

CAN UNCLASSIFIED



Mobile network field experiments

Summer 2017

Tricia Willink Susan Watson David Brown Mazda Salmanian Ronggong Song Ming Li Pierre-Luc Drouin DRDC – Ottawa Research Centre

Defence Research and Development Canada

Scientific Report DRDC-RDDC-2018-R032 April 2018



CAN UNCLASSIFIED

IMPORTANT INFORMATIVE STATEMENTS

This document was reviewed for Controlled Goods by Defence Research and Development Canada (DRDC) using the Schedule to the *Defence Production Act*.

Disclaimer: Her Majesty the Queen in right of Canada, as represented by the Minister of National Defence ("Canada"), makes no representations or warranties, express or implied, of any kind whatsoever, and assumes no liability for the accuracy, reliability, completeness, currency or usefulness of any information, product, process or material included in this document. Nothing in this document should be interpreted as an endorsement for the specific use of any tool, technique or process examined in it. Any reliance on, or use of, any information, product, process or material included in this document is at the sole risk of the person so using it or relying on it. Canada does not assume any liability in respect of any damages or losses arising out of or in connection with the use of, or reliance on, any information, product, process or material included in this document.

Endorsement statement: This publication has been peer-reviewed and published by the Editorial Office of Defence Research and Development Canada, an agency of the Department of National Defence of Canada. Inquiries can be sent to: <u>Publications.DRDC-RDDC@drdc-rddc.gc.ca</u>.

© Her Majesty the Queen in Right of Canada (Department of National Defence), 2018

© Sa Majesté la Reine en droit du Canada (Ministère de la Défense nationale), 2018

Abstract

This report describes field experiments undertaken in the summer of 2017 using network and data applications developed under the Tactical Network Operations project. The aims of the activity were to develop a field testing capability and to obtain insight into the performance of wireless networks in static and mobile operation. We completed a number of experiments, building on each other to understand the impact of the physical environment, the network topology and the applications. Preliminary data analysis has shown that path loss, even in relatively ideal unobstructed line-of-sight links, is highly variable over very small distances and across different devices: this impacts the potential for power-based emitter geolocation. Additionally, these experiments have revealed some of the complexities of application data transfer in mobile ad hoc networks, which involve data being relayed through other nodes en route from source to destination. In particular, we found that the traffic and MAC layer protocol messaging appeared to cause interference and collisions when routed via relays, resulting in unexpected degradations in throughput and packet loss performance. Future work to investigate this further will require more detailed over-the-air observations in addition to logging on the radio nodes. The data collected during these experiments will continue to support the project's R&D activities, and further experiments will be planned to gain deeper insights and to investigate performance in more challenging environments.

Significance to defence and security

Future tactical networks will support a variety of application traffic types, and operate in mobile environments with dynamic topologies. To overcome the challenges in making these networks robust and resilient in dynamic and contested conditions or when under attack, and to reveal opportunities for exploitation, the complex interactions of the dynamic properties of the radio, network and traffic must be understood. The experiments and preliminary analysis reported herein represent a continuation of the R&D within the Tactical Network Operations project to achieve this goal.

Résumé

Le rapport décrit des expériences menées sur le terrain pendant l'été 2017 à l'aide d'applications réseau et de données développées dans le cadre du Projet d'opérations de réseau tactique. Elles visaient à développer les capacités d'essai sur le terrain et de mieux comprendre le rendement des réseaux sans fil en fonctionnement statique et mobile. Nous avons mené à bien plusieurs expériences, en mettant à profit les résultats des premières pour les expériences ultérieures, afin de comprendre les répercussions sur le rendement de l'environnement, de la topologie du réseau et des applications. Une analyse préliminaire des données démontre que l'affaiblissement de propagation, même pour les liens en visibilité directe, passablement idéaux car libres d'obstructions, varie considérablement sur de très petites distances et d'un dispositif à l'autre; cela ne peut que jouer sur la géolocalisation des émetteurs en fonction de la puissance. Ces expériences ont aussi mis en lumière certains des aspects complexes des communications de données interapplications dans les réseaux mobiles spéciaux, car ces données sont relayées par des nœuds intermédiaires entre la source et la destination. Plus précisément, nous avons observé que les messages des protocoles de couche MAC semblent entraîner interférences et collisions s'ils sont transmis par relais; il en résulte la perte de paquets et une dégradation inattendue du rendement. Pour étudier cela plus avant, il nous faudra observer les radiocommunications plus en détail et journaliser l'activité des nœuds radio. Les données recueillies pendant ces expériences appuieront les activités de R et D du projet, et nous prévoyons d'autres expériences afin d'approfondir nos connaissances et d'étudier le rendement des réseaux en environnements plus difficiles.

