaining permission for three connections between NSFnet and CAnet, and convincing NSF to partly fund one of them. This expectation is reflected in the costs shown for these links.

Costs The projected costs of the various components of the network are tabulated below, for each of the first four years.

Projected CAnet Expenses

	$ Thousands per Year			
	1	2	3	4
Nodal Processors	100	225	–	–
Communications Links	605	628	591	700
NSFNET Links	150	150	150	150
Support Staff Net Cost	202	287	323	368
Maintenace	–	40	40	40
Lodging/Travel	20	10	10	10
Shipping/Phone	15	10	10	10
NMB Travel	5	5	5	5
Contingency/Capital Replacement Fund	110	60	35	35
TOTALS	1,207	1,415	1,164	1,318

A considerable contingency is built into the first year's budget to compenste for the uncertainty with Telecom Canada's recent filing for lowered inter–city data link costs. Since the timing and explicit details of the tariff reductions are not yeat known, the contingency fund may be consumed in implementing the intitial topology. When the line costs are reduced, the contingency fund can be applied to subsequent year's expenses or used to improve the topology, through faster links, or added ones for redundancy, as performance concerns warrent.

Inflationary cost increases are not explicitly included because they will be covered by reductions in communications costs resulting from greater market penetration by INSINC and expected tariff reductions. Savings remaining after inflationary cost increases have been covered will accrue to the contingency fund.

The costs shown for Nodal Processors, Communications Links and Support Staff are net of the contributions by the participants.

Initial Cost Sharing Formula The NMB has agreed to share the costs of CAnet operation among the member regional networks. Half of the costs are to divided equally among the members initially one for each province), with the other half shared in proportion to the poplulation of each region. The contribution profile of several provinces was adjusted, at their request, to reflect their particular fiscal situation.

The NMB has, however, agreed to make special arrangements for Prince Edward Island, which may have difficulty paying its share because of its very small population. These

arrangements may take one of the following forms: connect with a lower speed line, with the savings realized being used to reduce PEI's membership fee; connect at the same speed as the other provinces, but delay implementation for up to a year.

Project CAnet Income

	%	Year ($ Thousands) 1	2	3	4
NRC	–	800	900	300	–
Ontario	22.5	67	115	220	297
Quebec	18.5	110	110	110	243
British Columbia	11	33	56	107	145
Alberta	9.5	28	49	93	125
Manitoba	7	21	36	68	92
Saskatchewan	7	21	36	68	92
PEI	5.3	10	30	55	70
New Brunswick	6.5	19	33	64	85
Nova Scotia	6.8	80	20	20	89
Newfoundland	6	18	30	59	80
TOTAL	100.1	1,207	1,415	1,164	1,318

Schedule The initial topology shown below will be installed in three phases. Installation is scheduled to begin May 1, 1990, with the aim of completing all phases in August 1990.

Future Growth The technology used to implement CAnet allows future enhancements in several areas. The line speeds can be increased at any time if funding is available; in fact, the nodal processors will operate at speeds up to T3 (45Mbps). Additional links can be added to provide redundancy, which would allow continued operation in the event a nedtwork connection fails. Finally, the nodal processors will support the emerging OSI protocols, and this capability can be implemented at little incremental cost.

3. R&D NETWORKS IN THE UNITED STATES

3.1 SUMMARY

Although the U.S. is regarded as possessing relatively advanced data communication networks, there is a movement towards leapfrogging the world research community and to implementing a National Education Research Network (NREN). According to the National Research Network Review Committee (1988), the powerful desktop workstation and centralized supercomputers used by the research community are not currently being used to their full potential because the existing data networks are seriously inadequate for their data communication needs. Against this backdrop the computer networking goals of the research community can be identified as:

- Improved access to and interaction with research tools (supercomputers, databases, microscopes, telescopes, high definition television and specialized software);
- Improved interaction among researchers (large file transfers, real–time data processing, and electronic mail with graphic capability); and
- Acceleration of networking technologies and services. Industry is interested in the commercialization of these technologies whereas "classical" users of the network are interested in the productivity gains and new opportunities that these technologies can provide.

During the course of this review we have identified a number of important characteristics of U.S networks which are responsible for both the successes and failures of research networks in the U.S.. Major developments are currently in progress which will shape the future of the U.S. national network. This summary touches upon these major characteristics and elaborates on the direction currently being adopted by the U.S. Government. This summary is divided into the following five topic areas:

- Network development at universities;
- The success of NSFNET;
- The Evolution to a Cohesive National Research and Educational Network;
- The development of a National Telecommunication Policy; and
- Federal Government Support;

Network Development at Universities

Educational institutions are the primary users of networking technology in the United States. The creation of networks has been led by campuses across the U.S.. Universities have developed the most important networking technologies whereas industry have developed the most important communications technologies. To date the role of industry has been to provide communication links and produce equipment. Although commercial, for–profit use of research networks in the U.S. is currently forbidden, new acceptable use policies are

currently under review by the Office of Technology assessment. The modification of the acceptable use policies of these networks relate primarily to industry accessing the network for "research" purposes, even if commercial benefit may occur.
It is anticipated that further involvement by industry will increase technology transfer from basic research discoveries to usable products and services.

The Success of NSFNET

NSFNET and NSF as the lead funding agency has been extremely successful beyond original expectations in improving collaboration among researchers. This is seen through the explosion of articles in all branches of science over the last several years. In fact, traffic is tripling per year (20% to 30% each month).[7] With over 200 universities connected, it has become the principal channel of communication for computational researchers across the U.S. Despite the success of NSFNET and the web of research networks contained in the U.S. network, Internet, they are no longer adequate. It is somewhat ironic that in the early days of NSFNET there was criticism for creating too much capacity whereas at present, NSF is hearing that not enough was created (56 Kb was nice on NSFNET but it was soon filled up). The U.S. experience suggests that there is a strong latent demand for networking.

MERIT, IBM and MCI have established a totally independent research network linking Ann Arbour, Michigan; Reston, Virgina; York Town, New York; and Milford, Conneticut which has been successfully used to develop and test new technology for the backbone and provides a crucial facility for the anticipated migration of OSI protocals over the coming years.

The Evolution to a National Research and Educational Network

There is general consensus in the U.S for the establishment of a National Research and Educational Network (NREN), with NSF as its lead agency. This stems from the existing U.S. Internet being far from uniform in the type and quality of service. It also does not yet reach the entire research community. Agencies like DoD, DoE, NASA and NSF have created their own networks (or in some cases several networks) to support individual missions. The results being that interconnectability of these networks is poor, protocols are different and a network user may have 5 or 6 identification numbers. This redundancy has resulted in a typical science building having 4 or 5 connections to different networks. The traffic on NSFNET and other research networks in the Internet is primarily electronic mail. Real time connections to advanced computers are frustrating because of delays and unexpected disconnects. Breakthroughs in enhancing service requirements depend on cross–disciplinary research and, therefore, the need for one cohesive national network.

A three phase plan exists to build the NREN:

> **Phase 1:** Is complete and involved the NSFNET backbone being upgraded to full T1 service.
>
> **Phase 2:** Is currently underway and involves upgrading the service provided to about 200 research institutions using a shared backbone with T3 (45 Mb/s) capacity.
>
> **Phase 3:** Will provide one to three gigabit per second networking services to selected research facilities and 45 Mbs networking to 1000 sites. Deployment is not expected until middle to late 1990s. Research in gigabit technology is currently underway. Several key development areas include (1) switching systems for ultra–high speed

[7] The volume of traffic increased from 195 million packets per month (startup year) to one billion packets per month (1989) — an increase of over 500%.

communications links to move information at billions of bits per second; (2) network to computer interfaces making it possible to exchange data in multiple formats and media at high speeds; (3) software programs to manage the network; (4) telecommunication links to support the network.

As NREN evolves from NSFNET, the three level hierarchy will be maintained:

Level 1: Campus or research organization;

Level 2: Mid–level network or regional network (defined on a regional basis or discipline basis) linking individual sites; and

Level 3: The backbone itself which would physically link the different regional networks.

There is a general commitment in the U.S. to move to ISO–OSI as soon as possible. DEC and IBM have announced to their international clients and the U.S. government that they fully intend to implement OSI by this fall. A transport layer of OSI running along side the same machines with TCP/IP already exists. Most of the experts in this field have indicated that interoperability between these two domains is not an area of concern .

Universality of access is also a feature of NREN. It is anticipated that the NREN will give researchers, students and colleges of all sizes, as well as large and small companies in every state, access to the same high performance computing tools, data banks, supercomputers, libraries, specialized research facilities and educational technologies.

Federal Government Support

It is clear that the Federal Government supports a dual role in future computer networking:

- The support of networking research and technology development; and
- The support of operational networks.

NREN will receive $400 M over five years if the proposed National High performance Computer Technology Act of 1989 is approved. This amount will be divided evenly between solving the existing production problems on NSFNET (which will evolve into NREN), and research and development of future technologies (ie. gigabit networks). It is expected that NREN would lead to improvements is researcher productivity of 100% – 200%, which could be worth billions of dollars to the U.S. economy.

This act, if approved, would authorize $1.75 Billion over the next five years toward the development of supercomputing, artificial intelligence and other technologies designed to further U.S research efforts. On the other side of the coin is the commercial vendor community, which has shown a strong interest in the project. Northern Telecom, AT&T Bell laboratories and Cray Research have all expressed an interest because the network provides a test bed for new products and a potentially rich market for services.

It is also evident that the U.S. Federal Government would like to see the services developed for NREN available commercially at the earliest possible time, thereby reducing the funding burden. The assumption is that the commercial carriers will have a gigabit network in placeby the time the necessary hardware and software developments to the upper layers and the applications have been completed.

Although a detailed plan for transition to commercial service does not yet exist, the concept suggests federal government support for a high performance research network will drive technology forward and develop a market for advanced computer networking services. Furthermore, government research on networking will stimulate the development of products that can be used across the country to build other high speed networks. However, for this to happen, industry must be permitted to use these research networks to a greater extent than at the present time.

National Telecommunications Policy

NREN is a key element of a national telecommunications policy in the U.S.. NREN supports national goals with respect to international competitiveness and provides an infrastructure to tackle the grand challenges and other large science projects like the Hubble telescope and the World Seismic Network. These super science projects could not go forward without the existence of a high speed computer network. Arguments such as these are likely sufficient to justify funding over five years. NREN is supported by the fact that it meets a strategic objective of the U.S. congress. The network is not based on numerical arguments but rather it is supported by national objectives of defense, security and international competitiveness. To further support funding, comparisons are made between the projected funding for NREN ($400 M over 5 years) to the 1989 total government R&D spending ($62 B). Such comparison shows that less than 1% of the national investment in R&D would be allocated to fund the NREN infrastructure.

3.2 U.S R&D ENVIRONMENT

In contemplation of the future of R&D in the United States, it becomes apparent that the United States possesses impressive assets but that it must also confront some significant challenges.

National R&D Effort and Performance

In an article prepared by Douglas Van Houweling for the Educom Review (Summer 1989) he points out that the estimated national R&D effort in the U.S. is $140 Billion for 1989. The federal government is expected to fund 47%, industry 49% and other sources the remaining 4%. In performing the research, government will conduct 11% in its own laboratories, industry will perform 72%, and universities 17%. Approximately 63% of these funds will be spent on the development of new technologies. Basic research accounts for about 15% of these funds. These statistics are summarized in Figures 3.1 and 3.2.

In comparison to other countries the U.S. spends more than three times Japan's expenditure and more than two times the combined expenditures of West Germany, France and the United Kingdom. Only the Soviet Union has more active scientists and engineers. However due to recent acceleration in R&D investment, West Germany and Japan currently spend a larger percentage of their Gross National Product on R&D than the U.S.

In order to sustain its place as the world leader in R&D it is not as simple as spending more on R&D (as a percentage of GDP) but must also deal with other complicated issues in parallel. Some of these issues are discussed below:

FIGURE 3.1: National Expenditure on R&D by Country.

	1981	1982	1985	1987
United States	76.6	83.9	96.5	104.3
Japan	25.5	29.9	36.1	na
W. Germany	15.9	16.4	18.5	19.4
France	10.8	11.9	13.1	14.0
United Kingdom	12.2	11.9	12.9	na

Billions of U.S. dollars: constant 1982

FIGURE 3.2: The National R&D Effort $130.8 Billion, 1989 (est).

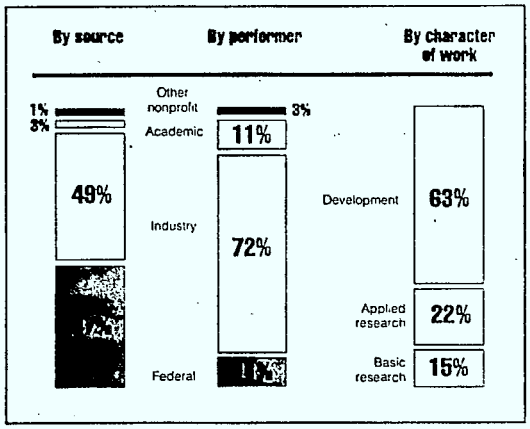

Changing Demographics

The underlying ethnic structure of the U.S. is significantly changing and projections are that it will change more radically in the future. For example, by the year 2000 about one third of post–secondary age U.S. residents, and almost half of all school–age children will be Black or Hispanic.

In addition, profound changes are anticipated in the age structure of the population over the next decade and beyond. In 1983 there were more Americans over age 65 than teenagers (for the first time in American history) By the year 2020, it is estimated that more than twenty percent of the population will be older than 65. This coupled with the increasing participation of women in the labour force, means that women over 40 will represent the most rapid growing segment of the U.S. labour force.

The U.S will not be able to utilize its full potential as an R&D power unless the changing demographic base is managed very carefully. These demographic changes have already caused a disparity between the population base required to run education and R&D institutions, and the population base that is available.

Shortage of Educated Researchers

Based on current trends, the U.S. educational system is not producing enough adequately skilled workers (approximately 8,000 science and engineering PhDs per year) for the information age of the future. Less than one quarter of one percent of all high school students will go on to attain a PhD degree in science or engineering, and the participation of minorities is significantly less than this. In 1986, only 100 Black American PhDs were produced by U.S. universities. This is not encouraging as the demographic forecasts suggests this to be one of the most rapidly growing segments of the work force. It becomes evident, then, that there will be increasing emphasis in using the limited researchers available in the most effective and efficient manner.

Shortage of Skilled Workers

The shortage of skilled workers is consistent with the changing demographic base and the shortage of educated researchers.

As Alan K. McAdams[8] points out, the U.S leads the world in patents drawn on pure science, and has one of the lowest lags between scientific article and patent (ie. 6 years).[9] According to an NSF–funded research project,[10] the U.S. leads the world in patented inventions which draw on pure science, and as patented inventions are increasingly drawing on pure science this becomes an important advantage. In order to maintain this position there is a growing need for the rapid diffusion of new scientific and technical knowledge, which requires support from those working in the information technology industry. However, time series on occupations in the U.S. reveals that the proportion of the workers in "new" information activities peaked in 1972, whereas in Japan the percentage keeps increasing every year.[11]

A national network would have the impact of "enlarging the pool of researchers" which is important in consideration of the projected shortfall of close to half a million scientists in the U.S. by the year 2000 [William Wulf Educom Summer Edition pg 26].

Managing Technological Evolution

According to a recent study by Cohen and Zysman (1988), "American difficulties in sustaining manufacturing innovation lie not in our machines and technology but in the understanding of how to exploit their productive promise." This is illustrated by the use of flexible manufacturing systems (FMS) for the production of comparable products in Japan and the United States. The number of parts made by an FMS in the United States is 10 versus 93 in Japan (almost 10 times greater), and the rate of new product introduction is 22 times as great in Japan as in the U.S.. The problems are not with technology but with the management of technology. It is argued that the generation and diffusion of "new" information would bring about a change in corporate culture required to attain more productive use of our technology.

Widely Decentralized R&D Program

As mentioned above, education, research and development activities are drawing on a changing workforce employed by a widely distributed and decentralized environment. With the exception of the proposed National Research and Education network, there does not appear to be any attempt to set a national agenda for research and development or to centrally coordinate efforts.

In addition, there is a real need for U.S. research establishments to be both competitive and cooperative. Basic research is becoming increasingly expensive to carry out without a joint partner to help share the cost. Technological changes are so rapid and costs so great that neither education, industry nor government can single–handedly provide all the high technology products and services necessary to a modern information based economy. Therefore, "strategic alliances" are very important.

---

[8] Economic Benefits and Public Support of a National Education and Research Network, Educom Bulletin, Summer/fall 1988.

[9] McAdams, Alan; Educom Bulletin, summer/fall, 1988.

[10] The research was carried out by Computer Horizons, Inc, of Haddon Heights, NJ.

[11] Morris–Suzuki, Tessa, "Capitalism in the Computer Age", New Left Review, Nov. 1986, pg 81–91.

3.3 SCOPE OF THIS REVIEW

There are hundreds of networks in the U.S. with many different characteristics in regard to size, region, purpose, administration, speed and capacity. This U.S. review focuses on private, long–haul (national), research–related, non–commercial networks.

Commercial and non–commercial networks are two types of networks that are readily distinguishable. Non–commercial networks are not necessarily free since there is a general move towards increased cost recovery. However, non–commercial networks are nonprofit. Most non–commercial networks support mail, but many do not support computer conferencing. Most of the ones that have dedicated links support remote login, but many of the dialup networks do not. File transfer support varies widely across these networks. Commercial networks, on the other hand, charge for their service in order to make a profit. These include systems in the U.S. like Compuserv, The Source, MCI Mail, AT&T Mail etc. There are few points of interconnection within these systems, and even fewer with these the non–commercial networks.

A network can also be characterized by the geography that it covers. A regional network is larger than a campus network but smaller than a national network, and often have backbones interconnecting smaller networks. National networks frequently interconnect regional networks with a wide area backbone. The U.S. has several national networks, each one specializing in a different academic discipline, protocol suite or funding source. Examples of these are NSFnet, MILnet and the proposed National Research and Education Network (NREN). International networks, like Bitnet, link hosts on different sides of a national boundary.

The networks selected for this review are presented in Table 3.1 and are categorized by the sponsoring federal agency. The selected networks are national networks used for ground and space research activities. As a group they serve researchers across the United States and have worldwide connections. Individually, these networks are in different stages of development, and vary in size, capacity, protocols, and services.

TABLE 3.1: NETWORKS SELECTED FOR THIS REVIEW

I EXISTING NETWORKS

1) THE INTERNET

2) DEPARTMENT OF DEFENSE NETWORKS:

- ARPANET and Defense Research Internet (DRI);

3) NATIONAL SCIENCE FOUNDATION NETWORK (NSFNET) — Three Level Network:

- Level 1: National Backbone
- Level 2: Twenty–one Mid–Level Networks:
- Level 3: Over 250 Campus Networks

4) DEPARTMENT OF ENERGY (DOE) RESEARCH NETWORKS:

- Energy Science Network (ESNET);
- Magnetic Fusion Energy Network (MFENET);
- High Energy Physics Network (HEPNET);

5) NATIONAL AERONAUTICAL AND SPACE ADMINISTARION (NASA) RESEARCH NETWORKS:

- NASA Science Internet (NSI);
- NASA Science Network (NSN);
- Space Physics Analysis Network (SPAN);

6) COOPERATIVE NETWORKS

- Bitnet; and
- CSNET.

II NEW FUTURE NETWORKS

1. National Research and Education Network (NREN)

3.4 SUMMARY OF NETWORKS IN UNITED STATES

Tables 3.2, 3.3 and 3.4 and 3.5 compare and contrast the characteristics of the U.S. networks. Size, purpose, protocol used, capacity/speed, media used, acceptable use, management and administration costs and benefits as well as a number of other discriminators are summarized by these tables. The main conclusions under each of these are indicated below:

Purpose: Networks have emerged from four major agencies in the U.S. government: NSF, DoE, DoD and NASA. Each network has come into existence for its own special mission or purpose. For example ARPANET was started when administrators within the Advanced Research Projects Agency realized that contractors required the same databases, graphic facilities and powerful CPUs. NSFNET evolved out of a need for improved access to super computers but file transfer and electronic mail account for 90% of the traffic. The original purpose of MFENET was to connect physics departments doing research in nuclear fusion. This mission–oriented approach to networking has resulted in poor interconnectability of these networks. A network user may have 5 or 6 different identification numbers. This redundancy has resulted in a typical science building having 4 or 5 connections to different networks. The purpose of NREN is to create one cohesive network, uniform in service to all potential network users.

An important component of NSFNET (and ARPANET to a certain extent) is research into network design and performance. An independent research network is used to develop and test new technology for the backbone and provide a facility for the anticipated migration of OSI protocols in the coming years and for the development of gigabit technology. The testing of products and services does not in any way affect the performance or reliability of the operational network — the two are kept separate. The Test and engineering network consists of a loop from Merit in Ann Arbour, Michigan to IBM Technical Computing Systems in Milford, Conneticut to IBM research in Yorktown, New York to MCI in Reston, Virginia back to Merit. Packet video is being experimented with on this network (potential application to video conferencing).

Size: NREN is projected to service over one million end users, which is by far the most ambitious undertaking in the United States (and in the world). BITNET also has a wide user base because of its international availability ie. 32 different countries. With NSFNET there are over 10,000 users and the amount of traffic generated by these users surpasses all other existing networks in the U.S..
Figure 3.3 summarizes the relationship between bandwidth and number of users for a selected number of networks.

Protocol: TCP/IP is by far the most commonly used protocol in the States. There is a general commitment in the U.S. to move to ISO–OSI as soon as possible. DEC and IBM have announced to their international clients and U.S. government that they fully intend to implement OSI by this fall. A transport layer of OSI running along side the same machines with TCP/IP already exists. Most of the experts in this field have indicated that interoperability between these two domains is not an area of concern. SPAN is an exception and is based on DECNET Phase IV and will migrate to Phase V for ISO–OSI compatibility. NREN will initially use TCP/IP because of its high reliability and will migrate to OSI as soon as it becomes available. However one of the failings of TCP/IP protocals is the lack of built–in features for network management. However, the need has been recognized by network vendors and substantial progress will be seen over the next few years.

TABLE 3.2 COMPARISON OF U.S. RESEARCH NETWORKS GENERAL INFORMATION

NETWORK NAME	CLASSIFICATION	DISCRIMINATING CHARACTERISTICS	FIRST YEAR OF OPEATION	SIZE (# of hosts/users/ sub-networks)
I EXISITING NETWORKS				
1.0 The Internet				
• ARPANET	Non-Commercial; Research [1]	• First Network to utilize packet switched technology currently being phased out and replaced by DRI	1969	Connects over 700 Networks, and over 500,000 host computers
• NSFNET	Non-Commercial; Research [1]	• Most Advanced and generally available National Backbone interconnecting regional networks and super computer facilities; three level inter network	1987	200 hosts
• MILNET	Non-Commercial; Military	• Unclassified, Production Network. (Does have classified segment)	1983	1,500 hosts
• CSNET	Non-Commercial; Research [1]	• Metanetwork to merge with BITNET, classified as a regional NSFNET network	1981	180 hosts (Dec. 1988)
• ESNET	Non-Commercial; Operational [2]	• Replacement of HEPNET and MFENet, backbone network for all energy research programs	1988 [3]	4,000 users
2.0 NASA Science Network				
• NSN	Non-Commercial; Operational [2]	• Supports TCP/IP Protocal, connects five major NASA Hubs across the nation	1987	Expected Population of 10,000 hosts
• SPAN	Non-Commercial; Operational [2]	• Supports DECNet Protocal, nationwide network	1981	2,400 hosts (Nov 88)
3.0 Cooperative Networks				
• BITNET	Non-Commercial; Academic	• Membership includes academic institutions and collaborating research institutions	1981	2,300 hosts at 450 sites in 32 countries
II NEW FUTURE NETWORKS				
Proposed NREN	Non-Commercial; Research, Operational and Educational	• Single-high speed streamlined and compatible backbone system. Envisions a three phase build up to 3 GBit/sec technology	Phase I to be completed by 1991 if Congress approves funding	Million end-users

NOTE: [1] The primary purpose of research networks is research into network technology.
[2] Operational networks do not necessarily promote the development of new network technology.
[3] ESNet currently under construction, will become the high speed backbone carrier for HEPNet and AEPNet.

Source: [1] The Matrix; Computer Networks and Conferencing Systems Worldwide; John S. Quarterman; 1990.
[2] HICKLING Research and Personal Interviews.

TABLE 3.3: COMPARISON OF U.S. RESEARCH NETWORKS

Technical Information

NETWORK NAME	MEDIUM	SPEED/QUALITY	SOFTWARE/PROTOCOL
I EXISITING NETWORKS			
1.0 The Internet			
• ARPANET	• N/A	• 56KB	• TCP/IP
• NSFNET	• leased lines & fibre optic circuit and digital microwave network	• T1 1.5 Mbps,	•TCP/IP
• MILNET	• leased land lines	• 56 Kbps, and	• TCP/IP but has not yet adopted DNS domain nameservers, DOT Internet Protocol System
• CSNET	• public telephone system	• 1200 or 2400 bps	• TCPIIP and a variety of other
• ESNET	• leased lines	• 56 Kbps	• IP backbone over X.25 supporters TCP/IP, DECNET and the MFEnet NSP protocols
2.0 NASA Science Network			
• NSN	• leased lines and microwave links	• T1 + 56 bps circuit to the European Space Operation Centre (ESOC) in GDR	• TCP/IP protocol
• SPAN	• leased lines	• 9600 bps as a baseline	• TCP/IP, hosts must run DECNET Phase IV
3.0 Cooperative Networks			
• BITNET	• leased phone lines	• 9600 bps	• Network Job Entry (NJE) protocol, VM or CDC
II NEW FUTURE NETWORKS Proposed NREN	Fibre Optic	Phase I - TI Phase II - T3 Phase III - Gigabit	TCPIIP initially with movement to ISO-OSI. Gigabit protocol under development.

Source: [1] The Matrix; Computer Networks and Conferencing Systems Worldwide; John S. Quarterman; 1990.
[2] HICKLING Research and Personal Interviews.

TABLE 3.4: COMPARISON OF U.S. RESEARCH NETWORKS

Management and Use

NETWORK NAME	MANAGEMENT AUTHORITY	MANAGEMENT OPERATIONS	ACCESSS	USE
I EXISITING NETWORKS 1.0 The Internet • ARPANET • NSFNET • MILNET • CSNET • ESNET 2.0 NASA Science Network • NSN • SPAN 3.0 Cooperative Networks • BITNET	• Several agencies including FRICC • Division of Network and Communications Research and Infrastructure (DNCRI) • Defense Communication Agency (DCA) • CSNET Executive Committee • NMFECC of LLNL • NASA Science Internet Program Office (NSIPO) at Ames Research Centre (ARC) • Data Systems User's' Working Group (DSUWG), project scientists, project and network managers and routing centre managers • Incorporated Board of Trustees elected by membership	• DDN NIC and the Network Operation Centre (NOC) at Bolt Beranek and Newman (BBN) in Cambridge, Mass. • Management and Operations - MERIT Inc, End User Support - NIC and the NSF NSC at BBN • DDN Program Management Office (PMO) • CSNET Coordination and Information Centre (CIC) at BBN and The University Cooperation for Atmospheric Research (UCAR) under contract to NSF • ESnet Steering Committee - a representative of each DoE energy research program • NISPO • DSUWG • BITNET NIC, network services	• non-commercial use • non-commercial use • determined by the U.S. Department of Defence, classified segment and readily accessible part • industrial, academic, government & non-profit institutions engaged in research and advanced development in computer science and engineering • restricted to projects sponsored by DoE • NASA and participating universities and laboratories • NASA and participating universities and laboratories • communication network among universities with no special requirements, restrictions or fees for membership, virtually unrestricted	• government agencies, educational institutions, private corporations for the purpose of collaborating among researchers and as a testbed for new developments in networking • national U.S. research network, research institutions - not-for-profit • not normally used for networking research, long haul military production network • facilitate research and advanced development in computer science or engineering • DoE Energy Research Programs • coordinate OSSA science requirements, network engineering using PSCN requirements, gateways and other internal requirements, network technology reserach and development • space related research, file transfer and mail only • non-commercial exchange of information
II NEW FUTURE NETWORKS Proposed NREN	• To be determined	• To be determined	• To be determined	• To be determined

Source: [1] The Matrix; Computer Networks and Conferencing Systems Worldwide; John S. Quarterman; 1990.
[2] HICKLING Research and Personal Interviews.

TABLE 3.5: COMPARISON OF U.S. RESEARCH NEWORKS

Funding and Future Plans

NETWORK NAME	EXPANSION	FUNDING BODIES	FUNDING MECHANISM
I EXISITING NETWORKS			
1.0 The Internet			
• ARPANET	• will be retired in March 1990.	• U.S. DOD, several government agencies and numerous companies	• Grants
• NSFNET	• T3 by 1990, moving to support ISO-OSI protocols	• NSF, State of Michigan , also study agreements with IBM & MCI on network research and development including service and facilities	• Mostly government grants, NSF funds UCAR & BBN for operation of NNSC (NSF Network Service Centre) NSF provides initial and partial funding to mid-level networks
• MILNET	• in the process of completing implementation of fully qualified domain names	• DOD	• mostly government grants
• CSNET	• merge into NREN	• CSNET members	• Initial funds from NSF, since 1985 self-supporting with annual dues from members with rates set according to classification
• ESNET	• merge into NREN	• Department of Energy	• n.a.
2.0 NASA Science Network			
• NSN	• merge into NREN and transition to ISO-OSI protocols	• NASA	• Code T - Program Support Communications funds the Program Support Communication Network (PSCN*), Code E - Office of Space Science and Applications + user fees
• SPAN	• merge into NREN and transition to ISO-OSI protocols	• NASA	• as above
3.0 Cooperative Networks			
• BITNET	• BITNET II, planned merger with CSNET (ONEnet)	• BITNET members	• membership fees dependent on the class of use and the member's budget
II NEW FUTURE NETWORKS Proposed NREN	• Would merge all lead agencies into one cohesive national backbone	• High Performance Computer Technology Act of 1989 (awaiting approval) dedicating $1.75 billion over the next five years toward the development of super computing, artificial intelligence and other technologies	• $400 over 5 years if HPCT Act is approved

* PSCN is the Physical Infrastructure for NSN and SPAN

Source: [1] The Matrix; Computer Networks and Conferencing Systems Worldwide; John S. Quarterman; 1990.
[2] HICKLING Research and Personal Interviews.

FIGURE 3.2: RELATIONSHIP BETWEEN BANDWIDTH AND NUMBER OF END USERS FOR SPECIFIC U.S. NETWORKS

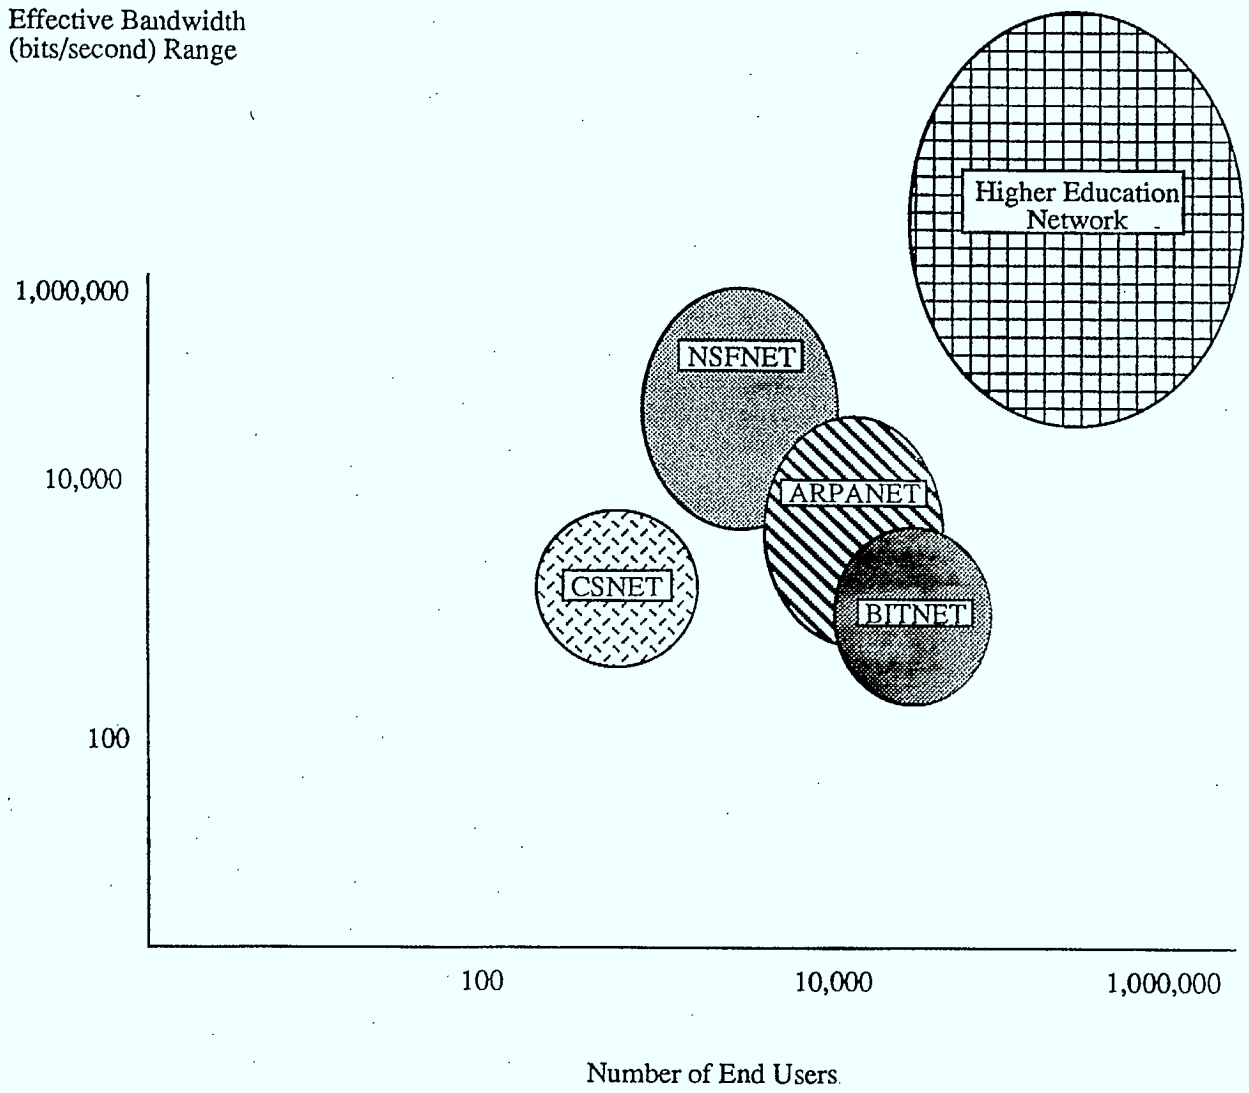

Source: EDUCOM Bulletin, Winter 1987.

An area of research that several government and military groups are exploring is networking protocals that are specifically designed for multigigabit. rates. There is general consensus that Open Systems Interconnect (OSI) and Transmissions Control Protocol/Internet Protocol will not be able to handle 200 M bit/sec across the country, so it is likely that Phase 3 of the NREN would use neither.

Capacity/Speed: The networks studied in this review are predominately T1 (1.544 Mb/s), or are 56Kb transitioning to T1. The T1 netwoprk infrastructure allows member institutions to move beyond electronic mail, file transfer and remote terminal access applications to more sophisticated interactive graphic work and multi–media mail. The second phase of NREN involves the implementation of a T3 shared backbone in the near future. The third phase of NREN will provide one to three gigabit per second networking services to selected research facilities and 45 Mb/s networking to 1000 other sites. Deployment is not expected until middle to late 1990s. Research in gigabit technology is currently underway. Several key development areas are (1) switching systems for ultra–high speed communications links to move information at billions of bits per second; (2) The development of network to computer interfaces making it possible to exchange data in multiple formats and media at high speeds; (3) Software programs must be written to manage the network; (4) telecommunication links must be configured to support the network.

Media: Leased land lines (fibre or copper) are the most commonly used medium in the States with some networks like NSFNET and NSN taking advantage of microwave links. CSNET utilizes the public telephone system as a dialup service. SPAN is composed over a number of different media, including dedicated lines, as well as both private and public X.25 packet switch networks. The migration to higher speeds (T3 and gigabit) will necessitate the use of fibre.

Acceptable Use: As the networks described in this review have emerged as mission–oriented projects funded by government agencies, it is not surprising that access to a network is usually restricted to those researchers who receive funding from the relevant sponsoring agency.

For example BITNET prohibits "commercial usage" and does not permit the sending of proprietary software products. BITNET originally would not permit two industrial members to communicate ie. had to involve a university. This has subsequently been relaxed because projects may involve multiple industrial and multiple acaedemic researchers.

CSNET was created by the NSF with a charter that its users had to be engaged in computer science or related research or engineering. CSNET has broadened itself to permit users from all science backgrounds. CSNET has specific limiations intended to restrict unnecessary transmissions like junk mail.

NSFNET is intended to support research or instruction at non–for–profit institutions in the U.S. The rules do not prohibit non–university research or non–profit research.

Acceptable use of NSFNET spans a wide range of scientific disciplines, which includes only those researchers engaged in non–commercial research or industrial laboratories receiving NSF funding. NSF guidelines leave considerable room for interpretation. NSF decides what is consistent with their mission. The other physical networks that make up the NSFNET are the regional and super computer networks which although partially funded by the NSF are really independent. This independance means that the regional and supercomputer networks have their own guidelines.

Fee Structure: Many faculty believe that utility–oriented cost–recovery accounting policies virtually guarantee underuse of networking resources. They argue that networking services should be funded centrally in the same manner as libraries.

One argument against user fees based on use is that if we believe huge benefits accrue from exchanging information, then we should encourage this interaction rather than discourage it. Also from an operational view (the above is a philisophical view) the overhead costs associated with metering the network are very large. Network hardware and software now in use in the Internet cannot assess charges.

Another viewpont expressed is that only after the growth, use patterns, and legal impediments are fully understood should the national science Internet be treated as a utility.

Management and Administration: With the exception of NSFNET most of the networks described in this review are managed by network "groups" within the sponsoring agency. However NSFNET is an example of a different management model which is professionally managed and administered. A combined system for administering NSFNET does not exist — administration is divided functionally.

The Division of Network and Communications Research and Infrastructure (DNCRI) part of the Computer Science and Information Science and Engineering Directorate within NSF is responsible for establishing policy for the network.

Merit Inc. is responsible for the management and operation of the NSFNET backbone. Merit is a nonprofit membership consortium of eight Michigan universities with technical support from IBM and MCI. IBM provides broad research experience and product offerings in systems and connectivity (including network management) and MCI's provides experience in telecommunications and state–of–the–art digital networks.

End user support for NSFNET is provided by a Network Information Centre (NIC), the NSF Network Service Centre (NNSC) at BBN Systems and Technologies Corporation in Cambridge, Massachusetts (cooperative agreement with the University Corporation for Atmospheric Research (UCAR)). Both BBN and UCAR are funded by NSF.

Funding: For the most part funding of these research networks is provided through the budgets of the sponsoring agencies. The budget for NSF has doubled over the last five years. With respect to NSFNET, IBM contributed the packet switching and network management equipment and software and MCI the long distance data transport facilities.

NREN will receive $400 M over five years if the proposed National High performance Computer Technology Act of 1989 is approved. This would subsidize roughly 40% of the cost of for Stage 1 and 2 (described earlier) and 60% for gigabit research (Stage 3). (Refer to the five year reference levels in Table 3.6). The act, if approved, would authorize $1.75 Billion over the next five years toward the development of supercomputing, artificial intelligence and other technology designed to further U.S research efforts. The commercial vendor community has also shown a strong interest in the project. Northern Telecom, AT&T, Bell laboratories and CRAY Research have all expressed an interest because the network provides a test bed for new products and a potentially rich market for services.

	YR 1	YR 2	YR 3	YR 4	YR 5
	IN MILLIONS OF CURRENT $				
Interagency Interim NREN	14	23	55	50	50
Gigabit R&D	16	27	40	55	60
TOTAL	30	50	95	105	110

Security: After the Internet virus went out over the national backbone, it is not surprising that security is very topical. It is a difficult problem to resolve because it must be integrated into the computer operating system and network software. Security steps taken with MIT's Athena project – a joint research effort between the school, IBM and Digital – produced the Cerebus protocol. It's features include encryption, authenication, and validation.

Costs and benefits: Bitnet took five years to be self sustaining and currently costs each institution about $25,000 per year. NREN is expected to be self sustaining in 10 years. Currently, NSFNET costs $50,000 to $100,000 per year per institution for communications to connect to a mid–level network (which does not count campus networking and support costs).

The shortage of high speed links has caused researchers to limit or even table certain projects that require major remote collaboration. As a result, many researchers "filter out" highly collaborative projects. This has resulted in a chicken and egg problem in which potential developers of a network have been unable to find enough interest to cost–justify the project.

It is generally believed that networks like NREN would lead to greater research productivity and improved access to specialized databases and facilities (supercomputers, accelerators, and telescopes), accelerate the transfer of results to the private sector and substantially benefit international competitiveness as new information plays a critical role in stimulating economic growth and international competitiveness.

For further information on any of these networks please refer to the next section which provides more detail on each network.

3.5 EXISTING NETWORKS

This section describes the characteristics of the networks contained within The Internet, The NASA Science Internet, Department of Energy Networks and briefly studies BITNET and CSNET as examples of a cooperatively run network.

3.5.1 The Internet

Background:

- Web of networks connecting U.S. government research and academic institutions;
- Represents primarily a federation of autonomous U.S. networks which run the TCP/IP protocol suite connected through gateways and sharing common name and address spaces;
- Largest number of participating networks are in the U.S. — although The Internet connects to TCP/IP networks in Canada, Mexico, Japan, Europe, New Zealand, Australia and New Zealand; U.S. Internet is widely connected abroad and just in the last year, T1 connections have been put in place to Japan and Australia. There are 6 links to European researchers;
- Main backbone networks of the Internet are NSFNET and MILNET;
- Constituent parts like NSFNET have many component networks on its own;
- Exists as a testbed for new developments in networking;
- Number of hosts connected to the Internet is in the order of 40,000 – 500,000 with close to one million users — and is growing exponentially each year;
- A more complete understanding of the Internet can be attained by studying its primary components; ARPANET, NSFNET, MILNET, CSNET, ESNET and MFENET;

Technical Design:

- Hosts in the networks of the Internet use the TCP/IP protocol (some networks run the TCP/IP protocol on top of X.25 on public data networks);
- A Government Open Systems Interconnection Procurement Specification (GOSIP) mandates eventual conversion to ISO–OSI protocols for all U.S. government networking. This will not have any strong effect on the use of TCP/IP on the Internet for about two years;
- Most hosts and networks on this system support TELNET for remote login, FTP for file transfer and SMTP for mail, date, system status and directory services;
- There are point to point connections over terminal lines, HYPERchannel links, dialup links, and T1 microwave links;
- The DNS domain system is widely used but not exclusively (top level domains in the Internet include COM, EDU, GOV, ORG, MIL, NET, ORG, and US);
- Routing is hierarchical — a route is found to the appropriate network through gateways by the Gateway to Gateway protocol (GGP) and the Exterior Gateway Protocol (EGP);

3.5.2 Department of Defense Networks

Background

- ARPANET implemented in 1968, as a small research project, when administrators with the Advanced Research Projects Agency (DARPA an arm of the U.S. Department of Defense) realized that contractors required the same resources (databases, graphic facilities, powerful CPUs etc.);
- Network demonstrated the viability of long–haul packet switched technology – – and eventually grew into a national U.S. backbone, resulting in the current Internet; Represents the first packet, connectionless architecture;
- In 1983 it split into MILNET, a production military network, and ARPANET a research network;
- MILNET is a long–haul, stable (few disruptions to service), unclassified military production network that is not used for networking research;
- MILNET has about 1,500 hosts (vast majority located in the U.S);
- A classified segment of MILNET exists in addition to the readily accessible part;

Technical Design

- MILNET uses the TCP/IP protocols except has not yet adopted DNS domain nameservers and uses static host tables for host name to address mapping;
- ARPANET and MILNET PSNs are connected by 56 Kbps leased lines;

Implementation Design

- ARPANET has been retired and is gradually being taken out of service (March 1990) because its long–haul links have been overtaken by new technology (which is cheaper and faster);
- ARPANET is to be taken over by the **Defense Research Internet (DRI).**
- Access to MILNET is determined by the Department of Defense (DoD)
- Management of MILNET is by the Defense Data Network Program Office of the Defense Communication Agency.
- Funding is by DoD;

3.5.3 National Science Foundation Network (NSFnet)

Background

- NSFNET is part of the Internet, and over 200 universities are connected;
- NSFNET is the most advanced and generally available network now in operation and has become the national U.S research network;

- In 1987 established to facilitate access to the national supercomputing centres;
- Initial purpose to provide access to supercomputers but file transfer and electronic mail account for about 90% of the traffic;
- NSFNET connects to ARPANET, BITNET, CSNET, BCnet, CDNnet (connected by a link from the University of Washington to UBC) and Onet (connection from Cornell to University of Toronto);
- Traffic is increasing at a rate of sixty–fold per year (20% to 30% each month); one billion packets of information were transmitted in June of 1989;
- Important component of NSF backbone effort is research into network design and performance.
- Independent research network is used to develop and test new technology for the backbone and provides a facility for anticipated migration of OSI protocols in the coming years and for the development of gigabit technology;
- Test and engineering network consists of a loop from Merit in Ann Arbour, Michigan to IBM Technical Computing Systems in Milford, Conneticut to IBM research in Yorktown, New York to MCI in Reston, Virginia back to Merit. Packet video is being experimented with on this network (potential application to video conferencing).

Technical Design

- NSFNET is hierarchical network connecting 13 network hubs with full T1 service (see Map). Shortly within the next fiscal year it appears likely the NSFNET backbone will be upgraded to T3.
- Most mid–level networks use are T1 or 56Kb/s leased circuits. Local nets generally use Ethernet (FDDI Lans are appearing) and 100Mb/s is becoming common at the local level;
- NSFNET is interconnected to ESNET, ARPANET, NASA Science Internet by designated gateways and dedicated links;
- NSFNEt backbone uses MCI's fiber–optic circuit and digital microwave radio network to carry data. Connections are point–to–point T1 links.
- Uses packet–switching technology (Nodal Switching Systems (NSS) supplied by IBM). Each NSS consists of nine IBM RT/PCs (runs IBM version of 4.3BSD UNIX), connected to two token ring networks for redundancy;
- Software controlled circuit switching equipment is an Integrated Digital Network Exchange (IDNX) is provided by IBM;
- Subrate T1 multiplexing, dynamic alternate routing and dynamic bandwidth allocation will eventually be provided by these 13 nodes;
- General support for ISO–OSI protocols exists. At Interop 89 Merit successfully demonstrated a prototype in wide–area OSI connections on NSFNET. This demonstration spanned coast–to–coast with computer links operating over a T1 connection;

- The prototype OSI implementation is designed to coexist with TCP/IP — fully dynamic routing is supported for both protocols. Intend to offer CLNP as production protocol on the backbone as soon as possible;

- It is possible to send IP–encapsulated DECNET over NSFNET backbone;

- Some of the NSFNET networks use internet and transport protocols other than IP and TCP such as DECNET or NSP protocol. Therefore NSFNET is actually a metanetwork;

- A Network Operations Centre (NOC) provides 24 hour coverage seven days a week;

- A 24 hour WATS line is available for trouble shooting and general status information. The program that provides online information and communication services is called GRASP and is a merging of IBM's GRAND with Stanford's production data management system SPIRES;

- Two IBM mainframes exist at Ann Arbor campus for the NOC and for Information Services;

Implementation Design

- Open to users from the higher education and research community across the nation. Use for commercial activities by for–profit institutions is generally not acceptable;

- Super Computer account can be set up for any researcher whose project is judged meritorious through a peer review process;

- It is a three level hierarchical network consisting of the following:
(1) **a backbone** which connects separately administered and operated mid–level networks and NSF funded supercomputer centres; (2) **mid–level networks** consisting of regional, discipline based and super computer consortium networks (see Table 3.7); (3) **campus networks** connected to mid–level networks.

- Combined administration for all three levels of NSFNET does not exist — Administration is divided functionally:

  – Division of Network and Communications Research and Infrastructure (DNCRI) part of the Computer Science and Information Science and Engineering Directorate within NSF is responsible for establishing **policy** for the network.

  – Merit Inc. responsible for the **management and operation** of the NSFNET backbone (Merit is a nonprofit membership consortium of eight Michigan universities).

  – **End User Support** provided by Network Information Centre (NIC), the NSF Network Service Centre (NNSC) at BBN Systems and Technologies Corporation in Cambridge, Massachusetts (cooperative agreement with the University Corporation for Atmospheric Research (UCAR)). Both BBN and UCAR are funded by NSF.

TABLE 3.7: NSFNET MID-LEVEL NETWORKS [*]

BARRNet	San Francisco Bay Area Regional Research Network [2]
JVNCNet	John von Neumann Center Network [1]
Merit	Merit Computer Network [2]
MIDnet	Midwest Computer Network [2]
NCSAnet	National Center for Supercomputing Applications Network [1]
NorthwestNet	Northwestern States Network [2]
NYSERNet	New York Educational and Research Network [2]
PSCnet	Pittsburgh Supercomputer Centre Network [1]
SDSCnet	San Diego Supercomputer Centre Network [1]
Sesquinet	Texas Sesquicentennial Network [2]
SURAnet	Southeastern Universities Research Association Network [2]
THEnet	Texas Higher Education Network [2]
USAN	University Satellite Network [3]
WESTNET	Mountain States Network [2]

NOTE: [*] See attached maps.
[1] Supercomputer consortium network
[2] Regional Network
[3] Discipline–oriented network

SOURCE: John S. Quarterman, Texas Internet Consulting, THE MATRIX: Computer Networks and Conferencing Systems Worldwide; Digital Press; 1990.

- NOCs exist for some of the regional networks as well as NICs (eg. NYSERnet), therefore questions or problems start at the most local center;
- NSF funded Merit $14M (over 5 years) to reengineer the NSFNET backbone. The state of Michigan Strategic Fund contributed $5 million;
- Merit established agreements with IBM and MCI on network research and development. IBM supports Merit with the Nodal Switching Subsystems and network management applications, and to the migration of ISO protocols over the life of the program. MCI provides the transmission backbone and participates in service enhancement and network management efforts;
- Under this agreement MCI provides leased lines for the backbone and expert services with regards to circuit switched technology;
- Hardware and software for the backbone nodes and for network management and information services, provided by IBM;
- Proposal to implement T3 by 1990 but funding has not yet been forthcoming;
- Recent international connections have been made which will potentially lead to a global research network.

NSFNET T-1 DATA NETWORK NEW TOPOLOGY

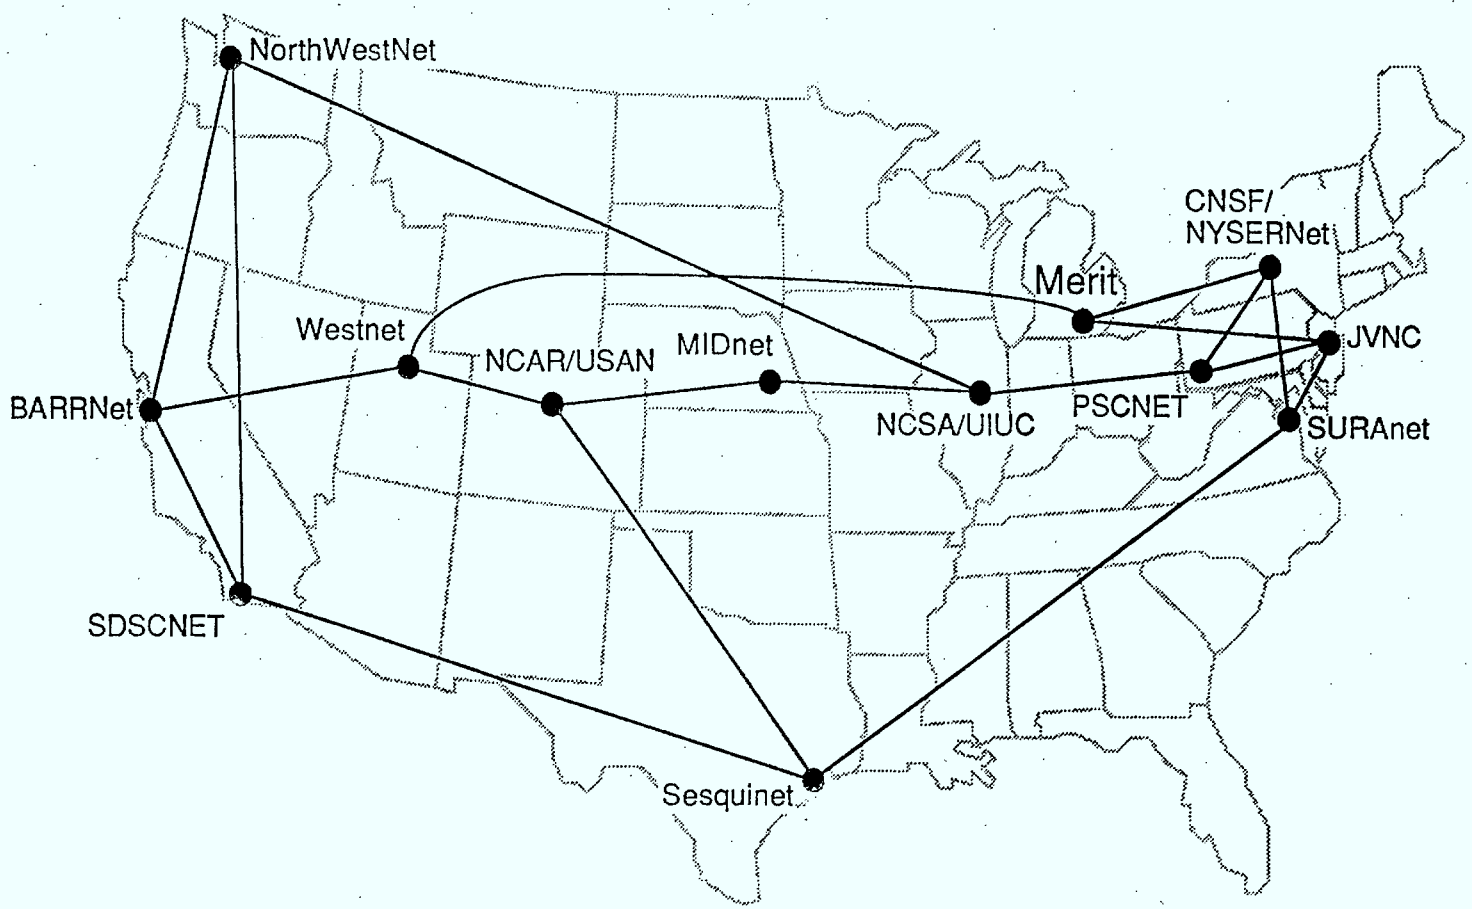

Source: National Science Foundation; Washington, D.C.; September 1989.

