Communications Research Centre

A DESCRIPTION OF THE ELAT RADAR DISTRIBUTED COMPUTING SYSTEM

by

Eloi Bossé, Jean Caseault and N.R. Fines

This work was sponsored by Department of National Defence,
Research and Development Branch under Project No. 011LA13.

CRC REPORT NO. 1425
OTTAWA, JANUARY 1988

of Canada
of Communications
Gouvernement du Canada
Ministère des Communications

Canada

COMMUNICATIONS RESEARCH CENTRE

DEPARTMENT OF COMMUNICATIONS
CANADA

A DESCRIPTION OF THE ELAT RADAR DISTRIBUTED COMPUTING SYSTEM

by

Eloi Bossé, Jean Caseault and N.R. Fines

(Radar and Communications Technology Branch)

CRC REPORT NO. 1425

January 1988
OTTAWA

This work was sponsored by the Department of National Defence, Research and Development Branch under Project No.011LA13.

CONTENTS

ABSTRACT

The Experimental Low Angle Tracking (ELAT) Radar distributed computing system which supports real-time data acquisition, processing, and distribution for radar applications is described. The underlying motivation behind its development was to build a general-purpose interconnect architecture for the real-time processing of high resolution algorithms to estimate the elevation of low flying targets. The common requirement of these algorithms is the collection of information from a number of different sources and its real-time processing and distribution to various destinations. The ELAT radar processing system consists of a preprocessing and a post-processing section which are connected via a distribution network. Aside from supporting the processing system, the network also interfaces with a large number of other subsystems, each with its own characteristics.

1. INTRODUCTION

This document describes the distributed computing system architecture currently under development at the Communications Research Centre (CRC) to support an experimental low angle tracking (ELAT) radar. The ELAT radar is being developed as part of a research program into high resolution techniques for estimation of the height of low flying targets such as sea-skimming missiles. New techniques under investigation such as the Correlation Height Analysis (CHA) technique [1,2] show promise of solving this problem. The ELAT radar system is being developed to test the CHA technique as well as other high resolution angle estimation techniques in a realistic maritime environment.

The unique feature of the CHA algorithm is that it utilizes separately detected samples of the received wavefront from N different antenna heights. The processing subsystem must combine and process digital samples from the separate elements forming the antenna aperture. The N-channel sampled-aperture system increases by N times the amount of data gathered and thus increases the complexity and size of the processor.

In the ELAT system, a variety of digital processing algorithms is carried out in order to extract the required information from interference and noise. This digital signal processing is characterized by high sampling rates and heavy computational loads.

The handling of the digitized data, namely, their collection, preprocessing, combining, distribution and postprocessing, requires complex automated procedures and special-purpose architectures. The proposed architecture is made-up of a pre-processor subsystem and a post-processor subsystem. The post-processor subsystem is designed to perform the high

resolution algorithms, logging, and graphical presentation of the results. It consists of off-the-shelf equipment such as minicomputers, array processors and general-purpose devices. The pre-processor subsystem is specially designed to provide the first level data processing.

Thus, the computer system architecture must support:

1. The acquisition of large amounts of information;

2. The processing (pre-processing and post-processing, with respect to the distribution network) of the acquired data; and

3. The distribution of data from the sources to the destinations.

A discussion of these three main functions is presented in the following sections with emphasis on data distribution.

2. THE ELAT RADAR DISTRIBUTED COMPUTING SYSTEM

The ELAT radar system uses multiple sensors, which provide both the radar and environmental information. Because the radar system receives interference and noise in addition to the desired signal, it is necessary to perform a range of processing algorithms to extract the required information. Consequently, the ELAT radar system is composed of multiple processing units which are optimized for specific operations (thresholding, integration, doppler filtering, and high resolution techniques).

Because the signal processing is made up of a set of different functions coordinated to perform the complete task, a distributed approach becomes very attractive. The distributed system can be decomposed into the following major modules (see Figure 1):

1. RECEIVER ELEMENTS - detects radar signals using a multi-channel, sampled-aperture, receiving system. Each receiver channel develops analog in-phase (I) and quadrature (Q) signals via coherent quadrature baseband detectors.

2. DATA ACQUISITION - simultaneously digitizes the sampled-aperture radar signals (I and Q) at video rate from all channels.

3. SIGNAL PREPROCESSING- lowers the effective data processing rate by processing the digitized data on a channel by channel basis.

4. DISTRIBUTION NETWORK - collects data and status from various sources and distributes these signals to multiple destinations.

5. SIGNAL POSTPROCESSING - carries out the real-time high resolution tracking algorithms and provides graphical representation and data storage for further off-line analysis.

6. SHIP MOTION MODULE - provides the necessary ship motion information required for the high resolution algorithms.

7. ENVIRONMENTAL INFORMATION - comprises temperature, wind speed and direction, sea-state.

By decomposing the system in the above fashion and providing the appropriate infra-structure, one can alter the operation of the system by changing either the front-end or back-end portion of the system. For example, the radar sensor could be replaced by one which operates in a different frequency band. Also, a new postprocessor could be added in parallel with the existing processor without the need to modify hardware or software.

The system must be able to process large blocks of radar data in real-time. Therefore it must be powerful, fast, and versatile. To meet the experimental requirements, it must also be capable of being easily modified or enhanced. Regarding its implementation, both single processor and multiprocessor-based system designs were investigated. The uniprocessor approach was found unattractive because of its cost-to-performance ratio, its inflexibility, its inability to accommodate future changes and/or enhancements and the low reliability factor associated with a single device [3]. The multiprocessor approach was identified as the most suitable choice for meeting the most critical requirements of the application.

In order to maintain the real-time processing capability and to provide system flexibility, a multilevel tree structure was chosen as the basic computer architecture (see Figure 1).

Some of the inherent benefits associated with this tree architecture are:

1. Reduction in the number of processing elements across the various levels in the tree.

2. Decreased software complexity, as units are self-contained dedicated processing elements each with its own modular software.

3. Potential increase in allocated processing time, due to the possibility of data reduction as data moves through the various levels.

4. Reduction in the amount of soft (recoverable) errors accumulated near the point of acquisition.

5. Fail-soft capability, as real-time diagnosis routines can be performed between adjacent levels. The system is capable of masking out bad data resulting from software errors and/or hardware malfunction.

6. Accommodation of new devices such as faster microprocessors and special-purpose LSI/VLSI hardware modules at a later date.