Importance pour la défense et la sécurité

Les réseaux tactiques de l'avenir prendront en charge bien des types de communications entre applications, et devront fonctionner en environnements mobiles et selon des topologies dynamiques. Afin de surmonter les obstacles qui nous empêchent d'assurer la robustesse et la résilience de ces réseaux en situations dynamiques et contestées sinon carrément hostiles, et en dégager les occasions d'exploitation, il faut comprendre clairement les interactions complexes entre les caractéristiques dynamiques des émetteurs-récepteurs radio, du réseau et des communications qu'il véhicule. Les expériences et l'analyse préliminaire décrites dans le présent rapport s'inscrivent dans la lancée des recherches et du développement du Projet d'opérations de réseau tactique visant à atteindre cet objectif.

Table of contents

Abstract	i							
Significance to defence and security	i							
Résumé								
Importance pour la défense et la sécurité								
Table of contents	iii							
List of figures	v							
List of tables.	vii							
Acknowledgements	viii							
1 Introduction	1							
1.1 Background concepts.	1							
1.1.1 Mobile ad hoc networks	1							
1.1.2 Transport protocols	2							
2 Experiment plans and equipment.	3							
2.1 Experiments overview	3							
2.2 Hardware	3							
2.3 Software	4							
2.4 Measurement support.	4							
2.5 Location	4							
3 Physical layer evaluations	7							
3.1 Objectives	7							
3.2 Plan and execution.	7							
3.3 Preliminary results.	8 0							
3.3.2 Effect of device variation	0							
3.3.3 Relationship between distance and RSSI	13							
4 MANET operation and link connectivity	15							
4.1 Objective	15							
4.2 Experimental setup and procedure	15							
4.2.1 UDP and TCP measurement between two static nodes	16							
4.2.1.1 UDP transport—packet delivery rate	16							
4.2.1.2 TCP transport—throughput	17							
4.2.2 UDP and TCP measurement between two mobile nodes	17							
4.2.2.1 UDP transport—slow mobility	17							
4.2.2.2 UDP transport—fast mobility	17							
4.2.2.3 TCP transport—slow mobility	17							
4.2.2.4 TCP transport—fast mobility	18							
4.2.3 TCP measurement between two nodes using a relay.	18							

	4.3	Result	s and disc	ussion .																			18
		4.3.1	UDP and	TCP mea	suren	nent	betv	veen	tw	o st	atic	no	des	—re	sul	ts							18
			4.3.1.1	UDP tran	sport-	-pao	cket	deli	ver	y ra	te.												18
			4.3.1.2	TCP trans	sport-	-thro	ougl	hput	•						•								20
		4.3.2	UDP and	TCP mea	suren	nent	betv	veen	tw	o m	obi	le n	ode	es—	resi	ılts							21
			4.3.2.1	UDP tran	sport-	-dro	op ra	ate .															21
			4.3.2.2	TCP trans	sport–	-thro	ougl	hput							•								23
		4.3.3	TCP mea	surement	betwe	en t	wo i	node	s us	sing	, a r	elay	/—	resu	lts								24
5	Gate	way-ins	pired expe	eriment .																			26
	5.1	Scenar	io and obj	jectives .																			26
	5.2	Varyir	ng the dista	ance betw	een re	lays																	27
	5.3	Varyir	ng the dista	ance betw	een de	estina	atio	n and	d so	ourc	e to	the	eir r	espe	ecti	ve	neig	ghb	our	rel	ays	5.	30
6	Traff	ic analy	sis data c	apture .																			32
	6.1	Object	ives.																				32
	6.2	Plan a	nd executi	on																			32
	6.3	Prelim	inary resu	lts																			33
7	Lesso	ons lear	ned and fu	iture work																			34
	7.1	Opera	tional less	ons learne	d.																		34
	7.2	Experi	mental les	ssons learr	ned.																		35
	7.3	Softwa	are develo	pment req	uirem	ents																	36
	7.4	Future	experime	ntal work																			36
8	Conc	lusions	-																				38
Re	ferenc	es .																					39
			· · ·		· ·	•	-	•	-	-	· ·	-	-	•	-	-	-	•	•	-	-	-	