NATIONAL SCIENCE FOUNDATION NETWORK (NSFNET) MONTHLY TRAFFIC (MILLION PACKETS)

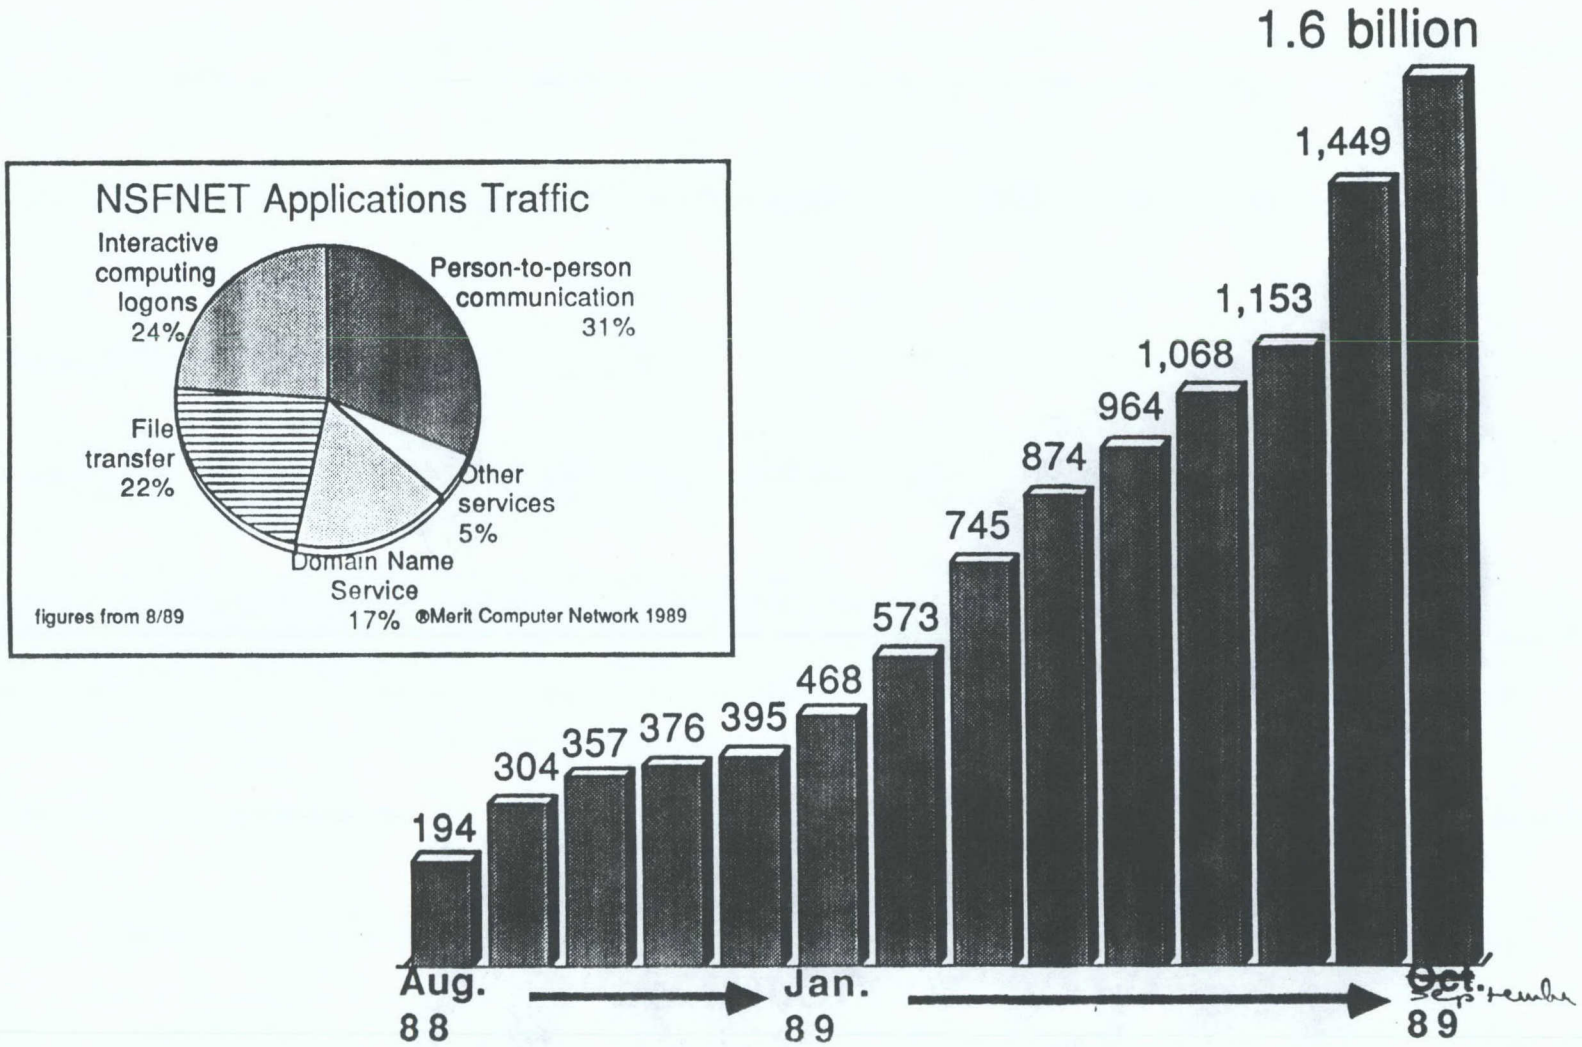

Source: National Science Foundation; Washington, D.C.; September 1989.

BARRNet

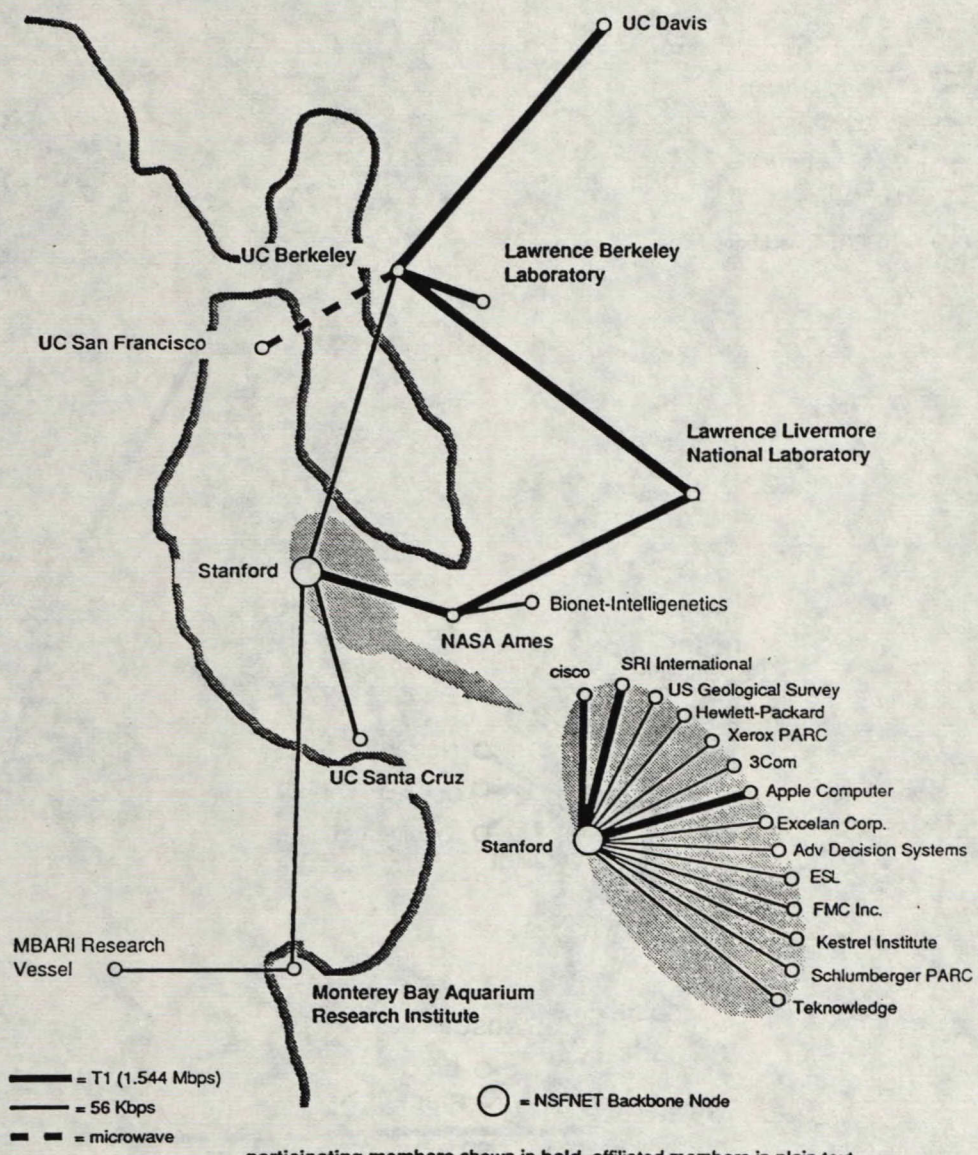

participating members shown in bold, affiliated members in plain text.

Sites to be added 6/89: DEC WRL, Lucid Systems, MIPS, mt Xinu, Rockwel PAL, Sun Microsystems, Tandem Computers, VLSI Technology.

Source: Users' Directory of Computer Networks; Compiled by Tracy Lynn LaQuey, University of Texas; July 1989.

CERFnet

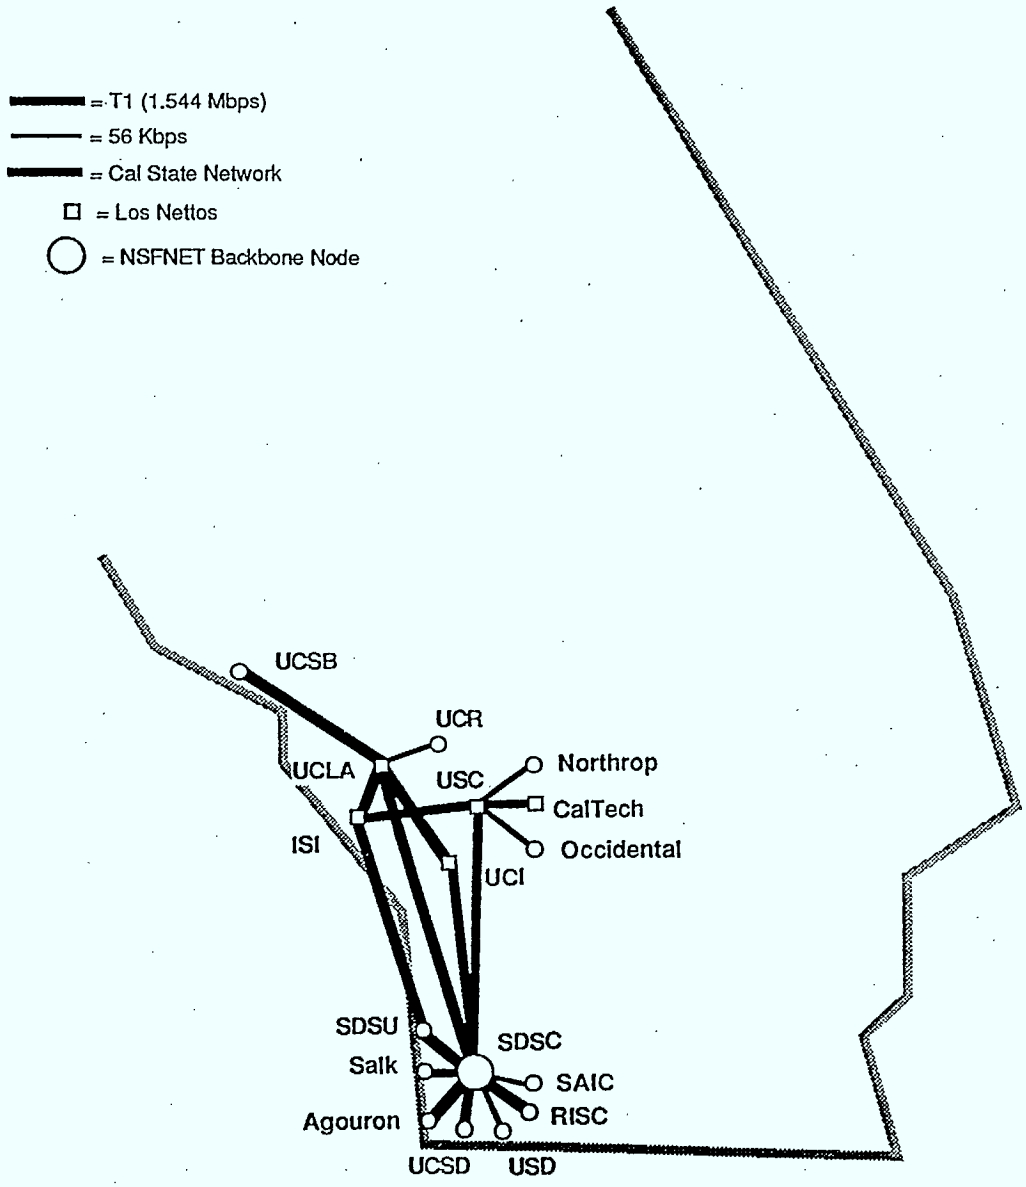

Source: Users' Directory of Computer Networks; Compiled by Tracy Lynn LaQuey, University of Texas; July 1989.

CICnet

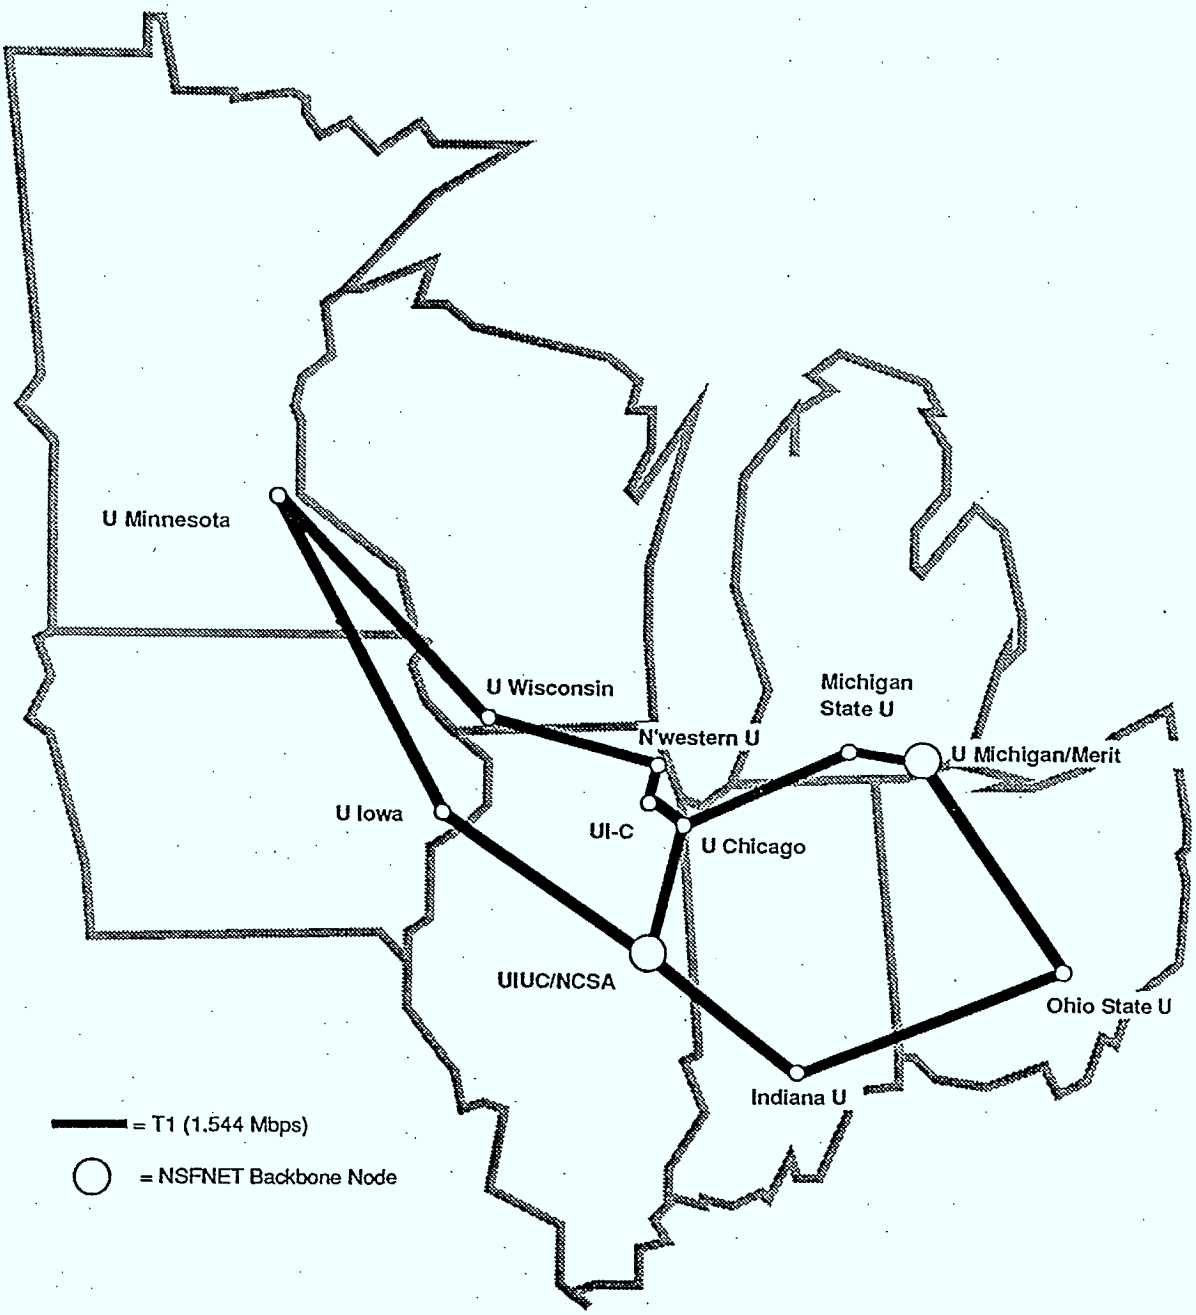

Source: Users' Directory of Computer Networks; Compiled by Tracy Lynn LaQuey, University of Texas; July 1989.

JVNCnet

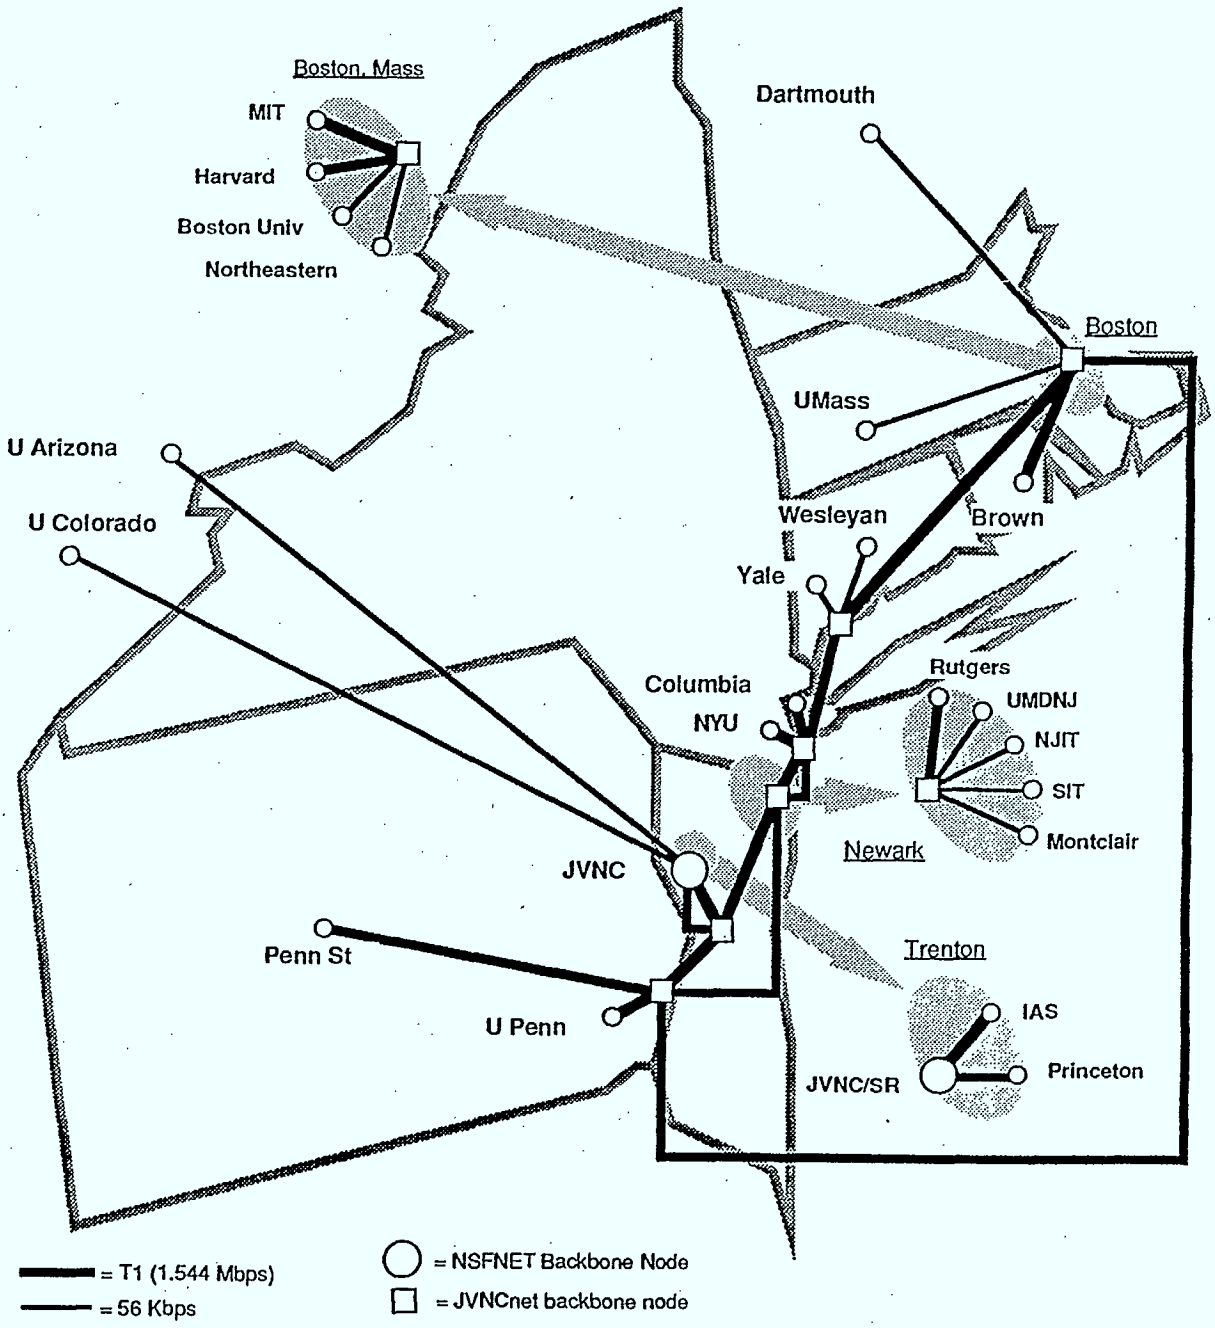

Source: Users' Directory of Computer Networks; Compiled by Tracy Lynn LaQuey, University of Texas; July 1989.

MIDNet

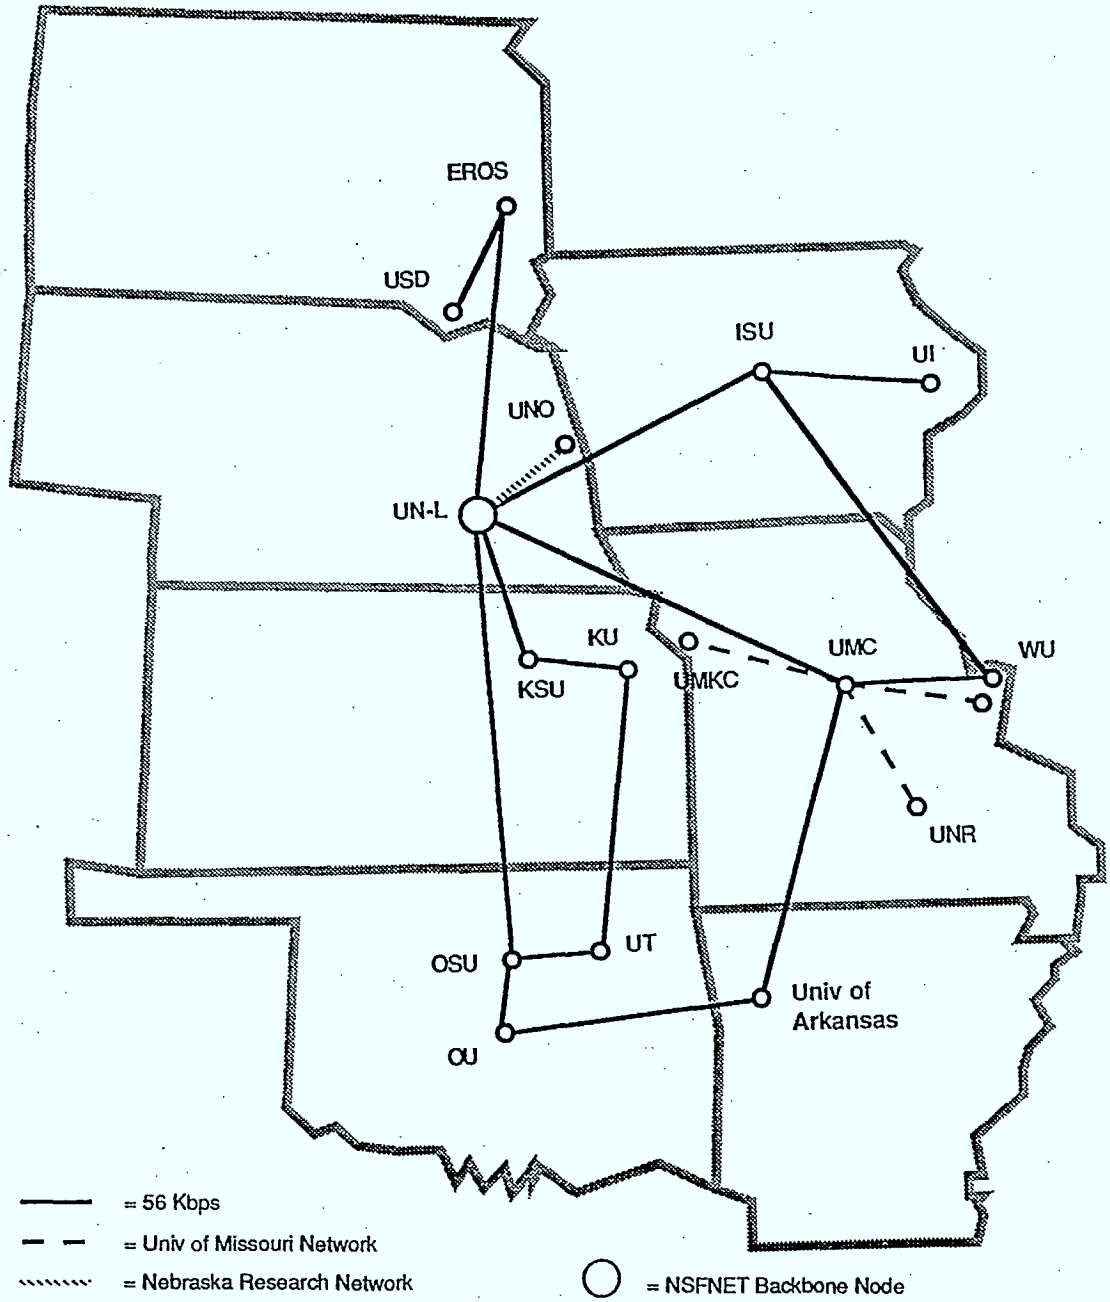

Source: Users' Directory of Computer Networks; Compiled by Tracy Lynn LaQuey, University of Texas; July 1989.

MRNet

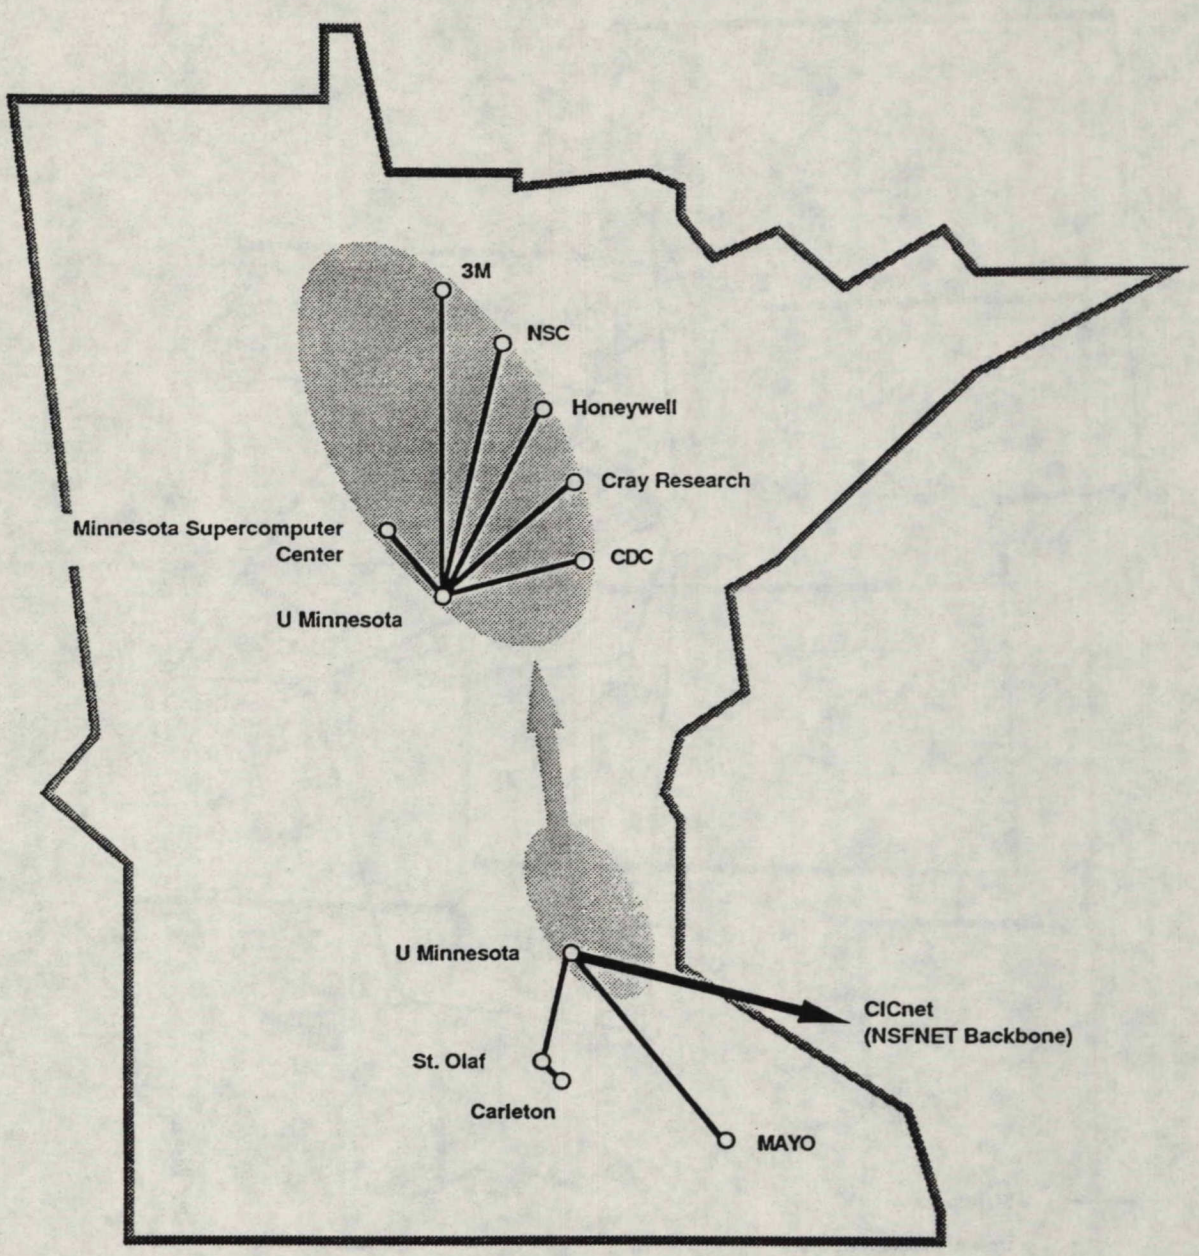

Source: Users' Directory of Computer Networks; Compiled by Tracy Lynn LaQuey, University of Texas; July 1989.

NCSAnet

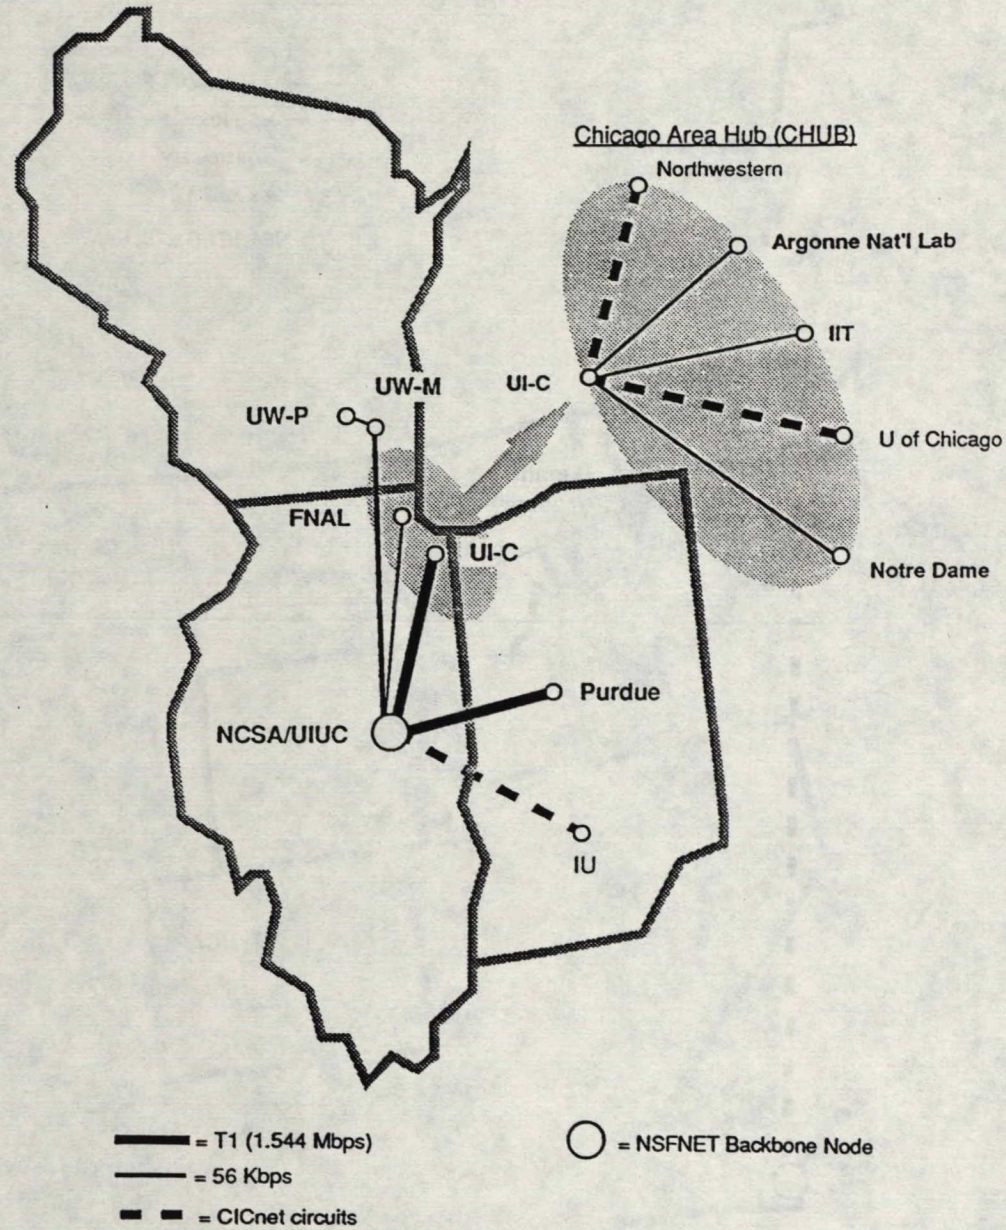

Source: Users' Directory of Computer Networks; Compiled by Tracy Lynn LaQuey, University of Texas; July 1989.

NorthWestNet

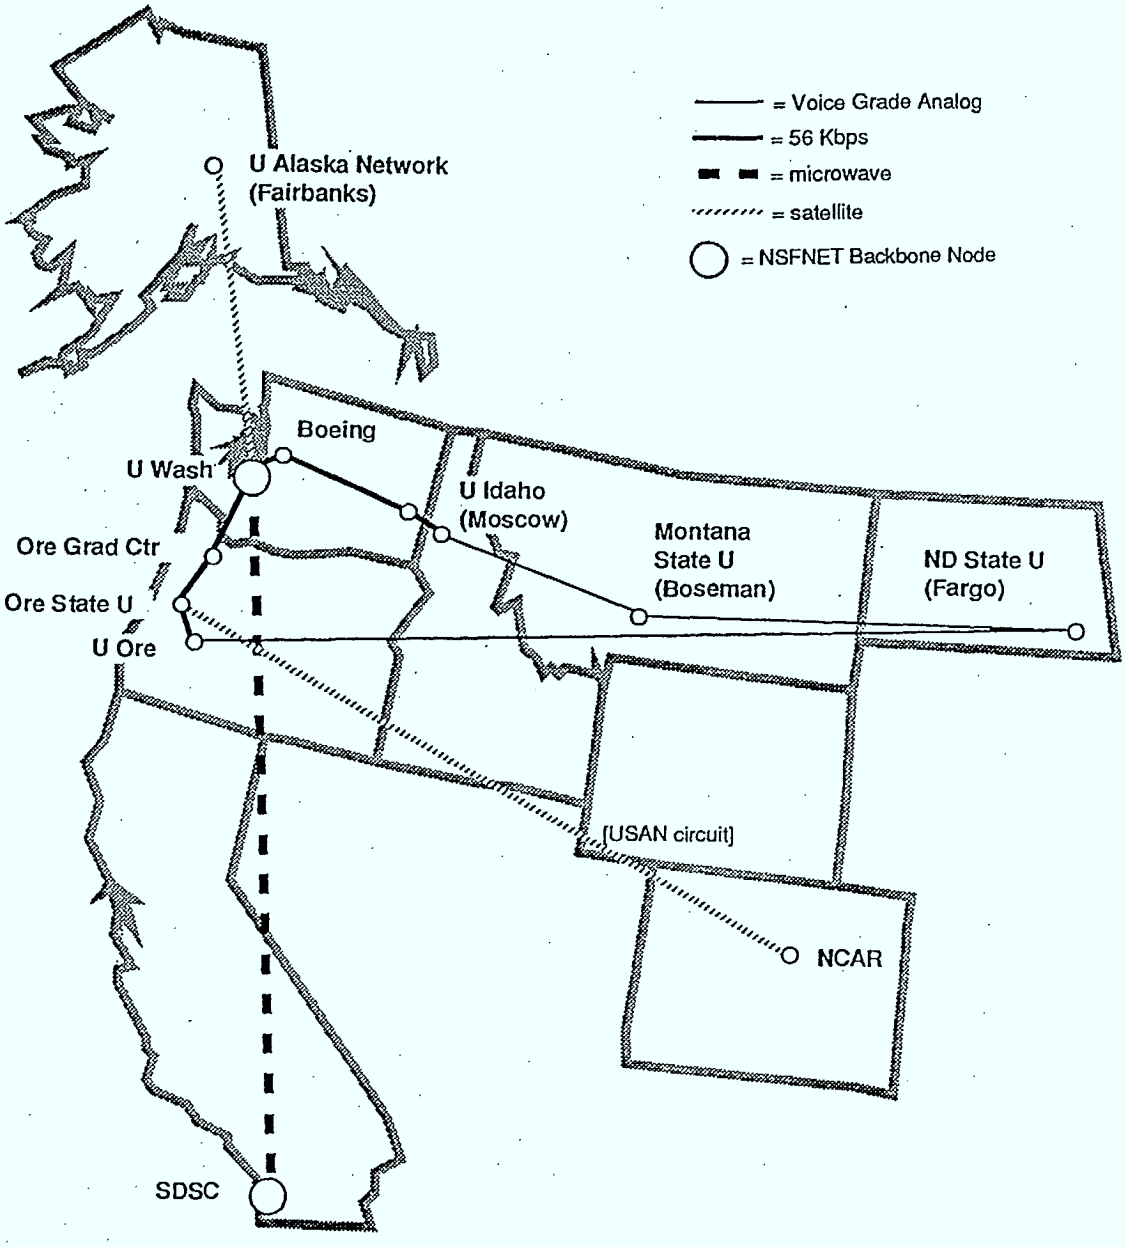

Source: Users' Directory of Computer Networks; Compiled by Tracy Lynn LaQuey, University of Texas; July 1989.

NYSERNet

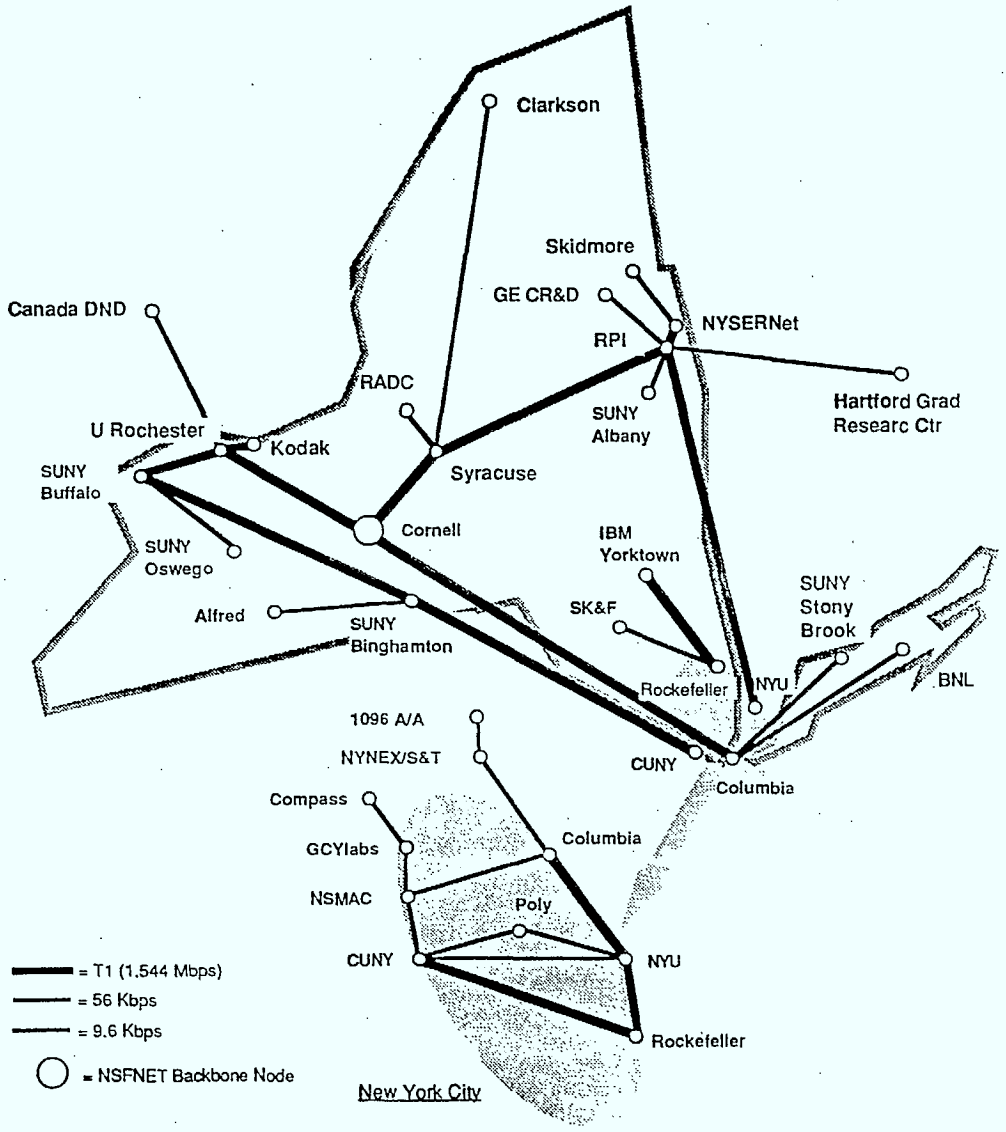

Source: Users' Directory of Computer Networks; Compiled by Tracy Lynn LaQuey, University of Texas; July 1989.

OARnet

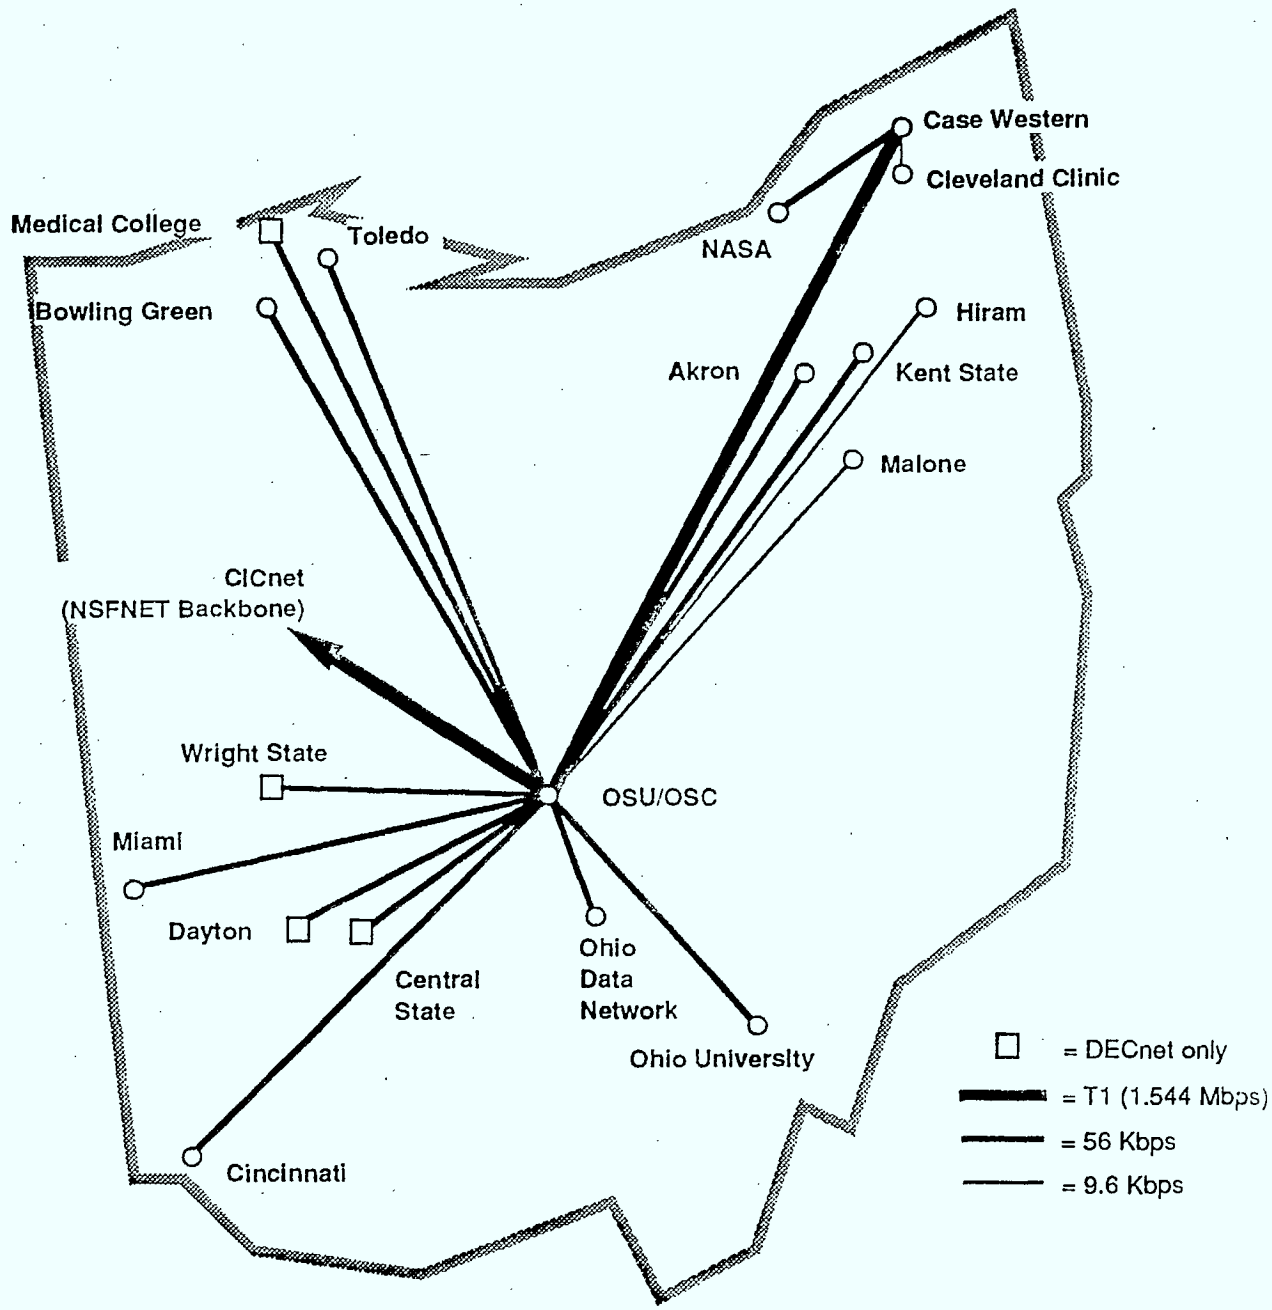

Source: Users' Directory of Computer Networks; Compiled by Tracy Lynn LaQuey, University of Texas; July 1989.

PREPnet

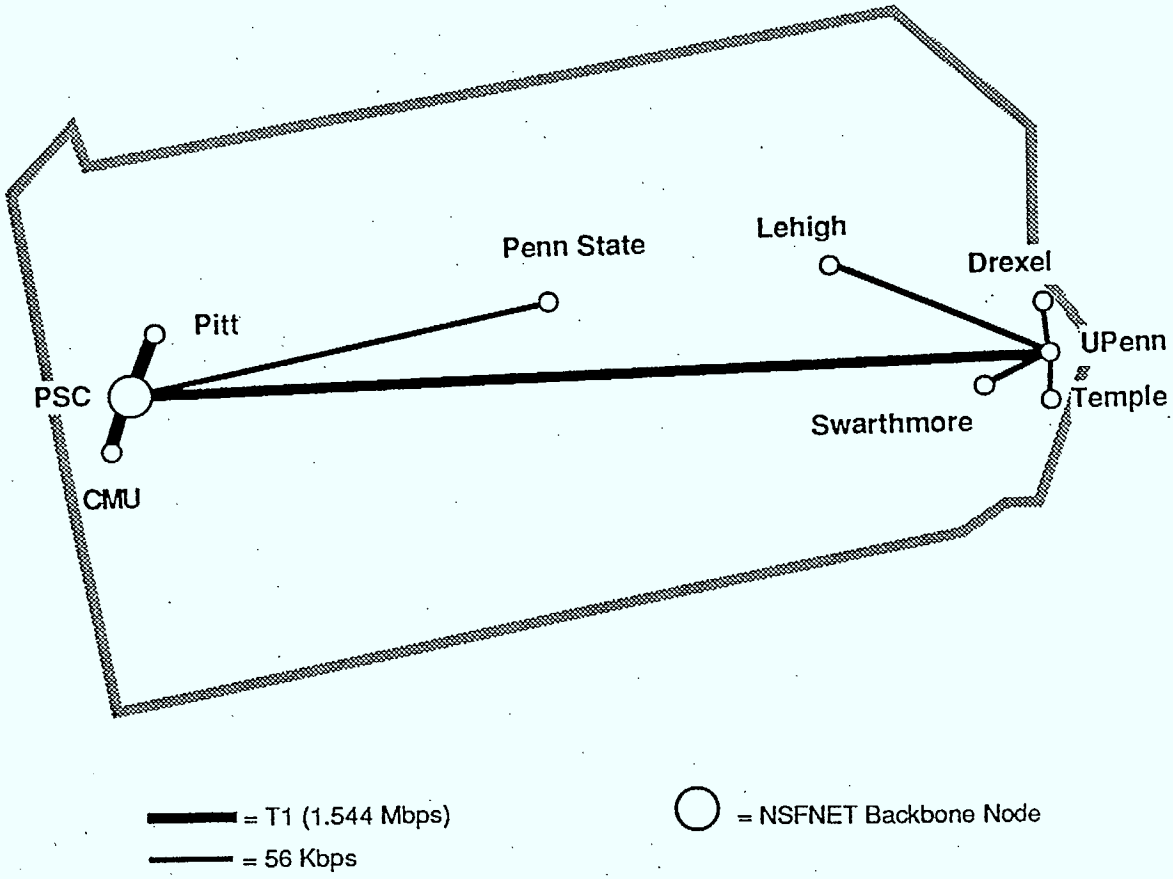

Source: Users' Directory of Computer Networks; Compiled by Tracy Lynn LaQuey, University of Texas; July 1989.