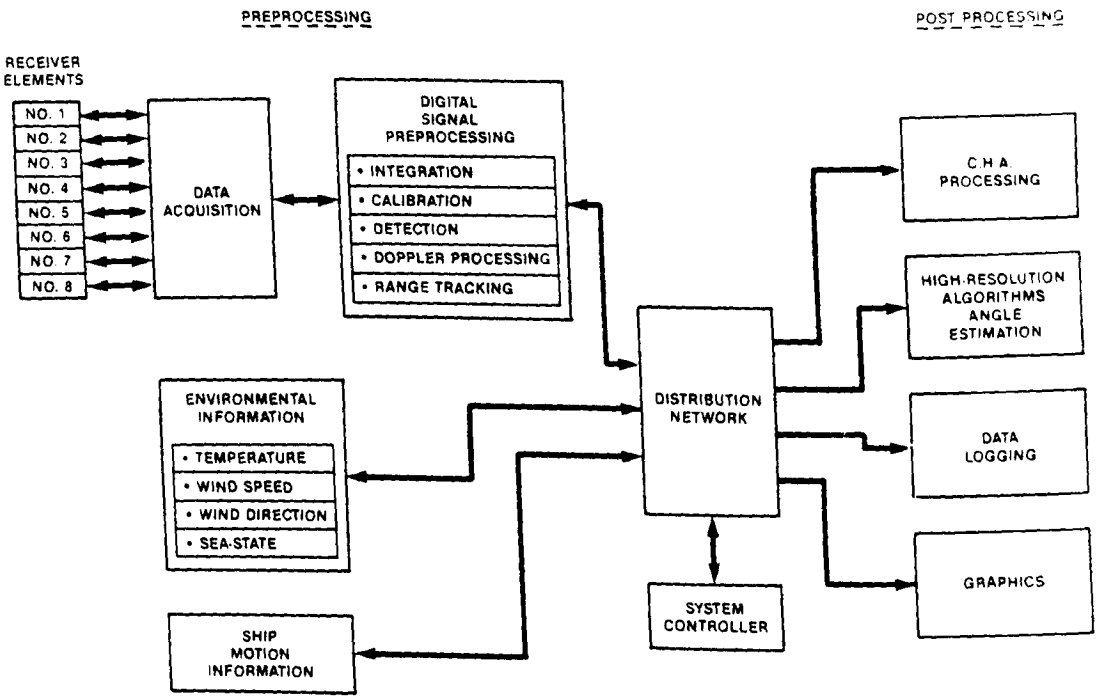

Figure 1 - The ELAT Distributed Processing System

3. DATA ACQUISITION

The ELAT radar is a monostatic system with separate but colocated receiving and transmitting antenna systems. The principle of operation is to illuminate (limited flood lighting) the target with pulsed output power from the transmitting antenna and to receive the reflected signal on a linear vertical array antenna. Each element is connected to an individual receiver for conversion of the microwave frequency information to the video frequency band. Each receiver includes quadrature (Q) and in-phase (I) detectors which are followed by analog to digital conversion, in each channel. Using this receiving arrangement, discrete samples of the electromagnetic field across the vertical antenna aperture can be converted to the video band for further signal processing. An array of this type is sometimes known as a "Sampled Aperture Antenna".

Another acquisition mode is also provided whereby the receiver channels are combined to form a single narrow receive beam. In this mode, only two analog-to-digital conversions are required on the combined I and Q channels.

The ELAT radar uses other sensors to provide environmental information and consequently other analog to digital conversions are also associated with it (Fig. 2).

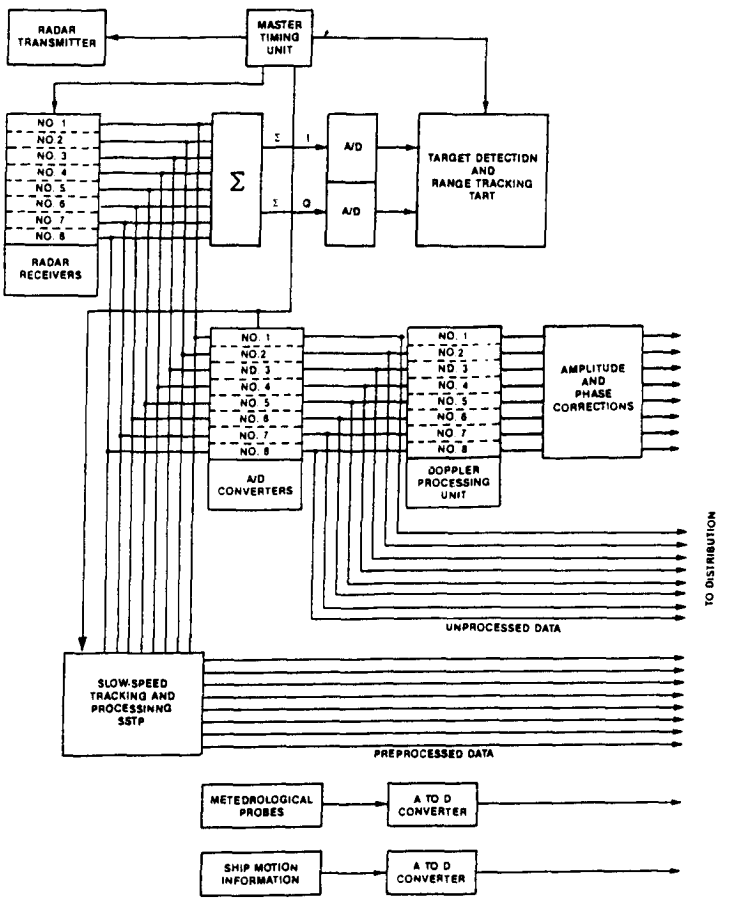

Figure 2 - ELAT Data Acquisition and Preprocessing System

4. DATA PROCESSING

To provide real-time processing and because of the multiplicity of complex signal processing algorithms, the overall processing task is split into a number of well defined functions. Optimized units can then be used to perform these functions. A multilevel tree structure was chosen to link all these units together. By using multiple processors,

the system software complexity is reduced significantly. Each processing unit may have its own operating system. The architecture is made up of a preprocessing and a postprocessing subsystem. Based on the environment where the system is to be located, namely aboard a ship, it is assumed that the separation distance between processing sections may vary from 15 to 60m.

4.1 Preprocessing Section:

Signal preprocessing is required for a number of reasons as listed:

1. The amount of data to be analyzed in the CHA processor must be reduced. This is necessary because the CHA algorithm is computationally demanding so that a new estimate of height cannot be produced after every pulse transmission.

2. The signal to noise ratio (SNR) must be high to obtain a good estimation of height using the CHA algorithm. This requires that the data be coherently integrated.

3. Sea clutter returns must be eliminated.

4. Calibration techniques must be implemented to compensate for systematic errors.

These requirements are satisfied by providing two units: the target acquisition and range tracking unit (TART) and the doppler processing unit (see Figure 2). In the present radar design, a pulse repetition frequency (PRF) of 2 KHz and a pulse width of 200 nsec are used. For the 200 nsec pulse the radar range cell resolution is 30 m. The major signal processing load occurs in the TART unit which requires scanning through 500 range cells to obtain a maximum unambiguous range of 15 Km. The signal processing operation in this unit involves filtering the data from each range cell in a continuous manner to remove clutter and detect the presence or otherwise of the target. The determination of the presence of a target is complicated by the fact that the radar itself is in motion. The doppler frequency shift of the clutter varies according to platform motion. A programmable digital filter bank whose responses can be shaped, depending on the prior knowledge of the velocity of the platform and the pointing direction of the antenna, has been designed [4]. The TART unit works on combined I and Q signals from all 8 channels and transfers the range information via the Master Timing Unit (MTU) to the 8 data acquisition channels for selective range gating. The TART unit must process at a real-time computation rate of 320 million operations/second [4] to filter each of the 500 range cell data.