List of figures

Figure 1:	Aerial view of Shirleys Bay campus	5
Figure 2:	Experiment zones.	6
Figure 3:	Histograms of collected RSSI measurements for a static receiver and transmitter, runs 1–3.	9
Figure 4:	Histograms of collected RSSI measurements for a static transmitter and a circularly moving receiver, runs 1–2.	10
Figure 5:	Collected RSSI measurements for two collection runs with unobstructed line-of-sight, and one run with scatterers (people) in the line-of-sight, moving from the transmitter toward the receiver.	11
Figure 6:	Histograms of collected RSSI measurements for a static receiver and different transmitting devices. (Numbers .xx identify the devices' final IP address digits).	12
Figure 7:	Collected RSSI measurements as distance is increased from 25 to 250 m	13
Figure 8:	Distance error, resulting from difference between propagation model and actual mean RSSI, vs. actual distance in metres.	14
Figure 9:	The experimental area for MANET operation experiments	16
Figure 10:	Packet delivery rate for UDP broadcast from Tx to Rx	19
Figure 11:	Evolution of total number of dropped UDP packets	19
Figure 12:	Average TCP throughput for a 120 s test.	20
Figure 13:	10 s moving average of TCP throughput.	21
Figure 14:	Nodes in relative motion, slow walk (175 s to walk 200 m), stationary Tx and mobile Rx.	22
Figure 15:	Nodes in relative motion, 33–40 s to drive 200 m, stationary Tx and mobile Rx	23
Figure 16:	TCP throughput 10 s moving average, nodes in relative motion, stationary Tx and mobile Rx.	24
Figure 17:	Average TCP throughput between nodes separated by 250 m, communicating through a relay.	25
Figure 18:	Experiment plan for a three-link route MANET, containing two relays	26
Figure 19:	MANET UDP throughput via a route with two relays vs. distance between relays	27
Figure 20:	MANET TCP throughput via a route with two relays vs. distance between relays	28
Figure 21:	Two instances of use case MANET UDP instantaneous throughput via a route with two relays vs. time.	29
Figure 22:	MANET TCP and UDP average throughput via a route with two relays vs. distance between relays.	29

Figure 23:	MANET TCP throughput via a route with two relays vs. distance from destination to its neighbouring relay (link c) and vs. distance from source and destination to their respective neighbouring relay (links a and c).	30
Figure 24:	MANET UDP throughput via a route with two relays vs. distance from destination to its neighbouring relay (link c) and vs. distance from source and destination to their respective neighbouring relay (links a and c).	31
Figure 25:	The experimental area and physical setting for traffic analysis data capture	33

List of tables

Table 1:	Mean and standard deviation of collected RSSI measurements, for stationary and moving receivers.
Table 2:	Mean and standard deviation of collected RSSI measurements for 5 different transmitters. Two runs were collected for each transmitter

Acknowledgements

The authors thank Lucas Rooyakkers for his development of the GPS software application that was used to synchronise the Android phones, and for his enthusiastic contribution to the experiment activities.

1 Introduction

Wireless networks operate in a medium that is time-varying and unpredictable; these variations in the physical links among radios have repercussions for higher level radio, network and application functions. To maintain adequate application performance, networks should be resilient to changes in the physical medium. This is conventionally achieved by automatically adapting the throughputs of individual radios and, in the case of Mobile Ad Hoc Networks (MANETs), the network's topology.