SDSCnet

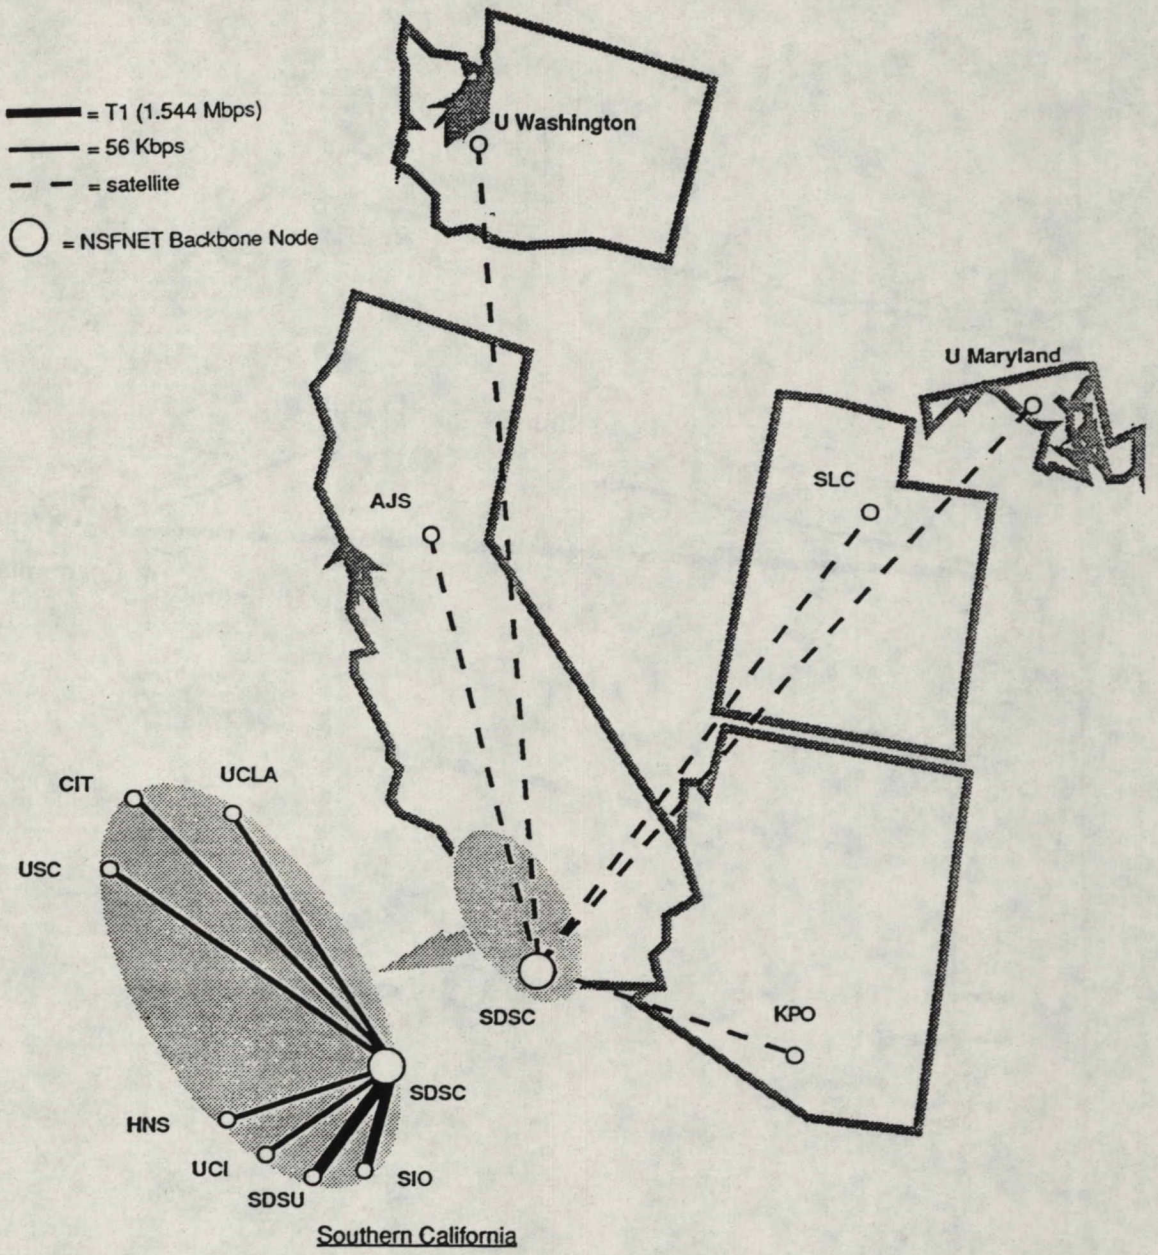

Source: Users' Directory of Computer Networks; Compiled by Tracy Lynn LaQuey, University of Texas; July 1989.

SESQUINnet

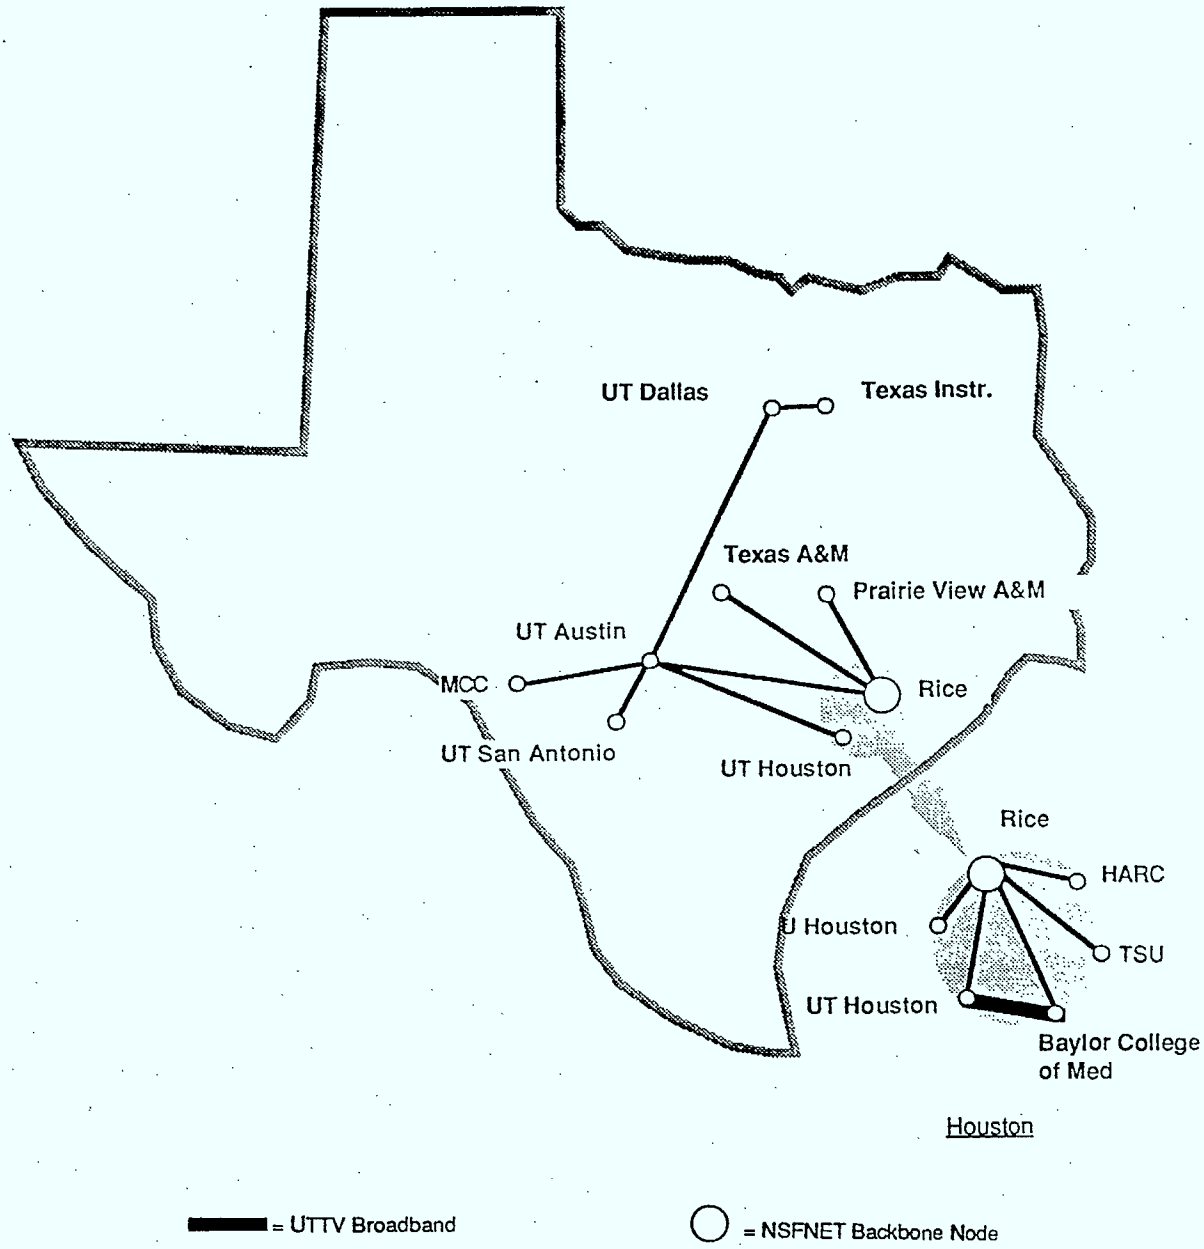

Source: Users' Directory of Computer Networks; Compiled by Tracy Lynn LaQuey, University of Texas; July 1989.

SURAnet

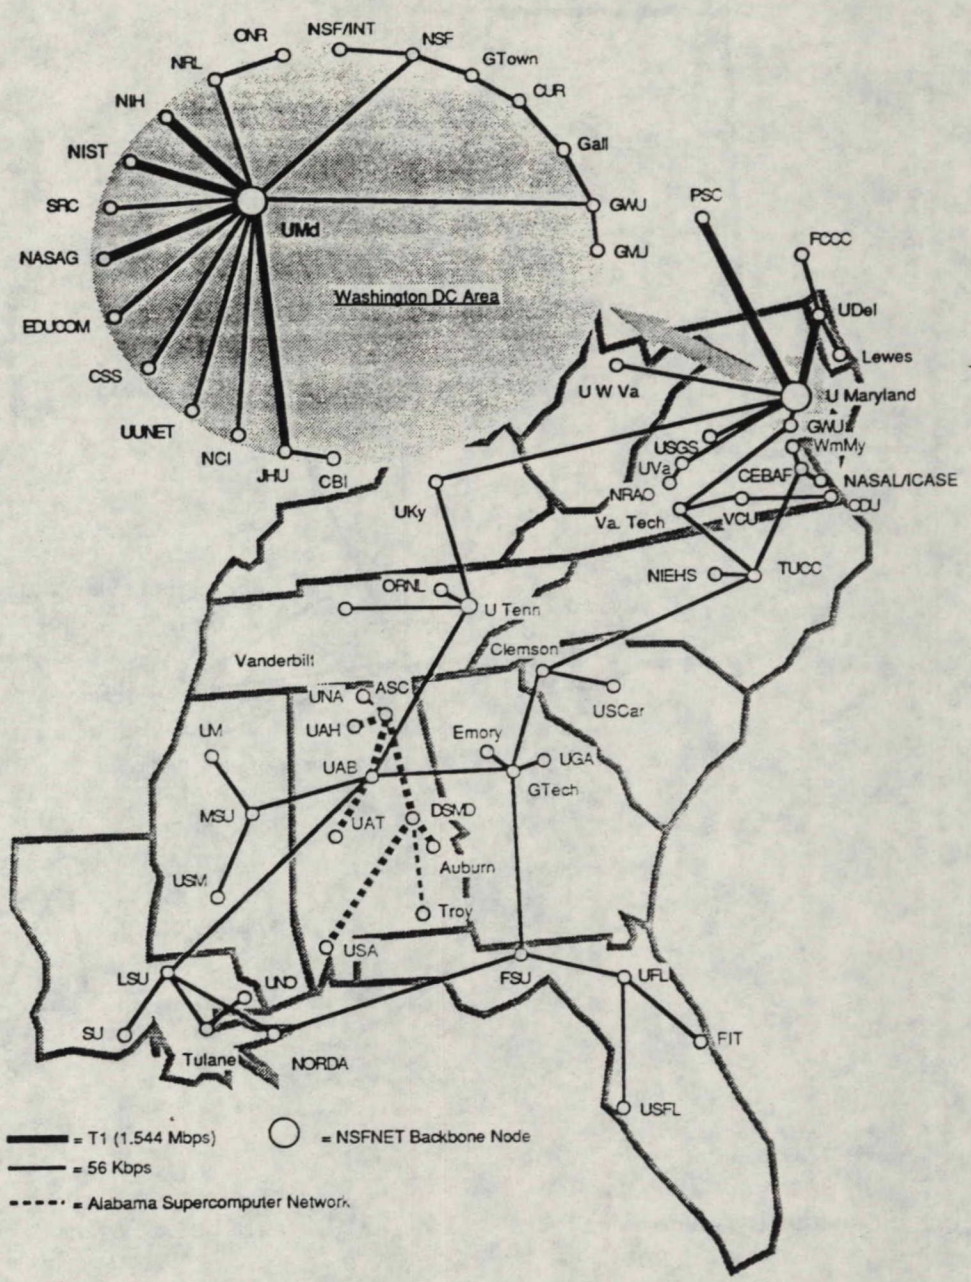

Source: Users' Directory of Computer Networks; Compiled by Tracy Lynn LaQuey, University of Texas; July 1989.

THEnet

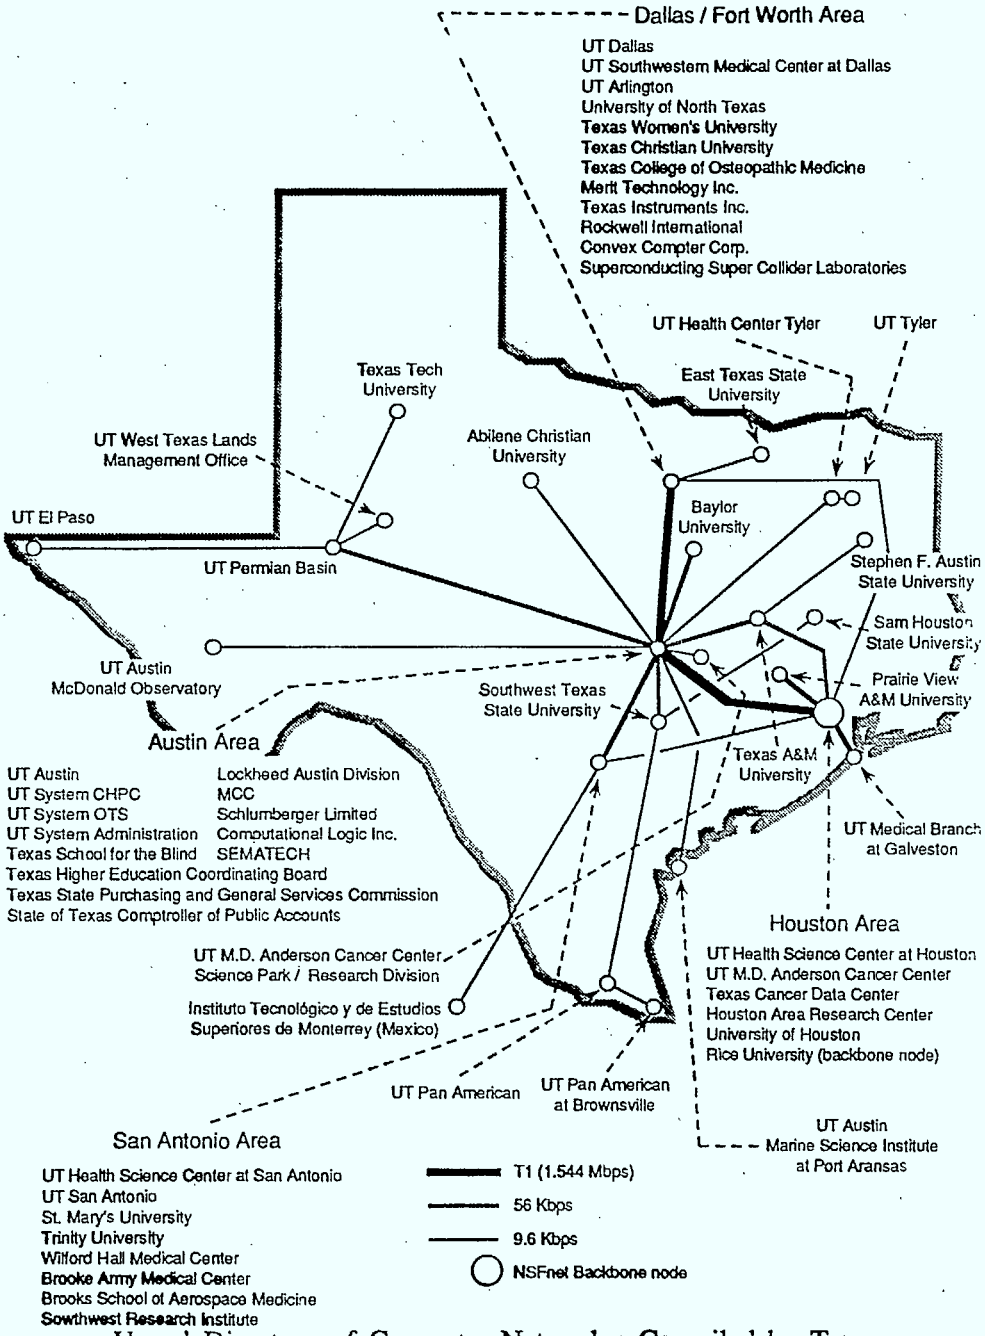

Source: Users' Directory of Computer Networks; Compiled by Tracy Lynn LaQucy, University of Texas; July 1989.

WestNet

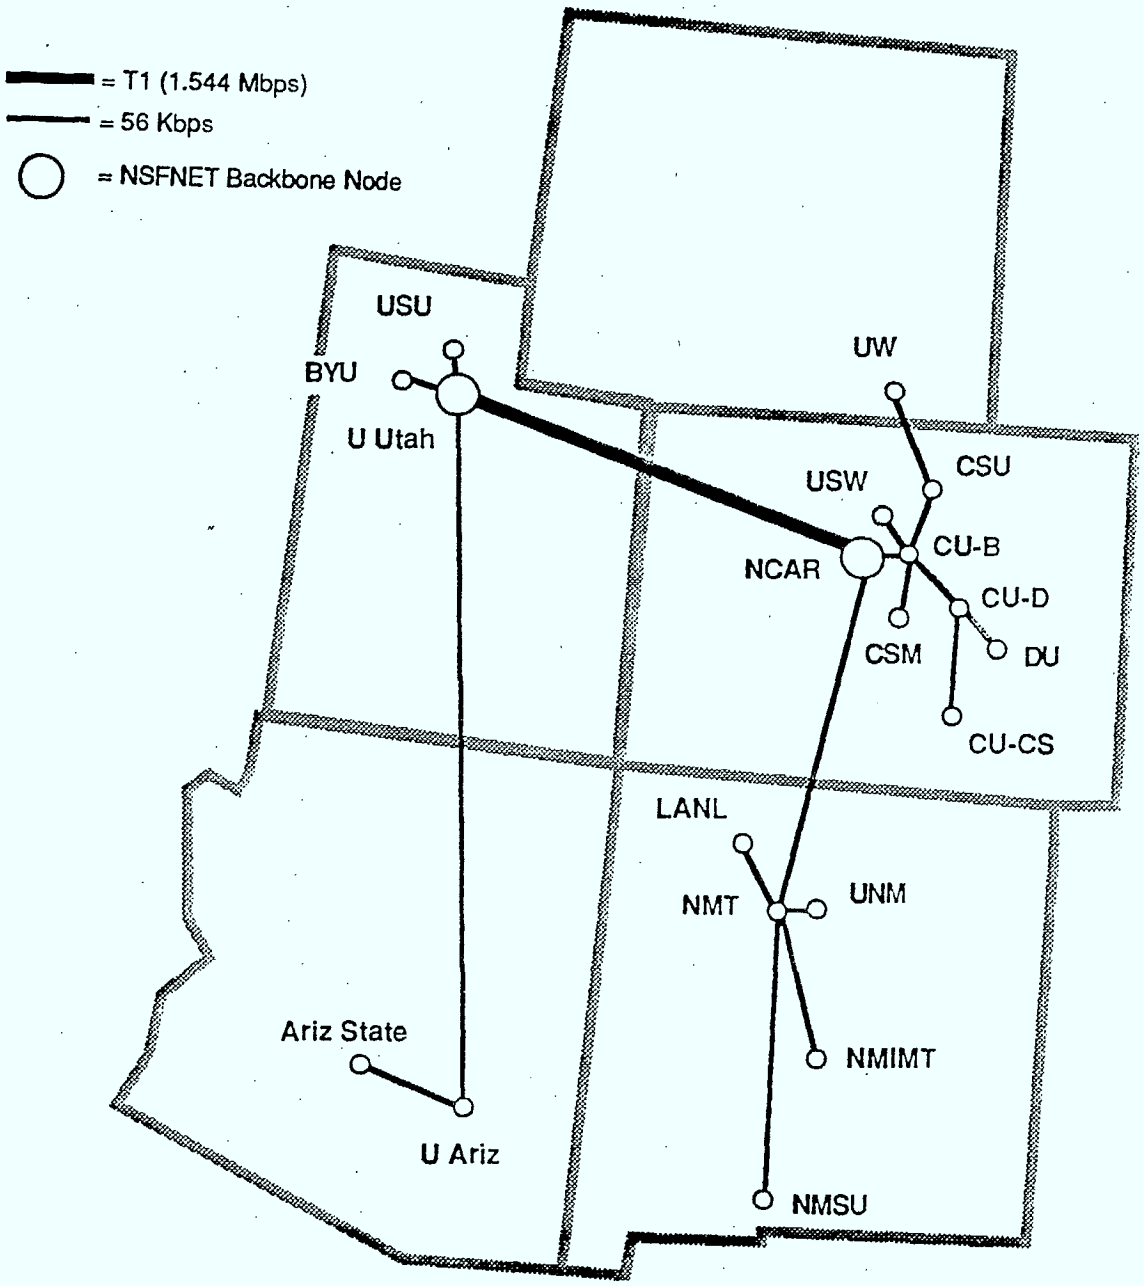

Source: Users' Directory of Computer Networks; Compiled by Tracy Lynn LaQuey, University of Texas; July 1989.

3.5.4 Department of Energy (DOE) Research Networks:

(a) Magnetic Fusion Energy Network (MFENET)

Background

- Original purpose to connect physics departments doing research in nuclear fusion (has since expanded to reach all DoE programs);
- Originated in mid 1970s;
- About 120 hosts on the network — five super computers can be reached;

Technical Design

- Mail, file transfer, remote command execution and remote login are all supported;
- Specialized remote procedure calls used for interactive graphic terminals using special purpose protocols (called NSP);
- The use of nonstandard protocols has lead to interoperability problems with other networks — changing NSP suite over IP protocol and eventually move to ISO–OSI protocols (such as CLNS). The replacement of MFEnet will be MFEnet II;
- Existing links range from 9600bps to 56Kbps leased lines to 112Kbps satellite links;
- Speed depends on intervening links;
- Reliability is high;

Implementation Design

- Funded and administed by DoE;
- Managed from National MFE Computer Centre (NMFECC) of Lawrence Livermore National Laboratory (LLNL);
- Access restricted to DoE funded researchers;
- ESnet is to provide the high speed backbone carrier for MFEnet and HEPnet as well as transition to MFEnet II;

(b) Energy Science Network (ESNET)

Background

- Energy Science Network (ESNET) is currently under construction but is operating in a limited way;
- ESNET is a backbone network for all DoE energy research programs;

HEPnet (ESnet-DECnet) TOPOLOGY
As of June, 1989

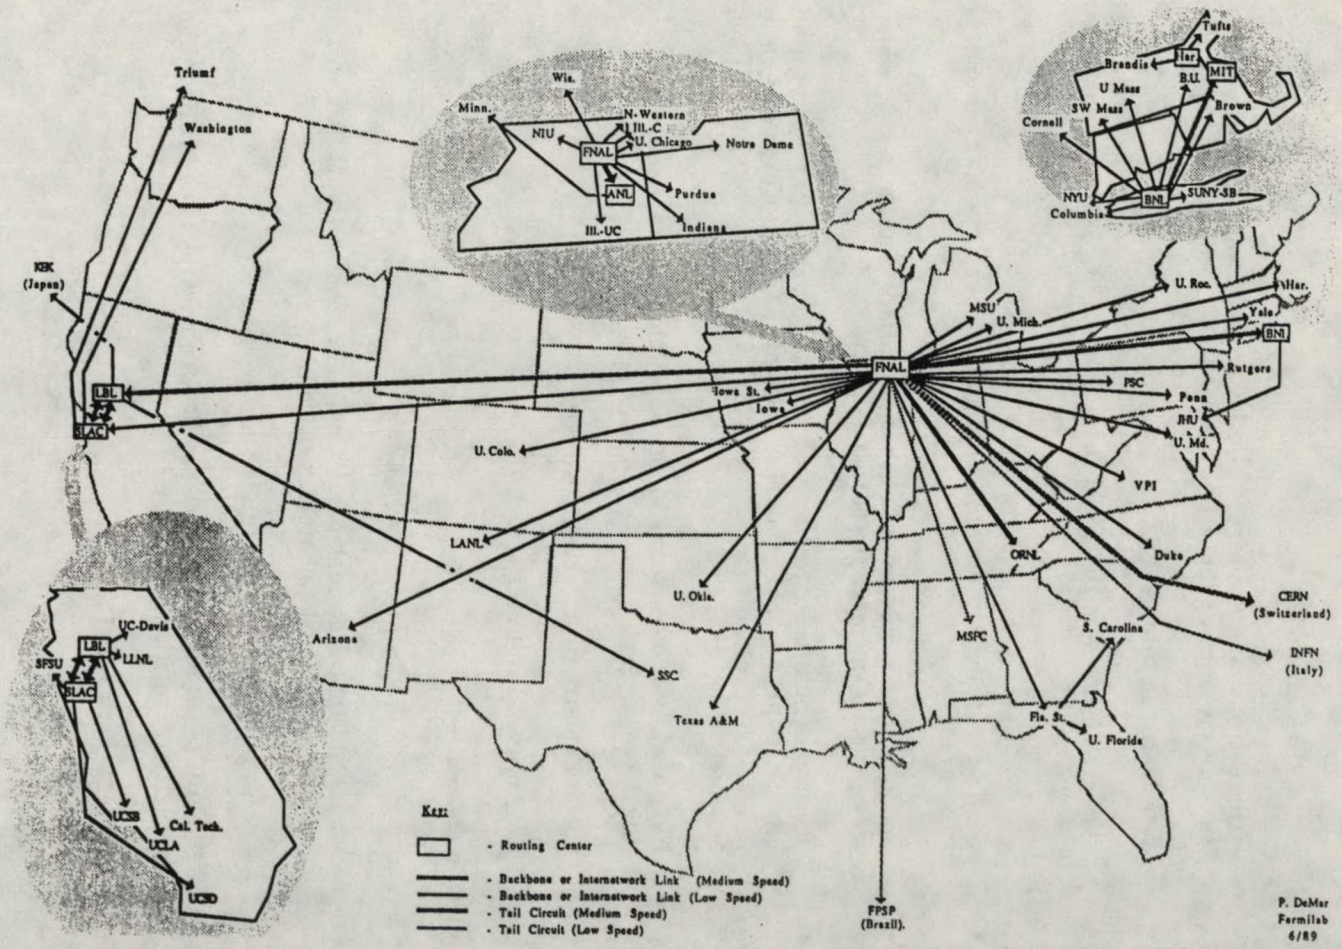

Source: Users' Directory of Computer Networks; Compiled by Tracy Lynn LaQuey, University of Texas; July 1989.

- ESNET is a replacement for the existing High Energy Physics network (HEPnet) and the Magnetic Fusion Network (MFEnet);
- The purpose of ESNET is to provide energy researchers around the world the computational power, interconnectivity and file storage to perform the immense simulations required for magnetic fusion and other energy related problems;

Technical Design

- MFEnet used the nonstandard NSP protocol which does not conform to the TCP/IP based protocols in use on the Internet. ESnet is designed to overcome these limitations;
- Both X.25 and IP will be provided in ESnet which will support TCP/IP, DECNET and the MFEnet NSP protocols;

Implementation Design

- Access restricted to projects sponsored by DoE;
- Installed and operated by the NMFECC of LLNL;
- Policy decisions are made by steering committee comprised of represented of each DoE research program;
- In the final phase of implementation ESnet will consist of MFEnet, MFEnet II and the X.25 backbone;

3.5.5 National Aeronautical and Space Administarion (NASA) Research Networks

- NASA Science Internet (NSI) is composed of two complementary parts: a TCP/IP component called the NASA Science Network (NSN) and a DEC–net component called the Space Physics Analysis Network (SPAN). NSI is connected to ARPANET, NSFNET, CSNET, BITNET and numerous other networks;

(a) NASA Science Internet (NSI)

Background

- Is connected to ARPANET, BARRnet, NSFNET, CSNET, BITNET, HEPNET;

Technical Design

- A PSCN T1 backbone installed in November 1988 with engineering provided by Boeing Computer Services (BCS);
- There is some redundancy built in;
- T1 backbone managed as pools of 56Kbps links allocated primarily to voice and data and switched dynamically as the need arises;

NASA Science Network
Gateway Interconnections

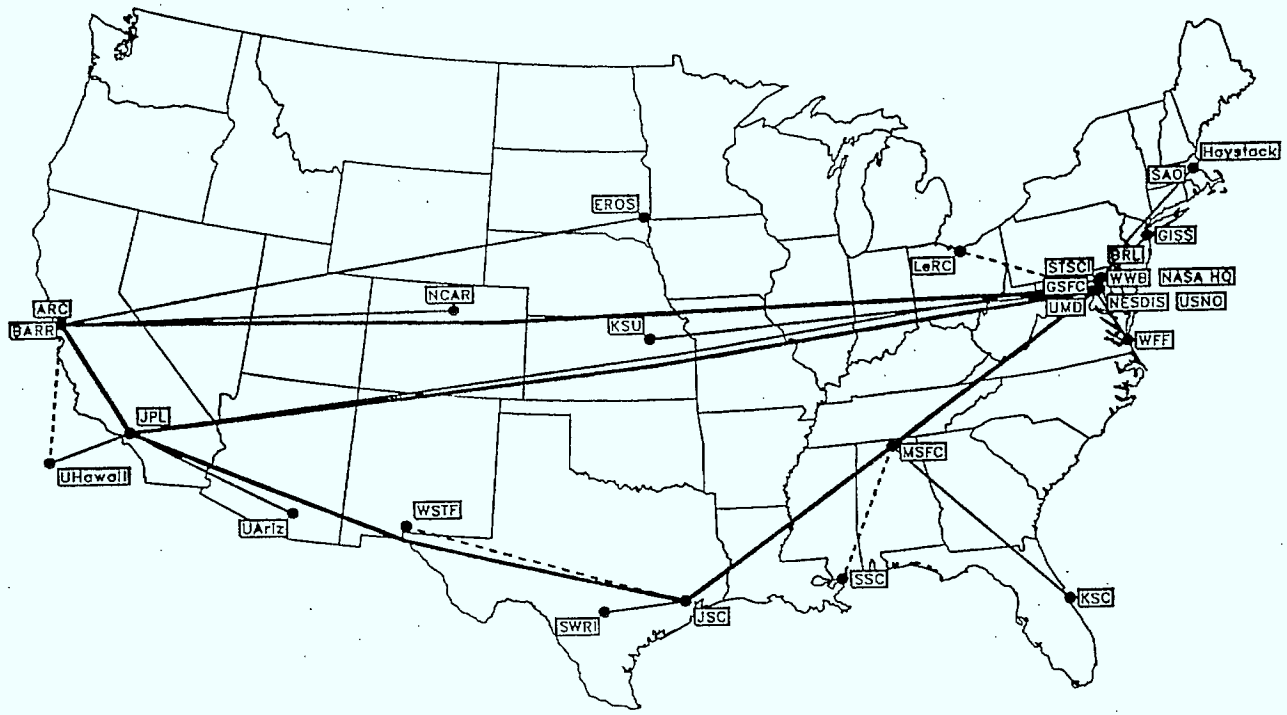

Source: Users' Directory of Computer Networks; Compiled by Tracy Lynn LaQuey, University of Texas; July 1989.

- Switching is provided by 14 roadrunner switches. There are eight earth stations;
- There are plans for eventual transition to ISO–OSI protocols;

(b) **National Science Network (NSN)**

Background:

- NASA Science Network (NSN) installed in early 1987, and grew out of a project called the Pilot Land Data System;
- Serves NASA flight projects and all NASA disciplines and about 30 backbone hosts;

Technical Design

- TCP/IP internetwork;
- Uses the PSCN infrastructure;
- There are plans for eventual transition of NSN and SPAN to use ISO–OSI protocols;
- NSN is configured with major hubs at five NASA centres, with each major hub connecting to at least two other hubs;

Implementation Design

- Managed by NASA Science Internet Project Office (NSIPO) on behalf of the Office of Space Science and Applications (OSSA);

(c) Space Physics Analysis Network (SPAN)

Background

- SPAN is a nation wide DECNET network which began operations in 1981
- SPAN interconnects space and earth science researchers.
- SPAN is an operational network as it is not intended to promote the development of new network technology, but it is a research network providing an infrastructure for space related research.
- Has connections to Japan, Canada, and many countries in Europe
- There are close to 3,000 SPAN host computers.
- SPAN is predominantly general purpose, but has been used to support specific missions (such as Haley's Comet),
- Access to SPAN is relatively inexpensive and it is easy to connect to.

NSI/SPAN
North American Sites

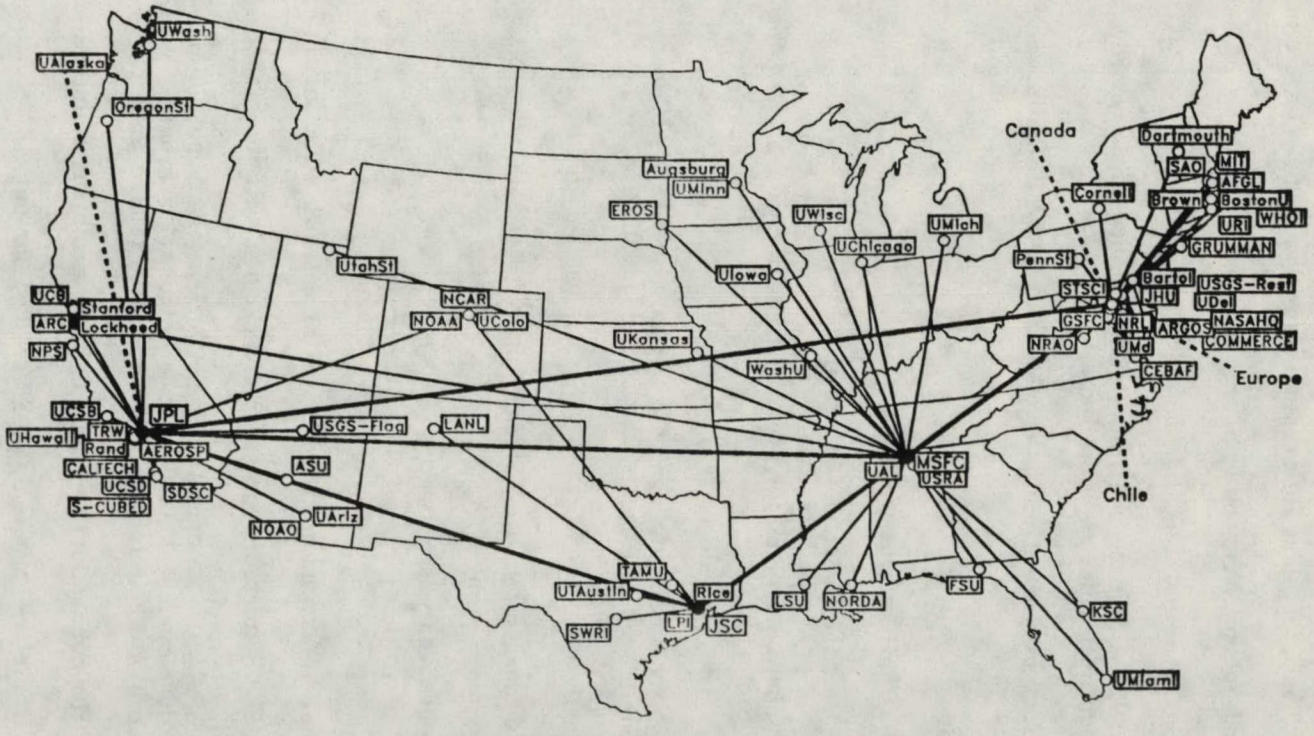

Source: Users' Directory of Computer Networks; Compiled by Tracy Lynn LaQuey, University of Texas; July 1989.

Technical Design

- SPAN is based on DECNET Phase IV and will migrate to Phase V for ISO–OSI compatibility (Lower layers provided by PSCN);
- Backbone consists of redundant 56Kbps links among five routing centres;
- There are many local area networks indirectly connected to SPAN;
- Direct gateway to NSFNET and BITNET exists;
- Migration to ISO–OSI protocols are planned.

Implementation Design

- Guidance for the network provided by users through Data System's Users Working Group;
- Administration done by project managers, network managers and routing centre managers;
- Various offices of NASA fund SPAN;
- NASA science communications requirements are usually met with 9600 bps circuits. Higher bandwidth or more extensive services require project–oriented funding.

3.5.6 Special Interest Groups

Special interset group networks are loosely defined and include many different administrative philosophies. However the common ground between these networks is that they are run by a subset of their own users and not by any outside agency. Most of the funding comes from the sites, hosts and users they connect.

(a) CSNET

Background:

- CSNET established in 1981 to facilitate research and increase collaboration in order to accelerate the development of computer science and engineering technologies;
- There are about 180 hosts overall (many serve as gateways into internal company networks or national networks);
- CSNET is considered to be a NSFNET mid–level regional network although it is not directly connected to the NSFNET backbone;
- The success of CSNET led to the development of NSFNET;

Technical Design

- Mail is only service supported on all parts of CSNET because developers of CSNET noticed that electronic mail was the most popular service on the ARPANET and thus their aim was to connect institutions that did not have access to ARPANET to those that did;

- As network evolved it was realized that mail was not enough, and thus CSNET has become a metanetwork composed of several parts that vary in their additional services (ie. some support remote login and file transfer), lower level protocols, speed and reliability;

Implementation Design:

- CSNET administered by CSNET Coordination and Information Centre (CSNET CIC) at Bolt Beranek and Newman (BBN) in Cambridge Massachusetts;
- CSNET CIC provides technical, operational, administrative and end user support;
- The University Corporation for Atmospheric Research (UCAR) under contract to NSF, oversee BBN's operation of the CIC;
- NSF provided initial funding for the establishment of CSNET. The total project budget was $5 million over the five years 1981 through 1985. There is no current funding from NSF and the network has been self–supporting since 1985;
- Annual dues are collected from member organizations with rates set according to academic or industrial classifications. Member organizations must also pay for hardware and communications costs directly;
- The CSNET executive committee and the BITNET board of trustees are in agreement that CSNET and BITNET should merge (using the name ONENET). This merger reduces redundancy in management and service and provides a symbol of a step forward towards a single national research network;

(b) BITNET:

Background

- In 1981 BITNET (Because It's Time Network) started as a small store–and–forward network of IBM computers centred at the City University of New York (CUNY);
- The membership includes academic institutions and collaborating research institutions and serves more than 2300 hosts at several hundred sites in 32 countries. (BITNET refers to the combined constituency of the U.S. and Mexican part, NetNorth (Canada) and EARN (the European Academic Network));
- BITNET established for facilitating non–commercial exchange of information, via electronic mail transfer of documents, programs and data and access to BITNET server machines which provide information retrieval services;
- Services provided are electronic mail; file transfer; and interactive massaging;

BITNET in ASIA TOPOLOGY
As of February 6, 1989

Source: Users' Directory of Computer Networks; Compiled by Tracy Lynn LaQuey, University of Texas; July 1989.

Technical Design:

- Gateways exist between BITNET and the Internet, CSNET and many other networks worldwide;
- BITNET uses Network Job Entry (NJE) protocol (came at no extra cost with IBM mainframe systems);
- This is a store and forward network;
- Hosts are interconnected by leased telephone lines at 9600bps;
- Gateways exist between BITNEt and CSNET and the Internet;

Implementation Design:

- Online directory, electronic newsletters, end user documentation, workshops, seminars and conference presentations are made available are services provided by BITNET Network Information Centre (BITNIC);
- The CSNET executive committee and the BITNET board of trustees are in agreement that CSNET and BITNET should merge (using the name ONENET). This merger reduces redundancy in management and service and provides a symbol of a step forward towards a single national research network;

3.6 FUTURE NETWORKS

This section describes the future network activities in the U.S.

3.6.1 The National Research and Education Network (NREN)

Background

- NREN represents the information superhighway of the future;
- The NREN is currently before congress;
- A report completed by The Computer Research and Applications Committee, in 1987, concluded that the U.S. was lagging behind other countries in establishing a national research network to connect every research network in the United States. According to the Federal Research Internet Coordinating Committee (FRICC) the current network does not adequately support collaboration or access to unique scientific resources, nor is the long-distance bandwidth available for real time system to system collaboration;
- According to a recent Computerworld article (August 1989) "the lack of high-speed links has caused researchers to limit or even table certain projects that require major collaboration across distance";
- The missions of NREN and NSFNET are quite different. The purpose of NREN is that it attempts to support the needs of all scientific networking whereas NSFNET is geared more to providing access to SuperComputers;

- The ultimate goal of NREN is to provide access to all scientific instruments, databases, and that the network have sufficient band width that users appear to be attached to a world local area network;

Technical Design:

- Presently research groups communicate with one another and with remote super computer centres through haphazardly, overlapping and interconnected regional networks each of which has its own user interfaces, high–level communication protocols and access methods;

- NREN would potentially iron out protocol differences and interconnect these networks over a single, high speed, streamlined and compatible backbone system;

- The TCP/IP protocol would be used at the outset to promote reliability;

- The main focus of development will be on gateway standards to facilitate interconnection of the networks;

- Physically the early phases of NREN would evolve from today's networks; whereas the ultimate phase would be a major departure from today's technology and would require advanced research to be successful;

- Project currently underway which involves a "multi–Gb/sec optoelectronic data transport network" which would allow packets to remain in optical form throughout the transmission period (including switching). At present data carried in optical form must be converted back to electronic mode for switching (two to five years away from testing this principle using 24 nodes);

Implementation Design

- NREN would interconnect networks owned by the various federal agencies to promote resource sharing and collaboration among researchers;

- The NREN proposal has two distinct parts: (1) An operational Network providing reliable service and rapid transfer of very large data streams; and (2) R&D on a gigabit network to be deployed in the latter 1990s. The NREN project envisions a three phase buildup to 3G bit/sec target, within ten to fifteen years;

  – **Phase 1:** Several 1.5 Mbs backbone lines with regional networks operating at 1.5 Mbs, and institutional access at 56 kbs. Selected institutions to have access at 1.5 mbs. This Phase would take approximately two years. Much of what is include in Phase 1 is already underway as substantial pieces are already operating at T1.

  – **Phase 2:** Would be an upgrade of the Internet to provide 1.5 Mbs backbone lines to 200 or 300 research facilities across the nation. Some of these facilities are anticipated to be industrial laboratories. This phase would require several 45 Mbs backbone lines and is projected to be complete in five years. Development and implementation costs are estimated at $5 million per year with operating costs of $40 – $50 million per year. Phase 2 is within the

capability of industry to provide.

- **Phase 3:** 1.5 Mbs service to nearly every research facility, 45 Mbs to 1000 sites and 1 to 3 Gbs to selected sites. Development costs are projected at $400 million per year spread over 10 years with operating costs of $200 million are anticipated. This phase could be deployable in 10–15 years. Phase 3 requires substantial research on interfaces and protocols, switching and advanced software applications that can utilize speeds above a billion bits/second.

- Funding for these two initiatives is projected at $400 million over five years; roughly 60% for gigabit research and the other 40% for Stage 1 and 2. Refer to the following five year reference levels in the table below:

	YR 1	YR 2	YR 3	YR 4	YR 5
	IN MILLIONS OF CURRENT $				
Interagency Interim NREN	14	23	55	50	50
Gigabit R&D	16	27	40	55	60
TOTAL	30	50	95	105	110

- Accounting methods and access restrictions are still under study;
- NREN could build upon nearly 24 existing low–speed agency networks and make use of several national high speed fibre–optic networks;
- NSF has assumed responsibility for supporting a backbone for the NREN and will coordinate collaboration among Federal agencies;
- DARPA lead agency for advanced networking technology research and development leading to advanced networking technology gor gigabit speeds;
- DOE will provide networking support for energy research community;
- NASA will provide networking support for the aerospace research community;
- NIST will participate by establishing networking standards, with emphasis on protocols and security standards;

4. R&D NETWORKS IN EUROPE

4.1 SUMMARY

In the words of New Scientist magazine, "Europe is a divided house of fragmented, small markets and small, mutually uncomprehending academic and industrial research communities". There are a number of factors unique to the European situation which complicate communications across national borders, such as language, national security concerns, and differing telecommunications systems, protocols, and equipment.

As a result, much of the network capability in Europe today links research communities within each nation, with Pan–European links between major nodes of each national network. These networks are generally small, low–capacity, and intended mainly for news and mail transfer.

However, the increasing importance of the European Community (EC) indicates that the current fragmentation and low performance of data communication networks in Europe will quickly become a thing of the past. By far the largest portion of the recently–proposed five–year "Framework" for EC–financed research is the $5 billion for "Technologies for Information and Communication"; a major part of this will be aimed at computing and telecommunications.[2]

A recent report to the European Community by RARE, an association of European researchers using supercomputers, proposes a European High–Speed Networking Initiative which would allow remote users better access to supercomputer power at various locations around the EC. This would improve European work in areas such as combustion analysis, astronomy and space physics, fusion research, the automobile industry, and meteorology.

4.2 EUROPEAN R&D OVERVIEW

Europe's overall R&D effort has lagged behind that of the United States and Japan in recent years, both in terms of financial support as well as productivity. This is more a reflection of the increase in R&D in those two nations rather than a decline in European effort. However, the fragmentary nature of national research efforts combined with significant communication barriers has resulted in much duplication and wasted effort.
The annual financial effort of the Western European nations devoted to R&D is in the range of $100 Billion. The European Community, financed by the EC member nations, directs approximately $2 Billion of this under its current "Framework" program, which expires in 1992. A recent article describes the EC effort as follows:

> "Brussels is the source ... of some of the most efficient scientific support in Europe. This is because the EC does not so much buy test–tubes as see that they are well used. The Community adds value to basic research money. Its money helps scientists to co–ordinate work, share data and often spark off each other. It finds the gaps and overlaps in European science, and tries to fill and prune them.
>
> Most of all, the Community overcomes the communications barriers that plague European science. It arranges meetings, studentships abroad and language training ... The Community brings industries together for research and, anathema on the Continent, even arranges joint research between academia and industry."[3]

Recently, a successor "Framework" has been proposed for the years 1990–1994, in which the areas of R&D have been cut down to six. Nearly one half of the budget of $10 Billion would be spent on "Technologies for Information and Communications", an area which includes programs in computing and telecommunications.

In most countries, a large proportion of R&D is government–funded and takes place in government–run research institutes and universities. Figure 4.1 shows the distribution of government funds for R&D in five important EC member countries. Generally, the high government involvement in R&D has produced a greater gap between successful research and implementation by industry with a major exception being West Germany. The EC efforts will help to break down many of the technology transfer barriers.

One of the major problems facing European countries, especially Britain, over the next decade is an inability to attract promising people to science and engineering careers, due mainly to greater opportunities and higher salaries in financial and service industries. Moreover, demographic change is resulting in ageing populations requiring greater government–provided social services. Providing these services will strain the resources of the European community and put further pressure on government financing for R&D. This may lead to a shortage of talented young scientists in research, and engineers in key industries.

In summary, the European R&D community faces many of the same problems as that of the United States and Canada. However, the greater importance of the EC will result in better coordination and communication of R&D efforts and more joint efforts across national boundaries, and the large markets available to EC–based companies after 1992 should enhance private industry's ability to profit from R&D. These factors will make high–technology research more affordable. National interests will ensure that European R&D is never highly centralized and, therefore, the need for high–quality R&D data communication networks will become critical in the next decade.

4.3 EUROPEAN NETWORK ENVIRONMENT

The environment within which data communications networks must function in Europe poses significant barriers to efficient networking. One of the most obvious is cost: charges for high–speed data transmission lines are much more expensive then costs in the U.S..

In most European countries, telecommunications is a monopoly of the government–run postal and telecommunications authority (PTT). Each PTT has its own lines and equipment, and their services are, in general, considered slow, costly, and unreliable. Protocols supported vary from country to country. In many countries, the PTT must approve the use of any piece of equipment connected to its lines; they are often slow about doing so, and a common complaint is that researchers are unable to use the latest technology because it is not approved for use by the PTT. The PTTs are also under political pressure and sometimes legal constraint to use equipment made by domestic manufacturers.

It is expected that the European Community will oversee the integration and standardization of telecommunications within Europe, according to the ISO/OSI standards, and will be able to apply pressure to the national PTTs to improve and standardize their operations.

The major European coordinating group for networking is called Reseaux Associes pour la Recherche Europeenne (RARE). The objective of this organization is to support national academic network operators in migrating to international standards. RARE was started by a group of enthusiastic experts meeting in Luxembourg in 1985. It has been assigned responsibility for the specification phase of COSINE (Cooperation for Open Systems

FIGURE 4.1

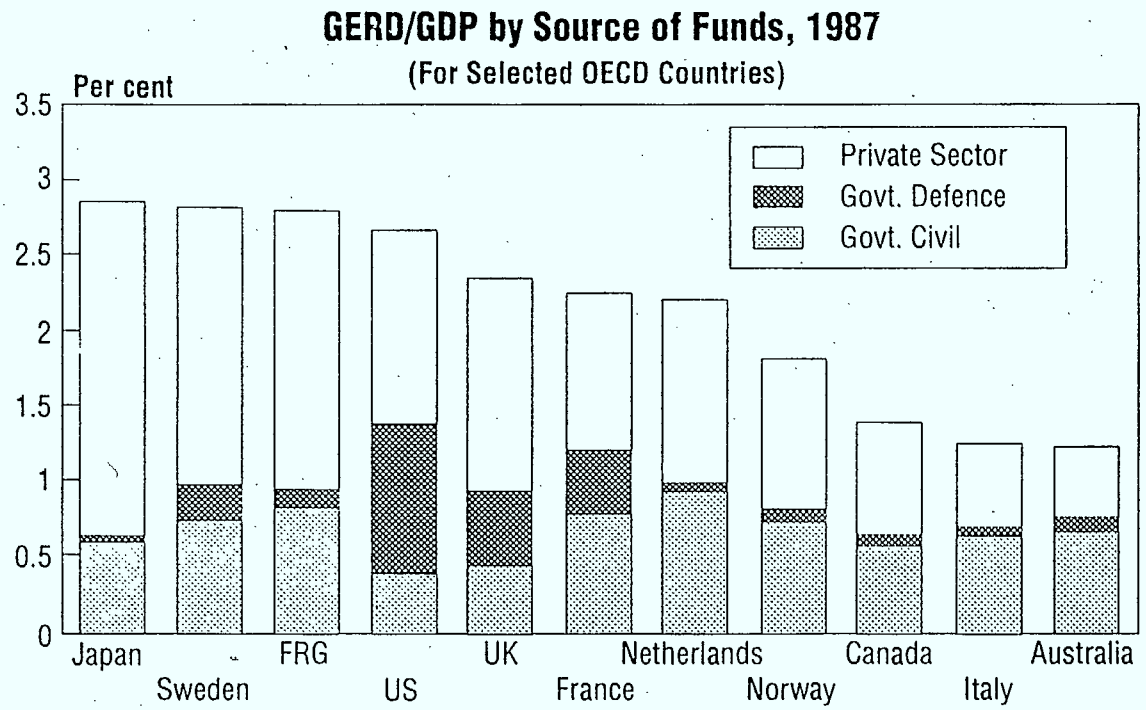

Source: ISTC; Selected Science and Technology Statistics; 1989

Interconnection Networking in Europe), a project of the European Economic Community (EEC) within the EUREKA program, which is a predecessor to its Framework R&D program. The COSINE project is expected to assist approximately 500,000 users in the industrial and academic research community. RARE has developed close relationships with two major Pan–European networks, EARN and EUnet. The RARE group has also been involved in harmonizing network links with the U.S., and contributed to the EEC Green Paper on the Common Market for Telecommunication Services and Equipment".

Current membership in RARE includes 18 of 20 possible European countries, and application requests from many non–European countries indicate the interest, acceptance, and trust placed in RARE's integration work. RARE has a permanent secretariat located in Amsterdam, where all documentation concerning its activities is available. The following is a short list of working groups within RARE:

- WG1: Message Handling System based on ISO X400; the group is charged with mapping current message handling systems into X400; and coordinate the operational R&D MHS service.

- WG2: File Transfer, Access and Management (FTAM) is involved in establishing conversion between existing de facto standards and ISO's FTAM.

- WG3: Information Services Exchange of Operational Information seeks to establish a general mail directory.

- WG4: Network Operations and X.25: This group examines current national network operations in the different countries, most importantly to coordinate international operation of these services. The most important decision taken so far is to upgrade the initial speed of 64 Kb/s to 2 Mb/s as soon as feasible.

- WG5: Full Screen Services: This group's objective is to provide full–screen terminal emulation on host systems. Their plans are as follows:

  - Short–term: Improve performance of CCITT Triple X
  - Medium–term: Apply ISO's VT Standard
  - Long–term: Provide open standards for windowing and bit–map oriented interactions.

- WG6: Medium and High–Speed Communications: This group supports the development of higher speed facilities by the European PTTs; for example X.25 networks faster than 64 Kb/s, switched circuit services, satellite links with 64 Kb/s or 2 Mb/s. Progress in this area has been faster than expected.

- WG8: Management of Network Application Services: Network management, security, authentication, and accounting are some of this group's concerns.

Most of the academic networks in Europe are not professionally operated and managed. This is considered a serious problem because of the wide variations in equipment and protocols in use. Industrial participation in academic networks has been minimal due to concerns about reliability, security, speed, and management. Instead, companies tend to use private value added networks (VAN) offered by large computer vendors or private carrier companies, or they build their own networks based on lines leased from the PTTs.

RARE is involved in talks with the three major pan–European networks (EUnet, EARN, and HEPnet) about coordinating and integrating services throughout Europe, including shared gateways and infrastructures.

4.4 SCOPE OF THIS REVIEW

There are a wide range of networks within the main countries of Europe. This study will be restricted to discussing the major Pan–European research networks, and the major research networks within each country. The following list summarizes the European networks described in this section:

I EXISTING NETWORKS

(a) Pan–European Networks:

1) EUnet;
2) EARN;
3) HEPnet;
4) EAN;
5) NORDUnet

(b) Major National Networks:

1) United Kingdom
2) Federal Republic of Germany
3) France
4) Italy
5) The Netherlands
6) Ireland
7) Spain
8) Denmark
9) Finland
10) Norway
11) Sweden
12) Switzerland
13) Austria
14) USSR

NEW FUTURE NETWORKS

1) VENUS
2) Germany: Forerunner
3) U.K.: Cambridge Fast Ring
4) Italy: GARR

Figure 4.2 summarizes the characteristics of the major continental and national networks in Europe.