Since the target range is found using the TART subsystem, it is only necessary to perform the doppler processing on a single range cell.

Therefore, the doppler processing unit works at a slower speed than the TART unit and can be done in software. A programmable array processor can satisfy the performance requirements [5]. However, the doppler filtering must be performed in parallel for each of the eight acquisition channels. Data are coherently integrated in blocks of N returns sampled at the selected range interval. Sampling synchronization is performed by the MTU. In addition, prior to doppler filtering, the same processing unit compensates for systematic errors in the in-phase and quadrature channels.

Finally, if the high speed channels fail (main system), a back-up system is provided but the 2KHz real time rate is not preserved. This back-up system, the slow speed tracking and processing system (SSTP), can implement all the functions such as data acquisition, target detection, range tracking and doppler filtering on its own (Fig. 2). All the remaining devices in the preprocessing section are made up of various dedicated hardware or small processing units to provide environmental conditions, ship motion and other useful information used in the high resolution algorithms for accurate determination of target elevation.

4.2 Postprocessing Section:

The postprocessing section is dedicated to the real-time implementation of the high resolution algorithms. These algorithms process a large amount of data coming from different sources. These sources provide radar and environmental information (sea-state,temperature, etc...) which demand many computations. In particular, the CHA algorithm requires about 50,000 operations/pass [6]. The postprocessing section provides graphical representation of results and also data storage for further off-line analysis.

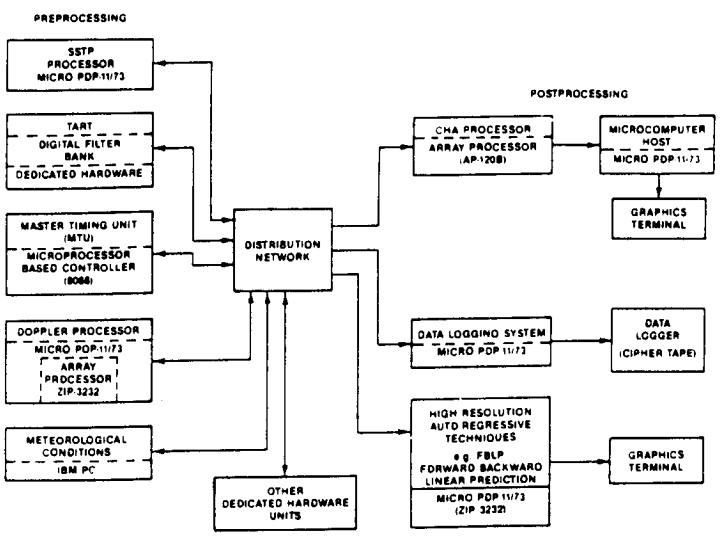

Figure 3 - ELAT Computer Systems

In the recent development of the ELAT radar system, two algorithms have been implemented: the CHA and the forward backward linear prediction method (FBLP) [7]. Both use general purpose computers, array processors, and other off-the-shelf devices (Fig. 3). The preprocessing and the postprocessing can be easily expanded allowing parallel processing of several algorithms.

5. DATA DISTRIBUTION

The main requirement for the system distribution network is to support fast transfers of large blocks of information (data, status, and control), from a large number of input devices to a variety of output devices, in a reliable fashion. Given this requirement, a unique structure had to be developed.

The design of the system distribution network is based upon a study by Anderson [8] of the interconnection transfer strategy, control method, and path structure. In particular, emphasis is placed on studying the system architecture as seen by consideration of the type of communication allowed between the various modules, rather than the structure of the modules themselves [9]. Specifically, the following areas are considered:

1) The physical characteristics of devices to be connected;

2) The data transfer rate requirement;

3) The type of buses to be used (electrical characteristics);

4) The rules governing the transfer of data and control (communication protocol);

5) The type of connection (interconnection structure system).

5.1 Devices Classification

To facilitate the physical definition of the bus architecture, a classification of system devices is carried out. The various devices in the system are classified into three categories: producing devices, consuming devices, and controllers. The definition of each of the device categories is as follows:

i) Producer/Talker/Driver

A device which is capable of generating data or control information and has the potential to put the data on the bus either directly or indirectly.

ii) Consumer/Listener/Receiver

A device which is capable of removing data or control information from the bus at the specified transfer rate.

iii) Controller

A device which is capable of changing the control parameters (e.g. transfer rate, number of words) and/or configuration (e.g. which producing device is in command) of the bus.

5.2 Data Transfer Rate Requirement

In the proposed system, the worst case condition occurs when a data sample is recorded from each channel every radar pulse repetition interval (PRI).

As there are eight channels, the total number of words from the data acquisition subsystem is 8 x 2 x 1 = 16 words. It is also assumed that all of the remaining sensors may generate up to a total of 112 words every radar PRI. There is also a need to transfer an additional 128 words of preprocessed data after N PRIs resulting in a 1 megaword/sec (1 Mwords/sec) transfer rate requirement for a PRI of 500 usec.

5.3 Type of Buses

In consideration of the anticipated data rate (1 Mwords/sec) and the distance (greater than 15 meters) over which data has to travel to the consumers, one can search for an appropriate standard which meets these requirements. In general, there are two basic methods for electronic communication between devices in a computer system: single ended transmission and differential transmission. Each of these methods will be briefly described.

i) Single ended transmission

This method utilizes a single transmission line and is only suitable for short distances and slower data rates since, as the line length increases, the signal quality is degraded by noise. This is acceptable for the producer bus since the interface units are connected together to the controller via a backplane.

ii) Differential transmission

This method utilizes two signal lines for transmission. Thus unwanted signals appear as common-mode levels and are rejected by the differential line receiver. This is more appropriate for the consumer bus which needs a high bandwidth (16 Mbits/sec) and long distance transmission capability (60

meters). Because of these requirements, the RS-422A is preferred over other standards which failed to meet at least one of the requirements (speed for Ethernet, IEEE488, and RS-232C; distance for RS-232C, and IEEE488; and expansion for IEEE488).

5.4 Communication Protocol

In order to transfer data and control information in a reliable and orderly fashion, it is necessary to establish a set of rules which govern the information transfers on the bus (a communication protocol). In the proposed system, there are two main global buses: an input bus with multiple producer interface units and an output bus (fig.4). In this section, the key features associated with both buses will be defined.

INPUT BUS

Data is collected and assembled in blocks into the producer interface unit (PIU). At each PRI, the bus control unit (BCU) interrogates, sequentially, every interface unit and receives information in a fixed format block. The BCU selects devices with these control lines:

- chip select (CS)
- address
- read signal (RD)
- control line

If necessary, after the fixed format block has been transferred to the consumers, the BCU can send back information (control, status) to the target device by providing:

- chip select (CS)
- address
- write signal (WR)
- control line

OUTPUT BUS

The suggested communication protocol for the output bus is a very simple one due to the fact that it does not require handshake signals. Whenever data becomes available, the bus control unit via the bus interface unit performs the following steps:

i) it activates the data block transfer ready line,

ii) outputs data on the bus followed by a write signal,

iii) repeats this process until all data are sent,

iv) reads the request lines from each of the consuming devices to detect a consumer request, and

v) resets the data transfer ready line.