The Tactical Network Operations (TNO) project, within the Director General (S&T) Joint Force Development (DGSTJFD) portfolio, started in April 2014. According to the Project Charter, the project "addresses the need for future wireless network security, management, and full-spectrum operations through a balance of R&D activities" [1]. A core facilitator of this project has been a MANET application, developed in preparation for the TNO project, which operates using WiFi at 2.4 GHz on Android phones and Linux laptops. This MANET implementation supports several network and cyber tools that we have developed, including Situational Awareness (SA) and Command and Control (C2) applications. Some of these tools were trialed at DRDC – Toronto Research Centre in 2014 [2].

To understand the performance of the SA and C2 applications in real conditions, it is necessary to test and evaluate the behavioural characteristics of the MANET over which they operate. While mathematical analysis and computer simulations are vital to the development of concepts for network use and cyber security, they are unable to reproduce the complete effects of the real world. Therefore, to understand the impact of the complex interactions within each radio and across the network as a whole, it is essential to perform live experiments.

This report describes a number of field experiments that we undertook in the summer of 2017 at the Shirleys Bay campus, with the aims of developing a field testing capability and obtaining insight into the performance of wireless networks in static and mobile operation. The objectives of this report are: to document the planning and implementation of the experiments; to provide preliminary analyses of the data collected; and to identify lessons learned. These will inform future experiment activities within TNO and subsequent projects.

In the remainder of this section, we give a summary of background concepts related to the experiments. In Section 2, a high-level view of the experimental plans and a brief description of the system components is presented. Details on each experiment phase and their outcomes are provided in Sections 3–6, and lessons learned and future experimental directions are presented in Section 7. Conclusions are summarised in Section 8.

1.1 Background concepts

1.1.1 Mobile ad hoc networks

Mobile ad hoc networks, or MANETs, are a form of self-organising, infrastructure-less network, in which messages may be relayed through peer nodes, over multiple hops, to reach their destinations. Routes connecting pairs of nodes may be determined in a proactive way, meaning that routes are regularly re-computed regardless of demand, or in a reactive way, in which routes are computed only when demanded by network nodes. In both proactive and reactive modes, probing messages are sent by nodes

to estimate the quality of links to their immediate neighbours, and to exchange local information that is required to compute routes.

In the MANET implemented and used in these experiments, the Optimised Link State Routing (OLSR) protocol [3] is used, which is a proactive routing protocol. We used the OLSRd implementation of OLSR version 1 [4]. Each node emits "Hello" messages at defined intervals, which are used to sense the links and exchange information about one- and two-hop neighbours, as well as Topology Control (TC) messages, which are used to inform each node about link states to support route computation. The intervals at which these messages are sent are configurable; it is known that these values should be adjusted to accommodate the speed of network nodes [5], but in our experiments, mobility was sufficiently low that we used the default values (2 s for Hello messages and 5 s for TC messages).

The Ad Hoc On-Demand Distance Vector (AODV) routing protocol uses route request and route reply messages to compute routes from specified sources to desired destinations [6]. Only routes for active communications are maintained, and no resources are used to compute routes that are not needed. This protocol was designed for dynamic conditions and low network demand. An implementation of AODV has recently been completed within TNO, which will be used in future experimentation.

1.1.2 Transport protocols

Transport protocols are used in IP-enabled communications to provide a mechanism to deliver data streams from an application on one node to a (possibly different) application on another. In addition to carrying the data stream itself, the protocol also transmits control data to support error detection and correction, and may have some form of Automatic Repeat Request (ARQ) to retransmit uncorrectable packets.

In these experiments, we used two of the most common transport protocols, the User Datagram Protocol (UDP) and the Transmission Control Protocol (TCP). UDP is "connectionless," i.e., it is a "send-and-forget" protocol that does not ensure data delivery. It has only a small amount of overhead, to communicate the application information and for packet segmentation. However, note that for unicast (addressed to a specific destination) UDP in WiFi, there is a Medium Access (MAC) layer connection, in which the source node transmits a "request to send" (RTS) message, which is responded to by the destination with a "clear to send" (CTS) message. Further, Acknowledgements (ACKs) are issued by the MAC layer, so UDP operating over WiFi is not truly connectionless.