FIGURE 4.2
EUROPEAN NETWORK CHARACTERISTICS

NETWORK	EXTENT	NO. OF NODES/ HOSTS	NO. OF USERS	LINE SPEED (Kb/s.)	PROTOCOLS	FUNDING	TOPOLOGY
Ellnet	Continent	1221 N	N/A	2.4-9.6	UUCP, TCP/IP	user fees	Modified Star
EARN	Continent	600 H	>30,000	variable	NJE	membership	Tree
HEPnet	Continent	~12 N	N/A	64	DECNET, others	N/A	Star
EAN	Continent	~140 H	N/A	9.6	Ean X.400	N/A	N/A
NORDUnet	Scandinavia	4 N	N/A	64	Variety	Nordic governments	Star
JANET	U.K.	20 N	N/A	9.6-512	Coloured book	government	Tree
Starnet	U.K.	19 N	~1,000	9.6-512	DECNET, coloured book	government	Tree
DFN	W. Germany	120 H	>5,000	9.6	X.400, proprietary	government	N/A
BELWU	W. Germany (Regional)	~5 N	N/A	140,000	TCP/IP	government	Star
Aristote	France	~10 N	N/A	4.8-9.6	X.400, X.25 proprietary	membership	N/A
Reunir	France	60 N	N/A	4.8-2,000	X.25, SNA	membership	Tree
SURFnet	Netherlands	375 H	N/A	9.6	DECNET, NJE	N/A	Tree
HEANET	Ireland	30 H	N/A	N/A	Coloured book	government	N/A
UEuroKom	Ireland	1 N	N/A	N/A	PortaCOM	European community	N/A
Ean	Spain	22 H	N/A	9.6	Ean X.400	Government	Star
DEnet	Denmark	N/A	N/A	64-128	DECNET, TCP/IP	Government	Modified Star
DKnet	Denmark	~150 H	N/A	1.2-9.6	UUCP	User fees	Star
FUnet	Finland	>1000 H	N/A	14.4-64	TCP/IP, others	User fees and government	Star
UNINETT	Norway	11 N	600	N/A	proprietary, others	Government	N/A
SUnet	Sweden	>1700 H	N/A	64	DECNET	Government	Tree
SWITCH	Switzerland	10N	N/A	48-2,000	X.400	Membership	Tree
ACONET	Austria	~10 N	N/A	N/A	DVS-NW	Government	N/A
UNA	Austria	13 N	N/A	N/A	DECNET	Vendor, government, user fees	N/A

Source: [1] The Matrix; Computer Networks and Conferencing Systems Worldwide; John S. Quarterman; 1990.
[2] HICKLING Research and Personal Interviews.

4.5 PAN-EUROPEAN NETWORKS

4.5.1 EUnet

Background

- A Pan–European co–operative R&D network for UNIX users. Its purpose is to provide users with modern communications facilities, particularily electronic news and mail, and interconnections to other networks. It is closely tied to the European UNIX Users Group (EUUG);
- Widely used by mathematicians, computer science researchers, and other R&D people generally in small or medium size companies and university departments having their own computing facilities (most of which run UNIX);
- An important means of technology transfer between industry and academia;
- Started at the April 1982 EUUG meeting by a group of interested users in the Netherlands, Sweden, Denmark, and the U.K., but quickly spread throughout Western Europe.
- 1221 registered sites in 19 Western European countries as of May, 1988. There is one backbone host per country which is the sending/receiving station for international traffic; there is also a central host through which nearly all international traffic is routed; this host is located at the Centrum voor Wiskunde en Informatica (CWI) in Amsterdam.
- In May, 1988, traffic through the central machine amounted to 2.5 Gb/month, and was increasing at the rate of 100 Mb/month.

Technical Design

- A mixture of X.25 links and international phone links, some using leased lines, for the international backbone.
- Recent technology advances have made phone lines more attractive than X.25, (mainly faster modems and higher capacity digital lines), but only where the PTTs have approved the equipment.
- Most X.25 links are nominally 9.6 Kb/s but effectively half that. Within the national backbones, there are still many dial–up telephone links using UUCP which are 2400 bps; however new modem technology allows 10 Kb/s or more. Where high–speed links are available, TCP/IP is sometimes used over X.25 with RFC877 encapsulation.
- Routing is managed by the backbone hosts, each of which knows the organization within its own country and which hosts are in which foreign country. This information is automatically exchanged on a daily basis between backbone hosts. Thus the user need only know the name of the host he wants to connect with; the network takes care of the routing.
- There are also a number of interconnections with other networks, both national and international.

EUnet - Topology in France

Source: Dr. Fred Casadei; Basel, Switzerland; January 1990.

Implementation Design

- EUnet is an informal operating agreement among the administrators of the backbone hosts. The entire network is essentially run by volunteer labour; this has provided satisfactory performance and avoids bureaucratic overhead costs and minimizes liabilities.

- The backbone host administrators meet to discuss operating strategy and planning; EUnet is represented by the EUUG (European UNIX Users Group) whenever a united front is required, such as in working with the CEC and RARE.

- Each user site pays a flat fee of about $4/month for access. Mail is charged to the originating host on a per message and per link basis. This is based on the PTT charges for the kind of link used. X.25 links are generally lower cost than equivalent North American phone links, although there is wide variation between PTTs. The host computer equipment time is donated by the host.

- Future plans include eventual migration to ISO–OSI protocols and co–operation with the CEC and RARE initiatives for network integration and standardization. In Spring 1989, EUnet began to set up a continental network using TCP/IP, using T1 links from Amsterdam to France and Stockholm, thus providing access to most of the transatlantic links to NSFNET in the U.S.

4.5.2 EARN

Background

- European Academic Research Network is essentially the European arm of BITNET. It links more than 600 hosts and 30,000 users in every Western European country, plus many North African and Middle Eastern nations.

- Formed in 1983, with much of the funding being supplied by IBM between 1983 and 1987.

- Most sites are in large computer centres at universities, government research centres, and large companies; many hosts are IBM VM or Digital VAX VMS mainframe machines.

- Main thrust is to build a uniform system extending throughout the EC which is independent of the national PTTs. For this, EARN has to receive special permission from each PTT to operate there; however this gives the advantage of being able to charge annual flat fees to its user institutions that are published in advance (allowing budgeting) rather than per use charges that depend on the particular PTT being used.

- Currently uses NJE protocols, like BITNET. However EARN will migrate directly to the ISO/OSI protocols, using Digital, IBM, and Northern Telecom supplied hardware, software, and technical support.

- International EARN backbone is all IBM machines for which ISO/OSI implementations exist, however hosts within countries will only require protocol converting gateways for continued connection.

EARN TOPOLOGY
As of February 6, 1989

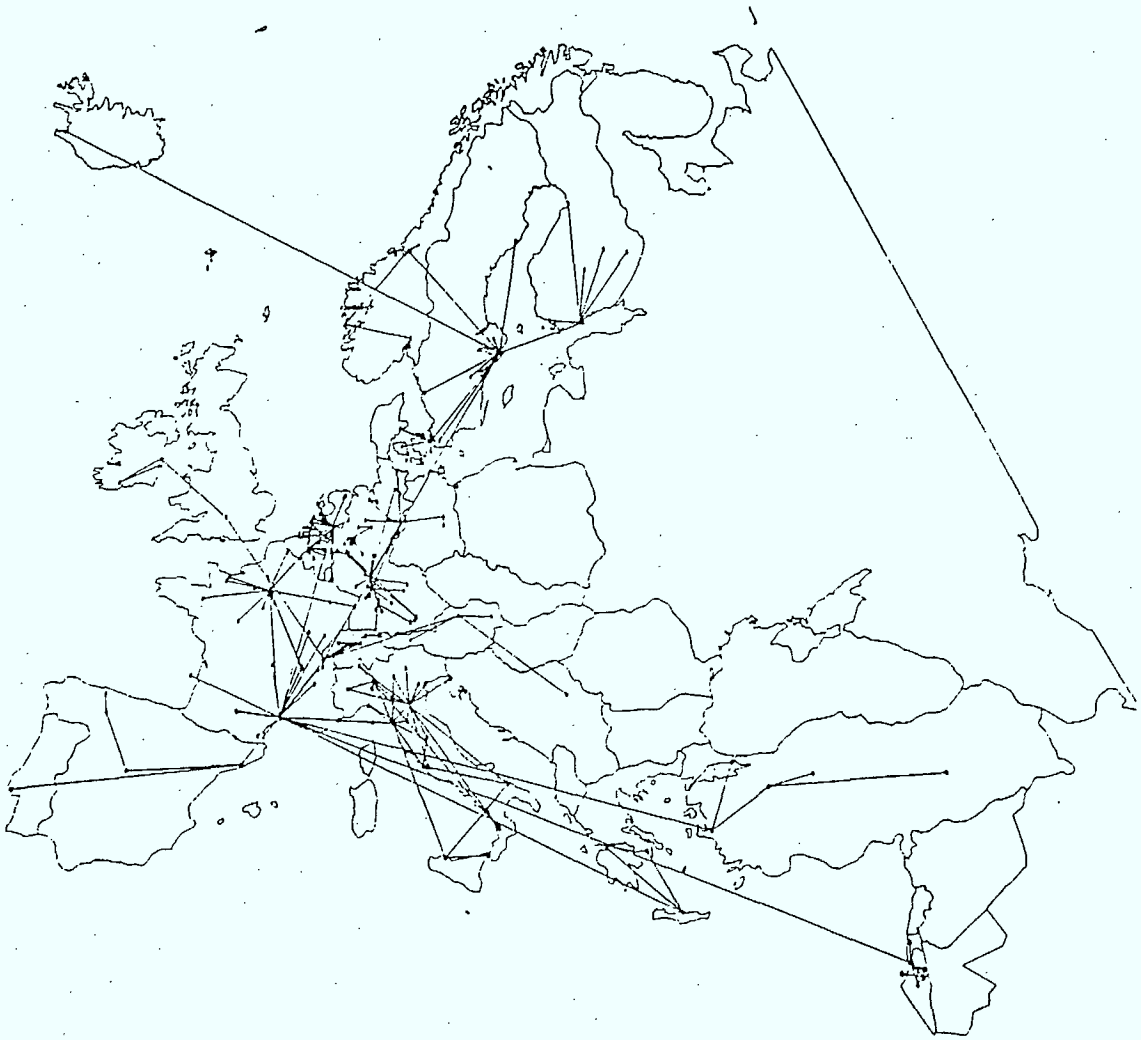

Source: Users' Directory of Computer Networks; Compiled by Tracy Lynn LaQuey, University of Texas; July 1989.

Technical Design

- Logical architecture similar to that of EUnet in that each country has a central or backbone host through which the international traffic is routed. Within each country, a number of other hosts may be connected to the backbone host.
- Uses X.25 links supplied by the PTTs in each country for its connections between national centres. Links within countries vary as to type and speed.
- Standards, protocols, and services supported are essentially the same as BITNET, ie NJE supplied with IBM mainframes.
- There are gateways to several national academic networks such as JANET in the U.K., and a transatlantic link between Montpelier, France and New York.
- In converting to ISO/OSI standards, EARN will initially supply a private network supplying the 1984 version of X.25 required for OSI; the PTTs generally support only 1980 X.25. The X.400 mail protocol will be the first OSI application supported.

Implementation Design

- EARN is an association registered in France. The board of directors comprises one representative from each member country, who may be elected by members from that country or appointed by the government. Planning and policy are decided by the executive committee elected by the board.
- Each country pays the cost of one international line to another country in Europe. The cost of the transatlantic link and administrative costs are collected from the national EARN members based on GNP. Internal national funding is the concern of the EARN members in each country.
- The total annual costs are about $1.8 million.
- National membership requests are voted upon by the board, but connections within a member country are the responsibility of that country's central administrative body.
- EARN offices are located in Paris and Dublin.
- IBM provided $15 Million of support for EARN between 1983 and 1987, including much of the equipment, support personnel, and funding transatlantic links.

4.5.3 HEPnet

Background

- The European section of the worldwide High Energy Physics network. In Europe it provides connections between CERN, the European Centre for Nuclear Research, and several national networks and physics laboratories. Transatlantic links are to MIT and a satellite link to FERMILAB in Batavia, Illinois.

Technical Design

- Most of the links are dedicated lines operating at 64 Kb/s.

- The DECNET X.25 protocol is most often implemented, although a wide variety of other protocols are used. To accommodate these, a number of application converters have been developed. One of these, GIFT (General Internetwork File Transfer Protocol) allows conversion from any of five commonly–used file transfer protocols to any other, without intermediate storage. MINT, short for Mail Internet Transfer Protocol, converts between any of seven common mail systems.

Implementation Design

- Run by an organizing committee chaired by the CERN representative
- Future plans include the introduction of 2 Mb/s dedicated links, and the gradual conversion to ISO/OSI protocols as they become available economically.

4.5.4 Ean Europe

- Essentially a collection of national networks which are using the Ean implementation of the X.400 protocols.
- Origins in CDNnet in Canada in 1983. As of October 1987, there were approximately 140 hosts supporting Ean.
- Most of the sites within countries use 9600 bps leased lines, while intercountry links are generally by X.25 PDN. There are gateways through CERN to EARN and EUnet.

4.5.5 NORDUnet

Background

- A joint undertaking of the Nordic countries (Denmark, Finland, Iceland, Norway, and Sweden); does not extend to the rest of Europe, although there are interconnections with other networks.
- Arose out of the annual Scandinavian university network conferences and was formally started in 1985.
- Goal is to provide harmonized network services to Nordic research and development users and to establish good inter–Nordic relations in networking.

Technical Design

- One backbone node in each country
- International links are 64 Kb/s leased lines
- Backbone is provided by SCANTEL, a consortium of the Nordic country PTTs which provides international communications and satellite links.
- Backbone is able to carry a wide variety of protocols, as major design factors were an ability to connect with U.S. networks and to have ISO/OSI compatibility.
- Only recently made operational; link to Iceland is still under discussion. Plans to move to 2 Mb/s links are also being discussed.

Implementation Design

- Administered by NORDUNET, a program of the Nordic Council of Ministers, and funding comes from NORDUNET and the four major countries.
- NORDUNET steering committee also implements development projects concerning practical infrastructure using existing protocols, and ISO/OSI pilot networks and experiments.
- These projects are generally assigned to one of the national member networks.
- Total approved funding for NORDUnet is 10 Million Norwegian Crowns (about $1.85 million Canadian) over five years to 1990. Proposal for follow–on funding is under discussion.

4.6 MAJOR NATIONAL NETWORKS

4.6.1 United Kingdom

There are two major research and development networks in the U.K. other than the British parts of Pan–European nets such as EUnet and Ean. These are JANET and Starlink.

4.6.1.1 JANET

Background

- Joint Academic Network (JANET) was established in 1984 to serve the needs of the scientific and research community for information exchange.
- Represents the integration of several independent networks developed by various organizations within the U.K. to serve their users.
- Supported by the Computer Board for Universities and Research Councils (CB), funded by the U.K. Department of Education and Science.
- Main users of JANET are the Universities, the Government Research Councils, and the Polytechnics. There were 915 registered users as of September 1986; however the true number of connected hosts (including those on LANs and PADs) was estimated as around 1500. Usage at the end of 1985 was around 24 Mb/day.

Technical Design

- Services include mail, file transfer, remote login, and remote job entry. Access to supercomputers is available through LANs at their locations.
- Currently uses the so–called Coloured Book Protocols, a set of locally developed software packages, but future plans call for migration to the corresponding ISO standards.
- Science and Engineering Research Council (SERC) has set up one–year postgraduate courses in networks and distributed systems, to alleviate a shortage of capable people. In addition, annual workshops on Networking are offered.

JANET - Topology in England

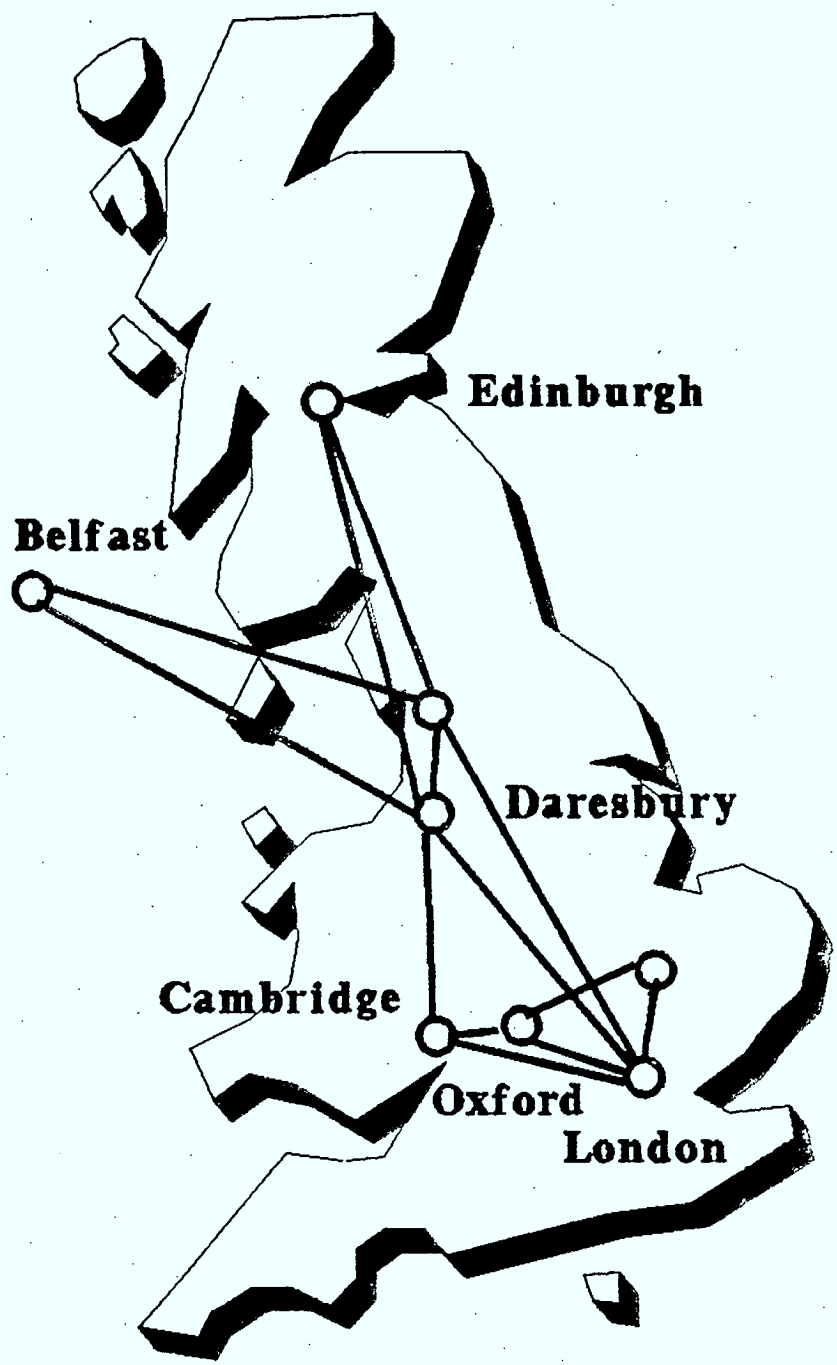

Source: Dr. Fred Casadei; Basel, Switzerland; January 1990.

- Network consists of 120 sites with 8 X.25 switches located at main nodes spanning Great Britain and Northern Ireland. Figure 3.2 shows the location of the main switches and the high speed backbone links between them.

- Main switches are GEC 4000 series running GEC–supplied software. Seven of the eight switches are housed at Network Operations Centres while the Network Control Centre houses a Network Management Unit comprising a switch and associated control hardware and software.

- Switching nodes have capacities of forwarding up to 1000 packets per second; the main trunk line speeds are 512 Kb/s, while other lines are 64 Kb/s digital or 48 Kb/s analog. Most subscriber lines run at 9.6 Kb/s.

- Some of the main trunk lines are 2 Mb/s lines time division multiplexed and shared with other applications. These lines are leased from British Telecom, the U.K. PTT.

- Coloured Book Protocols are used for the higher layers. Local networks connected to JANET use either X.25, Cambridge Ring (CR82), Ethernet, or IEEE 802.3, with the latter two gaining in popularity at the expense of CR82.

- Gateways exist for EUnet, EARN, Ean, the Internet, and JUNET, the Japanese Universities Network, although access to all gateways is controlled due to the cost.

Implementation Design

- Network Advisory Committee (NAC) is responsible for JANET planning and policy. It consists of two CB officials, two Research Council representatives, a representative from the Polytechnics, the JANET Director of Networking, and the chairman of the Users Group.

- Technical and administrative support is the responsibility of the Network Executive (NE) located at the Network Control Centre. It provides budgeting, planning, reporting, and coordinates user support activities, in addition to being the final level of arbitration for disputes related to faults and problems. It is responsible for decisions concerning overloaded lines, poor reponse times, and poor availability.

- The seven NOCs have an average operational staff of two each during working hours, and are responsible for day–to–day operation, reliability monitoring, and initial response to problems.

- Does not achieve its goal for MBTF (Mean Time Between Failures) greater than 2000 hours, but availability is better than 99.8%; this indicates the good organization which is one of the strengths of this network.

- No direct charges are made for usage to the Universities and Research Councils. Polytechnics are charged to join the network, but individual members of polytechnics who hold Research Council grants may be connected.

- Annual CB budget for both JANET and University LANs was approximately $7 Million in 1986; a further amount of $10 Million was approved in 1986 for capital expenses related to upgrading the switches and trunk lines to make use of new British Telecom links.

4.6.1.2 Starlink

Background

- A network for astronomers whose purpose is to provide interactive computing facilities for spectral and image work.
- Need arose in the mid 70s when new astronomical data sources such as satellites and radio telescopes began producing data faster than it could be processed.
- Distributed network using VAX 11/780s at each of six sites, connected by leased lines, was in operation by April 1980.
- There are now 19 sites with a total of 50 hosts and 950 users. Users consist of:

Research Astronomers	87%
Programmers	10%
Administrators	3%

Technical Design

- Originally used its own leased lines, but now uses JANET lines and so can be considered a special part of JANET.
- Both DECNET and Coloured Book Protocols are used.
- Programmers have produced a package of astronomy application software, a copy of which resides in each host. The network also provides a central database of astronomical data, and mail and news facilities.
- Original VAX 780s and 750s are being replaced by MicroVAXs.
- Not expecting much future growth in terms of users since nearly all astronomy sites are now connected.
- Starlink has been a pioneer in the use of graphics standard GKS.

Implementation Design

- Administered by a group at Rutherford Appleton Laboratory.
- Funding provided by the Science and Engineering Research Council, part of the U.K. Department of Education and Science.

4.6.2 Federal Republic of Germany

Other than the German portions of international nets such as EUnet (called Dnet in Germany) and EARN (AGFNET in Germany), there are two research and development networks operating at this time, the DFN and BELWU.

4.6.2.1 DFN

Background

- The national research network connecting every university, college, and research laboratory in West Germany.

- Purpose is to provide a communications and resource sharing system using the ISO/OSI protocols. To do this, appropriate software implementations of the standard are being developed. This project is directed and funded by the German Ministry for Research and Technology.

- Started operation in 1986 with seven hosts connected by public X.25 lines.

- At the end of 1988 there were 120 active members used by over 5000 researchers. Usage has increased at over 100% annually over the past few years.

- Membership is open to universities, public or private research institutes, hardware and software manufacturers, and even other industrial companies.

Technical Design

- Over some time the X.25 PTT supplied lines have proven inadequate for some applications, and higher speed T1 leased lines are being installed and put into operation. These will be operated in parallel with the existing links.

- The Deutsche Bundespost, the German PTT, is currently installing a fiber optic cable system, which will greatly enhance transmission bandwidths.

- The DFN project is heavily involved in the development of ISO/OSI protocols and has implemented ISO/OSI protocols such as TP0, FTAM, and X.400. It has also developed conversions at gateways allowing many other protocols such as TPC/IP over IEEE 802.3 as used in many LANs.

Implementation Design

- Network administered by DFN–Verein located in Berlin.

- MRT funding is expected to end in 1990.

4.6.2.2 BELWU

- Network is intended to connect campus LANs in the German state of Baden–Wurttemberg using 140 Mb/s optical fibre links.

- First link was demonstrated in February 1988.

- One major objective is to provide remote access to the CRAY 2 computer installed at the University of Stuttgart.

- Higher level protocols are currently TCP/IP but conversion to ISO/OSI is expected.

4.6.3 France

Aside from the French branches of EUnet (Fnet), EARN, and HEPnet (PHYnet), there are two national research and development networks, called ARISTOTE and REUNIR.

4.6.3.1 ARISTOTE

Background

- Major purpose is to act as a testbed and club for French research institutions to develop networking technology.
- Organized as a non–profit association, with three classes of members. Full membership is granted to French public R&D institutions, of which there were nine as of May 1988. Associate members are public agencies and groups interested in the goals of OSI implementation, such as CERN. Correspondent members are mostly from industry.
- Services provided comprise remote login, mail using the X.400 protocol, and a directory service called THORN which is under development.
- Transmissions are made using X.25 at 4.8 to 9.6 Kb/s over TRANSPAC, the French PDN.
- Future connections may use the TELECOM–1 satellite or the 64 Kb/s digital packet switched service provided by the PTT.

4.6.3.2 REUNIR

- A network connecting many French universities and research institutions.
- Purpose is operational support and communication of other research, rather than experimental research in networking itself.
- Services supported include mail, file transfer, remote login, and some calculation and archiving services.
- Contains a wide mixture of parts including private X.25 networks using PTT lines, SNA networks connecting mainframes, LANs using TCP/IP, and HYPERchannel links allowing access to supercomputers attached to some sites.
- Connections between these parts use either X.25 or leased lines, with speeds ranging from 4.8 Kb/s to 2 Mb/s.
- Administration is the responsibility of two directors assisted by the host administrators, a user group, and a technical team.
- REUNIR represents France in the CEC COSINE program.
- Some funding from the French Education Ministry and the National Science Research Institute (CNRS) is being used to develop and offer an X.400 mail service.

4.6.4 ITALY

Italy at present does not have any unique national R&D networks. A network called INFNET, originally connecting the sites of the Italian National Institute for Nuclear Physics, has spread throughout the country to 118 hosts by July 1987. INFNET can be considered as the Italian portion of HEPnet and connects to SPAN in the U.S.

4.6.5 THE NETHERLANDS

The two networks in the Netherlands are national branches of the European networks, but the national network SURFnet, besides being the national branch of EARN, has some features unique to the Netherlands.

4.6.5.1 SURFnet

- Network for research and higher education in the Netherlands
- Name is an acronym translating to Cooperating University Computing Facilities.
- As of 1988 there were 85 connected organizations, with 375 hosts connected by 9.6 Kb/s links operated by the Dutch PTT.
- One part of SURFnet uses DECNET protocols, while the other is the national branch of EARN donated by IBM to SURFnet and running NJE.
- Run by a registered company called SURFnet B.V. which was incorporated January 1, 1989. The SURF Foundation representing users holds 51 % while the PTT holds 49 % of the shares.
- By the end of 1989, SURFnet plans to have a backbone of twenty–five 64 Kb/s links in place using COSINE X.25 protocols.
- The first step in migration towards ISO protocols is conversion to X.400 mail protocols, and there are already several hosts using X.400 in a pilot project for RARE. A second project is introducing DNS domain names, and after this conversion, users will not need to know the names of gateways to other networks.

4.6.6 Ireland

4.6.6.1 HEANET

- Higher Education Authority Network is the national R&D network in Ireland whose purpose is to provide access to specialized hardware and software and networks in other parts of the world, and to promote information exchange and cooperation.
- Network currently connects 30 hosts at seven colleges throughout the Republic of Ireland.
- Coloured Book Protocols are primarily used, although others such as NJE are used in connecting nodes for mail service.
- Network uses both leased lines and Telecom Eireann's X.25 packet switching network.
- Managed and funded by the Irish government Higher Education Authority.

4.6.6.2 EuroKom

- The European Community supports a program called ESPRIT, the European Strategic Programme for Research in Information Technology. The ESPRIT program supports cooperative research into software technology, computer–integrated manufacturing,

microelectronics, and related areas.

- More than half the ESPRIT projects use EuroKom, an electronic mail and conferencing system based on a central machine in Dublin, Ireland.

- Accessed through existing gateways on other networks such as EARN.

4.6.7 Spain

Spain has national branches of the three main networks EUnet (Enet in Spain), EARN, and HEPnet (FAENET). In addition there is a RARE–supported experimental R&D network. National responsibility for networking is given to the Interconexion de recursos Informaticos (IRIS) which also provides funding for EARN, FAENET and the experimental Ean network. IRIS is part of a non–profit organization sponsored by the Spanish government, and represents Spain in EC and other European projects such as RARE and COSINE.

4.6.7.1 Ean

- RARE–supported experimental network is a national research network using X.25 over the national PDN IBERPAC.

- As of November 1988, there were 22 hosts at 14 sites.

- Project plan runs four years from early 1988 through 1991.

- Organized as a star, with a Digital VAX 11/750 having 500 Mbyte disk storage as the backbone machine. More recent X.400 implementations use a MicroVAX 3600 as the backbone.

- Most hosts use 9.6 Kb/s X.25 connections, with a few being 1200 bps X.28 dialup links.
- Links with heavy traffic are planned to be upgraded to 64 Kb/s, with a similar international link to CERN.

- Development continues on ISO/OSI applications such as X.400 and experiments with FTAM; eventually directory service and management protocols will be implemented to complete the OSI migration.

4.6.8 Denmark

There are two main networks in Denmark; the Danish branch of NORDUnet called DENet, and the national branch of EUnet called DKnet.

4.6.8.1 DEnet

- Network connects local Ethernets in University departments across the country using MAC level bridges and 64 or 128 Kb/s leased lines.

- Originally provided access to three regional computing centres sponsored by the Danish Ministry of Education; these centres have since been merged into one, containing a variety of hardware by different manufacturers.

- Both TCP/IP and DECNET are used as protocols.

4.6.8.2 DKnet

- Although this is the national branch of EUnet and shares much of the same characteristics, its administration and funding is handled directly by the Danish UNIX User's Group.
- As a result, useful information on operating costs is available.
- Physical layout is star–shaped with a Sun–2/120 as the main backbone machine, communicating over two X.25 lines and four low–speed modems.
- As of May 1988, the total machines connected were about 150 at 60 sites, with measured monthly traffic through the backbone of 400 Mb.
- In 1988, the budget was about $90,000.
- Each site subscription was charged about $180, with an initial hookup charge about the same amount. Each service is charged separately (ie. mail and news would cost $360 per site).
- Sites pay directly for communication line costs, and a surcharge of $3/Mb of news is imposed.
- The Danish UNIX User's Group has a half–time employee responsible for running the backbone.

4.6.9 Finland

While there are connections to all of the international networks in Finland, the major national network is FUNET.

4.6.9.1 FUNET

Background

- Finnish University Network (FUNET) is a star layout with a backbone machine at the Helsinki University of Technology (HUT).
- Established in 1984 to provide remote access to computer centre machines, and grew to the point that the Finnish Ministry of Education began the FUNIT Project to coordinate the various protocols in use.
- Now thousands of hosts ranging from PCs and Macintoshes using the TCP/IP protocols through the range of machines running UNIX to IBM mainframes using NJE, and a CRAY supercomputer.
- Connection to the network not limited to Universities; there is a requirement that use be restricted to research and education purposes.
- An increasing number of private companies and Government agencies are connected, and they pay for the cost of their links.

Technical Design

- Essentially a wide–area Ethernet connecting local area Ethernets with MAC bridges and network layer routers.
- Uses leased lines at speeds varying from 14.4 Kb/s to 64 Kb/s, finding that in general leased lines are fixed cost, usually faster, and often cheaper than PTT X.25 connections.
- Fixed cost is a special advantage because it eliminates any need to bill connection charges to users, because it simplifies budgeting, and encourages greater use of the system.
- Services provided include mail, conferencing, file transfer, remote job entry, remote login, and interactive graphics.

Implementation Design

- Managed by the FUNET project, whose staff is funded by the Ministry of Education.
- FUNET Steering Committee consists of seven members representing the universities, the Ministry, and the PTT.
- Collects fees from the universities to cover the cost of the national links; the international links and capital equipment are funded by the Ministry.
- Project also involved in ISO/OSI program testing at several test sites.

4.6.10 Norway

As well as national components of the international networks, there is a major national R&D network called UNINETT.

4.6.10.1 UNINETT

- Network traces its origins to 1978; it is one of the earliest users of the Ean X.400 implementation in Europe.
- Boasts 600 users on 11 nodes.
- Funding is provided by the Norwegian Science Research Council and the participating institutions.
- Developed some of its own protocols before the ISO/OSI protocols were completely specified; these are supported along with DECNET and Coloured Book.
- In 1987, the Ministry of Cultural and Scientific Affairs set up a new UNINETT specifically to use ISO/OSI protocols.
- Goals are to connect all institutions of higher education, with a strong emphasis on reliability (95% mail delivery within five minutes), speed (64 Kb/s links), and online support services during working hours.
- All the existing networks are to be merged into the new UNINETT.

4.6.11 Sweden

SUNET is the national research–oriented network in Sweden.

4.6.11.1 SUNET

- Swedish University Network interconnects local and regional university networks with a star layout centred on the Royal Institute of Technology in Stockholm.
- System started operation in 1980 with a basic service of remote login by X.25 over the PTT Datapak.
- In 1985 a national DECNET was installed with six 64 Kb/s links connecting local Ethernets, which parallels the existing X.25 links.
- National branches of all the major networks are carried over these physical links.
- Altogether more than 1700 hosts connected.
- Network used by researchers and teachers of all disciplines.
- Operation and management is the task of a few departments at different universities, headed by the Umea University Computing Centre.
- Swedish Science Foundation funded the cost of expansion, and as of 1988 the network is funded by the Swedish government branch in charge of higher education.

4.6.12 Switzerland

Switzerland has a national research network called SWITCH. It is also the location of CERN, one of the major centres for networking in Europe. CERN has obtained from the Swiss PTT the right to switch third–party transmissions, making it a major gateway connecting every continental network. It is also a testbed for LAN technology and network protocol testing due to its large number and variety of computers.

4.6.12.1 SWITCH

- Founded in 1987 by the national government and those of the cantons within Switzerland having Universities.
- Objective is to connect the eight universities and two Federal Institutes of Technology.
- Initial funding of $12 Million was provided, and operations began in mid–1988.
- Committed to implementing OSI protocols when available, but in the meantime uses whatever is convenient.
- Two services that are available now are as follows:
- SWITCHmail, a messaging service based on X.400 with access, through CERN, to all other academic and R&D–oriented networks.
- SWITCHlan, an Ethernet–based network with bridges and routers (CISCO products) using 2 Mb/s (T1) and 128 Kb/s links.

SWITCH - Topology Swizerland

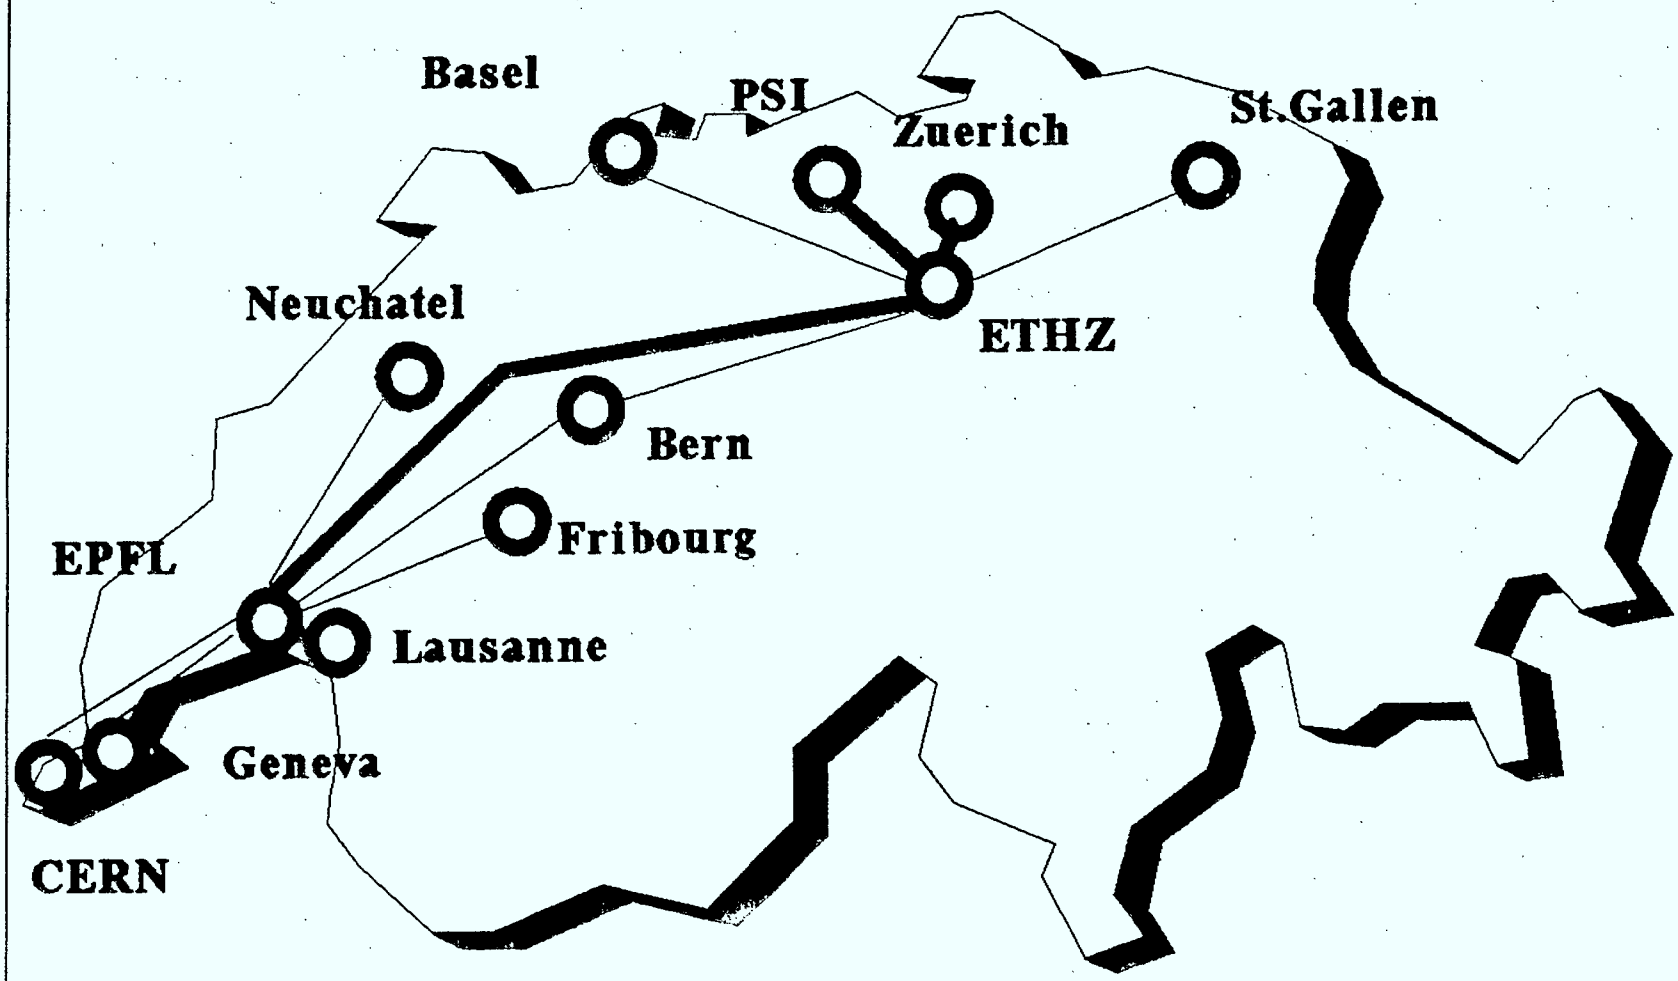

Source: Dr. Fred Casadei; Basel, Switzerland; January 1990.

- Mail service generally uses the Swiss TELEPAK 48 Kb/s links, considered adequate for mail service.
- The Ethernet part of the system is used mostly for access to two supercomputers.
- Several other services are planned for operation soon; file transfer, remote job entry, interactive terminal dialog, and a user and network information service.
- SWITCH is quite liberal about access to the network; research–oriented industrial companies are allowed to connect, although they pay a higher access fee than academic institutions.

4.6.13 Austria

Austria boasts two separate national research networks, ACONET and UNA, besides national branches of the continental networks.

4.6.13.1 ACONET

- Started in 1981 with the purpose of connecting existing LANs used for research at universities.
- Now reaches most institutes of research and higher education in Austria, and has interconnections with EUnet, EARN, and DATEX–P (the Austrian PDN).
- Services provided include file transfer, remote login, and remote job entry.
- Currently uses DVS–NW protocols developed in Germany, but will migrate to the ISO/OSI protocols as they become available.
- Funding provided by the Austrian Ministry of Science and Research
- Administered by an organization called ACONET located in Vienna.

4.6.13.2 UNA

- A vendor–specific network which connects Digital computers at all Austrian universities using DECNET over the X.25 DATEX–P links.
- All equipment donated by Digital; further grants were obtained from the Ministry of Science and Research
- 13 sites connected as of March 1987; connections to EARN and DFN were in place, and a connection to EUnet was being added.
- In the future, conversion to ISO/OSI will take place when appropriate DEC software is available (DECNET Phase V). These changes will move the network away from being vendor–specific. Eventually, UNA may merge with ACONET.

4.6.14 U.S.S.R.

In the USSR and the Eastern European countries, there are two R&D oriented networks. It appears that the major purpose of these networks is to provide access to Soviet and foreign databases. IASnet is a star network with a central host at the Institute for Automated Systems in Moscow and X.25 connections to leading institutes in the Socialist countries. As

of August 1988, IASnet was still in implementation.

Academnet connects institutes in the republics of the Soviet Union to a host at the IAS in Moscow. As of mid–1988 there were connections to institutes in five cities.

4.7 NEW FUTURE NETWORKS IN EUROPE

4.7.1 VENUS

A report submitted in Spring 1989 to the European Economic Commission by RARE, the European body for coordinating national academic networks, has proposed a European High–Speed Networking Initiative, to consist of a central backbone called VENUS, regional networks, and supporting networks for specific disciplines. Plans call for 20 to 40 nodes across Europe, serving from 200 to 500 sites. A phased improvement in network speeds is considered, with 2 Mb/s immediately, 100 Mb/s in a few years time, and gigabit speeds by the end of the century. Initial costs were stated as $6 Million annually in the near future, rising to $24 – $30 Million later.

4.7.2 Germany: FORERUNNER

In West Germany, a national fibre optic network is being installed by the Deutsche Bundespost, the national PTT. This "Forerunner" network, with capacities up to 140 Mb/s, is intended for video teleconferencing and other high–bandwidth requirements, and will be available for scientific networking including graphics transmission and access to supercomputers.

At the Regional Computing Center of Stuttgart University, applications are being developed and installed to make use of this capacity. Two projects are in progress to connect research centres located 1000 km and 150 km from Stuttgart with the Cray 2 supercomputer there at speeds of 100 Mb/s. Several years will be required to reach this speed, and costs of the high–speed link must still be negotiated with the PTT, which is in the midst of deregulation.

4.7.3 U.K.: Cambridge Fast Ring

The campus of Cambridge University is spread over relatively long distances within the city of Cambridge, U.K. A new network using about 200 kilometres of optical cable has recently been completed to connect existing local networks. Both the local networks and the connecting network use a design called the Cambridge Fast Ring, which permits relatively cheap hardware and low complexity. The new network, built with support from Olivetti Corp., has potential speeds in the Gigabit range and has been tested at 600 Mb/s.

4.7.4 Italy: GARR

An organization has been set up by the Italian Ministry of Scientific Research and Technology to develop a high–speed national backbone connecting all research institutes and universities in Italy. This organization is called "Gruppo Armonizziazione delle Reti per la Ricercha (GARR)". The system will use 2 Mb/s leased lines and offer four protocols: X.25, TCP/IP, DECNET, and SNA. Migration to full ISO/OSI protocols will be undertaken as soon as feasible. The backbone will also have a high–speed link to CERN, the European high–energy physics centre located near Geneva.

The network will be financed by the Ministry of Scientific Research and Technology at a estimated rate of 9.7 Million ECU/year (about $14 Million), with half to pay for leased lines and the remainder for control units and software.

GARR - Network - Topology Italy

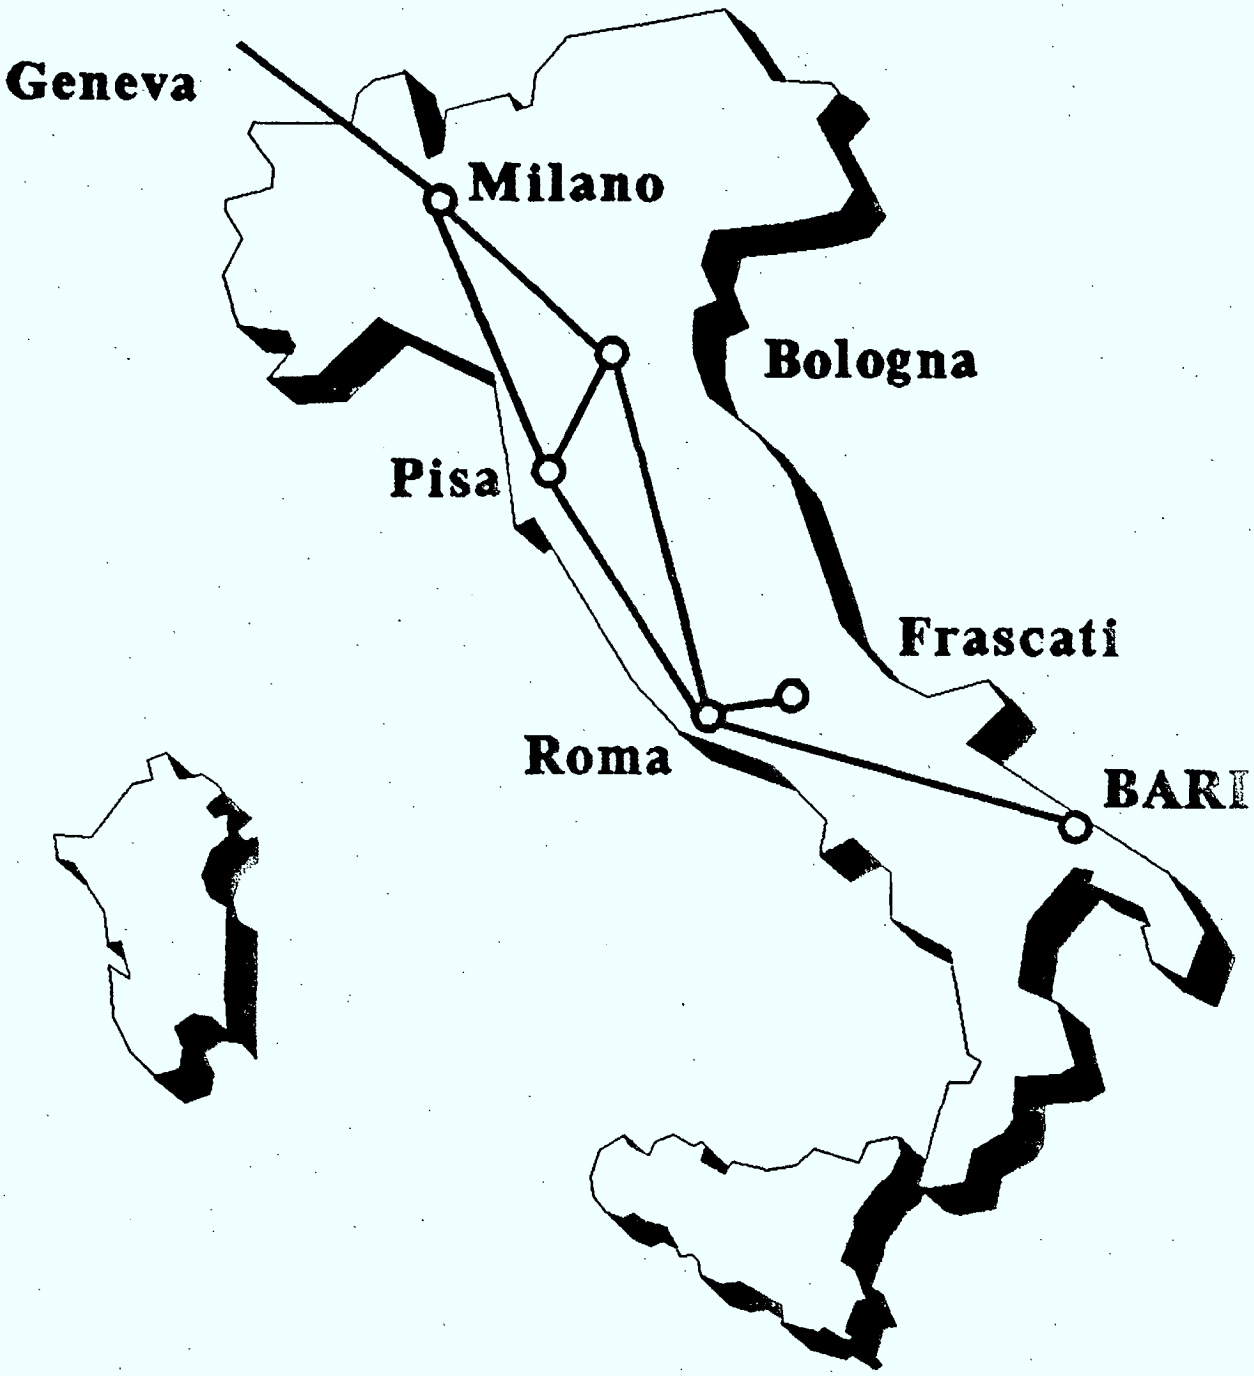

Source: Dr. Fred Casadei; Basel, Switzerland; January 1990.

5. JAPANESE R&D NETWORKS

5.1 SUMMARY

Prior to 1985, Japanese telecommunications were controlled by the government monopoly Nippon Telegraph and Telephone. The deregulation of the market in 1985 has promoted a boom in the number and variety of telecom services. One major component of these services is teleconferencing, which has seen significant growth fuelled by media publicity.
However, the number and use of R&D networks in Japan is surprisingly small. One reason is that the Japanese language and culture is not conducive to standard applications of networks, such as mail and news. A second reason is that there is, in general, fierce competition among universities which reduces the amount of cooperation and communication among them. Finally, the intense concentration of Japan's population and industry means that often wide area network services are not necessary because they can be easily reached in person.

The Japanese language can be written in five methods; only one of which is compatible with Western–style keyboards and network software. This accounts for the popularity of facsimile communication in Japan, which probably replaces much of the traffic for which a network would be used in other countries.

There are several standards for encoding Japanese characters. The basic standard is JIS X 0208, which provides for character representations in two bytes, compared to the standard ASCII which uses only one byte. Special conversions to much foreign software have been completed which allow their use in Japanese systems.

The Japanese government is encouraging and in some cases financing the installation of Local Area Networks at universities across Japan. In addition it is involved in planning and funding two new high–speed international links which will connect with U.S. and possibly European networks. These may form part of its initiative to increase basic R&D in Japan through cooperation with other companies.

There are three predominant, operational, research networks in Japan:

N–1: is the oldest domestic network and connects mainframes.

Science Information Network (NACSIS): is a packet switching network operated by common carriers and is used for super computer access and library access. This network appears to be the unifying force to realize a single network that serves the R&D interests of Japan. This network also interconnects different networks such as JUNET, HEPNET and BITNET as well as university and institutional LANS by using gateways and protocol converters.

Junet: This network is the most widespread noncommercial network in Japan, which also provides a testing environment for research into computer networking and distributed processing, Japanese character handling, resource name managing and speed improvements. In addition to these production networks, Sigma is an experimental network which is used as a testbed for product development.

The Japanese networks are quite different from European and North American networks as they are used only peripherally for mail. The major purpose of the NACSIS system is

fast, local access to one of the largest databases of scientific and technical information in the world; electronic mail and other networking services are provided but assume less importance in the system design. Japan's R&D networks are more oriented towards public access to information databases and news than networks in other countries. The JUNET system is intended as a testbed for networking development rather than an R&D production network; its traffic is estimated at 85% news and only 15% mail. This may be a result of the relatively high number of FAX machines in Japan and the difficulty arising from typing Japanese characters on a standard keyboard.

In order to offer an online information retrieval service and catalog information service to researchers at universities, a high performance computer system is installed to meet the needs of large–scale database services. An optical disc storage system of 83 gigabytes of memory has been implemented for this purpose.

By working closely with NSFNET in the U.S. the Japanese Government will fund two–high speed international links to the U.S.. JUNET is already connected to CSNET/NSFNET in the U.S..

5.2 JAPANESE R&D ENVIRONMENT

Over the past decade, Japan's Research and Development effort has expanded dramatically. In the financial year 1987–1988, the country spent about $85 billion dollars on R&D, a record 2.8 per cent of its GNP. Japan now boasts about 490,000 researchers, more than one quarter of the developed world population and more than Britain, France, West Germany, and Italy combined.

However, this effort is heavily weighted towards the Development side of R&D; only 14 per cent of this effort is directed towards basic research. A full 80 % of the R&D in Japan is carried out by private companies, whose scientists and engineers are in general better supported than researchers in universities. Much of this work involves product development, an area not conducive to sharing of ideas or use of public resources due to security concerns.

The universities, which in other countries are the source of most of the best original research, are relatively poorly supported from an R&D viewpoint, and funding per researcher has declined in recent years to less than half the average in industry. In Japan, university is largely considered a stopping point on the road to a career in industry, rather than a place for a young researcher to do creative work. Two factors that inhibit the universities' performance are the scarcity of resources and the strict hierarchy that prevails in Japanese universities (and other institutions). Bright young researchers must spend a large proportion of their time assisting their professors in chores which they are overqualified to do.

One result of this is that Japan does quite poorly in measures of scientific creativity and originality. For example over the past century Japanese researchers have received only five Nobel prizes, a remarkably low number considering the country's population and relative prosperity. Japan's most recent Nobel laureate did his prize–winning work in Switzerland and the United States, and stated that "he would never have had the freedom to carry out original research if he had stayed in Japan".

This statement is echoed in a recent survey in which more than half the respondents agreed that the Japanese environment is not adequate to foster unique and creative research. Officials are concerned that this situation perpetuates foreign perceptions that Japan uses the global scientific system for its own benefit and does not contribute its fair share. Such attitudes have given rise to concern that access to foreign basic research may be further

hindered or cut off in the name of "techno–nationalism" where governments jealously guard technology as a national asset.

As a result, Japan has recently taken several initiatives to increase the quantity and quality of basic research it finances. One major element of this is to increase the contacts of Japanese scientists with colleagues in other nations by inviting them to work in Japan; a second, taken by Japanese companies, is to set up R&D labs overseas employing foreign researchers. A major research program called the Human Frontier Science Program has been set up and financed by the Japanese government to coordinate research into the fundamental processes of life and intelligence. Although one objective is to make friends overseas, a more concrete one is to bring Japanese scientists into contact and competition with the best of their colleagues overseas.

The situation is very different in Japan's companies. Relatively free from concerns about short–term profitability, high interest rates, and takeover threats, they invest heavily in product and manufacturing process R&D. Two other uniquely Japanese factors favour this; companies generally have access to large amounts of capital at low cost and thus are able to quickly benefit from R&D results, and the traditional Japanese employee loyalty means that long–term R&D programs do not become fragmented by the loss of key researchers to competitors or to entrepreneurship, as is often the case in North America.

The companies also benefit from cooperative R&D programs coordinated by the government, in which favoured companies participate and all share the results. This allows much larger efforts than any single company could afford. Until recently, Japanese domestic markets were heavily protected (many still are) to allow manufacturers to perfect new products and recoup R&D investment free from foreign competition.

These advantages have produced arguably the most successful industrial economy in the world; one which is predicated upon a large industrial R&D effort coupled with the ability to transform the results into competitive products quickly.

The future will see Japanese companies maintaining their lead in development–oriented work, while government–financed initiatives in basic research are expanded both within Japan and in cooperation with other countries and researchers.