A list of the technical specifications of the output bus is given in the Appendix.

5.5 Interconnection Structure Topology

The interconnection structure, which has been selected, combines the best features of the direct method (such as is found in single bus architecture) with the best features of the indirect method, in which address transformation is an intervening operation (such as is found in star interconnection structures) [10]. The resultant hybrid interconnection structure employs a single bus for the input devices and another single bus for the output devices. The use of single buses is dictated by reliability considerations outlined by Powell [11], as well as speed of operation. Between the two buses, a bus controller effectively provides a star configuration [12] by supporting centralized routing (transfer control method) with dedicated paths (transfer path structure) to the two busses.

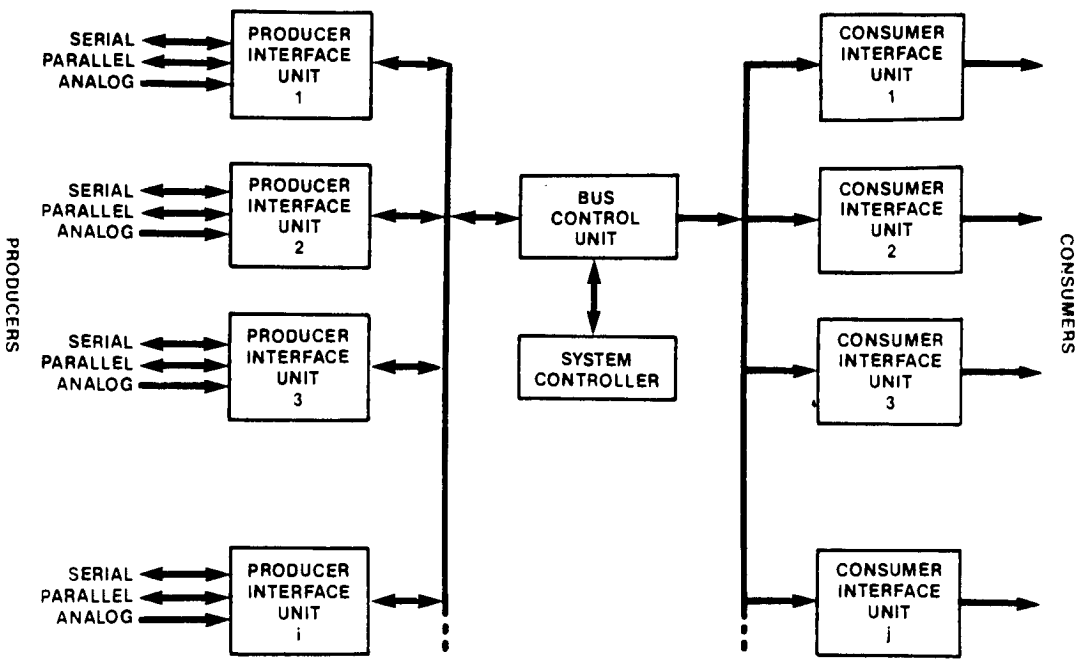

Figure 4 - the ELAT Distribution Network Architecture

The whole distribution network (Fig. 4) can be considered as a loosely coupled multiprocessor system with each processor having its own local program and data memory as well as I/O. In addition, some of the processors share a common communication memory in order to minimize the

intermediate transfers of data between the processors and facilitate the fast transfer of data. By using multiple processors, the system software complexity is reduced significantly. Each of the processors has its own simple executive which manages the resources and transactions within its domain. Thus, by increasing the amount of hardware, the need for a complex real-time multi-tasking executive (which would have been required on a uniprocessor system) has been eliminated.

The hub of the system is a dual-processor subsystem consisting of a master-microcomputer and a slave-microprocessor. This dual-processor configuration initiates, monitors and controls all the transactions occurring within the domain of the distribution network. The microcomputer master is referred to as the system controller (SC) whereas the specially designed microprocessor slave system is known as the bus control unit (BCU). Aside from the SC, the distribution network supports two other types of devices, namely producers which are devices that can output data on the bus and consumers which are devices capable of removing data from the bus.

Each of the producing devices has its own producer interface unit (PIU) to match its unique characteristics (serial, parallel, analog) to those of the producer bus. This permits simultaneous reception of data from all of the producing devices. The transfer of data from the producers to the consumers, and possibly between producers, is initiated and controlled via the bus control unit. In order to facilitate this transfer, the producer bus is bidirectional (Fig. 4).

The consumer bus is a unidirectional, high-drive bus with multiple receivers which are electrically isolated from the producing devices. Data on this bus is sent in block-mode DMA transfers. This is the most appropriate mode for a pulse radar. Any interested consumer has the ability to pick up any or all of the data associated with a given block via its consumer interface unit (CIU).

The main role of the distribution network is to facilitate information movement among producers, consumers and the system controller in a fast and reliable fashion. Therefore, it is necessary to identify the interface and control elements that are involved in this process. For each entity, a description will be presented, highlighting its major responsibilities, its functional block diagram, as well as a high-level description of its hardware and software components.

5.5.1 Bus Control Unit (BCU) -

The SC, being an off-the-shelf minicomputer, lacks the appropriate hardware modularity and software flexibility to solely manage the complete distribution network. Aside from the multi-interface requirement to multiple sources and destinations, the system is faced with a stringent requirement on the data transfer rate (a high data transer rate per device, coupled with data gathering from multiple devices). These condi-

tions made the performance unattainable with a commercially available device. Therefore, a custom-made subsystem had to be developed which could accommodate the various devices of the distribution network involved in the transaction, by interfacing to the SC on one side, and to the peripheral devices on the other. This led to the development of the BCU, which was designed to be a slave to the SC by executing its commands as well as managing the data packet transfers from the producers to the consumers.

As shown in figure 5, the BCU subsystem consists of the following main functional blocks: processing, DMA, memory, decoding, and I/O. A brief description of each of these elements follows:

1. The processing subsystem consists of the Intel 8086 16-bit microprocessor which provides the necessary intelligence for the BCU.

2. The DMA subsystem utilizes the Advanced Micro Devices 9516 for controlling direct memory access between producers and consumers in a block mode operation. Data located in various memory locations of different producers are collected and then transferred to the appropriate destination. In the present configuration, the maximum data block size is 256 words which are transferred directly from the various producers to the consumers at an effective transfer rate of 1 Mwords/sec.

3. The BCU memory is used to store message and control information (commands to the BCU or controls to external devices). It consists of 4Kwords of EPROM for program memory and storage of constants, and 2Kwords of static read/write memory for data and variable storage.

4. The decoding subsystem (memory select logic and I/O select logic) provides control logic for selecting the appropriate memory and peripheral devices by enabling either the input or output of the target device.