5.3 SCOPE OF THIS REVIEW

This review presents details on the three operational R&D–oriented networks (N–1, NACSIS and Junet) in Japan for which information is available, and an experimental network (Sigma) used as a testbed for product development. It does not include details on privately–owned networks which are not available to other users.

Aside from the networks that are national in scope, many universities have campus–wide LANS. Some serve widely scattered campuses or buildings of the same university as is the case of Kyoto University's KUINS and Tohoku University TAIN's. The KUINS network is unique in its extensive use of a digital PBX and X.25 packet switching system. The TAINS has a 100 Mbs fiber optic ring network. Since both are specific to a university they are not discussed in any detail in this document.

5.4 SUMMARY OF NETWORKS IN JAPAN

The Science Information System in Japan (see Figure 5.1) is very comprehensive with the National centre for Science Information System (NACSIS) as its nucleus. The telecommunication network interconnects over 500 universities, colleges, junior colleges and colleges of technologies is nation–wide in its coverage. This system covers science, the social sciences and the humanities; links university libraries, computer centres, information centres, information processing centres, national university research institutions, by computers and telecommunications networks; providing efficient scholarly information required by university researchers. It is anticipated that the system will expand links to include information systems of the private sector and from organizations that exist overseas. The major functions and services of this system are as follows:

- Collection of and search for primary information on science from various journals around the world;
- Search service of online databases consisting of 140 million books and 1.9 million journals held at 460 universities and institutions. This includes various types of information such as numeric and graphic/image; and
- Research and development into specialized computer hardware, software and database management methodologies;

Table 5.1 highlights some of the major characteristics of the national networks in Japan.

Participants: Only HEPnet permits industrial use with some 176 companies having access to this network. Clearly, JUNEt has the largest number of participants indicating the popularity of UNIX systems in Japan.

Services: The Science Information Network is by far the most ambitious network in Japan in terms of the scale of services that are being supported.

Protocols: A very wide range of protocols are supported by HEPnet. The remaining networks in Japan support a narrower range of protocols. For instance N–1 supports N1 and X.25. Junet supports TCP/IP and UUCP.

Bit Rates: The line speed of networks in Japan are relatively slow (ie. 9600 Bs – 48 Kbs) when compared to networks in the U.S. and Europe, which generally run at 1.5 Mbs.

5.5 JAPANESE NETWORK DESCRIPTIONS

5.5.1 EXISTING RESEARCH/UNIVERSITY NETWORKS

5.5.1.1 N–1

Background

- Earliest large computer network, first operational in 1981;
- By March 1987 connected about 64 mainframes at about 20 universities;
- N–1 is the name of the protocol used on the network;
- N–1 supports remote login and RJE but no mail service;

TABLE 5.1: MAJOR NETWORKS

	HEPnet	JUNET	BITNET	N-1 (8)
PARTICIPANTS				
UNIVERSITIES	30	40	27	100
COMPANIES	176			
COMPUTERS	>100	>1,000	48	200
PROTOCOLS	DECnet (1) TCP\IP (2) SNA (3) ASCII (4) FNA (5) HNA (6) X.25	UUCP (7) TCP/IP	RSCS	N1 X.25
BIT RATES	9,600 b/s 1,200 b/s	9,600 b/s 64 kb/s	9,600 b/s	48 kb/s

(1) DECnet: Digital Equipment Corporation Network Protocol
(2) TCP: ARPANET Transmission Control Protocol
(3) IP: INTERNET Transmission Protocol
(4) SNA: IBM System Network Architecture
(5) FNA: Fujitsu Network Architecture
(6) HNA: Hitachi Network Architecture
(7) UUCP: UNIX to UNIX CoPy
(8) Eventually replaced by the Science Information Network

Source: Dr. Elmer Hara, University of Regina; February 1990.

FIGURE 5.1: SCIENCE INFORMATION SYSTEM

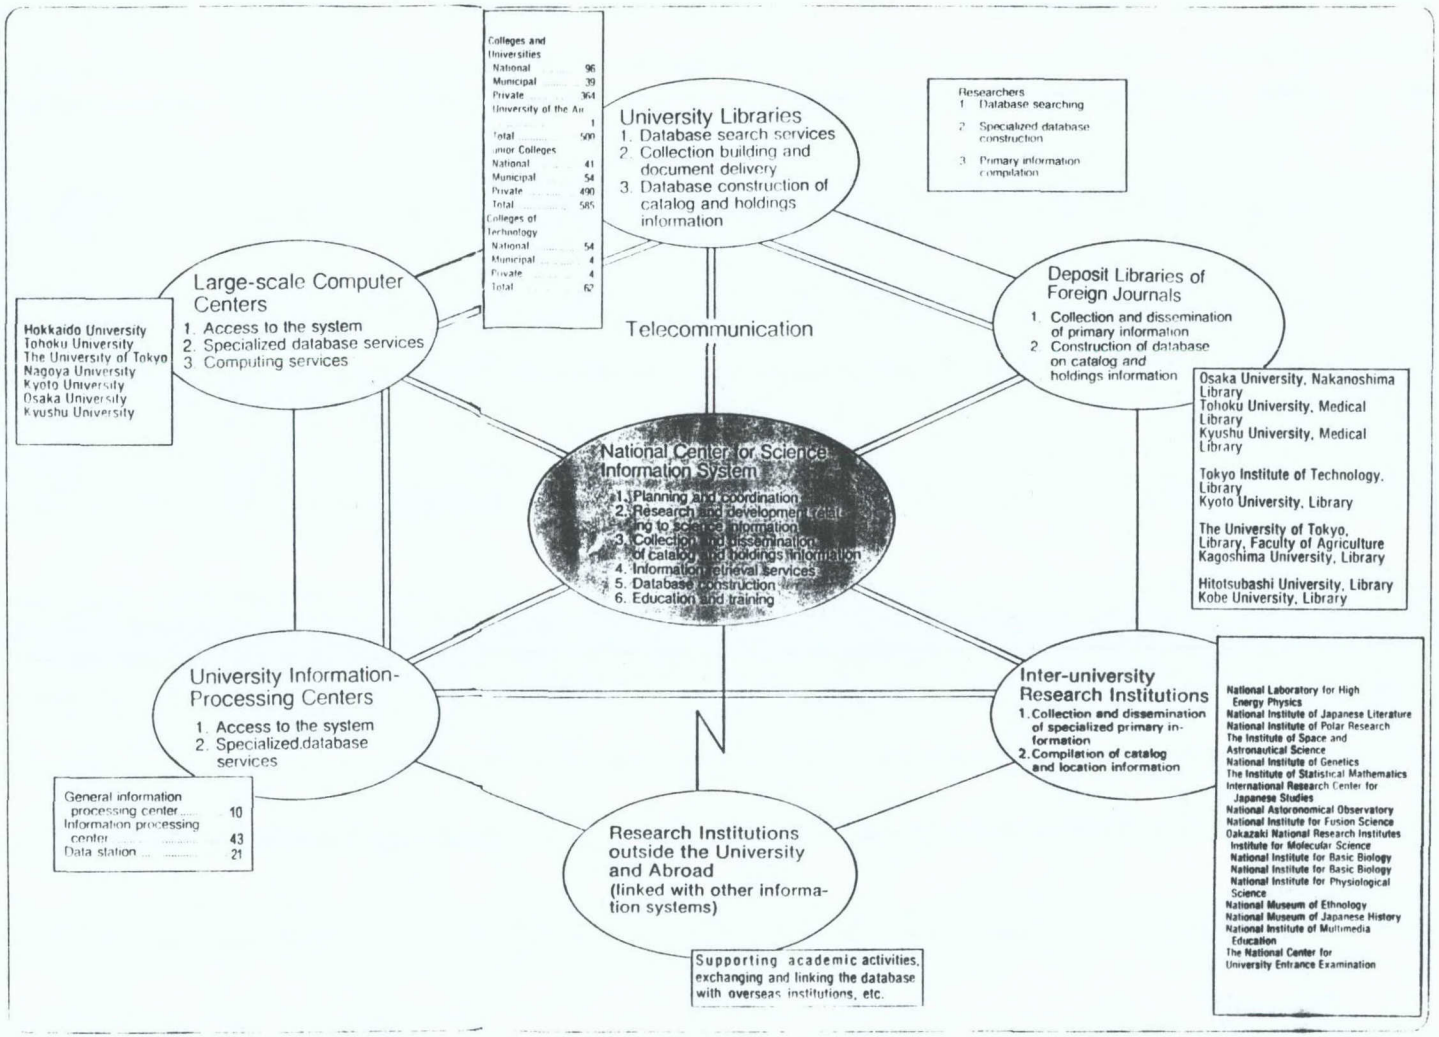

Source: National Centre for Science Information System (1989/90)

- Network used to share special hardware (array processors), software (PROLOG/KR) and databases.
- Line speed is generally the most common
- NACSIS established in 1987 can be considered as the successor to N–1.

5.5.1.2 Science Information Network

Background

- The National Centre for Science Information Systems (NACSIS) created in 1987, operates a network called the Science Information Network (also referred to as NACSIS).
- This network is the successor to N–1;

Technical Design

- Privately operated packet switched network employing high speed digital circuits to interconnect packet switching nodes;
- Tree–shaped network with its own packet switches;
- Line speeds range from 48 Kbs to 768Kps.
- The network uses N–1 and X.25 protocol and has the same services as N–1. Experimentation is currently underway to solve the lack of a mail protocol.
- The network will support network interconnection, closed group services, and virtual network operation as well as protocol conversion. In order to realize these network functions with adequate efficiency, network protocols will be adopted to OSI.

Implementation Design

- Most researchers belonging to national universities and national research institutions are eligible to transfer research–related information through the network.
- The incorporation of an integrated communication service into the network is planned. This will include full texts, graphics, voice and image communication.

5.5.1.3 Japanese University / UNIX Network (JUNET)

Background

- The major nationwide noncommercial network in Japan, started in 1984
- Purpose is to promote information exchange among Japanese researchers and with researchers outside Japan.
- Also provides a testbed for research on networking and distributed processing

SCIENCE INFORMATION NETWORK

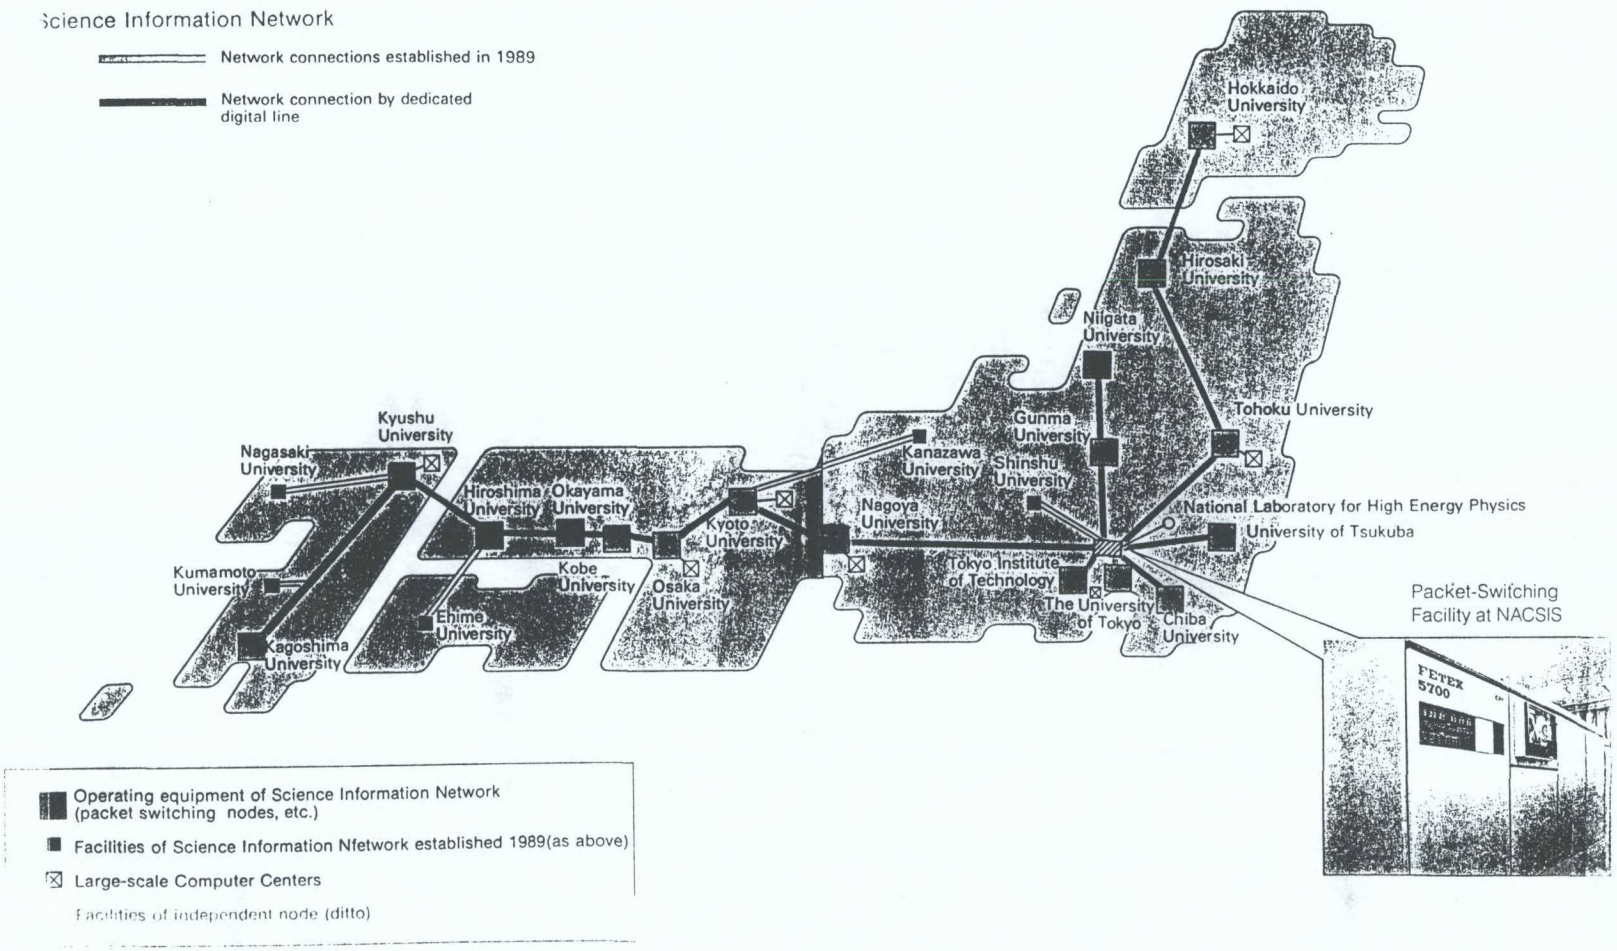

Source: National Centre for Science Information System (1989/90)

- As of August 1987 there were 87 connected organizations across Japan, with concentration in the Tokyo and Osaka areas.
- 85 % of traffic is news, 15 % mail. Connecting organizations are from a spectrum of universities, government research centres, and corporations

Technical Design

- Most JUNET links are 9.6 Kb/s; a major area of research has been maximizing utility of these links
- UUCP is main protocol, although there are facilities for TCP/IP connections
- JUNET uses the JIS X 0208 standard for communications between machines, which specifies 16–bit characters and uses escape sequences to indicate character set being used. ASCII text is default; this simplifies connections with other networks such as USENET.
- There are two major international gateways with connections to EUnet in Europe, UUCP and CSNET in North America

Implementation Design

- Administrators of backbone hosts meet monthly to administer network
- Connection costs are paid by each institution

5.4.1.3 Sigma

- An R&D testbed network which connects cooperating vendors in the Sigma project
- The Sigma project is intended to produce a standard workstation environment for use in Japan
- Network allows quick and simple distribution of developed software and other information concerning the project
- Sophisticated naming system insulates users from network operation details
- Underlying protocol is TCP/IP using X.25 over the PDN.

BITNET TOPOLOGY
As of February 6, 1989

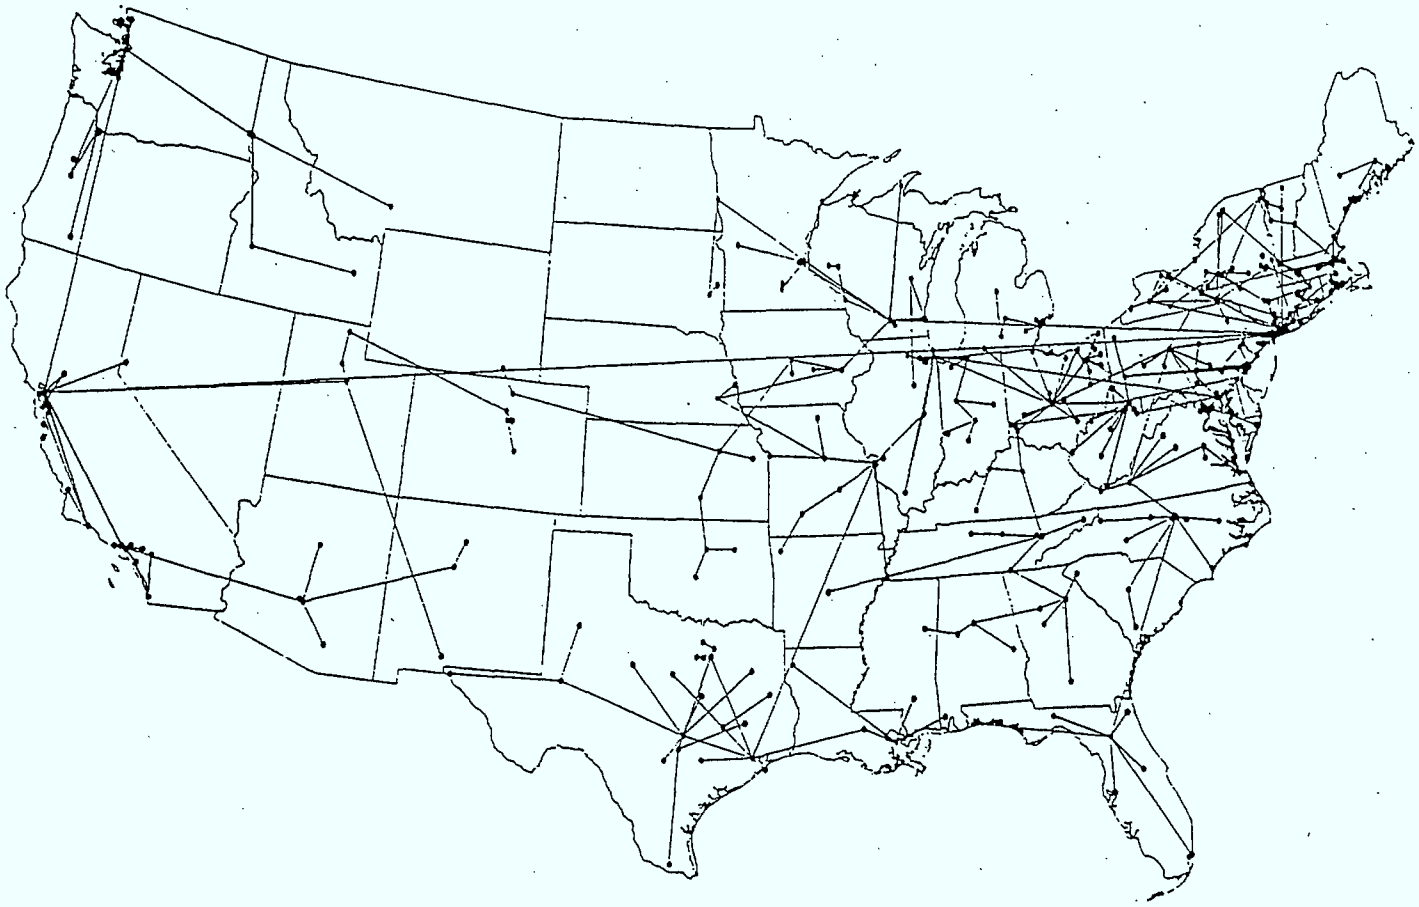

Source: Users' Directory of Computer Networks; Compiled by Tracy Lynn LaQuey, University of Texas; July 1989.

APPENDIX A: MINUTES OF EXPERT PANEL SESSIONS

SUMMARY MINUTES OF EXPERT PANEL SESSION 1
DECEMBER 18, 1989
DELTA HOTEL, MEETING ROOM ALGONQUIN C
9:00 a.m – 3:00 p.m.

I ATTENDEES

Panelists:

Ken Anderson (Mitel)
David McPherson (AECL)
Eugene Zywicki (Gandalf)
Mike Jager (EMR)
Mike Pascoe (Newbridge)
Dr. Probert (Ottawa University)
Dr. Kaye (Carleton University)
Dr. Bjerring (Western University)

Observers:

Pardeep Ahluwalia (NSERC)
David Astles (Centres of Excellence)
Bill Hutchison (Hutchison Research)
Dr. Digby Williams (ISTC)

Consulting Team:

Verne Chant (HICKLING)
Nick Gwyn (HICKLING)
David Arthurs (HICKLING)
Phil Kennis (HICKLING)
Dan Hara (HICKLING)
Roger Choquette (Comgate)
Tony Capel (Comgate)
Kal Toth (CGI)

II AGENDA

1. Introduction to the project and objectives of study (HICKLING).
2. Self introduction of all attendees.
3. User Needs and Technical Requirements
4. Provision of Test Facilities
5. Network Administration/Management/Implementation

III CURRENT STATUS OF THIS PROJECT

1. Digby Williams explained that CAnet has been initiated and organized by NRC and ISTC are advisors to the NRC. Vision 2000 is the second major piece of government involvement with an objective of personal networking capability by the year 2000. Direct thrusts are currently being worked out by an industry committee and will be completed within the next two months.

   ISTC needs to have a vehicle for communication R&D to improve international competitiveness and as a means of testing systems and performing collaborative R&D. The steering committee for this project is composed of representatives from DOC, NRC and TB, and chaired by Cliff Mackey (ADM). DSS, Public Works and DND have been informed about this network. If it is decided to get into this type of network it will be driven by ISTC. The ISTC point of view is the longer term. However the user costs that result from the implementation of this network would not be that much different than the user costs in place right now (acknowledging that rates are going to come down over the next five years). Rates have to be affordable or the network will not endure.

2. Dr Bjerring acknowledged the longer term vision of ISTC but reiterated that the research community has short term needs in terms of increased connectivity and improved university and industry interaction, and is concerned with the high tariff structure which makes it difficult to compete with American universities – Canet can be T1 in a week or so and OSI at the same rate as the American's NSFNET however the problem is paying for it. Dr. Bjerring stated that he hoped ISTC would respond to these shorter term needs by working closely with NSERC and by supporting the perpetuation of funding to both NRC, for the establishment of CAnet, and to the provinces to maintain existing regional networks. Industry is becoming encouraged to become involved in CAnet, at the regional level.

IV NEEDS OF CURRENT RESEARCHERS AND LIKELY IMPROVEMENTS TO EXISTING SITUATION

1. Dr. Probert, speaking as a representative from TRIO, suggested that the connectivity between universities was not that bad and personally does not have a problem with bandwidth (but no requirement for access to super computers and large file transfers). The difficulty is in trying to encourage industry participants to utilize E–mail. The obvious reason that industry has not participated more is that the user interfaces on existing networks is not adequate and requires further enhancements.

2. From the panel it was determined that there are a number of generic applications that are relevant to this study but have not been identified in the agenda ie. teleconferencing and video–conferencing, distinction between point–to–point vs. broadcast, and the requirement for voice data and multi–media work stations.

3. Dave McPherson illustrated that the user interface to research networks exists but it takes perseverance to make it do what you want it to. There is a real requirement for a standardized user interface — as under the existing situation people are forced to become experts in a whole variety of computing environments. It is possible to transfer files through NetNorth but it better not be a very large file — there is a requirement to send files orders of magnitude beyond existing capability.

4. Dr. Kaye explained that one way or another the research community can get through to almost anybody, anywhere — however this requires keeping track of many things.

Why is fax network totally different from E–mail network? Band width is not available for graphic image files and video–conferencing although it is used to a small extent by industry, it is not simple nor elegant to do it. With air travel becoming more dangerous and expensive there is increasing demand for video and media conferencing. A 1.5 megabit network must be seen as a short–term step and must move quickly to at least a 45 megabit network. Research problems and challenges are with a 100 megabit network therefore from a test–bed point of view T1 does not make that much sense.

5. Dr. Bjerring explained that NetNorth is a store–and–forward network, which means that bottlenecks may occur, and in this sense a user on one end really does not interact with users at the other end — bandwidth is not really a problem on such a network. In true interactive computing, remote login and large file transfers, band width becomes very important. 100 megabits is an implementable local area network speed. Long distance, long haul networks are a little away from this. Important to keep in mind the hierarchy from what is going on at the desktop and on down the line.

6. Pardeep Ahluwalia mentioned that there are super computer facilities that people cannot connect to because there is no effective way of getting into them. Many of these facilities are sitting relatively idle because of the lack of bandwidth. It is necessary to improve the utilization of these facilities soon before they become obsolete.

7. Dr. Probert responded to a question posed by Bill Hutchison that technology transfer is an area where large gains should be made — in particular distributed software engineering. This is a point to multi–point requirement. Tremendous productivity advantages exist if the expertise of a limited number of experts is broadcast around the country in an interactive way.

8. Dr Kaye stressed that he agreed with Bill Hutchison that there is a pressing need to keep up with other countries and to be looking ahead. American researchers are well on their way to implementing multi–media conferencing (which cannot run over T1 lines) and movements of hyper text which may involve voice, data, graphics and video. Dr. Kaye also mentioned that Canadian universities do not put forth adventurous, farsighted proposals because the infrastructure is simply not available to support them.

9. Dr. Probert noted that company's dealing with him are $25 million or greater in size. Dr. Kaye added that in his opinion size is not an issue but rather importance is placed on whether or not a firm is a forward looking company. Both panel members felt that the Centres of Excellence would become vehicles for small company involvement in R&D.

V THE CONCEPT OF A TEST-BED NETWORK

1. Dr. Bjerring stressed that "research" in the context of this network means two things – – research on the network vs. the network that is used to support a broad–based research community. The research community may be concerned that if they are using the network in a production sense, somebody else is using it to test and tweek certain aspects of it — which may crash the system.

2. Ken Anderson mentioned that there has to be some identified commercial advantage. Mitel has their own network and are well linked around the world. Mitel makes

considerable use of USEnet which was established in the States about 10 years ago. This network is used to post questions on a number of broad topic areas.

3. Dr. Kaye mentioned that the possibility of crashing the network is more likely with some types of research than others. Testing the physical path over which data flows is a problem. However there is a great deal of research that is required that has nothing to do with physical path — ie design of efficient gateways and protocols. No reason to see why this common use could not exist with users who require reliable service. If wish to test a given protocol stack that is quite different from the one in general use, and have a willing partner there is no reason to suspect that this will cause interference with other users. Therefore this network could and should be designed so that both of these objectives are served. Must be forward looking — has to be much more than just a cure of current problems. Has to grow and support future oriented research. (Networks of the future will need to eliminate intermediate store and forward nodes in networks).

4. Mike Pascoe highlighted that there are gains from a supplier's point of view as being a part of a world class project. The whole value to the manufacturer is not the products sold into the network but rather the gains in the intelligence and the optimization of that product. Suppliers should not worry so much about the cost of this network being prohibitive. In addition, in order to develop new products Newbridge has to see a market, and that market is directed largely by what is happening in the States. Although there are many value–added developments that could be put together, Newbridge is more interested in moving ahead with technology and trying to provide backward compatibility.

5. Eugene Zywicki commented that the common element between CAnet and NSFnet is the Internet protocal for routing and if you are building on top of that you are limited to the network that you can experiment with. However Dr. Kaye mentioned that there may be merit at testing at a lower level, different networking protocols for higher speeds which could not be tested on CAnet. (CAnet could well become a service provided by this network). The rest of the network could be used by IT testers and suppliers for the migration to OSI).

6. Eugene Zywicki's notion of testing on the network is that media access must be predefined, and the necessary conversion is done to ride on top of the network without affecting the network. To build a higher speed OSI network with all the protocol stacks embedded in the network is not going to suffice. (Don't want another X–25 type network). Dr. Kaye noted that the network should provide transparent, edge to edge data flow transfers with variable bandwidth access at the physical level in order to permit R&D on the network to its fullest capacity.

7. Dr. Bjerring suggested that researchers are more interested in talking south of the border and that technology that we build into our networks must be the technology that is being built into the networks south of the border.

VI IMPLEMENTATION/ADMINISTRATION/MANAGEMENT OF THIS NETWORK

1. Dr. Probert mentioned that the rate structure is very important. The only way to keep the network is if it can pay for itself. High speed computers sit idle because of the existing rate structure – ie. most of the users belong to a single institution. How should the rate structure be set up for these different types of users.

2. Eugene Zywicki mentioned that with ONET they pay a annual membership fee and then as long as they get value out of this they are happy — ie. do not cost justify every connection. Where is the line drawn at using the network from a research point of view or using the network for own business? Who are the users of this network?

3. Ken Anderson mentioned he would like to see this session influence government policy in some way. MITEL up against companies in the States like Intel, that are backed by the U.S Government, and up against Japan Inc. and the entity of the European common market. There is a lack of a national policy in Canada. Government should make it strategically important and perhaps be the benefactor of this type of network.

4. Dr. Bjerring pointed out that the benefit of the networks as they now stand is that universities pour significant amounts of free labour into designing and managing these networks. Universities have been involved in network management for 5 or 6 years in Canada. Bjerring does not see this as a necessary activity for the universities to be in. Although private sector bids were received to run and operate CANET, the University of Toronto won out.

5. According to Dr. Kaye large universities are not involved in proprietary research, and that from their perspective security issues have tended to be somewhat overstated. Intrusion of viruses, on the other hand, is a much more important issue. However, security is important for a university in protecting certain kinds of administrative type data like student marks etc.. Depending on the company or agency security is a big issue. Some activities carried by AECL requires enormous amounts of security and carefully controlled interfaces.

6. It has been Dr. Bjerring's experience that universities are dead set against metering and usage based fees. Don't want to get into the administrative problems that would result with a metering system. In addition, the window environment envisioned would not work with such a system — too restrictive and would force users to find "other" ways to do things.

VII ECONOMIC FEASIBILITY

1. Dr. Bjerring provided some detail on the costs and rates for current networks. In Ontario, Netnorth (which is primarily an electronic mail network) consists of 2400 – 9600 baud leased lines, with 30 members across Ontario. The cost is about $6,000 for each member. ONET is largely a 56 Kilobit network, which supports electronic mail, file transfer, and telenet–remote logon, and the membership fee is $18,000. There are 15 members right now, includes NRC, Gandalf, BNR and three university based groups, HEPNET, IRTC and ISTS (which are two designated centres of excellence with private sector involvement) and the Ontario Centre for large scale computation – Cray facilities. Access to CAnet in Ontario could cost up to $15,000 additional if CAnet were a 56 kb/s network.

2. Costs identified above are for leasing lines. Administrative activity for NetNorth (keeping routing tables up to date etc.) is a small fraction of the total. With Onet (a TCP/IP network) there are no administrative aspects at all. For CAnet some where between 40– 50% of the $1.2 million (proposed) per year which is the steady state cost of a 56 kilobit network, is operating costs and the rest is leasing the lines. The $18,000 per year is the total share of operating the network which does not include the startup cost of buying a router (about $8,000 each). AECL noted that operating

costs represent about 50% of their total costs.

3. Eugene Zywicki mentioned that it is important to address the opportunities that we have missed up until now because of the lack of proper band width — ie. go south because of bandwidth.

4. According to Dr. Bjerring, Canadian universities spend well over $100 million per year on library services and serials. One of the real benefits of higher band width is to provide researchers across the country with a more efficient service.

5. According to NSERC, opportunities missed are common. Cooperative efforts with the States because cheaper for the researcher in Montreal to connect to NSFnet and get access to the various super computer facilities or accelerators then go through Toronto. We can't do this in Canada because the band width costs too much. There is a problem with both connectivity and band width. May not be losing people as a result of this but having trouble attracting people trained south of the border.

VIII THE EVOLUTION OF CAnet

1. From Dr. Bjerring's perspective there are three objectives of a national backbone network (1) increase collaborative R&D; (2) increase university/industry interaction; and (3) the creation and use of a test–bed network. Dr. Bjerring summarized that given the above three objectives what is the rationale for doing anything other than maintaining CAnet and letting it evolve into meeting these three objectives? Perhaps management and structure of CAnet may not encourage private sector involvement as much as it should however this can be improved by further private sector involvement at the regional level.

2. The panel noted that the relationship between Vision 2000 and the network being studied is primarily through the third objective eg. R&D in the communications area including using the network as a test–bed. The panel agreed that the CANET model does not cater to such a user.

3. According to Dr. Bjerring, NetNorth's strategic plan emphasised that it would evolve into CAnet. CAnet to a large extent will evolve in a similar manner as NSFnet, in terms of protocols (TCP/IP) and in terms of the types of implementation. Compatibility with NSFnet is very important as CAnet and Onet are viewed as extensions of the Internet. If the largest market of private sector is the U.S, then there is nothing wrong with the network being used as a test–bed which is an extension of the products and services offered in the States.

4. Eugene Zywicki commented that the common element between CAnet and NSFNET is the Internet protocol for routing and if you are building on top of that you are limited to the network that you can experiment with. In fact CAnet could very well be a service provided on this network. The rest of the network could be used for IT testers and suppliers for the migration to OSI etc..

5. Dr. Bjerring added, that being realistic, the backbone speeds being developed on NSFnet will always be higher then the national backbone speeds in Canada. However the Local area networks protocols and the network standards that tie into this national backbone need not follow this rule. There is no reason however that a national backbone could not be used to test environments like the successor to FDDI2 and other products at this level ie. local area environments.

SUMMARY MINUTES OF EXPERT PANEL SESSION 2A
OTTAWA, ONTARIO
JANUARY 18, 1990
DELTA HOTEL, MEETING ROOM DELTA C
9:00 A.M - 12:00 P.M

I ATTENDEES

Panelists:

Ed Acheson (Treasury Board)
Bill McCrum (DOC)
Michael Brandreth (CISTI)
Barry Reed (CRTC)
Al Keddy (GTA)
Roger Taylor (NRC)
Michael Roberts (EDUCOM)

Observers:

Digby Williams (ISTC)
Joe Padden (ISTC)

Consulting Team:

David Arthurs (HICKLING)
Phil Kennis (HICKLING)
Dan Hara (HICKLING)
Roger Choquette (Comgate)
Tony Capel (Comgate)

II AGENDA

1. Introduction
2. Michael Roberts (U.S. Experience)
3. Priorities and Issues
4. Trends in Broad Band Services
5. Policy
6. Management
7. Other

III INTRODUCTION

The introduction consisted of the following:

- an introduction of the panelists;
- scope of this study; and
- background information to the U.S. experience.

An introduction of the panelists:

All panelists introduced themselves, and expressed their particular interests and point of view with respect to this project.

Ed Aceheson: Interested in setting policy with respect to the administration of information technology, and making wise investments in that context. Also very concerned with the adoption of information technology standards.

Bill McCrum: Manages the research program of DOC with regards to open systems research, involved in standards work and also engaged in new initiatives with regards to open systems. Previous involvement in CDNnet and some other research networks in Canada.
Mike Brandreth: Here to give the library point of view. Interested in compatibility, distributed databases and costs.

Barry Reed: Responsible for technical input with regards to commission decision making with regards to communication tariffs. Interested in harmonization of rules across the country.

Al Keddy: GTA has recently been assigned special operating status, designed to put the agency on a more business like footing. Interested in the possibility of GTA playing a role in the management of this network even though the concept at present serves a community larger than the Federal Government.

Roger Taylor: Director of the Division of Informatics at NRC, responsible for the computer needs of NRC and also interested in this project because of active role in the development and management of CAnet. Sees this network as a natural extension of what he is currently putting together. CAnet is a network which uses today's technology to connect research scientists and students across Canada.

Michael Roberts: Vice President for networking at EDUCOM which is a U.S. association of colleges and universities and has about 600 members. He is also Director of EDUCOM's networking and telecommunications task force which consists of about 60 U.S

universities. His primary role is to represent the university
constituency in support of the National Research and Education Network (NREN).

Scope of this study:

David Arthurs briefly outlined the ISTC initiative and the objectives and scope of this study. Please refer to handout for an understanding of the general concepts underlying this network.

U.S Experience:

Michael Roberts was asked to brief the panel on the U.S. networking experience and to bring the panel up to date regarding the NREN. The major points raised in this discussion are as follows:

- Looking forward to bringing the Canadian network along;
- There are about 8 million students in U.S higher education, with about 1 million professors with a budget of about 150 million dollars;
- What is going on can be characterised by both a bottom–up and top–down activity;
- About 20 years ago the defense research provided some funding for campus networking people to construct a network that would do something about the vulnerability of the old point–to–point circuits for permitting researchers to communicate by remote hosts. Out of this came ARPANET which is the first packet, connectionless architecture;
- DEC and IBM have both thrown in with TCP/IP and have announced to their international clients and U.S government that they fully intend to implement OSI by this fall. Already have transport layer of OSI running alongside the same machines with TCP/IP — most of the experts on this field do not see this as a problem. The real problem is to create a transport mechanism which has the technology available to push bits along fibre very quickly and that can packetise and depacketise these bits. The challenge then is to provide advanced services over the network to support end users;
- It is estimated that there are 150,000 host machines and about 1 million active users on the U.S. Internet. The U.S Internet is widely connected abroad and just in the last year, T1 connections have been put in place to Japan and Australia. There are half a dozen links to European researchers;
- Four or five years ago the U.S congress became convinced that the U.S was falling behind other industrialized countries in terms of quantity and quality of research and development. 75% to 80% of all basic research in the U.S is conducted on university campuses, primarily funded by the Federal government. Study after study showed that research institutions were becoming obsolete and the tools and facilities available to campuses were not state of there art. As a result, one of the major fallouts of this is that the budget for NSF has doubled over the last five years;
- With regards to high performance computer networking there is a federal initiative called the high performance computer program. The U.S Office of Science and Technology Policy has written several paper son this over the last two years and out of this work the NREN concept has emerged;
- Promotion of the NREN began in 1987 and the federal agencies decided to get behind this single initiative. Government and universities are behind the NREN. There are also alot of private sector support, picked up in the last year;

- One of the objectives of the NREN is to accelerate the development of products and services that are technology driven and that will dominate markets over the next 10 –25 years. The present status of the NREN is that there is legislation pending in the U.S congress, that will create a legislative mandate for the NREN;

- Approximately $100 M per year is earmarked for the next five years, divided evenly between solving the existing production problems on NSFNET (will evolve into NREN) and the second half would go towards the research and development of the gigabit network. The evolution of the NREN is divided into three stages. The first two stages will create T1 connectivity to 400 – 500 private and educational sites and upgrade that capability to stay ahead of the demand for T3 services. Simultaneously research into product development that would gravitate that out to gigabit levels will then take place;
- As the above is going on the telecom industry is really scrambling to get back on top of their own technology problems. Many have companies have coalesced along the broadband ISDN concept High Speed Transport Mechanism. The gigabit test bed projects in the U.S. are going to use SONET for the transport layer so that the 10–15 year picture of the national telecommunications fabric and activities of the carriers will evolve to an all digital network at gigabit speeds. The NREN will add its own layers of applications;

- The U.S. is beseiged by Super Science projects like the Hubble space telescope. The data rate on this telescope alone is a continuous megabit per second. Two to three hundred researchers require immediate access to this information. This as well as other projects like the world wide seismic network cannot go forward without the existence of a high speed computer network;

- There is enormous agreement in the U.S that the NREN is something that should be done and should begin as soon as possible;

IV ISSUES

Bill McCrum questioned the acceleration of the implementation of this network, at a time when Telecom Canada has a very definite plan with regards to the evolution of their services and introduction of broad band services across Canada. As these plans exist, if we go outside of the timing of these plans there are enormous costs associated with this acceleration.

Michael Roberts commented that carriers have to be very attentive to market concerns. The determining factor for U.S carriers upgrading to Sonet technology is the capability to make them competitive in carrying digital information. Video is a very big market and it is primarily carried in analogue form toady. Unbelievable possibilities if you have a fast enough chip for compression of video images. Right now the expansion factor on analogue to digital video is such that there isn't any digital video out there. Most carriers want to get out of the analogue and especially out of analogue microwave and out of satellite because all advanced research is in terrestrial fibre. If we want to plan forward in time like handling video there is a need for higher bandwidth.

Michael Roberts also pointed out that in the States there is no plan for dedicated facilities at gigabit rates. The assumption is that by the time the necessary hardware and software developments to the upper layers and the applications have been done the commercial

carriers will have gigabit transport in place.

Bill McCrum questioned whether 80% of the needs are being already met? There may be some very specialized needs for very high bandwidth that could be met by leasing some very high speed lines. Accelerating the pace beyond the development of standards may come with a very hefty price tag.

Roger Choquette pointed out that there is general consensus that the current installed Telecom Canada capabilities would certainly not satisfy the requirements of the network in the long run. In the short run even if it does satisfy a T1 capability it is going to be prohibitively expensive.

Bill McCrum directed a question to Barry Reed from the CRTC regarding the potential significant reductions in the long haul service tariffs. Mr Reed commented that these rate reductions may be in the order of 50% for T1.

Roger Taylor added that the problem in Canada is that outside of BITnet and CDNnet, which are very slow networks, one does not develop applications because there is nothing to run them on. Anticipating that with CAnet coming on, which is at 56kb, people can start using that to develop applications. Applications in computational fluid dynamics requires high bandwidth. We can learn from the U.S. experience — 56 kb was nice on NSFnet but it was soon filled up.

David Arthurs commented that one aspect of this study is to look at user needs and to get a picture not only of current needs but also of future needs. For today the needs are not all that critical. There are two things to look at with regard to these needs (1) the applications which just are not possible today without higher bandwidth and (2) general usage requirements are increasing thus multiple users requiring even medium speeds results in a backbone which must be of higher speed.

Al Keddy mentioned that to get an initial network up and operating the first thing you have to do is look at the user needs. What is interesting is how are these needs going to grow and how much are people willing to pay for these needs to be met. The user needs are probably a lot bigger than what any of us probably imagine sitting around this table. Not visible because people are not moving it that direction because there is nothing there to satisfy them.

Michael Roberts provided some insight on the costs T1 connections in the States. A campus connection to NSFnet (through regional network) is in the range of $100,000 per site (unsubsidized). This probably will not go up much for T3. This may seem like alot of money, but the top say 100 schools in America have institutional budgets of over 1 billion dollars. In addition schools like Stanford have toll charges in their voice systems that gets toll charges of $1.5 M a year. Putting this cost in context it is not proportionately alot of money.

Mr Acheson commented that in Canada today a certain amount of that need can be met today by carriers. But if the need is beyond that then it becomes a political decision as to whether Canada wants to be at the forefront of science, and wants to get access to the Hubble telescope and the like.

Bill McCrum mentioned that the telecom industry may be worried about wiring up the country with high band width facilities because of the danger of the value–added network people to come in and steal all the business.

V TRENDS IN BROAD BAND SERVICES

In response to Roger Choquette's question regarding the availability of broadband ISDN Michael Roberts had the following comment. Broadband ISDN is an intensely competitive market with lots of small companies where you get different answers regarding the availability of this technology if you talk to engineering or marketing people. At the Interop show it was demonstrated last fall in San Hose they were demonstrating running TCP/IP and OSI transport. A Plasma Physics workstation with a 100 mip machine, colour display with 3–D motion and connect to Cray requires a gigabit on its own. There is general recognition that we are not pushing the technology as fast as we should. There is considerable pressure to lend financial support and policy support to bring this technology along at a faster pace.

Barry Reed commented that Sonet may have blown narrowband ISDN out of the water, and can't see any body going ahead with narrowband ISDN. Sonet allows you to do alot of things that you cannot do now — therefore potential benefits for everybody. Narrowband ISDN does not allow you to anything that you cannot do now. It allows you to do two things on the same bit of wire but does not see this as a real benefit. Sonet type technology will be available in about five years through Northern. ISDN will be hard to justify on a tariff basis. The evolutionary architecture is to build on the existing infrastructure. This is not a change over process but an add–on process and bringing in boxes that can do things with what is already there.

Michael Roberts explained that Sonet is defined up to 3.4 gigabits. In addition optical fibre modulation goes up to a 1000 of gigabits, therefore we will not run out of steam with regards to pushing bits across fibre for a long time.

Barry Reed pointed out that the problem comes in switching it. Putting it in and out of fibre is not the problem. You have to eliminate all those upper layers in the protocals.

VI PRIORITIES

In terms of priorities for this network Ed Acheson is concerned most with standards — which go beyond the interconnection capability at the telecommunication level. This is not only a telecommunications problem but a total system problem. Therefore it is important to speak of standards at both the bit level and the services level. In addition this network should be designed to give Canada a market share of something. Because we are going after this market share on an international basis, we should build on some of the international standardization acts that are out there.

Roger Choquette mentioned that the network will be used as a vehicle for the development of these services, as opposed to necessarily saying the network must have these services. The concept up to now is to make the network as transparent as possible. Most of the functionality would be in the end systems rather than in the network.

Digger Williams added that with respect to standards we must relate government goals and government programs to the next level of standards which are being developed. This could become a priority in terms of government procurement.

Ed Atcheson added that the second priority is to get the rates down. Our inter–city facilities are out of line in terms of costs. There is an incentive to use NSFNET to transport data between two sites in Canada. Roger Taylor mentioned that this is currently going on, and that NSFNET is concerned about this.

Bill McCrum mentioned that DOC is promoting the use of standards. OSI is going to become a very prominent area of interest within the department. The department is supporting a number of organizations across the country to introduce OSI concepts into real systems. DOC is interested in this network to the extent that it could lead to products that take advantage of international standards. Its use as a catalyst for technology is important but is unsure of how it would work. Connectivity and availability are tied into how you can harness the evolution of technological standards.

Michael Brandreth is not sure whether the high speed is directly related to the library problems. There is an increasing interest in revenue generation — therefore have to get out to more than traditional clients. The existence of a high speed network would give this extra push that does not seem to be there with existing technology. Although there is no doubt that libraries could have done far more with existing technology tha what they have done. Connectivity and standards are very important issues. Libraries in general are going to be interested in different types of services and different types of clients, and revenue generation. Speed is also an issue. Interested in image transmission to construction sites. Colour photographs to operating rooms.

Roger Taylor added that libraries will continue to have paper on their shelves until there is a more efficient way of methods of delivering iiformation to people. The network is certainly one of these. Optical discs, will not happen suddenly.

Barry Reed commented that from the CRTC perspective interested in a consistent framework existing across the country in terms of regulation. For a network like this one might want to try to enlist the carriers and put it together on a trial basis to not have to file tariffs. This is a quick fix. If you want something on going then one needs a shift in the tariff structure. The Commission cannot take a driving role in this. Not sure who will even regulate in the future. If you want to bring voice onto your network get into a whole host of regulatory problems that you do not exist if the network is data only. Roger Taylor added that consistency must exist in Canada but also with U.S. and international peer countries.

Digger Williams highlighted the mandate of ISTC as to increase the international competetiveness of industry, and to create an environment which would increase the numbers of trained personnel. This network can improve the visibility of Canadaian technology and for scientists to carry out cooperative research.

Al Keddy commented that he is most interested in the management structure of this network and whether or not the government is getting value for its money. Interested in the flat fee structure as in the U.S

Roger Taylor idicated that costs cannot be ignored or this network will not be used. The NRC mandate is scientific research and therefore interested in tools for science. The network should be seen as an efficient tool for science. Should take advantage of new technology therefore test–bed is important but cannot interfere with the service. Connectivity, availability all of these thing are important. NRC also interacts with industry and thus the catalyst is also very important.

VI POLICY AND MANAGEMENT

Roger Taylor indicated that not every body will have access to CAnet — it will be accessed only by regional networks. The regional networks can set their own rules but will provide access to all institutions that have a reserach or education interest. Designed to get all these sectors working together. An essential feature of this network is that it must be connected to the entire research community. If connected to NSFnet must make a commitment to use

network for similar purposes.

David Arthurs pointed out that the underlying premise behind the flat fee is that you it will not limit the use of the network. However a flat fee which is beyond the means of smaller companies limits the access by these companies.

Roger Choquette added that some consideration should be given to a fee structure which accomodates low speed access dial up.

Michael Roberts pointed out that 30% to 40% of the addresses on the Internet have a commercial address. There are essentially 21 regionals now — and are non–profit centres. The board of directors of these entities have a industry and academic representation. In the U.S the backbone is also run by a non–profit organization — MERIT. In the U.S the fundinq question is one really of semantics. At least 3/4 of the cost of the federal backbone in the U.S. comes down to the question of whether the federal agencies want this as a line item in their budgets that go directly to the backbone or do they want it as a line item in their grants that has to circulate through the system before it comes back to the backbone.

Michael Brandreth pointed out that from the library point of view the more clients the better. CISTI has been operating an information retrieval system for about 15 years. — CAN/OLE. The early principle underlying the CAN/OLE introduction is one of equality of service. There cannot be a penalization based on geography. Therfore collect calls are accepted.

Roger Taylor mentioned that NYSERnet has various levels of connectivity and speeds and their fees are adjusted to reflect this.

Roger Taylor described CAnet as being operated by The University of Toronto, IBM and INSINC as a non–profit organization controlled by a management board. The management board is made up of representation from each of the regions. So the regions will be set up in a similar way. At the regional level there is a lot of implicit subsidy going on with respect to universities donating time etc.

David Arthurs asked if this might change if the regional networks became more widely used by networks

Roger Taylor indicated that it might and may become more formalized and legal liability may become more of an issue especially with respect to medical imagery (ie. piping information into an operating room — and the network crashes).

Michael Roberts commented that with respect to security the onus has to be shifted against proprietary because it is a public network — public gain and public cost. The trends in the U.S are non commercial and non proprietary. The primary objective of this network is to transfer knowledge in universitys to industry. If you set up all these rules usesr will be scared to death — and defeat purpose. At present it is dealt with as a gentleman's agreement to behave properly — and this has worked.

Michael Roberts commenting on the test–bed concept mentioned that at the transport level you need a robust service for a research network. There are some links within NSFNET, which are run by MCI, that are not open and are used for experimentation. These links are used for experimentation with protocols and the higher layers.

SUMMARY MINUTES OF EXPERT PANEL SESSION 2B
OTTAWA, ONTARIO
JANUARY 18, 1990
DELTA HOTEL, MEETING ROOM DELTA C
1:00 P.M - 5:00 P.M

I ATTENDEES

Panelists:

Robert Vaive (NRC)
Michael Brandreth (CISTI)
Peter MacKinnon (Cognos)
Steve Cockle (Connaught Labs)
Prof. Egelstaff (Guelph University)
Chris Hughes (EMR)
Ike Goodfellow (Northern Telecom)
Harvey Nelson (McMaster University)
Roger Taylor (NRC)
Michael Roberts (EDUCOM)

Observers:

Digby Williams (ISTC)
Joe Padden (ISTC)
Phil Enross (Science Council)
Bill Hutchison (Hutchison Research)

Consulting Team:

David Arthurs (HICKLING)
Phil Kennis (HICKLING)
Dan Hara (HICKLING)
Roger Choquette (Comgate)
Tony Capel (Comgate)

II AGENDA

1. Introduction
2. Michael Roberts (U.S. Experience)
3. Priorities and Issues
4. Applications
5. Anticipated Benefits
6. Implementation
7. Other

III INTRODUCTION

The introduction consisted of the following:

- an introduction of the panelists;
- scope of this study; and
- background information to the U.S. experience.

An introduction of the panelists:

Robert Vaive: As manager of the Canadian Space Station Program Technical Management Information System is responsible for handling all the networking, data storage, communications for Canada's involvement in the Space Station Program. This involves participants from universities and industry, across the country. His dilemma is how to create, store and distribute large quantities of data to the U.S and 13 other member countries. Not all activities are R&D but a very large segment is. The networking requirement is considered to be an operational requirement. He has a very strong interest as a user of this network.

Michael Brandreth: Large supplier of documents, supplying about 1500 documents per day in the form of photocopies. CISTI can provide services on this network in the form of database searches, bibliographic searches etc..

Steven Cockle: Researcher with Connaught which is the world's largest producer of human vaccines. The information retrieval potential of this network is very appealing. Connaught receives significant amounts of scientific material in written form through CISTI and does make use of database searches. Connaught requires access to high speed computers for molecular calculations.

Peter Egelstaff: A Physics professor at Guelph University has a number of interests. On Netnorth 150,000 files per week move through Guelph. The campus library is quite modern and is anxious to receive data from a high speed network. A requirement exists for a central library for North America — with a capacity greater than anything which exists today. The veterinary school has a real need for electronic image transfer — same as medical schools. In physics many high computational experiments going across the country and in international labs. What will happen in the future is not what has happened in the past. Education is very important and should be given higher exposure.

Chris Hughes: Director of Computing at EMR, and also member of Supernet Consortium which is making use of the Cray in Montreal. The objective of Supernet is to permit users to gain cost–effective access to high speed computer resources. This ISTC high speed network could facilitate this. Would like to see a more cohesive approach to networking. Tis high speed network should (1) provide a vehicle for cooperative research; and (2) enable the exchange of applications and data and software across the country. A particular requirement is high speed transfer of large data. Currently run in a batch environment and are moving towards an interactive environment. Interested in high speed networking as a tool for high speed computing.

Ike Goodfellow: Vice–President at Northern Telecom and interfaces with the BNR Labs. His interest is two fold (1) interconnection between the labs and (2) the transfer of data from labs to manufacturing plants.

Harvey Nelson: Department of Epidimiology and Biostatistics. Is very involved in the

evaluation of clinical trials, new drugs, vaccines, or surgical vs. medical treatments. Real need for image transfer and convenient access and storage of data.

Roger Taylor: Director of the Division of Informatics at NRC, responsible for the computer needs of NRC and also interested in this project because of active role in the development and management of CAnet. Sees this network as a natural extension of what he is currently putting together. CAnet is a network which uses today's technology to connect research scientists and students across Canada. Education is very important. The regional network in Manitoba — is being joined quite readily by the community colleges.

Peter Mackinnon: Manager of CASE technologies at Cognos and is the largest software company in Ottawa. Cognos conducts extensive use of communications internally and also with universities. Cognos is involved in distributed computing. In addition was personally involved in the original creation of Supernet, NRC networking, and worked initially with ISTC on the proposal, which has led to this study.

Michael Roberts: Vice President for networking at EDUCOM which is a U.S. association of colleges and universities and has about 600 members. He is also Director of EDUCOM's networking and telecommunications task force which consists of about 60 U.S universities. His primary role is to represent the university constituency in support of the National Research and Education Network (NREN).

Scope of this study:

David Arthurs briefly outlined the ISTC initiative and the objectives and scope of this study. Please refer to handout for an understanding of the general concepts underlying this network.