5.5.2 Producer Interface Unit (PIU) -

The producer interface is concerned with the interconnection between the producing devices and the distribution network. The data generated by individual producers are collected by a Producer Interface Unit (PIU) which is capable of simultaneously servicing up to four producers. Each PIU has the following functions:

1. To provide the appropriate interface to match the characteristics of each producing device,

2. To assemble data in a special format for the BCU,

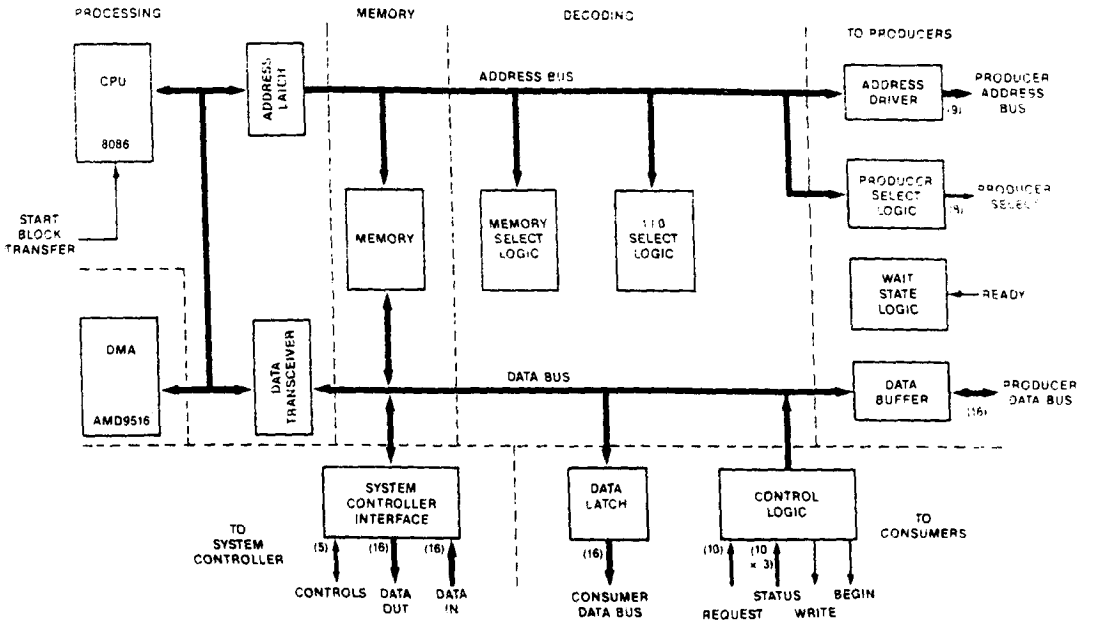

Figure 5 - The Bus Control Unit (BCU)

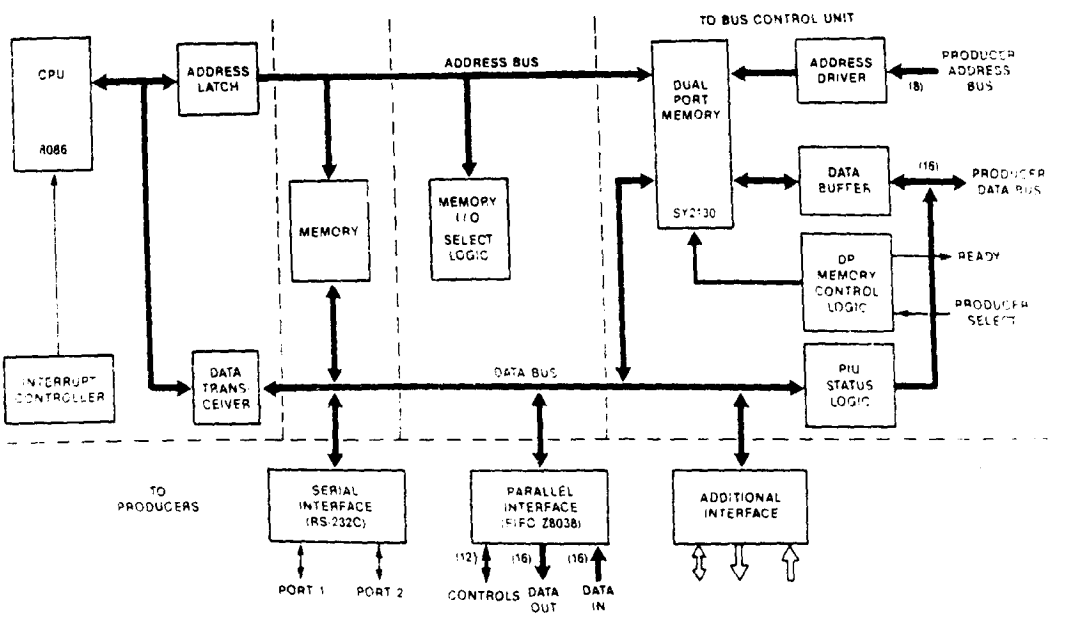

Figure 6 - Producer Interface Unit (PIU)

3. To provide a special interface to the BCU input bus,

4. To accept information from the SC via the BCU and send it to its producing device(s).

As shown in Figure 6, each PIU consists of the following main functional blocks: a processing section, a memory section, a decoding section (memory and I/O select logic), and an I/O section. The functionality of the first three sections is similar to that of the BCU. Most of the components used by the BCU, such as the Intel 8086 microprocessor, were also incorporated into the PIU design. The remaining section is unique to the PIU and is divided in two parts: the producer interface, and the BCU interface, which will be described next.

i) The producer interface

Each PIU is capable of servicing four producers, each of which may have different input/output characteristics. In the current design, the following interfaces are supported by the PIU: one multi-purpose parallel interface, two RS-232C type serial interfaces, and one producer interface unit (e.g. analog) which is accommodated via a reserved wire-wrap section. The multi-purpose parallel interface has a two-level structure. The first level connected to the outside world, which consists of 16-bit latched output lines and 16-bit buffered input lines. In addition a number of general-purpose control lines are used to synchronize data transfers. The second level, which is based on dual Zilog 8038 FIFO chips, multiplexes the first level input/output lines, and manages the transfer of control and data blocks to/from the producer from/to the PIU processor.

Each of the four producer interfaces is capable of interrupting the PIU processor and has a specific priority level associated with it. This priority level is hardwired in the current configuration and can be changed by modifying its location on the priority interrupt controller, which is part of the PIU processing section.

ii) The BCU interface

The main component of the BCU interface is a true dual-port read/write memory. This memory can be accessed simultaneously by both the BCU and the PIU and is used for transferring control and data information between them. The dual-port memory is logically segmented into two data blocks which permits the PIU to deposit data into one block while the BCU is removing previously written data from the other block. The control section on the other hand, has two fixed memory blocks, one for each direction. The PIU writes

while the BCU reads and vice-versa. In addition, the dual-port memory contains hardware semaphores for synchronizing access to the shared resource. Since multiple PIUs must be connected to the common bus, bidirectional tri-state buffers are provided on the outputs of the dual-port memory. In the event that both the PIU and the BCU attempt to access the same control memory location simultaneously, the BCU will be given priority and the PIU access will be delayed until the BCU has completed its operation.