U.S Experience:

Michael Roberts was asked to brief the panel on the U.S. networking experience and to bring the panel up to date regarding the NREN. The major points raised in this discussion are as follows:

- Looking forward to bringing the Canadian network along.
- There are about 8 million students in U.S higher education, with about 1 million professors and a budget of about $150 million.
- What is going on can be characterised by both a bottom–up and top–down activity.
- About 20 years ago defense research provided some funding for campus networking people to construct a network that would do something about the vulnerability of the old point–to–point circuits for permitting researchers to communicate by remote hosts. Out of this came ARPANET which was the first packet, connectionless architecture.
- DEC and IBM have both thrown in with TCP/IP but have announced to their international clients and U.S government that they fully intend to implement OSI by this fall. Already have transport layer of OSI running alongside the same machines with TCP/IP — most of the experts on this field do not see this as a problem. The real problem is to create a transport mechanism which has the technology available to push bits along fibre very quickly and that can packetise and depacketise these bits. The challenge then is to provide

advanced services over the network to support end users.

- It is estimated that there are 150,000 host machines and about 1 million active users on the U.S. Internet. The U.S Internet is widely connected abroad and just in the last year, T1 connections have been put in place to Japan and Australia. There are half a dozen links to European researchers.

- Four or five years ago the U.S congress became convinced that the U.S was falling behind other industrialized countries in terms of quantity and quality of research and development. 75% to 80% of all basic research in the U.S is conducted on university campuses, primarily funded by the Federal government. Study after study showed that research institutions were becoming obsolete and the tools and facilities available to campuses were not state of there art. As a result, one of the major fallouts of this is that the budget for NSF has doubled over the last five years.

- With regards to high performance computer networking there is a federal initiative called the high performance computer program. The U.S Office of Science and Technology Policy has written several paper son this over the last two years and out of this work the NREN concept has emerged.

- Promotion of the NREN began in 1987 and the federal agencies decided to get behind this single initiative. Government and universities are behind the NREN. Significant private sector support, has also picked up in the last year.

- One of the objectives of the NREN is to accelerate the development of products and services that are technology driven and that will dominate markets over the next 10 –25 years. The present status of the NREN is that there is legislation pending in the U.S congress, that will create a legislative mandate for the NREN.

- Approximately $100 M per year is earmarked for the next five years, divided evenly between solving the existing production problems on NSFNET (will evolve into NREN) and the second half would go towards the research and development of the gigabit network. The evolution of the NREN is divided into three stages. The first two stages will create T1 connectivity to 400 – 500 private and educational sites and upgrade that capability to stay ahead of the demand for T3 services. Simultaneously research into product development that would gravitate that out to gigabit levels would take place.

- As the above is going on the telecom industry is really scrambling to get back on top of their own technology problems. Many companies have coalesced along the broadband ISDN concept High Speed Transport Mechanism. The gigabit test bed projects in the U.S. are going to use SONET for the transport layer so that the 10–15 year picture of the national telecommunications fabric and activities of the carriers will evolve to an all digital network at gigabit speeds. The NREN will add its own layers of applications.

- The U.S. is beseiged by Super Science projects like the Hubble space telescope. The data rate on this telescope alone is a continuous megabit per second. Two to three hundred researchers require immediate access to this information. This as well as other projects like the world wide seismic network cannot go forward without the existence of a high speed computer network.

- There is enormous agreement in the U.S that the NREN is something that should be done and should begin as soon as possible.

IV ISSUES

Dr Egelstaff posed a question as to why this network should be used solely by researchers? If banks or other commercial institutions pay full cost for the service on this network why should they be excluded. Presumably this rate would be higher than the rate charged to a university researcher. It was mentioned by David Arthurs that the eligibility criteria for this network had not yet been defined — and at this point broad eligibility had not been ruled out. However it was noted that government subsidies directed to this network would not mix well with commercial use of the network.

With regard to the introduction of OSI in the U.S. Michael Roberts mentioned that it is really a software issue and not a hardware issue. At the transport layer software is available. Transition to the OSI transport layer protocols will begin over the next 24 months. The European situation is confusing in this regard. Enormously cost effective TCP/IP hardware and software has become available. Just within the last year many TCP/IP networks have sprung up all over Western Europe. However there is a general consensus in the U.S. that OSI is the direction to go.

Peter MacKinnon noted that the current state of OSI is very important consideration in the implementation of this network. David Arthurs mentioned that the study team is looking at a concept of OSI coexisting fully with TCP/IP.

V PRIORITIES

Robert Vaive mentioned that cost/benefit and services are two priorities that are coupled together. There is also a need to make this network reliable. Availability in the operational sense is extremely important and connectivity is required from coast to coast and throughout the U.S.. Very concerned with security as well.

Michael Brandreth indicated that the technology required is not yet there eg. still need to feed single pages into fax machine. Administraive and bureaucratic barriers stand in the way of new technology taking over. There is an increasing interest in revenue generation — therefore have to get out to more than traditional clients. The existence of a high speed network would give this extra push that does not seem to be there with existing technology. Although there is no doubt that libraries could have done far more with existing technology that what they have done up until now. Connectivity and standards are very important issues. Libraries in general are going to be interested in different types of services and different types of clients, and revenue generation. Speed is also an issue. Interested in image transmission to construction sites. Colour photographs to operating rooms.

Peter MacKinnon indicated that cost/benefit is overriding and ease of use must be considered — for end users it must be transparent. Security is also very important. Distributed computing is a movement towards real time manipulation of large volumes of data types be it graphics or motion.

Steve Cockle emphasized that cost/benefit is very important. Access to and storage of immense amounts of material is extremely important.

Peter Egelstaff commented that if universities are to use the network costs must be low. However if costs are not low it does not mean that universities will not use it.

If speed is increased then this will change the way researchers work. Connectivity to super computers is also important. Ease of use is extremely important if you want people to use it.

Chris Hughes mentioned that from the Supernet perspective, where there are industry and university partners, high bandwidth at low cost is priority number one. Regulation is holding us back in terms of the artificial costs placed on bandwidth in Canada. Security of the data is equivalent to integrity of the data. Data cannot be corrupted in anyway as it moves down the pipe. Connectivity is important in a more general sense. According to Mr. Hughes research in Canada is doing more and more of the same thing. He would like to see Canada levering up and leapfrogging over what the U.S is currently doing, in terms of applications. The researcher should really be able to get on the network and ask a question and the network would then direct him to the location where the answer resides. Intelligence built into the network is crucial. Don't need just another bigger pipe across the country, also need the applications and services.

Ike Goodfellow suggested that there is a requirement for large data files and CAD/CAM applications but also with low volume electronic mail. As a result band width on demand is very important. Northern Telecom is very concerned about security and refers to integrity as authenticity. Security is valuable. Software packages are expensive and want protection to be guaranteed. Ease of use is very important and not enough attention is paid. Must focus very heavily on standard protocols and transparency. Mobility is the other characteristic that this network must have. The concept of a fixed office and fixed desk is archaic.

Harvey Nelson mentioned that security is extremely important as handling confidential patient information and have signed papers saying that this will not be released. Clinical trials can make or break a company, therefore financial agencies etc. try to get this data, before it is officially released. Connectivity to all end users is critical. The network must also be very easy to use especially in the field. Continuing education very important for doctors as must keep abreast of modern technology ie. video education.

Roger Taylor commented that all of the aspects indicated are very important. In addition the environmental benefit due to electronic storage of data seems to be quite significant. Ease of use is important and so is cost. T1 from Toronto to Montreal is same as New York to LA — which is untenable. Security is important, secrecy of confidential documents. If data is corrupted this costs money. One spends substantial dollars on collecting data therefore protecting it is very important.

Peter MacKinnon emphasized that the institutional point of view expressed in terms of costs and benefits is important but this study should highlight national costs and national benefits. These cannot be overlooked.

VI APPLICATIONS AND BENEFITS

With better access to supercomputers it will be possible to push frontiers of science and push the frontiers of engineering. Realtime video and HDTV are applications that require high bandwidth.

IKE Goodfellow pointed out that high definition TV coupled with much better colour enhancement will lead to a whole host of things we see today but don't do because the quality is not quite right yet. For example still X–ray may become academic. Full motion video X–ray will evolve, but requires high bandwidth .

Michael Roberts noted that in the States, especially in the Physics community it is essential that the NREN offer end to end visualization services. Many validated scientific situations exist where the ability of the researchers to deduce what is really going on requires a three dimension real–time, colour picture.

Roger Taylor suggested that computational Fluid Dynamics has great potential for real–time 3–D colour graphics. With regards to the quantification of the benefits of this, time–to–market will be significantly reduced. Lead time to market depends on the product.

Robert Vaive suggested that teleconferencing and video conferencing would have huge benefits in terms of travel cost savings. NASA have their own satellite transponders, and 13 video conference rooms set up at about $ 1 million per room. According to NASA officials the payback on these rooms is less than one year. This also yields to more flexible interaction, calling people in at a moments notice. Would like to participate, but do not have the means. Travel costs are the biggest percentage of their operating budgets.

Steve Cockle indicated that through the Centres of Excellence Program — there is a need for interaction through videoconferencing. Too much time and money to bring people together.

Roger Taylor mentioned that new applications will emerge that are inconceivable at the moment. For example the introduction of a holographic screen would have a tremendous impact on the applications one could do.

Peter Mackinnon pointed out that the approach presented is a bottom–up approach. Top–down national goals for the country not the network should be focused upon. What is the role of science and technology in society? Areas in the environment, health care, transportation etc. are important, and part of the social fabric — manifested in technology. If this network is to be justified as a federally funded project these national goals must be looked at more closely. Unfortunately even with the Centres of Excellence projects the infrastructure is not there to carry out their mission. In fact this study was motivated by the fact that Centres of Excellence do not have this infrastructure, to effectively carry out there mandate.

Ike Goodfellow explained that there is no fundamental difference between T1 and T3 except the size of the pipe. Rollout of T3 is expected in 94/95. With respect to this timeframe fibre is not the problem but switching it poses a challenge. Bandwidth on demand is a year 2000 objective in a virtual public network. In a pure private network it could be done next year.

Chris Hughes explained that if you have a big pipe connected to a super computer you will use that facility in a different way. What we want to get into is visualization— interactive graphics on the Cray. You can't do that on a 56 KB line.

Robert Vaive pointed out that Library cost savings would be significant in terms of physically not having to go to the library to search for the information you want.

Dr. Egelstaff emphasized the need for gradualism. His suggestion was to immediately update the Canet to T1 and then to T3 when the need arises. He pointed out that the arguments made in the States for the need for such a network are not really numerical. The arguments made are qualitative — supported by national objectives of defense, security and international competitiveness.

Michael Roberts added that over the upcoming years more money will be spent on developing applications. Applications take a while to develop and then continue to exist for

a while. The next most expensive investment will be in workstations because performance factors are moving up so quickly. The network will be third on the list. In the U.S. it is expected that network improvements in the order of magnitude of three fold will occur between the 86 Arpanet and the end of the century.

Ike Goodfellow suggested a background report (Laboratory of Yorkton Heights) that suggested if anybody doing interactive computer work needed to have sub–second response time anything in excess of that the research productivity went down by 50%.

Dr. Egelstaff requested the dollar cost to implement CAnet at T1 bandwidth be included in these minutes. The costs are heavily dependent on the adopted topology. According to the U of T proposal a fully redundant T1 network with low hop count would cost $8.8 M annually. A completely non–redundant network, with low hop count, would cost $5M per year. A simple "beads on a string topology" would cost even less. If the proposed tariff reduction on megaroute service is implemented, these costs would be cut by 50%.

Robert Vaive commented that the NASA network group has a budget of $180 million per year and have just over 100 dedicated T1 circuits. They do not have any T3. With regards to video conferencing they logged 800 video conferences of which they would have liked his project to participate in 100 of them. In total 15,000 people attended these conferences and 5,000 trips were saved. Mr Vaive concluded that his agency could use T1 right now and T3 by 1995.

VII FEE STRUCTURE

Roger Taylor emphasised that the fixed fee scheme is the best way to go. There is enormous overhead with regards to administration of a usage based system, and in addition research would be inhibited.

Robert Vaive mentioned that people are comfortable with local vs. long distance charges. Would suggest a similar scheme based on usage (ie. $X for so many packets) otherwise abuse may occur.

Michael Roberts commented that in the U.S flat fee seems to work well. Some economists argue that only if there is limited capacity should a metered fee be implemented. However Michael Roberts argued that pricing and cost relationships in perfect market conditions only exist when you have a knowledgeable consumer and a well defined product and those things do not exist in the commodity we call research networks. This is not a classical market.

SUMMARY MINUTES OF EXPERT PANEL SESSION 3
MONTREAL, QUEBEC
JANUARY 29, 1990
AUBERGE DU GOVERNOUR, MEETING ROOM GOVERNOUR 2.
1:00 P.M - 4:30 P.M

I ATTENDEES

Panelists:

Peter Jones (Computing and Information Services, Dalhousie University)
Peter Garland (Spar Aerospace)
Christine Neilson (Department of Oceanography, Dalhousie University)
Bernie Turcotte (A.V.P Systems, CRIM)

Observers:

Joe Padden (ISTC)

Project Team

David Arthurs (HICKLING)
Phil Kennis (HICKLING)
Tony Capel (Comgate)
Roger Choquette (Comgate)

II AGENDA

1. Introduction
2. Priorities and issues
3. Applications
4. Anticipated Benefits
5. Implementation
5. Other

III PRIORITIES AND ISSUES

Peter Garland stressed that the network must have high standards in place for availability. Low response time is also very important. Flexibility to set up video conferencing from computer is important consideration. On demand requirements for bandwidth is a also a valuable feature. Security is a consideration and would like to be able to select various levels of security. Flexibility is the most important consideration.

Peter Jones indicated that the network has to be reliable. There comes a point however where diminishing returns exist with regards to reliability. For instance do not need the level of reliability required by the banking industry. The cost also has to be very low for the end user. Connectivity is also very important. There has to be one network — don't want duplication. Therefore it is important to see this as a follow–on or an extension to CAnet. Netnorth is moving into CAnet therefore a convergence into one high–speed national network seems attainable. The worst thing would be to have two separate networks. Security is a difficult subject and the more diverse the clientele the more this becomes an issue to deal with. User friendly transparency is ideal.

Bernie Turcotte commented that there is pressure from industry to make networks secure. However this is difficult to guarantee when working in an IP network with thousands of people. David Arthurs asked whether the private sector might implement this type of network if government doesn't. Mt Turcotte answered that researchers do not want to be at the mercy of commercial world. In the case of CAnet, administrators and researchers got tired of sitting around and decided to do something about this. Mr Garland also commented that such a venture is not particularly attractive from a private sector point of view. In his opinion entrepreneurs would not see this as an initiative to carry out.

IV APPLICATIONS

Mr Jones noted that although researchers in eastern Canada are not near super computer facilities, many require access to super computers. In many cases researchers have to go to super computer centres physically to carry out their research. The supercomputers used are in Dorval, Toronto or Calgary. Access to these facilities is pathetically slow. There is a requirement for 1 or 2 megabits per frame. At present, under CAnet a 56 Kb network, it will be possible to send one frame every 35 seconds. Although this is a beginning it is grossly inadequate.

Mr Jones added that there is a real need today among many of these research facilities for higher bandwidth. This need spans many disciplines — Oceanography, Climatology, Remote Sensing, Seismic Data Processing, Physics, Chemistry and Medicine. Researchers in these areas pass large amounts of data between each other. For many of them they require T1 in own short bursts.

Mr Jones also commented that in terms of data transfer there is tremendous potential in terms of productivity improvements at libraries. Most libraries have on–line catalogues, therefore it takes minutes to find where a resources resides but weeks to actually get it. Thus there is a real need to transmit data from library to library. Libraries are also under tremendous financial pressure, to avoid duplications with other libraries. Therefore this network provides an opportunity to reduce duplicative service.
Mr Jones summarized the new and innovative network concept which has recently emerged within Nova Scotia. With financial assistance from the Atlantic Canada Opportunity Agency and the provincial ministry of Science and Technology the Nova Scotia Technology Network is just getting off the ground. This network emphasises involvement from the private sector. Software Kinetics partnered with Dalhousie University have been chosen to implement and

manage this network — and have been set up as a for profit organization. The mandate of this network is research, education and technology transfer and should be operational in a couple of months. Systemhouse prepared a study four years ago which led to a recommendation to establish this network. The network will start out at 56 Kb (can't afford any more) and TCP/IP and will migrate to OSI when possible. The technology is the same as CAnet. The fee structure for this network is currently under discussion.
Peter Garland from SPAR indicated that within SPAR many of the older engineers required reeducation, and this network could provide this role. In addition SPAR would also be interested in the databases services that would be offered on the network.
SPAR is in the process of installing their own network using their own technology. This network would be used as a test–bed.

Christine Neilson described her centre of excellence as one with extremely strong industry, government and academic links. The centre uses Netnorth as much as possible. The actual data requirements for the centre in terms of using this network are limited. Data files are in the range of 1Mb – 10 Mb. Large amounts of file transfer are not required. The databases do not yet exist — and one of the goals of this network is to establish these high specialized databases. Although the data requirements at present are not that significant, she indicated that the members within the centre would find "new" and useful ways to do use this network.

OPEN is part of a globalized research effort in this field. OPEN is participating in two international projects which are very data intensive ie. the joint global area flux study and the world ocean circulation experiment. Data committees must be able to communicate with each other in different countries. Christine Neilson mentioned that at present she is not connected to the Department of Fisheries and Oceans (DFO), and thus communicates by telephone and fax.

Christine Neilson added that although remote sensing data is at the periphery of interest at the moment in several years it will become an important adjunct to her work. The data requirements in this research field is enormous and requires high band width communications networks.

V COSTS AND BENEFITS

Mr. Jones indicated that once such a network is in place there would be loud objection from the user base if it were ever taken away. However under the current circumstances it is extremely difficult to assess value. There is a latent demand out there for networking. Sharing band width is also an important issue to make economic for organizations using T1 ie. can amortize across a larger number of users.

Bernie Turcotte indicated that it is difficult to attract high calibre research scientists when the connection to the U.S Internet is only 9.6 Kb. In many respects Canada is a third world country in networking. He also mentioned that the highway comparison was a good one. An analogy was given that you can't build half a highway across Canada now and then decide to build the other half later. An iterative process just does not work.
Mr. Turcotte estimates that it takes about a year for a user to become fully educated and dependent on his network connection. Once they get it they figure out how useful it is.

Mr Turcotte also stressed that in Canada bandwidth is not affordable — and thus cost is a real constraint.

Mr Jones pointed out that this network would be the most cost effective way for government to spend money.

Ms. Neilson mentioned that within the centre of excellence program networking costs are eligible for funding. In the proposal made by her Centre a sizeable percentage of funding was allocated to networking costs. She also stated that the network would result in less duplication of work — and would facilitate project management.

Mr. Garland commented that the education benefits of such a network are very attractive. Instead of sending senior people to seminars across Canada and the States, seminars could be offered over the network. This would result in savings in both travel and time. Within SPAR it is estimated that engineers spend three days a year on this type of activity.

Bernie Turcotte suggested that with the network in place it would be possible to get hold of people more quickly — therefore getting answers more quickly. There is also tremendous potential to interactively correct "bugs" in software. To do this over the phone is a waste of time and is extremely costly and unproductive. Researchers are expensive, therefore cannot sit around idle.

Peter Jones stressed that quality and quantity of research would be improved. the bottleneck in terms of research productivity would rest with the individual and not with the tools and equipment that are used.

Mr Turcotte summarized by saying that it is difficult to quantify the benefits of such a network. It becomes necessary to look at the historical evidence — ie. NSFNET where usage is increasing by 20% every month.

VI POLICY, IMPLEMENTATION AND MANAGEMENT

Christine Neilson stressed that "universality" be kept in mind should this network be implemented. This network must reach out to non–traditional groups. Users do not think of fees they just logon. However, the type of usage may have to be taken into account. For example if a user only logons to the network once a week he should pay less than a user who uses it everyday. It is her opinion that the industrial partners involved with her centre would prefer a usage based fee.

Mr Garland also commented that it is important to offer equivalent levels of service to all users.

Peter Jones expressed his opinion that fees should not be metered by usage because of the high overhead and the fact that it inhibits use — which is counter productive. In the case of the Nova Scotia network (although not yet implemented) not all user will be equal — different types of access will be available as bandwidth and level of service will be variable depending on the needs of users. However fixed fees for fixed services is the general notion of this network. However this example is for a regional network. In the case of CAnet where regional networks are considered clients — not all regional networks will pay the same fee. 50% of the fee will be divided evenly and 50% will be based on population as a measure of likely use.

Peter Garland indicated that usage would not be a good parameter for a costing algorithm. He suggested that a hybrid approach might work best — and that another measurable parameter (other than usage) be used to meter use.

Peter Jones agreed with the need for "universality", and the need to make commitment to broad goals. He mentioned that promises down the line are not good enough. It was pointed out that in the U.S. NSF funds NSFNET and not individual researchers. There is a high

degree of subsidization in the States. Researchers in Canada must first catch–up before being able to promote themselves.

Bernie Turcotte suggested that it is important for Canada to have their own connections to Europe and Asia — and not have to depend on the U.S..

Peter Jones commented that conditions be imposed on regional networks — and that a regional network be defined in a meaningful way. Multiple networks in one province is discouraged, although this has not been specifically addressed.

CAnet is encouraging community colleges to join. Nova Scotia is interested in attracting high schools as members.

For this network hospitals and libraries have large potential as members.

Mr Turcotte commented that RISQ offers preferential rates to universities. If hospitals connect through universities then are eligible for this lower rate. If however hospitals want their own connection, a higher rate comes into effect. It was pointed out that the eligibility definition for access to this network be in line with the Internet because of wide spread interaction with U.S. researchers. the States.

According to Bernie Turcotte the most important barrier to this network is cost.

MINUTES FOR EXPERT PANEL SESSION 4A
VANCOUVER, BRITISH COLUMBIA
FEBRUARY 18, 1990
HOTEL VANCOUVER
9:00 A.M - 12:00 P.M

I ATTENDEES

David Taylor (INSINC)
Steve Baker (INSINC)
Jeff Berryman (ASI)
Alvin Fowler (UBC)
Jim Varah (UBC)
Saul Greenberg (ARC)
Ian Bardsley (MPR)
Chris Sibertz (Infosat)

OBSERVERS:

Paul Thiel (Ministry of Economic and Regional Development)
Digby Williams (ISTC)

PROJECT TEAM:

David Arthurs
Phil Kennis
Roger Choquette

INTRODUCTORY REMARKS:

Self Introduction of all panelists.

David Arthurs made a presentation concerning the concept of the network.

Roger Choquette made a presentation of the interim technical design of network.

Dr. Fowler asked for a clarification of whether the focus of this network was in fact research–based. David Arthurs reiterated that this network is not to be used for commercial purposes and will not compete with commercial services like Datapac or Envoy 100. The network will be used to support pre competitive research.

Ian Bardsley mentioned the need to explicitly consider video. In the future video will become very much the main medium of communication. The need to deal with video combined with audio is a high requirement and this must be factored into the design. The management of the flow of video is also very important. Image and video are quite different. Image transmission is being increasingly done today, whereas video cannot be done without special facilities. Thus this network would potentially have significant impact on a whole new research community.

Jeff Berryman added that if the projections of video come true any strategy that is laid down now must explicitly take into account video, and cannot be dealt with as an after–thought. The integration value (from a cost/benefit point of view) of video must be addressed.

Ian Bardsley commented that the easiest thing to do when designing a high speed network like the one proposed is to ask the user community what they would like to see in a high speed network. The user community will not be able to predict the benefits of the next generation technologies like video and therefore may not explicitly ask for it. This must be realized by the project team.

Saul Greenberg asked whether users of this network would contend with other users for bandwidth. Roger mentioned that the traditional design is to give a user community a channelized pipe. Another approach, bandwidth on demand provides capacity up to the limit of the pipe. In that sense you would be competing with people who are sharing that pipe with you, but in a more dynamic basis.

Sibertz mentioned that certain details with respect to the topology of the network are missing. International links have to be explicitly taken into account. The regional networks as they exist today do not have the money to connect to international networks.
Saul Greenberg mentioned that one reason the fax works so well is because of connectivity. Providing service to remote areas is also very important. If this is left up to the regional networks this will not happen. There has to be a federal presence to allow for this.

Steve Baker from INSINC mentioned that under the existing tariff structure DS1 lines cost $500,000 per month from Vancouver to Toronto. When looking at all interconnection pipelines, with multi users — the economics are staggering.

Jim Varah pointed out that before we get depressed about costs there are a number of things that we must look at. From his perspective speaking for academic researchers this network is absolutely essential. The U.S. has it, countries in Europe have it and the Japanese have it. Canada has to have it — and the money for it has to become available or tariff rates have to be modified.

Jim Varah added that Canadian researchers need this network for many reasons but video and image transmission are at the top of the list. The main stumbling block is lack of bandwidth. The Centres of Excellence have also put pressure on this network to come about. The Centres of Excellence are distributed across Canada and deal with leading edge technologies that require close collaboration.

Economic reality from INSINC's point of view is that the economics of providing new services under the current structure prevents participation from industry. It is very difficult to get into a network to do anything because the cost of access is staggering. Having a network that would allow industry to develop services and products in a cost effective manner is a very sound concept.

Sibertz mentioned that the entire U.S. is linked. In the U.S. they have bandwidth on demand.

Digby Williams pointed out that in the U.S there are many networks. The concept of this Canadian approach is to let one department begin to drive one network that can look at various aspects of the network. Over time intelligence and usage will develop on the network which hopefully will diffuse into the commercial world in the end. However it is impossible to build the car before building the highway.

Jeff Berryman commented that he views this ISTC initiative as a technology push rather than on a market pull or demand pull basis.
Jim Varah mentioned that from an academic point of view this is not just an operational network. Researchers at universities can contribute significantly to research on the network.

At Canadian universities there are a significant number of top class researchers in telecommunications.

Ian Bardsley inquired whether this network would be high speed narrow band or whether the concept is to build a broadband network. There is a very real basic difference between these two concepts. When you are talking about T1 speeds providing connectivity to a number of users this is a high speed narrow band network. When you start talking about the future and if we want to make a step up from where we are now we must really start talking about a broadband network.

Roger noted that the proposed architecture is in line with a migration to a broadband network. This would be available at the end of the decade. Ian Bardsley emphasized that broadband would be available in the next two or three years — however not ubiquitously. We must be realistic in providing a reasonable migration pattern. Do we want to wait until the end of the decade when this is ubiquitously there, or do we want to embark on an earlier migration pattern?

Ian Bardsley mentioned that with regards to economic feasibility the tariff structure is horrendous and imposes artificially high costs on users. The carriers do have an incentive to participate. If this venture can be done as a research and experimental vehicle the whole regulatory problem could be made to disappear. It is very important to get the carriers involved and to have them participate in the network. We must assume that we can defeat the existing tariff rules.

Jeff Berryman agreed with the point made by Ian Bardsley that we must go around the CRTC and not through it.

David Baker mentioned that this is not just a CRTC issue. We are talking about introducing pre–competitive aspects into Manitoba and Saskatchewan.

Saul Greenberg pointed out that if this network is for developing the future clearly 1.5 Mbs is not enough. The title of this network therefore is not descriptive enough.

Digby Williams agreed that T1 is not enough but never the less it is a big step from where we are today.

Jeff Berryman mentioned that people who empower Digby Williams are people around the table — we should then try to give him as much power as possible. Putting a network in place that could be replaced by a commercial network as soon as it is approved does not make very much sense. There is no point in selling the politicians a leading edged network and by the time it grinds through the bureaucracy it is no longer leading edge. We should transfer the power to Digger as much as possible. It better be leading edge and it better be looking at T3 and broadband and video, channelized vs. packet etc.. If this is not true, the descripter "leading–edge" better be taken out of all of these information packets. Politicians are willing and capable of understanding this concept as long as it is supported by researchers in the provinces.

Saul Greenberg agreed with empowering Digger for more bandwidth. Saul indicated that his network connections at present are an order of magnitude lower than his needs. If one looks into the future at multi–media, sound, real–time video over the network, very clearly the local area Ethernet also will not suffice. The next step with regards to multimedia is coming very soon. As a user he wants to extend his role.

Jeff Berryman went on to say that even at T1 this network would bring in a significant number of applications which cannot work in today's network environment. However we do

not want this network to stop there.

Dr Fowler asked whether the whole impact of ISDN on the telecom industry is being looked at. Roger Choquette mentioned the kind of ISDN technology that will become available over the next several years is not suitable for the types of applications being talked about in today's session — because it is not fast enough. However we would want it at as a low speed access mechanism. Roger reiterated that we would hope that Broadband ISDN — Sonet technology — would be available by the end of the century. Ultimately it would be ideal if this network could be provided using public services.

Steve Baker suggested that the greatest avenue for pre–competitive product development is in ISDN compatibility and ISDN products. In Canada the strength of our research resides with small companies that jump out of nowhere. For these companies there is a strong need to have a facility where they can test products in a pre–competitive environment. Otherwise, these products will not evolve.

Ian Bardsley suggested that the network not focus on narrow band, because by the time the network the network is implemented this will be passe. The focus has to be on the standards for Broadband IDSN because these are now emerging.

Jeff Berryman asked about the planning horizon for this study. How far are we looking into the future?

Roger Choquette answered that we are looking at this network in terms of evolution until the year 2000. In the long term where is Broadband ISDN going? Where are companies like MPR and BNR going? When will some of these products be available commercially?

Ian mentioned that to view this from the commercial perspective is incorrect. At the outset the products that go into this network will not be available commercially. Ian added that MPR has an ATM switch which they will be turning over in the next month.

If this network were to exist today, MPR could potentially use ceratin aspects of this network to help business development of its switch.

Jeff Berryman stressed that whether delivering raw bandwidth or delivering applications services, one service which should be provided on this network is a directory service. It should be X.500 based because that is the only game in town. One can look to the U.S. at the mess that they are in to see the importance of directory services.

Jeff Berryman asked whether this network might contribute to the development of a product that could be sold to the U.S. to help them escape the box they are in with respect to OSI. Or has the window of opportunity already passed by?

David Arthurs responded to this question by mentioning that there is a good chance that this could happen for two reasons: 1) In the U.S. networks are funded without having to justify the expenditure on an economic basis; and 2) industry is not encouraged to use the network to develop products. Therefore there is a good chance that Canada could well supply the U.S. with future network technology.

Jeff Berryman suggested that there is pent up desire among small companies or research organizations to get on the inside of the academic community network. This network should take advantage of this. Industrial access is important. Industrial users need more in terms of system integration. This is the only barrier.

II APPLICATIONS

Saul Greenberg mentioned that as a researcher in telecommunications great tools are available that run wonderfully over local area networks but tools are not available that work well over wide area networks. It is difficult to test these types of products. Therefore this test–bed feature of the ISTC network is very important.

Saul Greenberg gave an example of a multimedia mail bulletin board. People did not initially take advantage of this but now it is used extensively. At first there is an awareness problem.

Ian Bardsley would be interested in using this network from the telecommunication side as well as from the applications side. He would also want to use this network for research purposes and to provide hardware to a network like this. This would be an asset to sell to other networks. — design with real requirements.

Ian Bardsley also questioned whether the applications that users need today exist? The answer is maybe? This can be equated to a chicken and egg problem — must find people who have real communication problems today — and these people exist (ie. those who mail drawings and travel extensively to hold meetings). Ian mentioned that it is difficult to find users today that could stress this kind of high–speed network.

Saul Greenberg also mentioned that at present it is possible to develop prototypes, but there is no where to go beyond that. This network would let you go beyond that.

Alvin Fowler mentioned that part of his job is technology transfer to industry — ie. X.400 around the world and X.500 coming on stream — and the project team should be aware a protocol test centre in Alberta. A tremendous advantage could be obtained if protocols could be tested using this national network. Canada Post is doing value added onto X.400 — ie. produce electronic image then send image electronically to local post office. Once you have critical mass in the network it is possible to build in all of these types of things.

Jeff Berryman indicated that he had a problem with the outline pertaining to applications. Services are represented as software and products are represented as hardware this is not an accurate depiction. Products and services are orthogonal. Why are network management tools separated out? The classification shown really look at the network from a utilization point of view. Where does protocol development come under this classification? This slide must be considered in 3 dimensional matrix shown earlier ie. developing the things that network suppliers supply; network developers develop; and network users use.

Jim Varah added that some services are to be provided by the network and other services are to be provided by somebody else and run on the network. This distinction must be made.

Roger Choquette asked whether MPR's relationship with the telephone companies is the same as that of BNR. Ian Bardsley answered that MPR is charged with developing a vision of the future. MPR has a technology program and is involved in standards. It was mentioned that it is easier to be specific ie. supplying fibre in many of the buildings in the down town core and building a broadband system that will eventually connect the lower mainland and Victoria etc..

III BENEFITS

Jeff Berryman would like to see international competitiveness and national credibility as being specifically addressed by this report. It is very likely that this network would create markets. We should not focus so much on "stop watch type application benefits" because it will take a fairly specific proposal to come up with a quantification of these types of benefits. What we should focus on is what will happen if we don't proceed with this network.

Jeff Berryman suggested that enlarged collaborative research will lead to the formation of strategic alliances. This will not only be apparent at the research level but alliances will form all the way up the innovation chain.

Ian Bardsley commented that this network would help market MPR products internationally. Increased opportunity is very evident and international competitiveness is very important. There are benefits associated with having a more realistic test–bed. If this network doesn't exist opportunities will be lost.

It was mentioned by Chris Sibertz that this network would be of great value to smaller companies who do not have the dollars to create their own network. This comes under the theme of increased opportunity.

Alvin Fowler directed a question to Ian Bardsley as to how far laboratory testing can go in the development of new products. Ian indicated that the laboratory experiments are quite adequate for the development of hardware. However on the software side a live user community would help tremendously because the developer/tester would not have to speculate about the requirement and its use. Therefore the benefit would be realized much more from the software side then on the hardware side.

Saul Greenberg mentioned that if he is developing a piece of software he would achieve superior product improvement because the individuals testing it would be users — not the designer.

Dr. Fowler indicated that this network would accelerate the development of national databases under specific topics. However the first step is to collect the information and make that database available. This is really not a communications problem but rather a database problem. However the network would be the vehicle that allows the development to occur.

Steve Baker reiterated that this network would be of great benefit ie. to have access to the network and only pay for it when you use it. In a shared environment everybody shares the burden of cost.

IV MANAGEMENT ISSUES

Jeff Berryman mentioned that user fees will be paid by the regional networks and these types of questions should really be addressed to BCnet, who is not represented at the table.

Jeff Berryman indicated that as of yet he has not seen a specific proposal to address the question of willingness to pay. There is no proposal on the table.

Jeff Berryman indicated that he is apprehensive about volume dependent charges because he does not trust the existing costing algorithms. Distance charges should not be included in the fee. Users will not like the fact that if they send data further they will have to pay more. However users can live with the fact that if they send it faster or if they use more they should pay more.

Steve Baker emphasized that there is a dichotomy in using the network to develop products and also to use it in a production mode. Where does one draw the line? Standards are required. If separated too far a developer cannot do the kind of stress test that he may want to do. This dichotomy would have to be managed very carefully.

Alvin Fowler responded to this by saying that with respect to CDNnet researchers generally are willing guinea pigs.

V CONCLUSION

There was general consensus from everybody at the session that the concept of this network was acceptable in principle. An agreement in terms of user fees is a bit of a red herring at this point because a specific proposal is not on the table .

Chris Sibertz indicated that Infosat would be willing to invest in satellite hardware like communication stations in exchange for having access to the network.

Everybody in attendance was encouraged to write a letter to David Arthurs at HICKLING expressing their viewpoint on the concept of this network. These letters of support could then be used to build momentum for the network and to help sell the network politically.

MINUTES FOR EXPERT PANEL SESSION 4B
VANCOUVER, BRITISH COLUMBIA
FEBRUARY 18, 1990
HOTEL VANCOUVER
1:00 P.M - 5:00 P.M

I ATTENDEES:

John Boyce (UBC)
Mike Patterson (BCnet Manager)
Grenfeld Patey (UBC)
Andy Woodsworth (NRC)
Stephen Sutphen (University of Calgary)
Saul Greenberg (Alberta Research Council)
Rom Kieffer (University of Calgary)
John Demco (CDNnet Manager)
Renee Poutissou (TRIUMF)

OBSERVORS:

Digby Williams (ISTC)
Bill Hutchison (Ernst and Young)

PROJECT TEAM:

David Arthurs (HICKLING)
Phil Kennis (HICKLING)
Roger Choquette (Comgate Engineering — second half of session only)

II INTRODUCTORY REMARKS:

Self Introduction of all panelists.

David Arthurs made presentation of concept of the network.

Mike Patterson openned the discussion by mentioning that connectivity to the world is very important. Direct connections to Japan or Europe is something that should be addressed by this study. David Arthurs responed to this by saying that whether a direct connection to Japan is put into place or whether a connection is obtained through NSFNET has not yet been determined. However the study is explicitly looking at the international connectivity of this network.

Mike Patterson mentioned that NSFNET is currently TCP/IP but is moving towards connectionless ISO. It is his sense that CAnet is doing the same thing. He also questioned whether there were plans for X.25 or connectionless and whether IP would be carried on it, over it etc? Roger Choquette addressed this matter with Mike Patterson later in the session, after he arrived. Mike was supportive of the concept.

Andy Woodsworth mentioned that we have not made a point of including the North West Territories. Is this by design or by error? David Arthurs addressed the decision not to include the Nort West Territories because of the small population and the lack of Information Technologies companies that exist in these region. However the feasibility of remote areas of Canada having connections to this network are being addressed.

Mike Patterson brought up the relationship between CAnet and this network. CAnet is predominately an academic network but the regionals are actively seeking industrial R&D users. The test–bed is different but from the perspective of the regional network, there is no fundamental difference between CAnet and ISTCnet.
From the strict R&D user point of view there is no difference between this network and CAnet. A R&D user is looking for reliability and transparency and this will be provided by the ISTC network.

David Arthurs suggested that for todays discussion, which focuses on user needs, CAnet and ISTCnet should be perceived in the same way.

Rene Poutissou stressed that from a user perspective, network reliability is extremely important. It was added that if ISTCnet was used as a test–bed to develop and test new products and services the reliability of the network would be sharply reduced. It was pointed out by David Arthurs that the project team is very sensitive to the requirement for a reliable national network. The technical design of this network would be such that the operations of the network would not be affected (ie. the transport layer would not be accessible). It was added that the transition to OSI could be carried out without any loss in the reliability of the network. Some types of testing will not be permissable on the network.

Bill Hutchison indicated that one of the big differences between CAnet and this network is bandwidth. Rene Poutissou indicated that bandwidth is important for two quite different reasons: (1) to pass large data file from point A to point B; and (2) to ensure a quick interactive response. This is a bit of a problem in Canada. We will never have the volumes that they have in the States, but we do share the need for a very quick interactive response.

There are benefits of high bandwidth to bring the country together. How many companies can make a T3 switch — perhaps only a Northern Telecom for the moment. From an industrial point of view, the benefits of this network exist because of the value–added intelligence which is installed on the network (or at the edge of the network).

Rene Poutissou stressed that the fee schedule must be fixed for a given level of service (ie. speed of connection) and should not be based on usage.

John Demco added that some services should be provided on the ISTC network (ie. directory service). The management of CDNnet would like to help provide these services and contribute to the move towards OSI.

Rom Keiffer stressed that he wants this network to provide more than just electronic mail ie. wants this to become a true research network. He also stressed that a regional network (and the user community) cannot afford to connect to CAnet and ISTCnet. He would like to see this network provide cooperative processing. For this to happen significant amounts of bandwidth is required. T1 will not be enough to attract industrial partners — we must think in the longterm. Some of the leading future research is in software development and not hardware development. Realizes there are costs involved but this is the reality of todays situation.

Digger Williams mentioned that in reality it took 4 years to put a network in place which is 56 KB. We have to be able to sell this network and it would be very difficult to sell a T3 network at this time. It is easier to relate this T1 OSI network to what industry is currently providing. There are different communities that we have to look after. This is a dynamic situation.

Rom Keiffer also raised serious concerns with the cost of running lines. One of the problems with T3 and gigabit speeds is the cost of the lines. From the perspective of the carrier/supplier it doesn't cost any more to run gigabit than 56 KB cable (although the cost of hardware at each end is more).

Digby Williams added that no companies in this country have a T3 capability. Part of this study is to look at when T3 should be phased in. For the moment no protocols have been defined for gigabit speeds. As market projections for T1 is very large it is easier to sell this concept. Companies want to develop services on a network like this. T1 will suffice for 80% of the users.

III APPLICATIONS

David Arthurs mentioned that with respect to this discussion the project team wants to know what are the applications and demand for T3 and whether we need T3 or if T1 is enough.

Saul Greenberg spent a few minutes illustrating some new generic applications that might be possible if a high speed network were put into place.

- E–mail incorporated into documents. E–mail in machine readable form would be of great benefit.
- Real–time remote conferencing, see body contact — today video conferencing is really not that effective. There are other ways this can be done. For running a meeting tele–data can be used where physical data is physically imported and can point to specific data sets.
- Casual real–time interaction — video hallway. Many contacts are made informally.
- Asynchronous messaging — multi–media mail goes beyond electronic mail. Hundreds of people writing documents.
- Bulletin Boards today are all text–based and quite crude but could be extended to include hyper–text and hyper media.
- Access and operations of Computer Resources touches upon distributed processing — remote collaborative experimentation
- Digital Libraries has to do with information databases. At present there are some central depositories. Researchers have to know where information resides and have to be able to get at it quickly.

Andy Woodsworth summarized the functions, needs and network requirements of the Canadian Astronomy Data Centre (CADC). The CADC is responsible for acquiring and maintaining Canadian archives of data from the Hubble Space Telescope, which will be launched from the Space Shuttle on April 19, 1990. Archival research is a new area in astronomy and started with the space missions which generate huge amounts of data. The data is all calibrated and represents a very homogenous data set. The data archive rate of 2 GBytes per day is anticipated and the CADC has to be ready to accept this data. Another function of CADC is to support archival researchers with access to on–line interactive data catalog, delivery of screen data sets, and assistance with data analysis techniques an software packages. There is a requirement to ship frames to reseachers which are about 2 megabytes

in size. It would be possible to ship 2 megabyte frames quite quickly with a T1 network if this line was dedicated to meet this purposes. If shared with all the different users in Canada a T1 network would not be sufficient for distributing 2 megabyte images.

The needs of the Space Astronomy Community are increasing bccause of the Canadian Space Agency committment to spend $40 million over the next decade on hardware. The CADC needs (1) easy access to data catalogues of data using modern database query techniques; (2) high capacity network connections (CAnet is not adequate to meet their needs); (3) workstations with specialized data analysis and software packages; and (4) technical advice on processing techniques.

The network requirements of the CADC, in order of importance are as follows: (1) connections to peers in the U.S. (90% of the research is done south of the border); (2) must support a communications protocol that is available on the researchers computers and those of peers (TCP/IP or DECnet); (3) good interactive throughput — DSO is acceptable during nonpeak times but even T1 is not adequate if heavily loaded; (4) must be affordable and traffic–based charging is unacceptable (charging by packet is not realistic).

Rene Poutissou speaking as a researcher in physics indicated that remote login to super computers and the transfer of image (includes incorporation of graphics and E–mail) are two very important requirements. There is a need to communicate with American and European researchers. Unsure whether a dedicated 19.2 Kb line (which already have) would be better or worse then a shared T1 line. She added that real–time remote conferencing would save significant amounts in travel budget. Rene finds travel has been reduced tremendously. This kind of network would allow researchers within TRIUMF to seek out researchers around the world which share common research interests.

John Boyce mentioned that he uses E–mail quite frequently. In his research there is a real need for the transmission of colour images. The fax machine's current state of technology is not adequate to send image and yet it is the only device available which they can use. The reason the fax machine is used is because it is ubiquitous. The problem with E–mail is that the transmission of colour images and images in general cannot be done. It would be great to be able to freeze an image and contact a colleague and have him study the image in real–time. This dynamic interaction would be extremely helpful. There is also a need for the exchange of large data files which cannot be done today. He also emphasized the need for standard protocols and standards — that are consistent with European formats and standards.

Graham Patey is a very frequent user of electronic mail and relies heavily upon it. Connectivity is his real priority. Instantaneous response time is not that important, whereas connectivity is very important. There is also a need for the exchange of large data sets. Real–time transparent use of super computers is another useful application area for this network. His perception is that conferencing over the network could reduce the need to travel tremendously.

Rom Keiffer stressed the importance of truely distributed processing. Software is available to do this. Distributed queues are available.

There was general consensus that this higher speed network (at T1 bandwidth) could be used to carry out research activities that currrently cannot be carried out. The value added services provided on this network would be of great value.

The difference between T1 and T3 is a big jump, however two times T1 would not make that big a difference. Data transfer, real time visual, transmission of colour image, video conferencing, data access, conectivity, remote login and distributed processing are things that cannot be done at all or are cannot be done adequately today because of lack of bandwidth.

Saul Greenberg mentioned that if T1 is supplied services will improve. X–windows needs higher bandwidth. This network would allow users to do something that is value added.

For real–time conferencing 16 voice channels are required. This is expensive although there is a tremendous benefit.

Rene Poutissou added that the "value–added services" provided on this network must tie into the services provided internationally.

Rom Keiffer gave the example of the Lithoprobe project as an example of research being done today where the images that have to get transmitted back and forth is not possible with the networks available today. Datapac and leasing lines are not viable options. Chip manufacturers who wish to ship their designs to factories also find their hands tied.

The Centres of Excellence Project puts a real requirement for networking within the centre itself. The capability of two or more users to study the same graph/image at the same time would yield tremendous benefit.

With respect to the networking requirements within certain communities Rene Poutissesu pointed out that the High Energy Physics Community receives $3M of funding each year. 10% of this budget is spent on networking.

IV PRIORITIES

Saul Greenberg specified that availability and cost were very important priorities from the perspective of users. Development of services on the network and mobility will become more important once this network is in place. The tradeoff of cost with connectivity and bandwidth is important and it is difficult to prioritize within that (ie. pay more for greater connectivity and higher bandwidth). These must be taken together and is difficult to separate out.

Saul Greenberg indicated that although security is important a moderate amount of security goes along way.

Steve Sutphen indicated that users are most concerned with services. Security is very important but also creates real problems.

John Demco added costs must be predicatable. CDNnet is based on willingness to pay principles.

Privacy/reliability/reputiation/authenticity are all very important considerations.

V BENEFITS

The benefits of this network are difficult to measure because they are far into the future. Overcoming geographic barriers is a very visible benefit.

Until we discover the new applications many of the benefits are very difficult to see that far ahead.

VI MANAGEMENT AND OPERATION

Mike Patterson indicated that the annual rate for BCnet varies with speed. It is a graded fixed fee that is not metered. The fixed fee creates an incentive to use the network and is attractive from a budgeting point of view. Researchers feel uncomfortable with metered fees — this will be counter productive.

Mike Patterson also added that sometimes based on the type of traffic load that you expect to generate. Most people would feel comfortable with this.

Top level services should not be imbedded in fees. For example access to CAN/OLE, or other specialized services.

Rene Poutissou mentioned that management has to decide what gets put on the network and what user services get imbedded in the user fees.

Digger mentioned that one option to study might be a fee schedule which acts as an incentive to use the network in a non–traditional way.

Saul Greenberg mentioned that the network management should not decide what services to provide on the network (ie. university of Toronto Library system database). Users must have a say in the services which are included in their connection fee. Users want to make their own choices. All users who pay annual fee should have some vote (ie. through a representative) in setting policy and determining the services that are provided on the network. This is the only way users can remain happy.

Graham Patey indicated that if a company like DIALOGUE uses the network to make money, it should be charged to provide service.
NSERC should have some control and representation on the management of this network.

Having the carriers manage the network is one option, however would want to have a vote in how the carriers manage the network. The concept of Telecom Canada listening to clients is somewhat new.

This network has to be managed responsibly. One just has to look at the U.S. to see how a poorly managed network can become quite fragmented.

VII CONCLUSION

Everybody in attendance was encouraged to write a letter to David Arthurs at HICKLING expressing their viewpoint on the concept of this network.

APPENDIX B-1: ON-SITE INTERVIEWS WITH IT RESEARCH ORGANIZATIONS

NOTE: Records of these interviews have been maintained. They are included under separate cover due to the confidential nature of the information provided.

Company/Organization # 1	Alberta Government Telephone
Representatives:	Roger Pederson (Manager Research and Development) Peng Tan (Business Analyst)
Interviewers:	Phil Kennis (HICKLING) Roger Choquette (COMGATE)
Location:	AGT; Edmonton
Date:	90/03/07/9:00 P.M

Company/Organization #2	Alberta Telecommunications Research Centre (ATRC)
Representatives:	Ray Fortune (President)
	Brian Nicol (Government of Alberta; Director Telecommunications; Technology Research and Telecommunications)
	Catherine Gordey (Government of Alberta; Associate Director, Information Technologies, Technology Commercialization)
	David Antoniuk (Government of Alberta; Director of Electronics)
Interviewers:	Phil Kennis (HICKLING) Roger Choquette (COMGATE)
Location:	ATRC (Edmonton)
Date:	90/06/03/10:00 a.m

Company/Organization # 3	B.C Tel
Representatives:	Robert Alexander Vice President Corporate Development
Interviewers:	David Arthurs (HICKLING) Phil Kennis (HICKLING) Roger Choquette (COMGATE)
Location:	Terminal City Club; Vancouver
Date:	90/01/19/3:00 P.M

Company/Organization # 4	Bell Canada
Representatives:	Jim Shram
Interviewers:	David Arthurs (HICKLING) Roger Choquette (COMGATE) Tony Capel (COMGATE)
Location:	Bell Canada (Ottawa)
Date:	90/03/02/2:00 P.M

Company/Organization # 5	Bell Northern Research
Representatives:	David Barr (Director, NTD Operations) Kenneth Smith (Director, Network Services) Peter Carbone Marek Werenik
Interviewers:	David Arthurs (HICKLING) Roger Choquette (HICKLING) Tony Capel (HICKLING) Phil Kennis (HICKLING)
Date:	90/02/12/9:30 A.M
Location:	BNR (Ottawa)

Company/Organization # 6	Canstar Communications
Representative:	Douglas Mitchell (Vice President Business Development)
Interviewers:	Roger Choquette (COMGATE) Phil Kennis (HICKLING)
Date:	90/02/20/9:30 A.M
Location:	Canstar (Toronto)

Company/Organization # 7	The Canadian Bankers Association
Representatives:	Phillip Hogg (CBA) Manuel Silva (CBA) John Truman (Royal Bank)
Interviewers:	Phil Kennis (HICKLING) Roger Choquette (COMGATE)
Date:	90/02/22/2:00 P.M

Location:	CBA (Toronto)
Company/Organization # 8	CNCP
Representatives:	Gary Oliver (Vice President and General Manager Central region); Robert Armour Account Manager; Brett Butler, General Manager, Customer Service, Central Region; Casey Witkowitz, Director Customer Technical Support;
Interviewers:	Phil Kennis (HICKLING) Roger Choquette (COMGATE)
Date:	90/02/23/9:30 A.M
Location:	CNCP (Toronto)
Company/Organization # 9	Cognos Incorporated
Representatives:	Stanley Chan (Director Computing and Communication) Kenneth Aalders (Director Government Relations)
Interviewers:	Phil Kennis (HICKLING) Roger Choquette (COMGATE)
Date:	90/02/12/2:00 P.M
Location:	Cognos (Ottawa)
Company/Organization # 10	Develcon
Representatives:	Clement Gaudet (Vice President of Operations) Don Friesen (Product Manager)
Interviewers:	Phil Kennis (HICKLING) David Arthurs (HICKLING)
Date:	90/03/09/1:00 a.m
Location:	Develcon (Saskatoon)
Company/Organization # 11	Fonorola
Representatives:	George Best (President)
Interviewers:	Phil Kennis (HICKLING) Tony Capel (COMGATE)
Date:	90/02/26/10:00 a.m

Location:	Fonorola (Montreal)
Company/Organization # 12	Gandalf Technologies
Representatives:	Brian Penney Eugene Zwicki Duncan Glendinning
Interviewers:	Tony Capel (COMGATE) Phil Kennis (HICKLING)
Date:	90/02/23/9:30 A.M
Location:	Gandalf (Ottawa)

Company/Organization # 13	IBM
Representatives:	Paul Koch Dunc Howard
Interviewers:	David Arthurs (HICKLING)
Date:	90/01/24/1500
Location:	HICKLING (Ottawa)

Company/Organization # 14	McGill University(Centre of Excellence for Telecommunications)
Representatives:	Dr. Maier Blostein
Interviewers:	Phil Kennis (HICKLING) Roger Choquette (COMGATE)
Date:	90/01/29/10:00
Location:	McGill University

Company/Organization # 15	Microtel Pacific Research
Representative:	Mr. Bruce Hartwick, President
Interviewer:	Roger Choquette (COMGATE)
Date:	8 February, 1990/12:00
Location:	MPR (Vancouver)

Company/Organization # 16	Motorola
Representative:	Ted Strain (President)
Interviewers:	Phil Kennis (HICKLING) Roger Choquette (COMGATE)
Date:	90/02/19/2:00 P.M
Location:	Motorola (Toronto)

Company/Organization # 17	Memotech
Representatives:	Yvon Galarneau (V.P Sales Eastern Region)f Cheryl McKenzie (Account Manager) Jean Martel (Director Research and Development)
Interviewers:	Phil Kennis (HICKLING) Tony Capel (COMGATE)
Date:	90/02/27/9:30 A.M
Location:	Memotech (Montreal)

Company/Organization # 18	Mitel
Representatives:	Kēn Anderson Mike Afheldt Tony Bawcutt Dave Brown Don Dutton Bob Dyer Kent Elliott Mike Foster Howard Gray Bill Harris Hilary Smith Graham Thompson John Thompson Rick Zipes
Interviewers:	Phil Kennis (HICKLING) Tony Capel (COMGATE) Roger Choquette (COMGATE)
Date:	90/02/28/2:00 P.M
Location:	Mitel (Ottawa)

Company/Organization # 19	Mux Lab
Representatives:	Christian Legare (Vice President R&D)
Interviewers:	Phil Kennis (HICKLING) Tony Capel (COMGATE)
Date:	90/02/27/2:00 P.M
Location:	Mux Lab (Montreal)

Company/Organization # 20	Newbridge
Representatives:	Mike Pascoe (Vice President/General Manager Canada and International) Bill Murphy (Product Manager Local Networks Products) Peter Sommerer (Vice President Product Management) John Carosella (Director Information Systems products)
Interviewers:	David Arthurs (HICKLING) Roger Choquette (HICKLING) Phil Kennis (HICKLING)
Date:	90/02/05/9:00 A.M
Location:	Newbridge (Ottawa)

Company/Organization # 21	Northern Telecom
Representatives:	Ike Goodfellow (Assistand Vice–President Business Development) Richard Jung (Marketing Manager, Digital Switching Department) John Vice (Director, Strategic Marketing)
Interviewers:	David Arthurs (HICKLING) Roger Choquette (COMGATE) Tony Capel (COMGATE)
Date:	90/03/05/2:00 P.M
Location:	Northern Telecom (Toronto)

Company/Organization # 22	NSTN Inc.
Contact:	Mike Martineau (NSTN Inc., Marketing Manager) John Sherwood (Director; Communication Services; Dalhousie University)
Interviewers:	Phil Kennis (HICKLING) David Arthurs (HICKLING)
Date:	90/03/02/4:00 P.M
Location:	Dalhousie University (Halifax)

Company Name/Organization # 23	PRECARN
Contact:	Jean Claude Gabriel
Interviewers:	Phil Kennis (HICKLING) Roger Choquette (COMGATE)
Date:	90/02/06/9:00 A.M
Location:	Precarn (Ottawa)

Company Name/Organization # 24	Simware
Representatives:	Lew Shepherdson (Vice President, Technology) Robert Allison (Vice President, Development)
Interviewers:	Phil Kennis (HICKLING) Roger Choquette (COMGATE)
Date:	90/02/15/9:30 A.M
Location:	Simware (Ottawa)

Company/Organization # 25	Software Kinetics
Representatives:	Tim Symchych (Director of Telecommunications and Technology) Tony Moretto (Director Sales and Marketing)
Interviewers:	Phil Kennis (HICKLING) Tony Capel (COMGATE)
Date:	90/02/13/2:00 P.M
Location:	Software Kinetics (Ottawa)

Company/Organization # 26	SR Systems
Representatives:	Eugene Gaines (General Manager)
Interviewers:	Phil Kennis (HICKLING) Tony Capel (COMGATE)
Date:	90/02/27/2:00 P.M
Location:	SR Systems (Montreal)

Company name # 27	SR Telecom
Representatives:	Michael Morris (Vice President, Technical and Industrial Liaison)
Interviewers:	Phil Kennis (HICKLING) Tony Capel (COMGATE)
Date:	90/02/27/2:00 P.M
Location:	SR Telecom (Montreal)

Company/Organization # 28	Telesat Canada
Representative:	Linda Rankin (Vice President; Business Development) George Margita (Account Manager)
Interviewer:	Roger Choquette (COMGATE) Tony Capel (COMGATE)
Date:	90/01/10/13:30
Location:	Telesat; Ottawa

APPENDIX B-2: ON-SITE INTERVIEWS WITH PROVINCIAL GOVERNMENTS

NOTE: Records of these interviews have been maintained, and are included under separate cover, due to the confidential nature of the information provided.