There are two dedicated control lines (one for the BCU and the other for the PIU), which are used to signal a request for an action. The BCU utilizes its control lines (one for each producer cluster) in order to gain access to the particular PIU dual-port memory. In addition, each PIU uses its dedicated control line to signal to the BCU that it has a special message to send to the system controller.

5.5.3 Consumer Interface Unit (CIU) -

The consumer interface supplies the necessary logic to accept data from the BCU and transfer it to its consumer device. The major tasks of the consumer interface are:

1. To provide speed matching by buffering the incoming data,
2. To electrically isolate the consumers and producers,
3. To provide the decision-making capability required to verify whether information is needed by the particular device.

As shown in figure 7, the consumer interface comprises two sections: the transceiver section (driver and receiver boards), and control logic and storage section (control board). The transceiver section encodes and decodes the RS-422-A signals (data and control) via the HP-2602 opto-isolators, and transfers the information to the control logic and storage section (FIFOs).

The main consumer interface board consists of a control processor (8086, CPU) section which detects the beginning of a new block, removes the data from the bus, determines whether its consumer requires this block of data, verifies the validity of the incoming data, deposits the data into memory, and informs its consumer regarding the status of this data. In addition to the control processor section, the consumer interface also contains a memory buffer (FIFO) which temporarily stores the incoming data for later use by the consumer.

The communication protocol between the BCU and consumers has been set, such that the data is always available on the bus together with its clock. Any interested consumer may access this data. Thus, a consumer cannot interrupt the transfer. However, to verify consumer status, three status lines, as well as a single request line, are provided for the

consumer to inform the BCU about the system activities. Whereas each consumer has a dedicated request line, the three status lines are common to all consumers. If one or more request lines are activated, the BCU selects the appropriate consumer request condition by enabling its own status buffer.

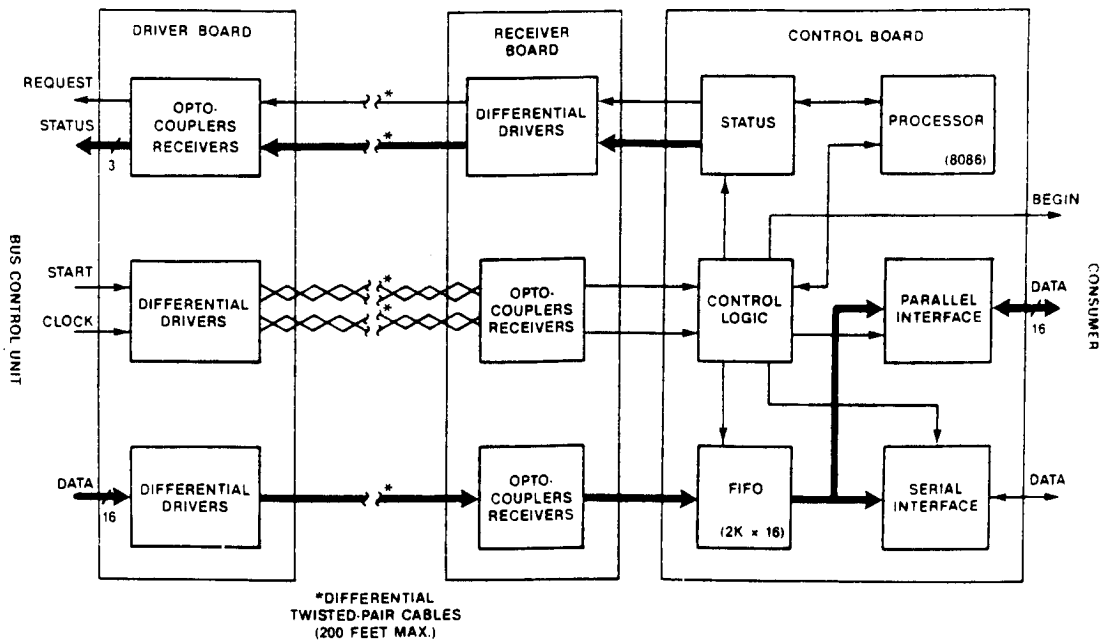

Figure 7 - Consumer Interface Unit (CIU)

5.6 Operation of the Distribution Network

There are two modes of communication available within the distribution network: a message and a packet. A message may include two parts: a single-word command and a multiple-word block of data information consisting of up to eight words (a message does not need to have a multiple-word structure data packet and may consist of the command alone). The command word defines the source and destination as well as the procedure required for the ensuing transaction. A packet, on the other hand, contains more information (up to 256 words) which includes data, status, and control information. This packet is assembled from the various producers and transferred to the consumers on each PRI.

The dual processor configuration SC and BCU, initiates, monitors and controls all the transactions occuring within the domain of the distribution network. The BCU software is made up of two major routines: the MAIN routine and the INTERRUPT routine. The former is invoked by software commands, and handles all functions which can extend over more than one PRI. The latter, on the other hand, is invoked by an external event

(hardware), and manages the data transfer which must be completed within the time of a single PRI (real-time).

The MAIN routine is responsible for the message communication mode such as failure detection, special-purpose data transfers, control and status transfers, and communication with the system controller. This routine consists of a test loop which attempts to detect an exception condition which will force an action to be carried out. After servicing the particular exception condition, a return is made to the MAIN routine.

The INTERRUPT routine performs all operations associated with the data packet transfer mode. These include: sending the start signal; transferring data block header information to consumers; reading and transferring the status of producers to the consumers; and collecting the data from the various producers, arranging it in a block format and transferring it to the consumers. In addition, the producer and consumer requests are read to detect an exception condition and set the appropriate flags, if needed, to inform the main routine, of this condition.

The two modes of communication mentioned above will be described in the following sections.

5.6.1 Message Communication Mode -

The message communication mode is achieved by the MAIN routine residing in the BCU. The SC via the BCU MAIN routine controls, and monitors the system. The SC communicates with the BCU via a dedicated parallel interface using a single encoded command word. The BCU accepts a command, decodes it, executes it, and sends the obtained information back to the SC.

The SC utilizes a basic command word to communicate with the BCU, whereas the BCU communicates with the SC using a special command word. The SC always begins a transaction with a basic command word containing seven fields which direct the BCU to carry out specific functions. The seven fields are:

BIT 15 — The SC is involved/not involved in the transaction,

BIT 14 — Direction of the transaction with respect to the SC (in/out),

BITS 13-12 — Other device type involved in the transaction (BCU, producers or consumers),

BITS 11-8 — Device number, or command type if the BCU is the other device involved in the transaction,

BIT 7 — Command extension bit. If set, more data is to follow with a subsequent command word,

BITS 6-4 Number of information words to be expected following the command word (from one to eight words),

BITS 3-0 Message type, if the PIU (i.e. producer) is the other device involved in the transaction.

Whenever the BCU wishes to inform the SC of the occurence of a special event, it uses the special command word. The special event represents either status information as a result of the SC inquiry or an urgent condition (request) occuring somewhere in the system. The special word is configured to inform the SC of the specific device status/condition. This word is divided into four fields:

BIT 15 This command has been originated by the BCU or not,

BIT 14 The source is either a producer type or a consumer type,

BITS 13-10 Status or request,

BITS 9-0 The producer's or the consumer's bit identifier. Every source that meets the conditions specified above will have its allocated bit set.