Province # 1:	Alberta
Representatives:	Peter Noden (Director Research & Statistics)
	Rob Taylor (Manager Data Communication Services)
	Kwan Wong (Senior Director, Network management Branch)
	Rom Kieffer (University of Calgary)
	J.W Hoffmeyer (Director, Information Technologies)
	Keith Switzer (University of Alberta, Manager, Network and Data Communication)
	Ian Simpson (Associate Director, Networking and Computing Centre Operations)
	Rod Yaehne (University of Alberta, Network Specialist)
	Walt Neilson (Alberta Reserach Council)
	Steven Sutphen (University of Alberta, Department of Computer Science)
Interviewers:	Phil Kennis (HICKLING) Roger Choquette (COMGATE)
Location:	Edmonton (Alberta)
Date:	90/03/06/2:00 P.M
Province # 2:	BC. Government
Representatives:	Allan L. Zemrau (Senior Advisor Office of Information Technology)
	Chris Wehrfritz (Senior Analyst Policy and Planning Branch, Ministry of Advanced Education, Training and Technology Science and Technology Division)
	David R.C. Miller (Communications Engineer, Ministry of Regional Development, Communications Industry Branch)
	Maurice Auger (Director Systems Policy, Office of Information Technology)
	Frank A. Vitek (Technology Advisor, Ministry of Regional Development, Communications Industry Branch)

	Paul Thiel (Senior Technology Advisor, Economic Diversification Branch, Ministry of Regional and Economic Development British Columbia Enterprise Centre)
Interviewers:	Phil Kennis (HICKLING) Roger Choquette (COMGATE) David Arthurs (HICKLING) Bill Hutchison (Ernst & Young)
Location:	Victoria
Date:	90/01/19/9:00 P.M
Province # 3:	Manitoba
Representatives:	James Reichart (Executive Director; Industry Trade and Tourism) Edgardo Gonzalez (Executive Director, Industry Trade and Tourism) Gerry Berezuk (Assistant Deputy Minister, Supply and Services) Andy Komher (Director; Government Services Telecommunications) Doug Birdwise (Policy Analayst; Telecommunications Policy Office) Dan Iwankow (Business Planner; Manitoba Telephone System) Markku Murtonen (Network & Technology Planning Manager; Manitoba Telephone System) Fred Cross (Customer Systems Engineering manager; Manitoba Telephone System).
Interviewers:	Phil Kennis (HICKLING) David Arthurs (HICKLING) Tess McLean (Ernst & Young)
Location:	Winnipeg
Date:	90/03/08/9:00

Province # 4	New Brunswick
Representative:	George Bouchard
Interview:	David Arthurs (HICKLING)
Location:	Conference Call (Ottawa)
Date:	90/03/12/9:00

Province # 5	Newfoundland
Representative:	Doug Moody (Department of Development)
Interview:	David Arthurs (HICKLING)
Location:	Conference Call (Ottawa)
Date:	90/03/13/9:00

Province # 6:	Nova Scotia
Representatives:	David Coleville (Senior Director, Communications Policy)
Interviewers:	Phil Kennis (HICKLING) David Arthurs (HICKLING)
Location:	Halifax
Date:	90/03/02/2:00

Province # 7:	Ontario
Representatives:	David Barr (Executive Director, Communications Division)
	Joan McCalla (Director, Telecommunications Branch)
	Merv Stevens (Manager, Telecommunication Industry Development)
	Phillip Dewan (Senior Manager, Technology Policy Branch)
	Dr. Maria Cioni (Director, Research Support and International Activities Branch, University of Colleges and Universities)

	Christopher Pringle (Sr. Management Consultant, Premiers Council Technology Fund)
	Michael Jenkins (Executive Director, Information Technology Research Centre)
	Peter Leach (President, Telecommunications Research Institute of Ontario)
	Greg Howell (Project Director Government Telecommunications Strategy)
Interviewers:	Roger Choquette (COMGATE) David Arthurs (HICKLING) Tess McLean (Ernst & Young)
Location:	Toronto
Date:	90/02/23/2:00
Province # 8:	PEI
Representatives:	Hugh Plant (Senior Director, Communications Policy)
Interviewers:	Phil Kennis (HICKLING) David Arthurs (HICKLING) Roger Choquette (COMGATE) Tess McLean (Ernst & Young)
Location:	Halifax (Conference Call)
Date:	90/03/02/2:00
Province # 9:	Quebec
Representatives:	George Corriveau (Industry, Trade and Technology)
Interviewers:	David Arthurs (HICKLING) Roger Choquette (COMGATE)
Location:	Quebec City
Date:	90/02/27/2:00 p.m

Province # 10:	Saskatchewan
Representatives:	Kent McKerlie (Director of marketing; SED Systems)
	Dr. Maguire (Department of Computer Science; University of Regina)
	Don Friesen (Product Manager; Develcon)
	Clem Gaudet (Vice President of Operations; Develcon)
	Dean Jones (Director Department of Computer Services; University of Saskatchewan)
	Jack Manns (Dean Graduate Studies; University of Saskatchewan)
	Gordon Pierce (Director Computer and Integration Services; Saskatchewan Research Council)
	David Fetch (Product Manager; SaskTel)
	Kim Wrigley (Executive Director; Sasktel)
Interviewers:	Phil Kennis (HICKLING) David Arthurs (HICKLING) Tess McLean (Ernst & Young)
Location:	Saskatoon
Date:	90/03/09/9:00
Territory 1:	North West Territories
Representatives:	Doug Heyland (Science Institute)
	Peter Dunn
	Emery Wilson
Interviewers:	Phil Kennis (HICKLING) Roger Choquette (COMGATE) Tess McLean (Ernst & Young)
Location:	Ottawa (Conference Call)
Date:	90/03/15/1:00

APPENDIX C: NETWORK USER SURVEY

To: All Network Users

NETWORK USER SURVEY

Today, networking in Canada supports the basic ability to communicate with sites around the world, but lacks the higher level facilities that would make it really valuable. Existing networks not only lag well behind the growing needs of the research community for access and capacity, they are too fragmented to develop unaided into a single, coherent system.

The Science and Technology Division of HICKLING has been asked by Industry, Science and Technology Canada (ISTC) to investigate the feasibility of establishing a National High Speed Communications Network for the Canadian research and development community.

The new network would have a greater capacity and functionality than existing networks. The primary objective of this study is to determine the feasibility of the network. Within the study two secondary objectives must be addressed: (1) the utility of the network as a means of enhancing collaborative research; and (2) the utility of the network as a means of promoting a more internationally competitive IT sector.

While the undertaking of this study is not to be construed as a commitment by the federal government to the establishment of this network, the study will provide a solid basis for such an initiative should it be found prudent.

User profiles must be developed and user needs and priorities must be determined before the feasibility of such a network can be assessed. The study team would greatly appreciate if you would take a few minutes to answer the following questions regarding your potential use of this network.

Please return your completed survey (email preferred) to:

InterNet:	(NetNorth/CDNNet)	SMail:	Network User Survey
	HICKLING@VM.NRC.CA		HICKLING
BitNet:	HICKLING@NRCVM01		6th Floor, 350 Sparks Street
FAX:	(613) 237-7347		Ottawa, Ontario
Voice:	(613) 237-2220		K1R 7S8

Any queries should be directed to Phil Kennis, HICKLING.

---------- CUT HERE ----------

Please respond (using the space provided) to the following questions:

(Q1) Your Organization/Department Name
(A1)

(Q2) Your Name
(A2)

(Q3) Your Position
(A3)

(Q4) Your Telephone Number
(A4)

(Q5) Your Network Address
(A5)

(Q6) Do you work for: (Please enter "Y" or "N" in the space provided.)

(a) yourself, as an independent	(A6A)	[ ]
(b) a university as a researcher or staff member	(A6B)	[ ]
(c) a university as a network manager	(A6C)	[ ]
(d) a university as a student	(A6D)	[ ]
(e) a government department or crown corporation	(A6E)	[ ]
(f) a publicly funded research institute	(A6F)	[ ]
(g) a private sector company specializing in contract R&D for others	(A6G)	[ ]
(h) a private sector company doing R&D for itself	(A6H)	[ ]
(i) other (Please specify)	(A6I)	[ ]

(A6I1) [ ]

(Q7) In what general area of research do you work?
(A7A)

(Q8A) If your general area of research is telecommunications, would you consider using the network to develop and test new products and new services?
(A8A) [ ] (Yes/No)

(Q8B) If "Yes" to the above question, please elaborate on the products and services you would develop and test. (Use as much space as necessary.)
(A8B)

(Q9A) Briefly describe the role of data communications in your research. (Use as much space as necessary.)

(A9A)

(Q9B) Please give examples of the applications you require. (Use as much space as necessary.)

(A9B)

(Q10) Indicate the manner in which you communicate with your research colleagues. Please specify what percentage of the total time you spend communicating with your research colleagues for each of the following: (Note: all percentages should add to 100%.)

(a) Verbal

(i)	in person	(A10A1) [ ]%
(ii)	on the phone	(A10A2) [ ]%
(iii)	other	(A10A3) [ ]%

(b) Written

(i)	mail	(A10B1) [ ]%
(ii)	fax	(A10B2) [ ]%
(iii)	electronic mail	(A10B3) [ ]%
(iv)	net file transfer	(A10B4) [ ]%
(v)	other	(A10B5) [ ]%

(c) Video

(i)	graphics	(A10C1) [ ]%
(ii)	conference	(A10C2) [ ]%
(iii)	other	(A10C3) [ ]%

(d) Other (please specify) (A10D1) [ ]%
(A10D2) [ ]

(Q11) Indicate which Research Network(s) and which Public Data Network(s) you currently access or plan to access in the near

future. Also please specify the relative percentage of your time spent accessing each network. (Note: all percentages should add to 100%.)

	Network		Time
(A9A1)	[ ]	(A9A2)	[ ]%
(A9B1)	[ ]	(A9B2)	[ ]%
(A9B1)	[ ]	(A9B2)	[ ]%

(Q12) How much of your time is spent communicating by network in the average week?
(A12) [ ] (hours)

(Q13) Which of the following aspects of a communications network are most important to you? Please rank in order of importance, using a scale of 0-10, with 10 being the highest.

		Ranking
(a)	Connectivity to whom you wish to communicate	(A11A) [ ]
(b)	Availability of the network when you wish to use it	(A11B) [ ]
(c)	Capability of the network in terms of capacity to support high speed applications	(A11C) [ ]
(d)	Cost of the network to use	(A11D) [ ]
(e)	Ease of use.	(A11E) [ ]
(f)	Communication services supported by the Network	(A11F) [ ]
(g)	Other (Please Specify.)	(A11G) [ ]
	(A11G1) [ ]	

(Q14) Would you like to have access to a more capable network?
(A14) [ ] (Yes/No)

(Q15A) If "YES" to Q14 please specify the applications that you would run on this network, that you cannot currently run using existing networks. (Use as much space as necessary.)
(A15A)

(Q15B) Also, please identify the benefits that you would gain from the existence of such a network. (Use as much space as necessary.)
(A15B)

(Q16) If "NO" Q14 please indicate your reason(s) for not requiring access to a higher capacity network. (Use as much space as necessary.)

(A16)

(Q17) Please provide any other comments you feel are relevant to this study. (Use as much space as necessary.)

(A17)

(End)

---------- CUT HERE ----------

Thank you for very much for taking the time to complete this survey.

APPENDIX D-1: CONTACTS MADE FOR PHONE INTERVIEWS

PHONE INTERVIEW LIST

ONTARIO INTERVIEWS

Company Name	Representative	City
1. TIE Communications Inc.	Paul Mitchell	Markham
2. Canada Wire Canstar Communications	Bob Gilmour	North York
3. Canada Systems Group	Maureen McCorry	Toronto
4. Valirate Electronics Ltd. Vlacom	Don Brandi	Guelph
5. National Business Systems Inc.	Richard Strafehl	Mississauga
6. Triple Crown Electronics	Tass Vamellis	Mississauga
7. Rockwell International	Terry McKinlay	Georgetown
8. Intelcan Technosystems	Clement	Ottawa
9. Rockwell International	Lou Mitrovich	Georgetown
10. Bel – Tronics Ltd.	David Sagl	Mississauga
11. Istec Ltd.	C.J. Cooper	Hamilton
12. Calniss Semiconductor Inc.	Adam Chowanico	Toronto
13. Canadian Marconi Company	Joseph Bedford	Kanata
14. Stetron International Inc.	John Ling	Markham
15. Teletone Ltd.	Donald Fraser	Markham
16. MDS Ltd.	Peter Brugger	Markham
17. GEAC CANADA LTD	Sam Ardijar	Markham
18. MDR Telemanagement Ltd.	Richard Malone	Mississauga
19. Gefanuc Automation Canada Inc.	Jill Yousie	Toronto
20. Foundation Instruments Inc.	Russell Smith	Nepean
21. Datem Ltd.	William Thomas	Nepean
22. Comtest Instruments Ltd.	Rick Brand	Nepean
23. Genesys Group Inc.	Gunter Hurz	Ottawa
24. Simware Inc.	Robert Allison	Ottawa
25. Zavitz	Ian Malcolm	Ottawa
26. Varian Canada	Mr. Tafler	Ottawa
27. Mirtone	Laurie Eisner	Downsview
28. Softkey Software Products	Greg Dynes	Toronto
29. Honeywell/Sperring Aerospace	Barry Thompson	Rockland
30. FRI CORP	Kevin O'Donnell	Toronto
31. International Verifact Inc.	Ron Kroll	Toronto
32. Teleride Sage	Dr. Louis Kates	Toronto
33. TIL Systems Ltd.	Robert Williams	Toronto
34. Telecompute Integrated Systems Inc.	C.M. Mehda	Toronto
35. Jonas & Erikson Software Tech. Inc.	Gary Jonas	Toronto
36. CMQ Communications Inc.	John Willim	Toronto
37. Rockwell International Canada Ltd.	Ed Tanaka	Toronto
38. Perle Systems Ltd.	John Feeney	Scarborough
39. Digital Management Group	John Dightam	Toronto
40. Leitch Video International	S. Goodship	Don Mills
41. Oracle Corporation Canada Inc.	Brian Courtney	Toronto
42. Sidus Systems Inc.	Henry Kalisky	Toronto
43. Netway Communications Inc.	David Burnett	Toronto
44. Leecraft Industries & Comm. Devices	Mr. Ainlay	Toronto
45. Canadian Educational Microprocessor	Mike Pejskar	Toronto
46. Edac Inc.	Tony Smith	Toronto
47. Xicom Technologies Corp.	John Taker	Ottawa
48. Telepanel Inc.	Christopher Howlett	Toronto

QUEBEC INTERVIEWS

Company Name	Representative	City
49. Micro Tempus	Richard Pelletier	Montreal
50. Micro Tempus	Rolf Veiler	Montreal
51. Canadian Marconi	Bernie Gale	Montreal
52. Canadian Marconi Co.	Raymond Cadieux	Montreal
53. SR Systems	Eugene Gaines	Montreal
54. Philips Information Systems	K. Wong	Montreal
55. SR Telecom	Mike Movies	St. Laurent
56. Apollo	Nick Vouloumanes	St. Laurent
57. Systerm	Claude Michel	Montreal
58. Marconi Avionics	Peter Ganvi	Montreal

BRITISH COLUMBIA INTERVIEWS

Company Name	Representative	City
59. MDI–Mobile Data International	Steve Flannery	Richmond
60. Epic Data Inc.	Duncan Smith	Richmond
61. Integra Systems Inc.	Ole Jensen	Vancouver

NOVA SCOTIA INTERVIEWS

Company Name	Representative	City
62. I.M.P. Group	H. Conner	Halifax
63. Cabco Research Ltd.	Craig Meredith	Porters Lake
64. Novatron Information Corp.	D. Potter	Halifax

NEWFOUNDLAND INTERVIEWS

Company Name	Representative	City
65. Unified Systems	Brian Hurley	St. Johns
66. Paragon Information Systems	Ken Hawn	St. Johns
67. SEA Ltd.	Art Garland	St. Johns
68. Newtech Instruments	Don Nickerson	St. Johns

ALBERTA INTERVIEWS

Company Name	Representative	City
70. Myrius Research Corp.	Kenneth Gordon	Calgary

SASKATCHEWAN INTERVIEWS

Company Name	Representative	City
71. Westbridge Computer Corp.	Chris Larson	Regina

MANITOBA INTERVIES

Company Name	Representative	City
72. GEC Plessy Telecom	Peter Riches	Winnipeg

APPENDIX D-2: SURVEY FORM FOR PHONE INTERVIEW

INFORMATION TECHNOLOGY USER SURVEY

PART I ORGANIZATION DESCRIPTION

Please Provide the following information:

I-Q1 Organization Name: _____

I-Q2 Your Name: _____

I-Q3 Your Position: _____

I-Q4 Do you work for:

a) yourself, as an independent;
b) a university as a researcher or staff member;
c) a university as a student;
d) a government department or crown corporation;
e) a publicly funded research institute;
f) a private sector company specializing in contract R&D;
g) a private sector company doing R&D for itself.
h) other (please specify) _____

I-Q5 Briefly identify the types of communication products and/or services that your organization provides:

Products	Services
a) Multiplexer	a) Public Video (television);
b) Packet Switch	b) Data (packet switching service eg. datapac,CNCP);
c) PBX	c) Voice (telephone only, CANTEL);
d) Data PBX	d) Personal communications???
e) Lan Server	e) Network management service (manage networks for others eg CNCP)
f) Modems	f) Value added database service (eg. IP Sharpe)
g) Carrier Switching	g) Value added data service (packing data, error detection, etc)
h) Key Systems	h) Mobile radio services
i) Satellite Systems	i) Other; _____
j) Communications Software	j) Other; _____
k) Fiber Optic Systems	k) Other; _____
l) Other; _____	
m) Other; _____	
n) Other; _____	

I-Q6 Briefly describe the nature of your research, current and future:

_____

PART II NETWORK EXPERIENCE

II-Q1 Refer to the list below and indicate which Network(s) are accessed by your organization/department. Also please indicate the relative percentage of time spent accessing each network. (NOTE: all percentages should add to 100%)

RESEARCH NETWORKS

%		%	
_____	CAnet	_____	NSFnet
_____	Netnorth	_____	CSnet
_____	CDNnet	_____	UUnet
_____	BCnet	_____	HEPnet
_____	Onet	_____	SPAN
_____	CRIM	_____	Intraorganizational
_____	DREnet	_____	Other (Please specify) _____
_____	ARPAnet	_____	Other (Please specify) _____

PUBLIC DATA NETWORKS

_____ Datapac _____ Infoswitch
_____ Other (Please specify) _____
_____ Other (Please specify) _____

II-Q2 Would you like to have access to a more capable network (Y/N) _____.

What features would you like to see, and how would you use such a network?

Please Explain: _____

PART IV TEST-BED USE

This section attempts to identify the current and future needs of developers and suppliers of communication products (hardware, software or combination of both). It is envisioned that the proposed high speed R&D network could provide a test-bed for the investigation and initial development of such products.

If your company provides communication products (examples listed in Q5 of Part I). please respond to Q1 - Q2; and/or if your company provides communication services (examples also listed in Q5 of Part I), please respond to Q3 - Q4.

DEVELOPING AND TESTING NEW PRODUCTS (Only if I-Q5 include products)

IV-Q1 If a national high speed R&D network were available, would you consider using it as a testbed for the investigation and initial development of communication products (Y/N)? _____

IV-Q2 If "NO" to Q1, why not?

Possible answers:

a) We have our own networks for this purpose and do not foresee a change to another network.

b) We do not feel that a high speed network is particularly useful for the types of products we develop.

c) Competitive reasons would preclude us from using such a Network.

If "YES" to Q1, why would it be used?

Possible answers:

a) We do not currently have access to such a network and believe it would be useful for the type of products we develop.

b) The volume of information exchanged routinely across the network and the number of users that interact daily provide a test environment that we can not simulate in a laboratory.

DEVELOPING AND TESTING NEW SERVICES (Only if I-Q5 include services)

IV-Q3 If a national high speed R&D network were available, would you consider using it as a testbed for the investigation and initial development of communication services (Y/N)? _____

IV-Q4 If "NO" to Q1, why not?

Possible answers:

a) We have our own networks for this purpose and do not foresee a change to another network.

b) We do not feel that a high speed network is particularly useful for the types of services we develop.

c) Competitive reasons would preclude us from using such a Network.

If "YES" to Q1, why would it be used?

Possible answers:

a) We do not currently have access to such a network and believe it would be useful for the type of services we develop.

b) The volume of information exchanged routinely across the network and the number of users that interact daily provide a test environment that we can not simulate in a laboratory.

PART V COSTS AND BENEFITS OF NETWORK

V-Q1 What percentage of your work is undertaken in collaboration with colleagues/branch offices in remote locations? _____ .

V-Q2 How would this percentage change, if at all, with a convenient means of contact through *The Network*? _____ .

V-Q3 What would be the benefits of *The Network* to your organization/department? (productivity improvement, marketing and increased competitiveness, new products and services opportunities) link with Q5 and Q6. _____

V-Q4 What is your organization's total annual budget for R&D?
$_____

V-Q5 What would you expect the percentage improvement in your staff or faculty productivity from the following features of the network? Please specify a high and low estimate. (Productivity improvement may be measured either in increased value of output, or reduced time necessary to do what they do now).

		High	Low
(a)	Access to new facilities	____%	____%
(b)	Better access to facilities you currently use	____%	____%
(c)	Ability to collaborate with remote colleagues through E-mail	____%	____%
(d)	Ability to transfer large computer files	____%	____%
(e)	Ability to address questions at large through bulletin boards	____%	____%
(f)	Ability to use the network as a test-bed;	____%	____%
(g)	Ability to use network to test new information services;	____%	____%
(f)	Other	____%	____%

V-Q6 Are their facilities or services that your organization might like to offer or market through *The Network*? Possibilities include computing facilities (super computers), specialized software, information data-bases, information services, library services. Please list.

_____

_____
_____
_____
_____

V-Q7 What is the average annual value of sales of your organization?
$_____.

What percentage of annual sales are exports? _____%

Would you expect access to *the network* to have an impact on your gross sales in the long run? _____
_____

If so, could you give a high and a low percentage impact on sales that you might expect? **High** _____% **Low** _____%

V-Q8 What percentage of your annual costs go to wages and salaries?
_____%

V-Q9 In light of the benefits of *The Network*, and assuming it is non-profit, how would you prefer to support the development and operation of the network? -- by user fees, or would you prefer to support the network through other means such as donating staff time, computing facility time, and/or equipment?

If user fees, what is the most your organization would be willing to spend on an annual basis for use of this service?

a) < than $10,000;
b) $10,000 to $15,000
c) 15,000 to $20,000
d) $20,000 to $25,000
e) More than $25,000

If other means of support, by providing

f) staff,
g) computing facilities,
h) hardware,
i) other _____

Would your organization be willing to support or fund for the right to use *The Network* as a test-bed? (Only if answer YES to IV-Q1 and/or IV-Q3)

V-Q10 How should user fees be administered?

a) Fixed charges and on what basis (eg, size of firm, type of firm);

b) Variable Based on Usage; and

c) Access to the network should be free

VII OTHER COMMENTS

VII-1 Please provide any other comments you feel are relevant to this study.

_____
_____
_____
_____
_____
_____
_____
_____
_____

End. Thank you for taking the time to complete this survey.

APPENDIX E-1: LETTERS FROM PROVINCES AND UNIVERSITIES

Province of
British Columbia

Ministry of
Advanced Education,
Training and
Technology
"SKILLS FOR LIFE"

Parliament Buildings
Victoria
British Columbia
Canada
V8V 1X4

RECEIVED
APR 02 1990

March 28, 1990

Mr. David Arthurs, MBA, CMC, P. Eng
Principal
James F. Hickling Management Consultants Ltd.
350 Sparks Street, 6th Floor
Ottawa, Ontario
K1R 7S8

Dear Mr. Arthurs:

Thank you for the presentation, given by members and associates of your firm, on the feasibility study you are undertaking for Industry, Science and Technology Canada.

The British Columbia Government is very interested in the proposal to establish a National High-Speed Communications Network for the Canadian research and development community and the benefits that will be derived for the research and academic communities. It will complement the research network which this province has established under the Advanced Systems Institute. This network will assist the work being undertaken by the various organizations participating in the Federal Government's Networks of Centres of Excellence.

As you know, the Province is already heavily committed to the support of these operations and to research at the three British Columbia universities. We will look forward to seeing the recommendations in your report, particularly your proposals for sharing of costs and utilization by both the business and academic communities.

../2

Since the Ministry of Advanced Education, Training and Technology is coordinating the science and technology activities of the Government of British Columbia, the Ministry of Government Management Services has asked us to take over the lead role in this development.

We would appreciate your indicating our interest to the Ministry of Industry, Science and Technology Canada, and our wish for an ongoing dialogue.

Yours sincerely,

Tim

Tim C. Padmore
Director
Policy and Planning Branch
Science and Technology Division

cc: Mr. Terry A. Prentice
Telecommunications Consultant
Office of Information Technology
Ministry Government Management Service

Manitoba

Deputy Minister of
Industry, Trade
and Tourism

Winnipeg, Manitoba, CANADA
R3C 0V8

RECEIVED
APR 02 1990

March 28, 1990

Mr. Phil Kennis
Senior Consultant
Hickling Management Consultants Ltd.
6th Floor, 350 Sparks St.
Ottawa, Ontario
K1R 7S8

Dear Mr. Kennis:

I would like to thank you and your colleagues for the presentation made recently to some of my planning and technology staff on the subject of the proposed High Speed Communications Network. Unfortunately I was unable to attend, however, Dr. Reichert has briefed me on the meeting and provided me with the presentation information. I understand that there was a useful discussion on the subject and that a more complete understanding of the concept has developed.

The notion of facilitating access to a high performance network as a means of stimulating the development and application of communications technologies for industrial and economic development seems worthwhile. I understand that no fundamental technical difficulty exists in creating the network, especially since it appears to be based on existing telco installations. However, the questions of costs and responsibilities for implementation and operation have not yet been sufficiently explored. Of course these issues are critical to gauge what the impact of participation would be on provinces. They must be fully assessed before any general endorsement or comment on the merits and prospects of provincial participation can be offered.

We look forward to the promised copy of your final report.

Yours truly,

for H. Eliasson
Deputy Minister

Nova Scotia

APR 04 1990

FAX

Department of Industry, Trade and Technology

Office of the Deputy Minister

World Trade and Convention Centre
1800 Argyle Street
PO Box 519
Halifax, Nova Scotia
Canada B3J 2R7

902 424-3231
Telex 019-22548
Fax 902 424-0514

Our file no:

March 30, 1990

Mr. David M. Arthurs
Principal
Hickling Science and Technology Division
James F. Hickling Management Consultants Limited
350 Sparks Street, 6th Floor
Ottawa, Ontario
K1R 7S8

Dear Mr. Arthurs:

I am writing to provide you with Nova Scotia's comments on the concept of a National High-Speed Communications Network.

As you know, the Government of Canada and the Province of Nova Scotia are jointly supporting the development and implementation of the Nova Scotia Technology Network with the objective of encouraging collaboration between and among research organizations and to stimulate technology transfer from the research community to the private sector.

It appears that the objectives of the proposed High-Speed Communications Network are similar to those of NSTN and would allow NSTN to link with other similar networks and organizations across Canada. Therefore, we support, in principle, the concept of a National High-Speed Communications Network.

We are aware that there is another current initiative to build a national research network that would link regional networks such as NSTN. We believe we should have a national research network but we do not believe we have the resources to provide two networks. The National High-Speed Communications Network must incorporate C.A.Net (originally NRCNet) in some manner. We cannot have two competing national research networks. In our view, the National High-Speed Network could evolve from C.A.Net by offering higher speeds, a wider range of services for a wider range of customers and thus be a national complement to our Nova Scotia initiative.

We would appreciate being kept advised of developments with respect to this proposal and, in particular, receiving a copy of the feasibility study when your work is completed.

Yours very truly,

David Conrad for

Thomas G. Merriam
Deputy Minister

Ontario

Ministry of Culture and Communications
Ministère de la Culture et des Communications

Communications Division
77 Bloor Street West
20th Floor
Toronto, Ontario M7A 2R9

Your File:

Our File:

March 1, 1990

Mr David M. Arthurs
James F. Hickling Management
Consultants Limited
350 Sparks Street, 6th Floor
Ottawa, Ontario
K1R 7S8

Dear Mr. Arthurs:

Further to your presentation on February 23, I wish to advise that from the perspective of the Ministry of Culture and Communications the concept of a national high-speed communication network for the Canadian research and development community clearly has merit and should be vigorously pursued.

You related to the group that a statement of support for the concept would be welcomed. As I stated at the meeting, with the exception of direct funding, I am prepared to consider other types of support identified for the project.

I wish you well in bringing this project to a successful conclusion.

Yours truly,

D. A. Barr
Executive Director
Communications Division

RECEIVED
MAR 08 1990

cc: J. McCalla
M. Stevens
Circ.
File

Ontario

Ministry of
Colleges and
Universities

Ministère des
Collèges et
Universités

790 Bay Street
11th Floor
Toronto, Ontario
M5G 1N8

790, rue Bay
11e étage
Toronto (Ontario)
M5G 1N8

RECEIVED
MAR 08 1990

March 7, 1990

Mr. David Arthurs
HICKLING
350 Sparks Street, 6th floor
Ottawa, Ontario K1R 7S8

Dear Mr. Arthurs:

Thank you for your presentation regarding a National High Speed Communications Network.

May I take this opportunity to summarize the points that I made during that meeting.

The mandate of my branch is to support the research capabilities of Ontario universities. Therefore, I support the prinicple of a national high speed communications network to the extent that it will further this goal. May I remind you that any response from the province of Ontario that would commit resources would be a corporate one on behalf of the various ministries concerned.

If there was a high speed network it should tie in the Ontario Centres of Excellence, the universities in Ontario, the Ontario Centre for Large Scale Computation (located at the University of Toronto) and have the capability to provide connections to networks internationally, at least to the U.S., Europe and Japan.

There is one other aspect of the network that I feel has great potential. That is, if the network would have the capability to deliver training opportunities to various sectors including business, industry labs, government, government labs, univerisities and the twenty-three Ontario Colleges of Applied Arts and Technology. Since our colleges are closely linked to

the business community and undertake applied research, their activity would complete the range of research-to-training uses for a high speed network.

I would be interested in receiving a copy of your final report.

Thank you for interest.

Yours sincerly,

for R. Lodin

Maria L. Cioni
Director

Industry, Science and Technology Canada
Industrie, Sciences et Technologie Canada

P.O. Box 1115
Confederation Court Mall
134 Kent Street, Suite 400
Charlottetown, P.E.I.
C1A 7M8

C.P. 1115
Confederation Court Mall
134, rue Kent, Suite 400
Charlottetown (Î.-P.-É)
C1A 7M8

March 6, 1990

Mr. William G. Hutchison
Partner
Ernst & Young Consulting
P. O. Box 251
Toronto, Ontario
M5K 1J7

Dear Mr. Hutchison:

Re: National High Speed Communications Network
Our file: 6450-2

I greatly appreciate being brought up to speed during the conference call of Friday, March 2. I have, since that time, been able to meet and speak with provincial government representation, private sector and Science and Technology Advisory Committee members.

Although all that I have discussed with you, and I will reiterate this herein, is accurate the Province feels that they are not ready to tie into the high band with 1.5Mbs. It seems that arrangements are well underway towards the installation of a 56Kbs line, linked with NRC into the grid. This will be at a cost of $100,000 over a three-year period.

I perceive this network maturing and the need for expansion as inevitable at some future date. The growth will encompass both federal and provincial requirements and will include linkage into the University of Prince Edward Island including the Atlantic Veterinary College, Holland College (community and technical training) and certain small businesses within the province.

.../2

The developmental benefits are appreciated for a more sophisticated system; however, this will come through evolution rather than revolution. We must undoubtedly be permitted to be kept informed of the growth and changes to the National High Speed Network and be made aware of how and where this province might dovetail into this system.

Again, thank you for your continued interest. If further information is required, please do not hesitate to contact me again.

Sincerely,

H. W. Plant
Executive Director
Prince Edward Island

Our File:
Your File:

Yukon Development Corporation

March 7, 1990

William G. Hutchison
Ernst & Young Consulting
Royal Trust Tower
Toronto-Dominion Centre
P.O. Box 251
Toronto, Ontario
M5K 1J7

Dear Mr. Hutchison:

Thank you for your recent letter which outlines your work on investigating the feasibility of a National High-Speed Communications Network for the Canadian research and development community.

Your project has little specific relevance to the activities of the Yukon Development Corporation because our operations relate almost entirely to business economic development rather than to innovative technology or research. However, I appreciate your apprising us of the steps contemplated in this area by Industry, Science and Technology Canada.

Yours truly,

Alex Raider
President

Mailing Address: Box 2703, Whitehorse, Y.T., Y1A 2C6
Location: 304 - 204 Lambert Street, Whitehorse, Y.T., Y1A 1Z4
Telephone: (403) 667-5028, FAX: (403) 668-3327

Information
Technology
Research
Centre

777 Bay Street
Suite 401, P.O. Box 125
Toronto, Ontario M5G 2C8
Tel (416) 978-7203
Fax (416) 978-7207
E-mail: mjenkins@itrchq.itrc.ca

March 9, 1990

Mr. David Arthurs
James F. Hickling Management Consultants
350 Sparks Street
6th Floor
Ottawa, Ontario
K1R 7S8

Dear Mr. Arthurs:

This letter responds to the presentation of the proposed ISTC sponsored National High-Speed Network Study.

TRIO and ITRC fully support the Federal initiative to create a high-speed backbone research network and commends ISTC in trying to create the focus necessary to implement such a significant undertaking.

There are, however, a number of points which we would like to make which we believe to be pertinent to the issue and which you must satisfy if this undertaking is to be relevant and useful.

1. Nature of Planned Service:

   There are two distinct needs in the research community and further needs in the product development (industrial) and educational environments.

   In research, the universities and industrial research communities need to be linked by an operational service which will enable services of voice, text and image for the purposes of furthering research in whatever areas disparate communities wish to study. This network must connect to the US Internet. An advanced network of this type must eventually subsume CA Net.

   In the narrower area of communications and information technology research, an experimental network capable of research into new services is necessary. For example, our centres intend to do research in multi-point/multi-point video communications together with voice, data and

.../2

image. As this network will require dedicated access for some research investigators, it must be administratively and functionally separate from any operational service network.

To support the trial of new developmental products, the experimental network must also be capable of tolerating physical changes and service breakdown. Such a capability is incompatible with an operational service network but may be operable in conjunction with the research network.

The final concept discussed was that of an educational network. Such a network is again of the service type but will necessarily involve video on a point/multi-point basis.

We conclude, therefore, that we are discussing an entity which cannot be a single network conceptually, although some physical plant may be common.

2. Network Cross-Section

We are unsure how you are estimating the potential network traffic, but a national "thin-route" (i.e. daisy chained single T1) is not a practical proposition if all these services are to be combined. With service and experimental network requirements, different cross-sections at different points will be necessary to allow experimentation with some network configurations.

We would conclude that the Quebec/Montreal/Ottawa/Toronto/ Waterloo/Western corridor would need multiple routes. These should thread through the research universities.

3. Network Integration

There are a number of current service networks including Netnorth, ONet, CA Net etc.. in Canada and other North American primarily academic networks, e.g. Bitnet and NSFNet. Gateways currently exist so that, at least for electronic mail, it is possible to access researchers across North America. The Ca Net initiative is bringing about a merger of the the regional Networks into a consortium. In Ontario, the regional Netnorth group has plans to merge with Onet and use TCP/IP services. If the ISTC initiative is to have credibility, it must integrate the Canadian networks and significantly extend their capabilities, without a major real increase in cost. The integration of the current Canadian networks must be agreed to by their organizers and must be part of the final proposal.

.../3

The important thing is to get the bandwidth needed between the collaborating researchers, whether they be inside one region or working on opposite sides of the country. Thus, the proposal must address the mechanism for bringing the regional network connections up to at least T1.

4. Timing

Too frequently, we have seen papers and proposals go into a prolonged evaluation procedure, only to be followed by a further study request and evaluation. Canada's capabilities are already significantly behind most other developed countries and progress elsewhere is proceeding at an exponential rate. We need a major effort to put a world class capability in place today. In 2 to 3 years, the opportunity for industry to take advantage of such a facility may well be lost, as we will be too far behind other international research networks. We will probably be able to buy this network technology, and the opportunity for research and early implementation will have been lost.

The fastest way to get the ISTC proposal in place would be to fund the CA Net and regionals to go to T1 immediately and leave them to administer it initially. The co-operation of the provinces in this effort, either by making their excess capacity available for many of the links, or by providing additional funds to the universities, would be an essential part of this approach.

A plan to evolve this setup into the sort of network envisioned by ISTC could gradually be put in place taking over the CA Net service as appropriate.

Considering the narrow perspective of TRIO and ITRC, it is our intention to establish a limited capability experimental network between Ottawa/Carleton/Queen's/Toronto/Waterloo this year for a joint TRIO/ITRC research initiative. Industrial support is being sought for this experimental network. It would be eminently sensible for the ISTC capability to be compatible with the network we wish to establish. An evolutionary approach would make this feasible.

5. Network Termination

The utility of the network is determined by the use which is made of it. Therefore, it is conditional on a real viable use to the end user. Planning the trunking, while essential is insufficient and it is necessary to extend the concept to the point of use in a professor's or student's terminal, a

.../4

research lab, a lecture room/theatre, and administrator's computer. This linkage to and functional capability of the campus network must also be world-class. The provincial governments must assure that the universities have the funds to complete the networking of the campuses.

6. TRIO and ITRC usage

TRIO (and other Ontario Centres of Excellence) pay overhead for university infrastructure support and we would not normally consider paying separately for a service network. To the extent that an ISTC network supported the needs for an experimental network as envisioned by TRIO and ITRC, the funding raised to support such research communications could be applied to the cost of the ISTC effort.

ITRC was an original founder of ONet and has joined this regional network as a full member. Although the administrative work done using the network does not justify the cost of a full membership, the ITRC Board believes support of the network is worthwhile because the ITRC researchers depend heavily on the provision of the network services. Indeed, there is a need for substantially better service than we are currently getting. The distributed administration of ITRC has been greatly facilitated by the existence of ONet and hence we expect to continue as a full member of the consortium.

TRIO has not taken such a position with respect to Onet, but its researchers use a combination of BitNet, NetNorth and Onet depending on their university affiliation and research community. Unless any network subsumes Onet, TRIO anticipates moving over to ONet administratively when this becomes the predominant researcher communication network.

We wish you success in bringing forward your proposal, but would stress the importance of early decisions as other network environments have or are being established and we cannot wait if we want to be competitive in this field of current Canadian strength.

Yours sincerely,

TELECOMMUNICATIONS RESEARCH INSTITUTE OF ONTARIO

Peter Leach
President

INFORMATION TECHNOLOGY RESEARCH CENTRE

Michael A. Jenkins
Executive Director

/mg

Saskatchewan

Saskatchewan Telephones

Communications Policy

3rd Floor, Saskatchewan Place
1870 Albert Street
Regina, Canada
S4P 4B7

RECEIVED
APR 05 1990

April 2, 1990

Mr. Phil Kennis
Senior Consultant
James Hickling Management
Consultants Limited
350 Sparks Street
OTTAWA, Ontario
K1R 7S8

Dear Mr. Kennis:

Thank you for your letter of March 14, 1990, in which you inquire whether I am prepared to support the concept of a High-Speed Communications Network in Canada.

Please appreciate that I cannot commit the Government of Saskatchewan either in principle, or from a practical point of view, to participating in such a network. An endorsement of this nature would, I think, have to have the sanction of the most senior levels of government and no-one from this department has sought such approval at this time.

Speaking only as someone who has been involved in a number of informational sessions respecting high speed data, and as one who personally is of the view that this medium holds significant promise both in terms of its capabilities as well as associated economic potential, broadly interpreted, I do support the idea of such a network being established. I agree with a number of the points that were raised at our informational meeting in Saskatoon, and would like to see such a network used in terms of providing quality public services such as health care and education in conjunction with other services that appropriately are provided by the public sector. I also agree that the establishment of this network could be used to foster and encourage certain private sector initiatives such as the development and marketing of certain hardware and, more significantly, software.

... 2

It would be particularly good for the country if these developments could occur in the less well established regions as well as those that are very well established. It is in this context that I support the concept and establishment of a high speed data network and I am hopeful that others too will be equally as supportive. I am sorry that I cannot provide any official status to my support, but I expect that many, including yourselves, are in a similar position at this time.

Yours truly,

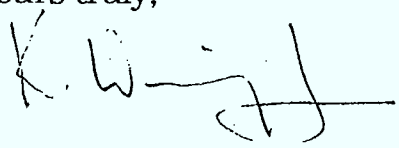

Kim C. Wrigley, Ph.D.
Executive Analyst

Saskatchewan
Research Council
15 Innovation Blvd. Saskatoon, Sask. Canada S7N 2X8 Phone: (306) 933-5400 Telex: 074-2484 Fax: (306) 933-7446

March 9, 1990

David Arthurs
HICKLING Management Consultants Ltd.
350 Sparks Street, 6th floor
Ottawa, Ontario
K1R 7S8

Dear Mr. Arthurs:

Re: **High Speed Communications Network**

We, at the SRC (Saskatchewan Research Council), have a definite interest in the proposed network. We are interested in using the network for:

- delivering research assistance to industry and better developing ways of doing this
- delivering specialized remote sensing products
- providing access to specialized data bases (geological data, SRC expertise and reports)
- providing access to pilot project use of manufacturing software

I am enclosing several documents that will give an overview of the SRC. It is possible that APRO (Association of Provincial Research Organizations) would be a good organization to coordinate the establishment of the regional networks. Jim Hutch, President of SRC, is also the current President of APRO and would be willing to talk to you directly about this. The provincial research councils have one foot solidly in industry and the other in the research community. For operation of the network, the telephone companies should be involved.

Please contact me regarding any follow-up discussion on this issue.

Yours sincerely,

Gord Pierce

GP/ms
Enclosure

cc: Jim Hutch
Ravi Maithel

Telecommunications
Research Institute of Ontario
TRIO

RECEIVED
MAR 16 1990

8 March 1990

Mr. David Arthurs
James F. Hickling Management Consultants
350 Sparks Street, 6th Floor
Ottawa, Ontario
K1R 7S8

Dear Mr. Arthurs:

TRIO and ITRC have prepared a joint response which is attached.

This response is the agreed position of both organizations with respect to the proposed High Speed Backbone Research Network.

We hope your study will conclude with strong recommendations to proceed immediately to implement a network which will meet the variety of service and research needs which we believe to be vital.

Yours sincerely,

TELECOMMUNICATIONS RESEARCH INSTITUTE OF ONTARIO

Peter Leach
President

PL:lw
Encl.

cc: D.A. Barr, A. Bjerring, M. Cioni, P. Dewan,
J. McCalla, D. McGeown, C. Pringle, M. Stevens,
M. Walmsley, D. Williams

Administration and Technical Program
340 March Road, 4th Floor *phone (613) 592-9211*
Kanata, Ontario K2K 2E4 *fax (613) 592-8163*

Member Services and Licensing
627 Lyons Lane, Suite 100 *phone (416) 842-7100*
Oakville, Ontario L6J 2Y2 *fax (416) 842-7110*

High Speed Communications Network Panel Session - Position Paper

Steven Sutphen

Department of Computing Science
University of Alberta
Edmonton, Alberta T6G 2H1

March 4, 1990

This paper is brief as I can see no need to repeat the standard comments and arguments that are given in support of a high speed network (e-mail, file transfer, remote logins). If the study wants to see proof of the usefulness of a medium speed network (T1 - 1.544 Mbps) they can examine the "Merit-NSFNET Link-letter." There are however a few things which I would like to comment on.

The most important point I want to make is procedural rather than functional, although it will effect the long term function of the network. I feel that ISTC should start funding a research arm to help guide them in the initial proposal and ongoing as the network evolves. The proposal mentioned that the network should conform to OSI standards. I agree that following standards is vitally important to the usefulness of any network. CDNnet has been an active participant in the OSI standards and was the world's first network to use the X.400 messaging standard. They have been participating in and following the evolving OSI networking standards closely and would be a valuable asset to provide guidance to the proposed network.

The second point is along similar lines. We must work together! Canada does not have a large enough population to support several disjoint efforts in networking. We do however have the large land area and long distances that make networking even more necessary than in other more densely populated technically advanced nations. CA*net, arguably the first Canadian backbone providing a wealth of services, has just started into implementation; NetNorth will largely be subsumed by CA*net; CDNnet is re-examining their business plan in light of the limited network dollars. All this points to the conclusion that we have to find a way to work together in establishing this national resource. I feel, from my naive place in the world, that a good scenario would be to have ISTC support the growth of CA*net into a more generally available resource than it can now afford to be (with its limited bandwidth and manpower). They should do what ever they can to encourage use of this resource. They should enlist the CDNnet as a research arm to guide the protocol evolution to OSI. This would then give Canada a substantial resource.

The third point concerns the speed of the initial plans for the network (voice, data, and images) verses the bandwidth (T-1, 1.544 Mbps). It would seem to me that few individuals could be accommodated in say a video conference at this speed. I would think that "real time" image transfer (for video conferencing) would be limited to a few feasibility experiments at best. It may be constructive to look to the USA to put the ISTC proposal into perspective. ARPANET was constructed initially in 1969, consisting of 10 to 20 nodes interconnected with 56 Kbps links. This is nearly identical to what CA*net is putting in this year. ISTC is proposing an initial bandwidth of 1.544 Mbps. This is the same as NSFnet implemented in 1988/9. I do not know of anyone able to do more than experiment (during quiet times) with interactive voice or video on NSFnet. The USA is looking toward the future also, as they have discovered the benefits of wide area networking (NSFnet is carrying over 2 billion packets per month now). Commercial intermediate level Internet connections are being offered by at least two companies—Performance Systems International (PSI, Reston, Va.) and UUNet Communications Services (Falls Church, Va.). These connections are demand driven (Data Communications, February 1990, pp.46). Looking into the more distant future, President Bush has endorsed a $2 billion initiative to engineer a "data super highway," with a 3 Gbps

bandwidth. Canada should be looking forward to this kind of bandwidth too.

The final point touches on applications. I feel that the experiences in the USA, Europe, and Japan have proven that 1.544 Mbps can be effectively used to help researchers interact and produce and disseminate more results to improve the world condition. Having just gone through the experience of convincing people about the usefulness of a WAN when they didn't know what they were missing, I know how hard it is to predict what the future will bring. To show how easy it is to be "so close and yet so far away" it is instructive to read the essay by John McCarthy entitled "Networks Considered Harmful for Electronic Mail" in the December 1989 Communications of the ACM. In this essay he looks at why faxes are much more widely used than e-mail even though e-mail has many advantages, and a much longer history. One begins to wonder at what the near future networking requirements will be when they start to consider the potential of delivering information digitally at an individual level.

I was just reading an article about the University of Alberta Library leading a fight for lower journal prices ("Joint Serials Renewal Negotiation Project"). They gathered information about the academic subscriptions from 50 Canadian libraries. They found that the major commercial publishers supply 14,628 journals to Canadian libraries costing $9,517,313 per year, or $651 per year per journal subscription. I wonder how many of the papers in those journals were in machine readable form when written, and could have been distributed directly electronically for much less cost. As more material is stored electronically one can easily imagine the "electronic National Library," the "National Jutebox" (which could be easily implemented using DAT right now), or the "National Video store" as contemplated by McCarthy. Certainly one area for development is in the area of information filters. In the March Scientific American there is an article on trends in communications ("The Road to the Global Village", Karen Wright, pg.83) that indicates that the human can acquire knowledge at 50 bits/second. Information providers easily exceed this rate, so the information must be filtered to individuals requirements.

In conclusion, I definitely feel that Canada must develop a High-Speed Digital Communications Network. This network would be a national resource and will become as vital as the telephone and fax are today (and may in time replace them for the general population). Because the low population density of Canada distributes colleagues over long distances, that are expensive in time and money to traverse, it is especially vital to the Canadian research community. To implement this network we must all work together, doing the research to incorporate the "right" protocols in a timely fashion, and working toward a common goal, hand in hand. This period of time has been called the "Information Age," and digital networking provides the most effective delivery vehicle for that information.

by John McCarthy

viewpoint

Networks Considered Harmful For Electronic Mail

Electronic mail over computer networks has been in use for almost 20 years. The widespread use of telefax is more recent. However, unless e-mail is freed from dependence on the networks, I predict it will be supplanted by the telefax for most uses in spite of the fact it is more advantageous. Information is transmitted more cheaply as character streams than as images, and multiple addressees are readily accommodated. Moreover, messages transmitted as character streams can be readily filed, searched, edited and used by computer programs.

Unless e-mail is separated from special networks, telefaxing will prevail because it works by using the existing telephone network directly. To become a telefax user, it is only necessary to buy a telefax machine for a price between $1,000 and $5,000 (depending on features) and to publicize one's fax number on stationery, on business cards and in telephone directories. Once this is done anyone in the world can communicate with you. No complicated network addresses and no politics to determine who is eligible to be on what network. Telefax is already much more widely used than e-mail. Japan estimates that 5 percent of the Japanese household population will have telefax machines by 1995. By the year 2010, the devices will cost about $200 and be found in 50 percent of the homes.

E-mail could work the same way at similar costs, but because of a mistake by DARPA about 20 years ago, i.e., making a special-purpose, special-politics ARPANET network the main vehicle for e-mail, it was combined with other network uses that require higher bandwidth and packet switching.

Another mistake was UUCP. It uses the telephone network, but three features inherited from its use within Bell Telephone Laboratories made its widespread adoption a blunder.

1. It assumes that both parties are using the UNIX operating system rather than using a general mail protocol. This is only moderately serious, because some other systems have been able to pretend to be UNIX sufficiently well to implement the protocols.
2. It requires that the message-forwarding computer have log-in privileges on the receiver. This has resulted in a system of relaying messages that involves gateways, polling and complicated addresses. This results in politics in getting connected to the gateways and often causes addresses to fail.
3. Today, forwarding is often a service provided free and therefore of limited expandability.

There has been a proliferation of networks and message services on a variety of time-sharing utilities. Some of them are commercial and some serve various scientific disciplines and commercial activities. The connections between these networks require politics and often fail. When both commercial and non-commercial networks interact, charging complications arise. A whole industry is founded on the technologically unsound ideas of competitive special-purpose networks and storage of mail on mail computers. It is as though there were dozens of special-purpose telephone networks and no general network.

The solution is to go to a system that resembles fax in that the "net addresses" are just telephone numbers. The simple form of the command is just MAIL <use>@ $<telephone number>, after which the user engages in the usual dialog with the mail system.

The sending machine dials the receiving machine just as is done with fax. When the receiving machine answers, the sender announces that it has a message for <user>. Implementing this can involve either implementation of protocols in a user machine or a special machine that pretends to be a user of the receiving machine or local area network. The former involves less hardware, but the latter involves less modification to the operating system of the receiving machine.

I have heard various arguments as to why integrating e-mail with other network services is the right idea. I could argue the point theoretically, but it seems better to simply point out that telefaxing is already far more widespread outside the computer science community. Indeed, it is often used for communicating with someone who is thought to have an e-mail address when forwarding connections seems too complicated.

The World of the Future

Eventually, there will be optical fiber to every home or office supplied by the telephone companies. The same transmission facilities will serve telephone, picturephone, telefax, e-mail, telnet, file transfer, computer utilities, access to the Library of Congress, the "National Jukebox" and maybe even a national video jukebox. In the meantime, different services require different communi-

cation rates and can afford different costs to get them. However, current telephone rates transmit substantial messages coast-to-coast for less than the price of a stamp. Indeed the telefax success, not to mention Federal Express, indicates that people are willing to pay even higher costs.

What will happen to e-mail over the next 20 years?

There are two kinds of problems, technical and political. Guess which is easier.

The main technical requirement is the development of a set of point-to-point telephone mail protocols. Any of several existing network mail protocols could be adapted for the purpose.

Presumably the same kinds of modems and dialers that are used for fax would be appropriate but would give better transmission speeds.

Organizationally speaking, perhaps the simplest solution would be to get one or more of the various UNIX consortia to add a direct mail telephone protocol to UUCP. Such a protocol would allow mail to be addressed to a user-id at a telephone number. The computer would require a dialer and a modem with whatever characteristics were taken as standard and it would be good to use the same standards as have been adopted for telefax. It must not require pre-arrangement between the sending and receiving computers, and therefore cannot involve any kind of log-in. Non-UNIX systems would then imitate the protocol.

Fax has another advantage that needs to be matched and can be overmatched. Since fax transmits images, fully formatted documents can be transmitted. However, this results in the loss of the ability to edit the document. This fax feature can be bested by e-mail, provided there arises a widely used standard for representing documents that preserves editability.

The political problem is more difficult because there are enormous vested interests in the present lack of system protocol. There are the rival electronic mail companies. There are the organizers of the various non-profit networks. There are the engineers developing protocols for the various networks.

I have talked to a few of them, and intellectual arguments have remarkably little effect. The usual reply is, "Don't bother me, kid, I'm busy."

It would be beneficial if the ACM were to set up a committee to adopt a telephone e-mail standard. However, I fear the vested interests would be too strong, and the idea would die from being loaded with requirements for features that would be too expensive to realize in the near future.

Fortunately, there is free enterprise. Therefore, the most likely way of getting direct e-mail is for some company to offer a piece of hardware as an e-mail terminal including the facilities for connecting to the current variety of local area networks (LANs). The most likely way for this to be accomplished is for the makers of fax machines to offer ASCII service as well. This will obviate the growing practice of some users of fax of printing out their messages in an OCR font, transmitting them by fax, whereupon the receiver scans them with an OCR scanner to get them back into computer form.

This is probably how the world will have to get rid of the substantially useless and actually harmful mail network industry.

In general, suppose the same need can be met either by buying a product or subscribing to a service. If the costs are at all close, the people who sell the product win out over those selling the service. Why this is so I leave to psychologists, and experts in marketing, but I suppose it has to do with the fact that selling services requires continual selling to keep customers. This, in turn, keeps the prices high.

I hope my pessimism about institutions is unwarranted, but I remember a quotation from John von Neumann that in effect said expecting institutions to behave rationally is like expecting heat to flow from a cold place to a hot place.

I must confess that I do not understand the relation between this proposal and the various electronic communication standards that have been adopted like X25 and X400. I only note that the enormous effort put into these standards has not resulted in direct telephone e-mail or anything else as widely usable as telefax.

I am grateful for comments from many people on a version distributed by e-mail to various boards.

John McCarthy, recipient of ACM's 1971 A.M. Turing Award, holds the Charles M. Pigott Professorship at Stanford University's School of Engineering.

CENTRE FOR INTEGRATED COMPUTER SYSTEMS RESEARCH • C • I • C • S • R •

THE UNIVERSITY OF BRITISH COLUMBIA

March 16, 1990

MAR 22 1990

Mr. Phil Kennis
Hickling Management Consultants Ltd.
350 Sparks Street, 6th floor
Ottawa, Ontario
K1R 7S8

Dear Phil:

Thank you for my copy of the minutes of the network panel discussion, held February 18, in Vancouver.

I have discussed the proposed network with researchers here at UBC, and I can report that support for such a network is very strong, particularly if it allows for a substantial research component. I can even say that the researchers are not averse to a user fee mechanism to cover expenses based on volume of usage.

I look forward to further developments with respect to the network, and hope that you can keep me informed on its progress.

Yours sincerely,

J.M. Varah
Director, CICSR

JMV/sp

Telephone (604) 228-6894
Fax (604) 228-7006
E/mail: varah@cicsr.ubc.ca

Office 2053-2324 Main Mall
Vancouver, British Columbia
Canada V6T 1W5

THE UNIVERSITY OF BRITISH COLUMBIA

Department of Computer Science
6356 Agricultural Road
Vancouver, B.C. Canada V6T 1W5
Tel: (604) 228-3061

February 8, 1990

High Speed Communications Network for Research and Development
Expert Panel Session
Hotel Vancouver

Subject: position paper for ISTC network feasibility study

These notes offer recommendations to Industry, Science and Technology Canada on the topic of a Canadian high-speed communications network for research and development.

1. The network should provide a wide range of services to the entire Canadian research and development community on a continuing basis.

- The network should coordinate and deliver application services, not just connectivity. The latter can be provided adequately by a number of commercial service providers. A national network should explicitly accommodate those not yet network-literate, and those who require more than simple connectivity.
- The lead department must take a strong, central leadership role, in addition to a coordinating role. To maximize its utility and impact, both technological and political, the network must be much more than a backbone to which research-oriented suppliers and consumers connect.
- Because of the long-term strategic importance of national networking, a continuing influence over the evolution of the network should be maintained. As well, the issue of immediate benefits versus costs should not be an overriding concern.
- ISTC should act very quickly, and should be prepared to be flexible in its planning and execution. Modern networking is possible because of advances in communications and computer technologies, and in both of these areas the rate of advance is astounding. Thus while long term projections are useful, it should be recognized that the landscape is changing so rapidly that the best plans made as recently as a few months ago may not be the best plans if made today. The network should be organized to encourage, accept, and incorporate change rather than to react to it.
- The network should encourage participation from small and remote centres, rather than penalizing them through high costs for connectivity. This is to give a more even footing to those researchers not at large institutions, to provide infrastructure to educate the next generations, and to encourage the growth of high technology organizations throughout the country.

2. The network should specifically support research and development in the area of telecommunications.

- The network should take advantage of existing Canadian experience with Open Systems Interconnection (OSI). In particular, the expertise built up by CDNnet, the national OSI network for research and development, should be incorporated into the network.
- Some very specific services should be targetted to provide new functionality in the network, and to raise its visibility. Examples of this are: exchange of multi-media documents via X.400 electronic mail, distributed X.500 directory services, and inter-library applications.
- The R&D community should be encouraged to become directly involved in the ongoing planning and engineering of the network. This includes researchers in the area of computer communications, as well as those in other areas such as multi-media and human factors. A good model here is the U.S. Internet Activities Board (IAB) with its Internet Research Task Force (IRTF) and Internet Engineering Task Force (IETF). Another good model is the Association of European Research Networks (RARE) and its Working Groups.
- The concept of a test-bed is an excellent one. In particular, the University of British Columbia presently has a number of OSI-related R&D projects which could benefit from the availability of a test-bed.

3. Close ties should be forged and maintained both nationally and internationally with bodies engaged in networking.

- A well-developed national network will have a significant positive effect on Canada's ability to compete and to prosper in the future. Thus the network should be a national initiative which extends beyond the confines of a single government department. There should be close cooperation with other departments and agencies involved in networking. This includes departments which currently operate networks or fund operational networks, and agencies which fund related work such as individual network-related research projects. A possible model here is the U.S. Federal Research Internet Coordinating Committee (FRICC), whose members sponsor most of the U.S. research on internetworking. FRICC members include the Defense Advanced Research Projects Agency (DARPA), the National Science Foundation (NSF), the Department of Energy (DOE), and the National Aeronautics and Space Administration (NASA).
- There should be cooperation with similar initiatives in other countries. This cooperation will involve pragmatic concerns (e.g. sharing of trans-continental bandwidth), policy and planning issues (e.g. appropriate international use and technical compatibility), and opportunities for international R&D.

John C. Demco, CDNnet Manager
Paul C. Gilmore, Professor
Gerald W. Neufeld, Assistant Professor

THE UNIVERSITY OF CALGARY

ACADEMIC COMPUTING SERVICES

2500 University Drive N.W., Calgary, Alberta, Canada T2N 1N4

Telephone (403) 220-6201

To: Phil Kennis
Senior Consultant
James F. Hickling Management Consultants Limited

From: Rom Kieffer
Manager, Datacommunications
Academic Computing Services
The University of Calgary

Date: March 22, 1990

Dear Phil:

I would like to make a few comments regarding the ISTC network initiative.

First of all, the time has come to build a Canadian network that goes beyond the technology that other countries and networks used in the past. The United States officially terminated the ARPAnet at the end of February, while we are planning to implement the CAnet using the same architecture and technology the ARPAnet introduced in the 1970's.

If Canada is to retain the reputation for excellence in communications, and to develop a reputation for innovation, now is the time to take decisive action, and to create the opportunities that will help to export technology rather than raw, unprocessed resources,

In my opinion those opportunities can only be realized if we are prepared to take the risks of a quantum step in the infrastructure of the network. I just finished reading about the effects of high speed railway networks in Europe. It used to take me two days to drive from Luxembourg City to Venice, the new trains will do it in a few hours. The authors of the article were discussing the benefits that this new culture would bring for business, lets go a step further, lets provide a communication infrastructure that will eliminate the need for some of the travel, plus add considerable advantages to the businesses using the network services. The future will bring fiber optic communications at a grand scale, lets build a vehicle for it to arrive early in Canada. There is no glory, or profit, in reinventing the wheel.

At the University of Calgary we are discussing the future provision of high speed fiber optics network services to the desktop, we already use ethernet to support those workstations today. A comparatively slow speed

national network would be flooded with traffic once high speed regional networks are connected.

I realize that there are many obstacles, not least the tariffs regulated by the CRTC. Besides the technological efforts, considerable work must be dedicated to a new pricing structure for communications services. What is the use of having Gigabit/sec capabilities if they become unavailable for cost reasons?

But what would stop the prospective partners in the research from petitioning the CRTC and the long lines providers for a research exemption from tariffs? It is my understanding that only commercial exploitation is regulated, not R&D activities. This would give us plenty of time to develop services, to perform research, and to evaluate the results, while tariffs are restructured. Most of the research I envision could be done without any national networks at all, laboratories all over the world are doing it already, but none of them have the national and large scale exposure that will germinate the unforeseen benefits and applications that only usage can provide.

There are many areas of research that this network could support, here are just a few:

- dynamic bandwidth allocation, both in time and in capacity.

- new error correction and avoidance protocols. Network nodes handling traffic at hundreds of Megabits/sec would spend more time calculating error correcting codes than forwarding messages. We could end up needing a supercomputer at each network node.

- Cooperative processing is at the beginning of its existence. Soon computer users at their workstations will be able to share computer power across networks, to access distributes databases transparently, and to access resources that no one entity could afford for itself. Some standards have been developed, more are in the design stage, but none are widely installed, especially in wide area networks.

Applications:

- One Canadian network that interconnects voice services, video transmissions, and data networks.

- Access to information stores, data libraries.

Beneficiaries:

There are direct beneficiaries who will be enabled to perform leading edge research, and indirect beneficiaries who will take advantage of the

facilities and services provided by the R&D community. The following paragraph lists only a few that come to mind immediately.

- Researchers can communicate effectively with their peers anywhere in Canada and in the world.

- Organizations can attract researchers to any location, not limited to by local availability of resources like supercomputers for example.

- Communications companies will have an engineering testbed, and a market for their products, as well as a demonstration plant for their foreign customers.

- Critical mass for commercial R&D. Success will attract more organizations and lay the foundations for more locations like Silicon Valley.

- Interconnection of vendors, contractors, and suppliers with each other and their customers.

- Opportunities for chip designers to produce electronics for the new networking environment, and for software developers to produce the communications packages that will be used in these networks.

Concerns:

- That the implementation contract not be given to one contractor who would drive it towards a company specific interest. A page out of the Japanese manual would come in handy: Companies there develop products in one common laboratory, then add their own added-value components and compete as hard as they can against each other and the rest of the world.

Yes, I do support the initiative, but as you see, I want more than just another network. Lets try to make this one count.

If I can be of any assistance, please do not hesitate to contact me.

Sincerely

Tom Kieffer

COLLEGE OF
GRADUATE STUDIES
AND RESEARCH
UNIVERSITY OF
SASKATCHEWAN
SASKATOON, CANADA
S7N 0W0

March 28, 1990

Mr. Phil Kennis
Senior Consultant
James F. Hickling Management
Consultants Limited
350 Sparks Street
6th Floor
Ottawa, Ontario K1R 7S8

Dear Mr. Kennis:

RE: Feasibility Study of a High Speed Communication Network

Subsequent to our meeting in Innovation Place, Dean Jones, Director of Computing Services, and Bob Kavanagh, Associate Vice President, Office of Information Technology Services, met to discuss the issue and consider an appropriate response.

Bob Kavanagh will, I believe, be writing a separate letter; his level of understanding of the technical issues is far superior to mine so I will not attempt to dwell on the technical issues.

Our response can be summarized as follows:

1. We recognize the fundamental importance of high speed networks and we have already taken steps to become part of CA net. We would welcome a more comprehensive network which includes industry because it should enhance our ability to work with industry.

2. We are concerned about financial matters and would not be able to make financial commitments beyond what are currently made, or anticipated, for the CA net.

3. We have some concerns about the ability of the telephone companies to manage the network. Bob Kavanagh may give more detail on that point.

Overall, we are supportive; the only significant difficulties we see are if the cost becomes prohibitive or if the project is delayed beyond the realistic time frame of about 2 years.

Sincerely,

Jack Manns.

Jack Manns, Ph.D.
Dean of Graduate Studies and Research
Associate Vice-President (Research)

RECEIVED
APR 02 1990

OFFICE OF THE DEAN
PHONE 306-966-5751
FAX 306-373-6088
TELEX 074-2659

JM:sb

cc: Bob Kavanagh

UNIVERSITY OF SASKATCHEWAN

COMPUTING SERVICES
56 PHYSICS

SASKATOON, SASKATCHEWAN
S7N 0W0

RECEIVED
MAR 30 1990

1990 03 27

Mr. Phil Kennis
Senior Consultant
Hickling Management Consultants Limited
350 Sparks Street, 6th Floor
Ottawa, Ontario
K1R 7S8

Dear Mr. Kennis:

Re: Feasibility Study of a High-Speed Communications Network for Research and Development

Thank you for inviting me to the information session which you held in Saskatoon, and for the opportunity to comment on this initiative.

I have discussed the concept of the network with my manager, Dr. R.N. Kavanagh, Associate Vice-President (Information Technology Services), and with Dr. J.G. Manns, Associate Vice-President (Research). Our discussions were positive, and each of them will be responding to you with statements of the University's position in respect to the concept of this network.

Yours sincerely,

Dean C. Jones
Director

DCJ/he

UNIVERSITY OF SASKATCHEWAN

INFORMATION TECHNOLOGY SERVICES
E282 ADMINISTRATION BUILDING

SASKATOON, SASKATCHEWAN
S7N 0W0

RECEIVED
MAR 30 1990

Tuesday, March 27, 1990

Mr. Phil Kennis
Senior Consultant
James F. Hickling Management Consultants Limited
350 Sparks Street, 6th Floor
Ottawa, Ontario

K1R 7S8

Dear Mr. Kennis;

Re: Industry, Science and Technology Canada Initiative for a High-Speed National Data Network

I regret that I was unable to attend the presentation on the ISTC network initiative in Saskatoon on March 9. I understand that the presentation was well attended and that it stimulated good discussion of the concept of a high-speed network to support industry-government-university interaction in Canadian research and development.

The purpose of this letter is to confirm that the University of Saskatchewan is supportive, in principle, of such a network. We believe that such an undertaking is going to be critical to Canada's ability to compete with countries that have very aggressive support for this type of infrastructure for research and development.

You will also be hearing from Dr. J. Manns, Associate Vice-President for Research and Dean of Graduate Studies & Research at our institution. My responsibilities include Computing Services at our campus, and from that context I wish to indicate the issues that we have identified that qualify our support for the proposed network.

First, on an economic level, our institution is already straining to move to CA*net membership. If and when the ISTC network is available, it will not be possible for us to support memberships in two networks at that level of cost. Therefore either our CA*net fee will have to also provide access to the ISTC network, or CA*net will have to be subsumed within the ISTC service, at no increase, and preferably a decrease, in costs.

Second, our support is dependent on the ability for <u>all</u> segments of our institution to participate: research, teaching applications and administration. Current "academic" networks allow universities to use the networks to support all of these areas (at our institutional discretion), and we would hope that the ISTC service would do likewise.

Finally, we would strongly support the concept of a management structure that is based on a 'consortium of users' model. Throughout North America, the most effective network management structures have the characteristic of relying on the Telco industry for bandwidth, but upon a user management board for policy and planning. We strongly believe that this would be appropriate for the network that ISTC is considering. The management board should have representation from all regions of the country.

In summary, the University of Saskatchewan needs to be connected to a national network that provides the richest possible communication with industry and government agencies and other academic institutions, to support our research, teaching and administrative functions. At the present time we intend to support CA*net as long as we can afford to. I believe that CA*net is itself only a step towards networks of even higher speed and functionality. Perhaps the ISTC network will be the next step. Canada is supposedly the nation of experts in communications. For the lack of coordination of the many national needs for networking, we currently have a hodge-podge of networks, resulting in a dilution of our collective ability to support them. I often regret that the DOC did not go ahead with CANUNET in 1971; perhaps we could have achieved a much earlier national focus and support structure in this critically important area.

I would appreciate being kept informed of developments in the ISTC network. If you wish any additional information about our interests, please do not hesitate to call.

Yours sincerely,

Robert N. Kavanagh

Robert N. Kavanagh
Associate Vice-President
Information Technology Services
(306) 966-6623

cc:
- **Blaine A. Holmlund, Vice-President, Planning and Development**
- **Dr. Jack Manns, Associate Vice-President and Dean of Graduate Studies and Research**
- **Dean C. Jones, Director, Department of Computing Services**

APPENDIX E-2: LETTERS FROM COMPANIES

Posted: Fri Mar 30, 1990 12:10 PM EST Msg: NVQA-8155-5000
From: AE.WALLE/MARKETING/AGT.BUS.SYSTEMS/AGT
TO: PKENNIS(FAX:6132377347), JHOULST(FAX:6132377347)
Subject: NATIONAL HIGH-SPEED COMMUNICATIONS NETWORK

To: ENVOY --CAL02V

SUBJECT: NATIONAL HIGH-SPEED COMMUNICATIONS NETWORK
THIS MESSAGE HAS THE FOLLOWING NOTE ATTACHED
For Your Information
Attention Phil Kennis
Attn Judy Houlstone

AGT

Albert Walle
Marketing Manager

Business Systems

Floor 3 East
3030 2nd Ave. S. E.
Calgary, Alberta
T2A 5N7
(403) 235-8380

March 30, 1990

Dr. D. Williams
Director - Information Technologies
Industries, Science and Technologies Canada
Government of Canada
235 Queen Street
Ottawa. Ontario.
K1A 0H5

Dear Dr. Williams:

National High-Speed Communications Network

On behalf of A.G.T., I would like to thank you for the opportunity to become involved in this major networking endeavor. We have no basic disagreement with the concept of establishing a national high-speed communications network.

However, A.G.T. is not in a position to participate directly in the proposed 'shared' network. Regulatory policies governing us do not permit A.G.T. (a common carrier) to have an ownership or equity position in any reseller or shared network.

A.G.T. is precluded from deviating from our tariffed rates for any network service. However, A.G.T. would like to assist ISTC in the development of this network. We will provide technical expertise to your organization to configure the Alberta portion of this network. Once a formal design document is available, A.G.T. Business Systems will provide a competitive quotation for the services required.

.../2

A.G.T. looks forward to dealing with your organization in the future. Please have your staff contact Albert Walle, at (403)-530-3247, to set up the appropriate direct contact. Budget restriction will require that all consultative work to be carried out at our home office.

Yours truly,

Albert Walle
Marketing Manager

AEW

cc: Phil Kennis

SR TELECOM INC.
8150 Trans-Canada Hwy.
St. Laurent, Québec
Canada H4S 1M5
Tel: 514/335-1210
Facs: 514/334-7783
Tlx: 05-824919

SR TELECOM

MAR 05 1990

February 28, 1990

Mr. P. Kennis
Senior Consultant
Hickling
350 Sparks Street, 6th Floor
Ottawa, Ontario
K1R 7S8

Dear Mr. Kennis:

This is further to our meeting on February 27, 1990 at SRT where we discussed some of the possibilities that might be offered by the proposed new High Speed Communications Network.

While it is always difficult to forecast the use of such a network until it really becomes available, it is clear that it could help us in a number of ways, such as:-

- easier and more user friendly access to existing data bases,
- access to activities in our field taking place at the universities,
- promotion and initiation of joint activities with companies with interests similar to ourselves,
- access to test facilities that may exist in the future. For example, the ability to thoroughly test ISDN interfaces and operation between our products and a simulated ISDN network would be of great interest to us if it were available on the network.

In addition to the above, it is possible that our point-to-multipoint TDMA microwave products could provide some local connections (64KB/s rate) in some parts of the network.

One area to be careful about, of course, is security. While activities might be 'publicly' initiated on the network, we would obviously not make the fruits of our own R&D, or possibly even its direction, freely available information!

Finally, as in all such matters, use of the network will be proportional to its perceived value and the costs involved to the user.

I hope the above gives a satisfactory picture of our interests.

Yours truly,

M.J. Morris
Vice President
Technical & Industrial Liaison

MJM/sm

B.C.Tel

BRITISH COLUMBIA TELEPHONE COMPANY
3777 KINGSWAY BURNABY B.C. CANADA V5H 3Z7 PHONE (604)432-2151 FAX (604)434-9916

90 03 27

Mr. Phil Kennis, Senior Consultant
James F. Hickling Management
Consultants Limited
350 Sparks Street, 6th Floor
Ottawa, Ontario
K1R 7S8

Dear Mr. Kennis:

Re: Feasibility Study of a High-Speed Communications Network for Research and Development

Further to your letter of March 14th and this morning's telephone conversation with Judy Houlston, the concept of such a high-speed communications network is indeed feasible from technical, economic and implementation points of view, and complimentary to the National Advisory Board on Science and Technology (NABST) objectives of a Canadian coordinated R&D process. It would be consistent with the Halifax Declaration and supportive of the theme of the next Conference in Edmonton (1990) of Science and Technology Advisory Committees.

Yours very truly,

RFA/tgl

Robert F. Alexander
VICE-PRESIDENT -
CORPORATE DEVELOPMENT

Bell-Northern Research Ltd.
P.O. Box 3511, Station C
Ottawa, Ontario
Canada, K1Y 4H7
Telephone: (613) 763-2211
Telex 053-3175

BNR

20 March 1990

RECEIVED
MAR 23 1990

Mr. David M. Arthurs, P. Eng.
James F. Hickling Management Consultants Ltd
350 Sparks Street, Floor 6
Ottawa, Ontario K1R 7S8

Dear Mr. Arthurs:

It was a great pleasure to meet with you and Mr. Kennis, and also with your associates from Comgate Engineering Ltd., on February 12th, and then again on March 16th, to discuss the Feasibility Study for a National High-speed Communications Network which your organization is carrying out on behalf of Industry, Science and Technology Canada (ISTC).

We do believe that it is appropriate to selectively invest public funds to promote "technology pull," to accelerate the introduction of leading edge technology when proven economic viability seems tenuous. The clear intention and expectation should be that the enhanced functionality would enable new business services which would achieve commercial viability quite promptly.

As was evident from the discussion at our first meeting, it appears that existing product plans (of Northern Telecom) are capable of fulfilling the requirements which you envisage, within the time frames which you indicate. NT's intention is to evolve its network product line at a pace which matches that of available CPE.

Telecom Canada evolves its national network in step with business service needs. It may be that the high-speed service needs which your study addresses could be the trigger to induce Telecom Canada to evolve its national network to provide this capability. Your meeting with Jim Schram and Doug Carruthers no doubt considered this possibility.

In Canada, due to Federal government and C.R.T.C. policies, cost and price (telco tariffs) are often rather unrelated, in both the long distance and the local services market, with the former subsidizing the latter.

...2

As we discussed quite extensively on March 16th, the evolving BNR/NT Corporate Wide Area Network provides a generally appropriate but somewhat simplistic "model" of the proposed Network which you are considering. We noted that the BNR/NT WAN, which serves a single corporate organization, must contend with many opinions as to how best the network and functionality should evolve. You can expect an even greater diversity of view from your clients/users, with their inherent diversity.

We understand that your proposed Network will initially link the "Regional" networks which currently link Universities at the provincial level. At BNR, we have found that linking existing networks in this way leads to substantial and unpredictable increases in network usage as new functionality enables new applications. Our internal approach has been to put the network into place and then monitor service levels and add capacity as required, as the traffic load builds.

The experience of the BNR/NT Corporate WAN, as it has evolved, suggests that traffic tends to expand to fill the capacity available (a network adaption of Parkinson's Law!). We noted the importance of a price-cost algorithm that enables the proposed Network to add capacity, at no net cost to itself, as usage builds. Conceptually, we understand and support the view that the proposed Network will provide an orderly evolution path for CA Net. Initially, we see the proposed Network as linking the already established provincial "Regionals", evolving quite quickly to have a life of its own.

Another concern which we discussed was the possible impact of a few large-bandwidth users on many small-bandwidth users. The network traffic implications of providing services with varied bandwidths and variations in other attributes is the subject of much activity in the Standards forums which are addressing B-ISDN. The BNR/NT Corporate WAN supports the embedded base of TCP/IP. Looking to the future, an orderly and appropriately timed evolution to OSI will be necessary.

Low volume locations must of course be addressed. The BNR/NT Corporate WAN makes use of switched services (e.g. voiceband data, Datapac, Datalink, Centrex Data) to provide appropriate transport where dedicated connections are uneconomic.

As we discussed both on Friday last and on February 12th, Network Operations are of critical importance to the success and usefulness of the proposed Network and are a complex and costly item.

...3

We noted that BNR/NT would likely not make extensive use of the proposed Network, since we have our own somewhat similar facility. We do now interconnect to the provincial "Regionals" (e.g. O Net), and we would use the proposed Network where it has unique capability (for example, to reach specific users). As we discussed, one area which could be substantially useful to BNR would be that of conformance testing, confirmation, and certification, for example OSI. We see the need for a national capability to which the proposed Network could provide access. The certification would have to carry the authorization of the appropriate agency, for example CIGOS.

Your proposal to put in place a National High-Speed Communications Network should help to improve the competitiveness of Canadian industry in world markets. We at BNR will make some use of it, and smaller organizations will find it even more valuable. We hope that the Project continues to the next phase, the detailed planning of the Network. We would welcome the opportunity for further discussion and to offer our views, based on our experience, of your proposals.

Yours sincerely,

David F. Barr

D.F. Barr
Director - NTD Operations
(613) 763-5568

cc: Mr. Roger Choquette, Comgate Engineering

CNCP

CNCP Telecommunications
3300 Bloor St. West
Toronto, Ontario M8X 2W9

Gary B. Oliver
Vice-President & General Manager,
Central Region
(416) 232-6836 Telex 06-218362
Fax: (416) 232-6180
Dialcom: 22:CNP314
Cellular: (416) 543-4023

MAR 01 1990

February 22, 1990

Mr. Phil Kennis
Senior Consultant
James F. Hickling
Management Consultants Limited
350 Sparks Street, 6th Floor,
Ottawa, Ontario
K1R 7S8

SUBJECT: ISTC Network Initiatives & CNCP

Dear Phil:

It was a pleasure to meet with you, Roger and Digby on February 20th in Toronto. As we discussed, CNCP is interested in this initiative. There looks like several opportunities to involve ourselves and our suppliers. There also seems to be a good fit for us in involving several educational institutions, not only from a joint project point of view (R & D), but also to give them more of an understanding of our industry and provide some opportunity for future hiring.

If you need anything additional from CNCP or if you have any additional communication concerning this project, please feel free to call or address it to me. I will await your comments to let me know what the next stage of the initiative is.

Sincerely,

Gary B. Oliver
Vice-President & General Manager
Central Region

GBO/ck

c.c. R. Dainty C. Witkowicz
B. Butler

Develcon

Develcon Electronics Ltd.
856 - 51st Street East
Saskatoon, SK Canada S7K 5C7

Telephone (306) 933-3300
Telecopier (306) 931-1370
Telex 074-2780

3/12/90

MAR 1 0 1990

Attention: David Arthurs
James F. Hickling Consultants Ltd.
350 Sparks Street, 6th Floor, Ottawa, K1R 7S8

Subject: National High Speed Communications Network

Dear David,

Develcon would like to express interest in your plans for a National Network. We can see benefits for our national research & development requirements as well as for Develcon. This program can benefit Develcon in three ways:

1) Develcon as a user: We maintain research facilities in Saskatoon and Toronto. As a user of the low speed public data networks, we could benefit greatly from a set fee, high speed access network for research.

2) Develcon as a supplier: As most western educational and research institutions have standardized on Develcon's networking product as a backbone carrier, Develcon is in a good position to be supplying an inexpensive gateway to the the network. In this respect, Develcon is very interested in providing product and if necessary, discussing the research and development requirements to design and manufacture the hardware / software required to facilitate the network. This is consistent with Develcon's strategic direction.

3) Development of new market opportunities: Develcon have excelled in the development of products which address the bottom four layers of the OSI model. If we were a part of national network, we would be more inclined to generate products which would address the applications layers of the OSI model. This could open up a new marketplace for Develcon.

Develcon supports your initiative and will gladly work with ISTC and others to see the network come to fruition.

Sincerely,

Clem Gaudet, P.Eng.
VP of Operations
Develcon Electronics Ltd.

cc: Don Friesen, Develcon Marketing.
Ron Dalenay, Develcon Sales Ottawa.

30

GEOVISION

RECEIVED FEB 07 1990

February 5, 1990

Mr. Phil Kennis
Science and Technology Division
James F. Hickling Management Consultants Ltd.
350 Sparks Street
6th Floor
Ottawa, Ontario
K1R 7S8

Dear Phil:

In response to your request, GeoVision's research needs would not be materially affected by the availability of an enhanced communications system such as the one described.

All of our essential information (computer systems technology) originates and is sourced from the United States -- primarily from hardware and software providers. Likewise, subject matter specific to the applications we serve.

As we are not an information provider, we could not have any contribution to make from this perspective either.

Sorry I can't be of assistance.

Yours truly,

Mary Stapleton / per

D.W. Seaborn
C.E.O.

DWS:mes

GeoVision Corporation
1600 Carling Avenue, Suite 350,
Ottawa, Canada K1Z 8R7
(613) 722-9518 Fax: (613) 722-5385

IBM

IBM Canada Ltd.

105 Moatfield Drive
North York, Ontario
M3B 3L9

RECEIVED
FEB 12 1990

February 6, 1990

Mr. P. Kennis
Senior Consultant
Hickling Associates
Ottawa, Ontario

VIA FAX

Dear Mr. Kennis

Thank you for your letter dated January 19, 1990, inviting IBM Canada to participate in the Expert Panel Session on High-Speed Communications Network for Research and Development.

Prior to receipt of this invitation, Paul Koch and Duncan Howard of our Ottawa office met with David Arthurs to discuss the feasibility of establishing this network.

It was agreed at that time that the Research and Development area of IBM Canada would gain no benefit from the access to such a network but, as a vendor of Information Technology Solutions, we would be prepared to assist Hickling Associates in this feasibility study and to share our experience and expertise with you.

Duncan Howard has been asked to be the coordinator of the IBM Canada resources that will be needed for this project. Dunc can be reached at (613) 788-6077.

We look forward to working with you in the near future.

Sincerely

W.N. Palm:cjs
Vice President
Telecommunications &
Business Development

cjs8-4/1

204

Memotec Data Inc.
600 McCaffrey St.
Montreal, Quebec
Canada H4T 1N1

Tel.: (514) 738 4781
Telex: 05 824228
Fax: (514) 738 4436
Envoy 100: Memotec.Sales

RECEIVED
APR 03 1990

memotec

Data Communications Division

March 30, 1990

Mr. Phil Kennis
Senior Consultant
HICKLING CONSULTANTS
350 Sparks St.
6th Floor
Ottawa, Ontario
K1R 7S8

Dear Mr. Kennis,

I would like to convey to you our appreciation in having chosen Memotec as part of the companies who have participated in the information gathering for the feasibility study of a high speed communications network for research and development.

Memotec has no doubt that such a network is needed and that would also be feasible.

Memotec would like to commit to a participation in any future phases of this project.

Yours truly,

Yvon Galarneau
National Sales Director
Canada

YG/tg

MITEL CORPORATION
350 Legget Drive
Kanata, Ontario
Canada K2K 1X3
(613) 592-2122

RECEIVED
MAR 30 1990

28 March, 1990

Phil Kennis
Senior Consultant
James F. Hickling Management Consultants Ltd.
6th Floor
350 Sparks Street
Ottawa, Ontario
K1R 7S8

Dear Phil,

Thank you for your presentation to Mitel Corporation in February regarding the National High Speed Communications Network for R&D.

Mitel supports the efforts of the Canadian Government to establish interest in such a network across Canada.

Mitel Corporation is interested in this project from two perspectives:
- as a user
- as a supplier of network components such as telephone sets, data sets, switching systems and semiconductor components.

Mitel envisages an R&D environment in which it could exploit the network not only through developing and testing commercial products, but through co-operation with Canadian research organisations, including government and the academic community, and with Canadian businesses interested in exploiting the international telecommunications market.

We would be pleased to remain involved in this process, from analysis through implementation. Please advise if there is any further information you require.

Yours truly

Hilary E. Smith
Vice-President
Business Development
Mitel Corporation

Mar 28/90.

Kenneth B. Anderson
Assistant Vice-President
Business Systems
Semiconductor Division

MOTOROLA

Information Systems

February 20, 1990

Mr. Phil Kennis,
Senior Consultant
James F. Hickling Management Consultants Limited
350 Sparks Street
6th Floor
Ottawa, Ontario
K1R 7S8

Dear Mr. Kennis,

Thank you for the opportunity to discuss the National High-Speed Network with you and Mr. Choquette last Monday. It is an interesting concept that has the potential of integrating Research and Development on a national basis.

Motorola would be interested in any or all of equipment supply, network installation, operation, and in working with the research community in evolution of network capabilities. We are committed to OSI Network Management, a key element of your requirements, and are developing the capability of extending management to non-Motorola devices and systems. You may also be aware that Motorola, as a "Rationalized" company, enjoys Class 1 status for Federal Government and related supply services.

I look forward to hearing the results of your technical and financial evaluation of the network.

Yours very truly,
MOTOROLA INFORMATION SYSTEMS

E.M. Strain
President

cc: R.C. Willcocks

9445 Airport Road, Brampton, Ontario L6S 4J3 Tel: (416) 793-5700 Twx: 610 492 5308 Tlx: 06-97565

165 Graveline
St. Laurent, Quebec
Canada H4T 1R3
514-735-2741
Toll-free U.S.: 1-800-361-1965
Fax. 514-735-8057

MUX LAB

Thursday, March [illegible] 1990

JAMES F. HICKLING MANAGEMENT
CONSULTANTS LIMITED
350 Sparks St, 6th floor
Ottawa, ON
K1R 7S8

Attention: **Mr. Phil Kennis**
Senior consultant

Sir,

Since 1988, Mux Lab Inc. has been following the definition of the concept of a national high-speed communications network.

As effective high-speed communications is becoming recognized as an integral and indispensable catalyst for advanced research in communications technologies, we think that the implementation of a canadian high-speed communications network can provide a mean for canadian industries to accelerate the development and testing of next generation information technology.

Mux Lab is already addressing the necessity to design and manufacture new solutions for the exchange of information in accordance with internationally accepted standards by aligning the development of new products on the OSI 7 layers model. These emerging standards will characterize future global markets for information technology products.

In the definition of the concept of a network that would link regional and international networks, Mux Lab's interest would be in the development of internetwork routing products, including gateways and bridges.

The implementation of a canadian high-speed network is very attractive to Mux Lab Inc. and you can continue to count on Mux Lab's interest and cooperation.

Sincerely,

Christian Légaré, ing.

Christian Légaré, ing.
Vice-president R. & D.

northern telecom

Northern Telecom
Canada Limited

2920 Matheson Boulevard East
Mississauga, Ontario
Canada L4W 4M7

Tel. (416) 238-7000
TWX 610-492-6311
Telex 06-960462
Fax (416) 238-7350

RECEIVED
MAR 19 1990

March 12, 1990

Mr. David M. Arthers
Science and Technology Division
Hickling Management Consultants Ltd.
350 Sparks St., 6th Flr.
Ottawa, Ontario
Canada
K1R 7S8

Subject: ***National R&D Communications Network***

Dear Mr. Arthers:

Thank you for the opportunity to review your conceptual proposal to develop and implement a high speed National Communications network for Research and Development applications. We certainly support your view that such a program would significantly accelerate the development of future oriented Canadian technology/services and act as an international showcase for worldwide R&D communications networks.

Northern Telecom's vision of an integrated, standards based, fiber optic communications network appears to be synergistic with your own. Our recent FiberWorld announcement is predicated on delivering these same attributes to the public and private switched telecommunications market. Our objective is to establish a state-of-the-art technology platform upon which a whole new generation of sophisticated high speed data and information services will be developed.

As you know all standards associated with such a network have not yet been resolved, and as such final specifications are somewhat of a "moving-target". Therefore, as your proposal develops we would welcome on-going dialog regarding implementation strategies and service requirements.

We look forward to meeting with you as your program evolves.

Regards,

B.B. Goodfellow
Assistant Vice-President
Business Development

30 Colonnade Road
Suite 300
Nepean, Ontario
K2E 7J6
Tel: 613/727-9576
Fax: 613/727-5672

PRECARN Associates Inc.

March 19, 1990

Mr. Phil Kennis
Senior Consultant
Hickling Management Consultants
350 Sparks Street, 6th Floor
Ottawa, Ontario
K1R 7S8

RE: Feasibility Study of a High-Speed Communications Network for R&D

Dear Mr. Kennis:

I enjoyed meeting with you and Mr. Roger Choquette a few weeks back, and I hope that our discussion on PRECARN's needs and perceived usage of a High Speed Communications Network will have been useful to your study. As you requested, I am sending you this letter to confirm our support for the concept of a National High Speed Communications Network for R&D.

PRECARN is an industrial consortium for pre-competitive research in Artificial Intelligence and Advanced Robotics. Our research program consists of a number of projects carried out jointly by industry, government and university laboratories. As we are an "Institute without walls", communication between our different projects, researchers and members is very important to our activities and success. As such we are in general support of the concept of a High-Speed Communications Network which could be used by our researchers and members. We would perceive using such a Network for exchanging very large volumes of data/programs etc, for two or more remote groups to work cooperatively online, and possibly even to hold remote meetings of researchers/committees.

Of course one very important limiting factor for using such a Network, will be cost. While it is difficult, at this point, for us to put a precise value on such a service, we generally feel that the pricing of the Network should be compared to similar services or to services it can replace. We also believe that usage of the Network would be initially low and increase gradually as people become familiar with the capabilities of the Network and new services/software become available to take advantage of its power.

I believe that this study is an excellent initiative and I look forward to reading the final report.

Yours sincerely,

Jean-Claude Gavrel
Director of Research Programs

Simware
Partners in the world's largest networks:
Solutions through connectivity software

MAR 27 1990

Simware Inc.
20 Colonnade Road
Ottawa, Ontario, Canada
K2E 7M6
Tel.: (613) 727-1779
Fax: (613) 727-9409
Fax-Sales: (613) 727-8797
Telex: 053-4130

23 March 1990

Mr. Phil Kennis
James F. Hickling Management Consultants Limited
350 Sparks Street
Ottawa, Ont
K1R 7S8

Dear Phil,

Thank you for including Simware in your research on the feasibility of establishing a Canadian R&D network.

Simware strongly endorses the concept of a national Canadian network to facilitate communication and interchange among the R&D community. I was particularly pleased that the current vision includes the "applied research and development" being done by companies such as ours. For years now, we have made "read only" use of the educational networks such as BITNET, but have always been limited because of our "outsider" status. We need a network in which, despite being a "for profit" company, we can still participate as full member subscribers.

Simware would benefit from such a network in several ways:

1. Timely access to information. Information about new products, new technologies, new applications, new competitors and numerous other subjects is almost invariably available "on the network" before it ever makes the trade press or we discover the information in the field. Such information is obviously critical to our success and competitiveness in the international marketplace.

2. Communication with customers. A large portion of Simware's existing and target customer base is already part of the global "internet" community. Being part of a network would allow us to be significantly more responsive to customer queries and problems, particularly for our international customers and distributors. (*e.g.* contrast E-mail versus Canada Post). Service quality is a key battleground in the 1990s, and network access is one way to achieve a competitive service advantage.

3. Access to multiple environments. Simware's business is connectivity, and there is no "standard" environment that matches all our customers' configurations. A small company such as ours can only hope to apply the "80/20" rule, and hope that the development and testing environment within our own organization matches that of most of our customers' networks. Having network connections facilitates co-

operation with customers or other vendors who have environments which we need to support.

4. Access and exposure to international networking standards. The proposed use of OSI protocols for the network would provide Simware with an opportunity to participate in OSI networking that we would be very unlikely to pursue on our own initiative. Furthermore, the very presence of such a network will spur demand for value-added products and create a market for OSI-related products.

In addition to using the network, Simware sees potential as a supplier to the network community. Simware's expertise is at several different levels, one of which is in providing easy-to-use, transparent user interfaces to complex networks. Since many of the primary users of the network will not be "computer jocks" by avocation or occupation, their ability to use all the facilities offered will be a key measure of the success of the network. The "presentation" and "automation" capabilities of Simware's workstation products can be used to create extremely simple interfaces for every type of application. Ease of use also directly impacts the type and amount of technical support that will be required as part of the network's infrastructure.

I am most willing to participate in the next step–please don't hesitate to call me for any more information or clarification.

Sincerely,

Lew Shepherdson
Vice President, Technology

LS:mw

Software Kinetics

06 March 1990

Mr. Phil Kennis
Senior Consultant
James F. Hickling Management Consultants Ltd.
350 Sparks Street, 6th Floor
Ottawa, Ontario
K1R 7S8

Subject: IST Net

Dear Mr. Kennis:

Thank you for visiting with us recently to review the IST Net and for soliciting our views on the subject. As mentioned during our meeting, we have a number of interests with respect to this network:

- First of all, as users, we see tremendous benefits for companies such as ours. We currently have an Internet connection which is most useful in communicating with clients and associates around the world.

- Software Kinetics also has an interest in designing and implementing the network. Our previous implementation experience with DRENet (Defence Research Establishment Network) and the Nova Scotia Technology Network qualifies us for this task.

- As well, we believe Software Kinetics can market, operate and manage the business of the network. We are doing this now in whole or in part for other clients and believe private sector involvement is a key to the success of the network. Note also that marketing the network to users is most important particularly if the network is to be self financing. Our experience as users and network designers qualify us for this task as well.

- Finally, the company is planning the development of products for which a network test bed is required. For example, we have an IP Data Router that we believe could be adapted for future use as an ISO router.

Software Kinetics Ltd.
65 Iber Road, Stittsville, Ontario, Canada, K2S 1E7
(613) 831-0888
Fax # (613) 831-1836

It is also important to have a market available for new products and IST Net should provide a market for new Canadian products.

Thank you again for your time and we look forward to receiving a copy for your final report.

Sincerely

T. A. Moretto
Director
Sales and Marketing

TAM:bl

Spar Aerospace Limited

Satellite and
Communications
Systems
Division

21025 Trans-Canada Highway
Ste-Anne-de-Bellevue, Quebec H9X 3R2
Telephone: (514) 457-2150
Fax: (514) 457-2724
Telex: 05-822792

SPAR

March 29, 1990

Fax 613-237-7347

RECEIVED
APR 02 1990

James F. Hickling Management Consultants Ltd.
350 Sparks Street, 6th Floor
Ottawa, Ontario. K1R 6S8

Re: Feasibility Study of a High Speed Communications Network for Research and Development

Dear Mr. Kennis

Thank you for your March 14th letter to Peter Garland. We were pleased to meet with you and the project team in January and we look forward to receiving a copy of your final report.

As we indicated during the day's discussions the concept of this network is attractive from many aspects. In discussion with others in the research and academic communities we have concluded that there is a real need for a high speed mesh network linking powerful computing resources with potential applications. The success of similar facilities in the US has demonstrated the level of utilization which could be expected with a Canadian system.

In addition to this day to day operational requirement, the facility would serve two other purposes, each critical to the Canadian telecommunications and information processing communities. As a test bed for new developments and applications, this network would allow the smaller Canadian entrepreneurs in these sectors to confirm the operation of new innovations in an actual operating environment before approaching the market. As a show case for Canadian and international network users, the High-Speed Communications Network will highlight the initiatives, technologies and capabilities of Canadian companies, and will demonstrate the technical and economic benefits of specific applications to potential customers.

In summary, the Network will serve to foster and display Canada's ongoing role as a world leader in communications technology.

Spar has acquired the experience required to take a lead role in such a project from our successful network implementations in North America and the Asia-Pacific region. We are looking forward to support you in both the planning and the implementation of the network. Please do not hesitate to call me for any additional information or further discussions.

Very truly yours

W.S. Corless
Director of Marketing
Canadian Business Networks

APPENDIX F: GLOSSARY

APPENDIX F: GLOSSARY

Address	The prefix of a coded message that identifies either the sender or receiver of the message. Any group of bits that identifies a network node as a separate, identifiable location.
AM	Amplitude Modulation
ANSI	American National Standards Institute
ARPANET	A wide–area network of packet switching nodes, operated by the U.S. Department of Defense for the support of research.
Asynchronous transmission	A method of transmitting computer data one character at a time. The length of time between characters is variable. Each character establishes a synchronization pattern for the receiver by means of start and stop bits.
ATM	Asynchronous Transfer Mode, a switching technique
B channel	Bearer Channel in ISDN
Bandwidth	The range of frequencies of an electromagnetic energy form measured from the lowest to the highest frequencies. The frequency range of a transmission medium such as the 3,000 Hz bandwidth of a voice–grade communications channel.
BBN	Bolt, Beranek and Newman; the originators of packet switching nodes (PSNs).
BBRA	Broadband Basic Rate Access
BISDN	Broadband ISDN
BPRA	Broadband Primary Rate Access
BPS	(bits per second) The measurement of the rate at which digital information signals are transmitted.
BRA	Basic Rate Access
BRI	Basic Rate Interface in ISDN
Bridge	A device that acts as a connector between similar local area networks. Bridges operate at OSI Level 2, the Data Link Layer.
Broadband link	A transmission channel with a wide frequency range that is divided into separate communication channels.
BSC	Binary Synchronous Code
CB	Channel Bank
CCITT	Comite Consultatif International Telegraphique et Telephonique
CCTA	Central Computer and Telecommunications Agency (agency of UK Treasury)

CD–ROM	Compact Disc Read Only Memory. Computer form of audio compact discs.
CEN	European Committee for Standardization
CENELEC	European Committee for Electrotechnical Standardization
Centrex	A Common Carrier service
CEPT	European Conference of Postal and Telecommunications Administrations
Channel	A path used for the transmission of electrical signals. One of several independent communication links available on a broadband link.
CID	Caller ID
CIGOS	Canadian Interest Group on Open Systems
CN	Call Negotiation
CO	Central Office
COM	One of the top–level domains. Stands for commercial and includes commercial enterprises.
Common carrier	An organization licensed to provide tele–communications facilities to the public.
Contention	A method of control that determines how the separate nodes of a network can access a shared transmission medium.
COS	Corporation for Open Systems; provides testing services for networking products.
COSAC	Canadian Open System Applications Criteria
CPE	Customer Premise Equipment
CR	Call Routing
CRTC	The Canadian Radio–television and Telecommuni–cations Commission. The agency responsible for regulation of the broadcast industries in Canada and the federally incorporated telecommunications common carriers.
CSA	Canadian Standards Association
CSA	Circuit Switched Access
CSMA/CD	Carrier Sense Multiple Access/Collision Detection
CSU	Channel Service Unit
CTS	Conformance Testing Services (Europe)
CTV	Cable Television

CUG	Closed User Group
CVSDM	Contineously Variable Slope Delta Modulation
D channel	Data Channel used for signalling in ISDN
DACS	Digital Access and Cross–connect System
DAS	Datapac Access Software
DASS	Digital Access Signalling System
Datagram	A self–contained package of data that carries enough addressing and routing information so that it can travel from source to destination without reliance on earlier exchanges between the source or destination and the transporting network.
DCS	Digital Cross–connect System
DEC	Digital Equipment Corporation
DECnet	A proprietary network for Digital Equipment Computers.
Digital signal	A discrete or discontinuous information signal that represents a numerical value; it is separated from other signals by an identifiable period of time.
Distributed processing	The movement of information processing functions from a central computing facility to separate locations equipped with independent systems.
DLA	Dedicated Link Access
DM	Digital Multiplexing
DOC	Department of Communications (Canada)
DS	Digital Signal
DSX	Digital System Cross–connection
DTE	Data Terminating equipment
ECMA	European Computer Manufacturers Association
EDU	One of the top–level domains. Stands for educational and includes educational institutions.
EGA	Enhanced Graphic Adapter
EGP	External Gateway Protocol
EIA	Electronic Industries Association (USA)
ESF	Extended Superframe Format, format of a T1 bitstream

Ethernet	A kind of network cable, or a network which conforms to IEEE standard 802.3.
FAX	Common abbreviation for facsimile; a page transmission service using scanners and copiers.
FDDI	Fibre Distributed Data Interface
FM	Frequency Modulation
FPS	Fast Packet Switching
FR	Frame Relaying
Frame	A self–contained package of data at the link layer.
FT1	Fractional T1
FTAM	File Transfer, Access, and Management (ISO)
FTP	File Transfer Protocol; a user–level protocol and program that you can use to transfer files over the network.
FTP	File Transfer Protocol (MIL–STD–1780)
Gateway	A device that acts as a connector between two logically separate networks. It has interfaces to more than one network and can translate the packets of one network to another, possibly dissimilar, network.
Gigabyte	One thousand million bytes.
GOV	One of the top–level domains. Stands for government and includes government organizations.
H channel	Higher speed channel (>64 Kbps) in ISDN
HDLC	High–level Data Link Control
HDTV	High Definition TV
HG	Hunt Group
IA	Isolated Access
IBL	Individual Business Line
IBM	International Business Machines
IDN	Integrated Digital Network in ISDN
IEC	International Electrotechnical Commission
iNET	A Bell Canada service based on Datapac, that provides Intelligent Network functions.

INTAP Interoperability Technology Association for Information Processing (Japan)

Internet Protocol Suite Alternate name for TCP/IP Protocol Suite.

IP Internet Protocol

ISAL Information System Access Line, a Bell Canada tariff for a data IBL

ISDN Integrated Services Digital Network

ISO International Organization for Standardization; responsible for publishing the Open System Interconnection Reference Model.

Kb Kilobyte: 1 Kb = 1024 bytes (accurately), a thousand bytes (colloquially). Also used as an abbreviation for Kilo–baud.

Kbps Kilobits per second. A measure of transmission speed.

KP Kilo Packet

LAN Local Area Network. A term used to describe the dedicated networks used to link computers and peripherals together within a relatively confined area. The area is usually an office, but may be as large as a building or campus.

LAP Link Access Protocol

LAPD Link Access Protocol D

LDDS Limited Distance Data Set

Leased line A communications channel reserved for the sole use of the leasing customer. Filtering elements are usually added to leased lines to support high data transmission rates. This process is called conditioning.

Line switching (circuit switching) A method of completing a direct physical communications path between two communicating devices. This contrasts with message switching where no physical circuit is established.

LPVS Linear Predictive with Variable Slope

MA Metropolitan Access

MAN Metropolitan area Network

MAP (General Motors) Manufacturing Automation Protocol

Mb Megabyte: 1 Mb = 1,045,576 bytes (accurately, 1024 x 1024), one million bytes (colloquially, and more usually).

Mbps Million bits per second. Another measure of higher transmission speed. If one were dealing with computer mainframe channel speeds, one would talk of Mbps and mean Megabytes, not megabits.

Message switching	A communications operation in which messages are received by a switching center and re–transmitted to their ultimate destinations.
MIL	One of the top–level domains. Stands for military and includes military organizations.
MILNET	A wide–area network of packet switching nodes, operated by the U.S. Department of Defense for the operational support of military communication.
Modem	A device that impresses digital signals onto a carrier wave for transmission over an analog transmission path. At the receiving end, the modem converts the analog signals back to digital pulses. Modems are used in pairs, one at each end of an analog communications line used for data communication.
MSAT	Mobile SATellite
Multiplexing	The division of a communication line into two or more separate channels either by separating it into independent frequency bands (frequency division multiplexing), or by assigning the same channel to different users at different times (time division multiplexing).
NBS	National Bureau of Standards (now called NIST)
NCP	Network Control Protocol or Program; the original host–to–host protocol for the ARPANET. In 1983, it was replaced by TCP/IP.
NCSCI	National Center for Standards and Certification Information (USA)
NET	One of the top–level domains. Stands for network and includes network service centers, network informations centers, and other organizations that have a hand in network management.
NISDN	Narrowband ISDN
NIST	(United States) National Institute of Standards and Technology
Node	A network location where communication links begin, end, or intersect. A node can be a user device with a direct attachment to its network, or it can be a communication processor that performs network functions for user devices.
NT1	Network Termination 1 in the ISDN reference model
NT2	Network Termination 2 in the ISDN reference model
NTSC	National Television Standards Committee
NUI	Network User Identifier
OC	Optical Carrier
ORG	One of the top–level domains. Encompasses nonprofit organizations.
OSF	Open Systems Foundation. A grouping of major computer industry companies to promote the adoption of microcomputer standards.

OSI	Open Systems Interconnection. A model of communicating processes established for use in the definition of inter–machine protocols.
Packet	A self–contained package of data at the network layer.
PAD	Packet Assembler/Disassembler
PBX	Private Branch Exchange
PCM	Pulse Code Modulation
PLP	Packet Level Protocol (as per X.25)
PRA	Primary Rate Access
PRI	Primary Rate Interface in ISDN
PRN	Packet Radio Network
Protocols	Communication software that establishes bit, character, or message synchronization between communicating devices; enables devices to recognize and correct errors; and determines how devices can access a network.
PSA	Packet Switched Access
PSN	Packet Switched Network.
PSTN	Public Switched Telephone Network
PVC	Permanent Virtual Circuit
QPSK	Quaternary Phase Shift Keying
R&D	Research and Development
RA	Remote Access
RC	Reverse Charge
RNR	Receive Not Ready
Router	A hardware and software device that connects hosts on different networks. Routers operate at OSI Level 3, the Network Layer.
RR	Receive Ready
SCC	Standards Council of Canada
SCOT	Steering Committee on Telecommunications
SDLC	Synchronous Data Link Control
Server	A provider of network service.

SNA	Synchronous Network Architecture
SNAcP	Subnetwork Access Protocol
SNDCP	Subnetwork Dependent Convergence Protocol
SNICP	Subnetwork Independent Convergence Protocol
SONET	Synchronous Optical Network
STM	Synchronous Transfer Mode
Store and forward	A data communication technique where switching node accepts messages from a communicating device and stores them until they can be passed on to their destinations.
STS	Synchronous Transport Signal. Electrical equivalent to OC in SONET
SVC	Switched Virtual Circuit
T1 Carrier	A designation of a transmission circuit, referring to its capacity. T1 Carrier has a bandwidth of 1.544 Mbps and T3 Carriers operate at 45 Mbps.
T1	A digital transmission system operating at 1.544 Mbps
TA	Terminal Adapter in the ISDN reference model
TCM	Time Compression Multiplexing
TCP/IP	Transmission Control Protocol/Internet Protocol
TCP	Transmission Control Protocol
TDM	Time Division Multiplexing
TE	Terminal Equipment
TE1	Terminal Equipment 1, ISDN compatible
TE2	Terminal Equipment 2, not ISDN compatible
Telcos	An abbreviation used to refer, collectively, to the telephone companies.
TOP	(Boeing) Technical Office Protocol
UCL	Underwriters Laboratory of Canada
User–level protocols	Protocols, such as TELNET, SMTP, and FTP, which allow you to perform operations or applications on the network.
UUCP	Unix to Unix Copy
V.24	A Physical layer standard
VGA	Video Graphic Adapter

VPN	Virtual Private Network
VSAT	Very Small Aperture Terminal
VT	Virtual Tributary in SONET
WAN	Wide area network; a network that spans great distances.
X.21	A Physical layer standard
X.25	A commercial packet network access protocol that specifies three levels of connections. The X.25 physical level, link level, and packet level correspond to the first three layers of the ISO/OSI model. CCITT standard protocol for network access
X.400	CCITT standard message handling service

TK5105.7/.A7/v.2
Arthurs, David.
Feasibility study of a
national high speed
AZCV c. 1 aa ISTC

DATE DUE - DATE DE RETOUR

AUG AOUT 23 1991	

ISTC 1551 (8/88)

INDUSTRY CANADA/INDUSTRIE CANADA
68785