Using the basic and special command words, the complete transaction communication protocol can be developed. This protocol is based on a master-slave relationship in which the SC initiates the command and the BCU executes it. The basic command word is used by the SC in every transaction initiation. Once started, the transaction is normally carried to completion. The BCU then returns to its wait state.

If during the execution of the basic command, the BCU detects an exception condition, it can interrupt the current command execution and inform the SC regarding this event using a special command word. The current transaction is then suspended until the special command word has been received and acknowledged by the SC. If the SC determines that an urgent action is required regarding this special command, it can suspend the execution of the basic command by issuing a new command. This command has a higher priority than the previously unfinished basic command and must be completed before resuming the execution of the unfinished basic command. By using a two level command structure, the SC can rapidly respond to an urgent request without destroying the data associated with the previously initiated command. Since all commands must either originate from the SC or be approved by it, the SC, at an one time, has the global picture of the whole system. The system operator, through use of the system monitor, can easily determine the status of the system and control its operation by initiating the appropriate command via the operator keyboard.

5.6.2 Packet Communication Mode -

The packet communication mode is invoked by an external interrupt

(PRI) and is managed by the BCU INTERRUPT routine. In order to be able to transfer large blocks of data coming from different producers, the design strategy was to build a distribution network with a variable memory organization and having data banks shared between the BCU and thePIUs. Many special features were incorporated into the design to accelerate address allocation to each data word: bank switching, shared addressing, programmable selection of the PIUs, and the DMA "latch while reading" option.

First, the data are collected from a producer when the PIU is interrupted by it. Data are immediately stored at their allocated addresses, inside the data block n.a (n is the PIU number and a or b designate the data bank) of the dual port RAM (shared by the BCU and the PIU). Once the entire group of data, designated for a particular PIU, has filled the allocated addresses inside the data block, the PIU switches the data block flag bit, and begins writing into the data block n.b, allowing the BCU to read block n.a. Designing the system so that a block of data is only available to be read by the BCU after it has been completely written by the PIU, decouples the BCU or PIU operations so that they can be carried out independently of each other, at maximum speed and without the use of complex logic control.

After the data has been sorted by the PIUs, the BCU can initiate data block transfers. This is done in real-time at 2 KHz. When the BCU receives a start signal, the DMA takes control of the bus, by supplying addresses, latching data words directly into the consumer port during the read cycle, thus allowing the bus to work at the maximum speed.

The data sent to consumers contains two types of information: the header information which contains the block type, status, and information coming from the BCU memory, and the data block itself, collected from various PIUs. The data block size is completely programmable; it contains a variable number of control words (BCU memory) plus a variable number of data words (PIU DPRAM). Because the BCU memory is adjacent to the PIU data, the BCU data, and later, the PIU data are transferred easily by the DMA. A special feature has been added to the BCU to increase the flexibility of the block transfers: the PIU data selection is done on a word basis instead of a block basis. Therefore, a data block is a selection of independent words, each one obtained from any PIU. The block, as a whole, becomes a group of PIU data packets. To make this possible, every BCU data block configuration must provide its own sequence of PIU selection. Up to 14 block configurations can be programmed, with immediate access to two of them by the DMA. Access to the other blocks is obtained by changing the starting address inside the DMA.

Finally, when the DMA has put the data on the bus, and has latched it onto the consumer port, the consumer driver board amplifies the signals and sends them to the consumer interface unit, 60 meters away at 16 Mbits/sec (or 1 Mwords/sec). The CIU, after receiving the block, decodes the block type word and the status word (status of PIU data) and

decides whether it should keep the block or not. If the test is positive, the data can be distributed to its allocated consumer, at the consumer speed.

6. SUMMARY

This report describes a distributed computing system which facilitates reliable high-speed data transfers and processing for systems with multiple sources and destinations. To support this requirement, an elegant system bus structure was developed which is called a multi-level tree interconnection structure. Such a structure provides many benefits derived from both the single bus as well as the star type of routing mechanism. It shows great promise for expansion to multibus star structures for other applications which could be characterized by the producer/consumer model described here.

In summary, the distribution network can be characterized by the following features:

1. Block size of 256 words (each word is 16 bit wide) with a block repetition interval of 500 microseconds,

2. An effective transfer rate of 1 Mwords/sec.,

3. The input interfaces must accommodate various devices and data structures (serial, parallel, analog),

4. A distance between devices of up to 60 meters,

5. The ability to support a large number of input devices (up to 32),

6. The ability to support a large number of output devices (up to 10),

7. A highly reliable distribution network characterized by graceful degradation; the distribution network being capable of operating in a degraded fashion after sustaining a single-point failure,

8. Modularity; common elements to reduce development costs, to increase reliability, and to decrease system mean-time-to-repair through a board level replacement,

9. The ability to control the operation of the input and output devices remotely,

10. A communication protocol which is independent of the physical characteristics and operating status of the various devices.

11. Neither an input device nor an output device can delay or interrupt the transfer.

Improvements in the system interconnection structure will ensue as technology develops. The system bandwidth is directly related to these improvements, as is the speed of operation of the Bus Control Unit. The BCU, at present, handles all of the DMA control, but in later improvements, it is conceivable that the other modules on the buses will assume the DMA function. In addition, increased data and processing rates, data quantity and distances over which data are transmitted can also be considered.

7. ACKNOWLEDGEMENTS

This work was undertaken by the Communications Research Centre under the Department of National Defence Project 011LA13. The authors wish to thank the following personnel who participated in the implementation of this system: Eli T. Fathi, Denis Lamothe and I.P. Sharp Assoc.

8. REFERENCES

[1] Litva, J., "A New Low Angle Tracking Technique", CRC Report No.1335, Ottawa,May 1980.

[2] Rook, B.J., Litva, J., "An Improved CHA Algorithm for Tracking Low Angle Targets", CRC Report no.1356, Ottawa, January 1982.

[3] Fathi, E.T., Krieger, E.D., "Multiple Microprocessor Systems: What, Why, and When", IEEE Computer, Vol. 16, No. 3, pp. 23-32, March 1983.

[4] Roy, D., Menard, P., Da Sylva, P., "Development of a Programmable Finite Impulse Response Filter for Sea Clutter Rejection", final report, Interactive Circuits and Systems Ltd., Ottawa, March 1986.

[5] Bosse, E., "Performance Requirement Analysis of a Multichannel Doppler Processing Unit (MDPU) for the ELAT System", CRC report to be published in 1988.

[6] Bosse, E., "Real-Time Implementation of the CHA Algorithm Using an Array Processor", CRC report 1422, December 1987.

[7] Ulrich, T.J, Clayton, R.W, "Time Series Modelling and Maximum Entropy", Phys. Earth Planet. Inter., Vol. 12, pp. 188-200, Aug. 1976.

[8] Anderson, G.A., Jensen, E.D., "Computer Interconnection Structures: Taxonomy, Characteristics, and Examples", Computing Surveys, Vol. 7, No. 4, pp. 197-214, Dec 1975.

[9] Thurber, K.J., et al, "A Systematic Approach to the Design of Digital Bussing Structures", AFIPS Conf. Proc., Vol. 41-II, 1972 FJCC, pp. 719-740, 1972.

[10] Fathi, E.T., Fines, N.R., "Real-Time Data Acquisition, Processing and Distribution for Radar Applications", Real-Time Symposium, Austin, Texas, Dec 1984.

[11] POWELL, D.R., "Dependability Evaluation of Communication Support System for Local Area Distributed Computing", Proc. 12 Annual International Symposium of Fault Tolerant Computer,Santa Monica CA,June 1982.

[12] Weitzman, C., "Distributed Micro/Minicomputer Systems, Structure, Implementation and Applications", Prentice Hall, Englewood Cliffs, NJ, 1980.

[13] Fathi, E.T., Caseault, J., Bosse, E., "A High-Speed Distribution Network for Multiple Sources and Destinations", IEEE MINI/MICRO WEST 85 CONFERENCE, Los Angeles, Feb 85.

APPENDIX

OUTPUT BUS TECHNICAL SPECIFICATIONS

A list of the technical specifications for the output bus is given below:

i) Electrical data communication standard: RS-422A

ii) Data transfer rate: up to 1 Mwords/sec or 16 Mbits/sec

iii) Length of transmission line: 60 meters maximum
Note: Based on new system requirements, the above two specifications could be modified according to the transfer rate/transmission length graph shown in figure A.

iv) Connector type: 64 PIN DIN connector allowing 32 differential pairs with the following designation:

A- 16 data pairs
B- 2 control lines
C- 4 status lines going back to the BCU
D- 10 extra pairs for future needs

v) Number of consuming devices: 10 maximum

vi) Bus width: parallel 16-bit wide

vii) Number of control lines: 2 minimum and 12 maximum

viii) Transfer method: DMA block transfer mode

ix) Block format:

A- Word 1
block type identifier:
1. initialization
2. raw data
3. preprocessed data
4. system controller software information or controls to consumers
5. target acquisition information
6. others

B- Word 2
update status word: old/new data (data validity)

C- Words 3 to 8
6 control words

D- Words 9 to 24
raw data (16 words)

E- Words 25 to 128
104 additional information words coming from other producing devices:

1. range
2. antenna height
3. inertial measurements

The format is similar for the preprocessed data (128 words), except that preprocessed data is 32 words wide and additional information is 88 words wide. Usually, for each PRI, 128 words of raw data (unprocessed data) have to be transferred, and based on the integration period, an additional 128 words of preprocessed data have to be carried by the bus.

UNCLASSIFIED

SECURITY CLASSIFICATION OF FORM
(highest classification of Title, Abstract, Keywords)

DOCUMENT CONTROL DATA

(Security classification of title, body of abstract and indexing annotation must be entered when the overall document is classified)

1. ORIGINATOR (the name and address of the organization preparing the document. Organizations for whom the document was prepared, e.g. Establishment sponsoring a contractor's report, or tasking agency, are entered in section 8.)

Communications REsearch Centre
3701 Carling Ave., P.O. Box 11490, Station H
Ottawa, Ontario, K2H 8S2

2. SECURITY CLASSIFICATION (overall security classification of the document including special warning terms if applicable)

UNCLASSIFIED

3. TITLE (the complete document title as indicated on the title page. Its classification should be indicated by the appropriate abbreviation (S,C,R or U) in parentheses after the title.)

A DESCRIPTION OF THE ELAT RADAR DISTRIBUTED COMPUTING SYSTEM (U)

4. AUTHORS (Last name, first name, middle initial. If military, show rank, e.g. Doe, Maj. John E.)

BOSSE, E., CASEAULT, J., and FINES, N.R.

5. DATE OF PUBLICATION (month and year of publication of document)	6a. NO. OF PAGES (total containing information. Include Annexes, Appendices, etc.)	6b. NO. OF REFS (total cited in document)
	25	13

7. DESCRIPTIVE NOTES (the category of the document, e.g. technical report, technical note or memorandum. If appropriate, enter the type of report, e.g. interim, progress, summary, annual or final. Give the inclusive dates when a specific reporting period is covered.)

CRC Report

8. SPONSORING ACTIVITY (the name of the department project office or laboratory sponsoring the research and development. Include the address.)

Defence Research Establishment Ottawa
Department of National Defence
Ottawa, Ontario, K1A 0Z4

9a. PROJECT OR GRANT NO. (if appropriate, the applicable research and development project or grant number under which the document was written. Please specify whether project or grant)	9b. CONTRACT NO. (if appropriate, the applicable number under which the document was written)
011LA13	
10a. ORIGINATOR'S DOCUMENT NUMBER (the official document number by which the document is identified by the originating activity. This number must be unique to this document.)	**10b. OTHER DOCUMENT NOS. (Any other numbers which may be assigned this document either by the originator or by the sponsor)**

11. DOCUMENT AVAILABILITY (any limitations on further dissemination of the document, other than those imposed by security classification)

(X) Unlimited distribution
( ) Distribution limited to defence departments and defence contractors; further distribution only as approved
( ) Distribution limited to defence departments and Canadian defence contractors; further distribution only as approved
( ) Distribution limited to government departments and agencies; further distribution only as approved
( ) Distribution limited to defence departments; further distribution only as approved
( ) Other (please specify):

12. DOCUMENT ANNOUNCEMENT (any limitation to the bibliographic announcement of this document. This will normally correspond to the Document Availabilty (11). However, where further distribution (beyond the audience specified in 11) is possible, a wider announcement audience may be selected.)

UNCLASSIFIED

SECURITY CLASSIFICATION OF FORM

13. ABSTRACT (a brief and factual summary of the document. It may also appear elsewhere in the body of the document itself. It is highly desirable that the abstract of classified documents be unclassified. Each paragraph of the abstract shall begin with an indication of the security classification of the information in the paragraph (unless the document itself is unclassified) represented as (S), (C), (R), or (U). It is not necessary to include here abstracts in both official languages unless the text is bilingual).

The Experimental Low Angle Tracking (ELAT) Radar distributed computing system which supports real-time data acquisition, processing, and distribution for radar applications is described. The underlying motivation behind its development was to build a general-purpose interconnect architecture for the real-time processing of high resolution algorithms to estimate the elevation of low flying targets. The common requirement of these algorithms is the collection of information from a number of different sources and its real-time processing and distribution to various destinations. The ELAT radar processing system consists of a preprocessing and a postprocessing section which are connected via a distribution network. Aside from supporting the processing system, the network also interfaces with a large number of other subsystems, each with its own characteristics.

14. KEYWORDS, DESCRIPTORS or IDENTIFIERS (technically meaningful terms or short phrases that characterize a document and could be helpful in cataloguing the document. They should be selected so that no security classification is required. Identifiers, such as equipment model designation, trade name, military project code name, geographic location may also be included. If possible, keywords should be selected from a published thesaurus, e.g. Thesaurus of Engineering and Scientific Terms (TEST) and that thesaurus identified. If it is not possible to select indexing terms which are Unclassified, the classification of each should be indicated as with the title.)

Radar
Distributed Computing System