TCP is a "connection-oriented" protocol, i.e., it uses extra overhead and re-transmissions to ensure that transmitted packets are received correctly, and in the correct order, (these are in addition to the MAC layer connection messages). Testing using the UDP protocol shows the number of packets that are lost due to poor link quality, while testing using the TCP protocol reveals the net application throughput, taking into account packet retransmissions due to poor link quality. Both these parameters are important in understanding how a network may support its applications.

2 Experiment plans and equipment

2.1 Experiments overview

As noted above, our objective with these experiments was to obtain insight into the performance of wireless networks in operational conditions, i.e., outdoors with nodes separated by tens of metres or more.

The experiments were designed to build on lessons learned, starting at the physical layer, using performance results and observations to fine-tune experimental conditions for higher layers and more challenging tests.

The first set of experiments, described in Section 3, were designed to evaluate the variation in received signal strength over small and large distances, and to determine whether there was a significant difference in measured received signal strength in different devices. This set of experiments directly informs the question of whether power-only measurements can support reasonably accurate geolocation of target nodes.

Building on the knowledge obtained from the physical layer evaluations, in the second phase of experiments we considered the link quality and its impact on MANET operation, described in Section 4. These experiments were designed to gain insight into throughputs at different node separations, both without and with the aid of a relay node. Both static and mobile scenarios were considered, as a first step to understanding the impact mobility has on overall MANET performance.

During the TNO project to-date, an application to implement a gateway concept connecting two or more MANETs, described in **Error! Reference source not found.**, has been completed, building on previous work. A full understanding of the performance of this complex architecture in a real environment requires repeated experimentation, to develop and build upon insights gained. In this phase, the behaviour of a MANET using a three-hop route was investigated: this is a preliminary proxy to the use of gateways connecting the edges of two distinct MANETs, where gateways are essentially complex relays using two radio interfaces, each with its own network stack. The experiments and results of this phase are described in Section 5; future experiments are being planned to enhance our understanding of gateway performance.

To support ongoing R&D in traffic analysis, on-air data packets were collected for different user applications, in single- and multi-hop scenarios. The applications used included blue force tracking, file transfer and TCP; in this experiment, the OLSR messages themselves were also part of the on-air data collected. The experimental setup and data sets are described in Section 6.

2.2 Hardware

For the experiments, we used Nexus 5 smart phones [8], which run the Cyanogenmod 13 operating system, which is an Android variant. This phone can be operated using its internal WiFi interface, and an external USB WiFi adapter can be used instead of, or in addition to, the internal interface. The external WiFi adapters we used were the Panda Wireless (PAU06) [9] and TPLink (TL-WN722N) [10] products.

2.3 Software

We used the DRDC MANET application (app) IRN MC (for improvised radio network—MANET controller), based on the OLSR protocol, to provide the test networks used in our experiments. For measurements not using this app, a GPS synchronisation app developed at DRDC was used to establish precise location and timing. These apps were developed prior too, and during, the TNO project.

To measure the received signal strength, we used the IRN Sensor app, which was developed at DRDC as part of a distributed network surveillance activity. It provides power measurements, averaged over short, definable windows associated with signals transmitted from selected nodes.

To test the performance of UDP and TCP, we used the IRN Test app, which provides support for measuring packet delivery rates for UDP and throughput for TCP in configurable tests between specified nodes in the MANET.

For data capture to support traffic analysis, in addition to the IRN Test app, we used the IRN blue force tracking app [11], which was developed as a de-risking prototype for the Tactical Edge Cyber Command and Control (TEC3) technical demonstrator [12]. The Serval Mesh app, an open source file sharing application [13], was also used for traffic generation.

The Android phones have a Nethunter terminal function installed [14], which is open-source software developed as a penetration testing platform. This includes the "tcpdump" functionality, which is a packet analyser used to display and count data packets observed on a radio interface [15].

2.4 Measurement support

The primary tool for supporting WiFi measurements was Wireshark [16], which is a network protocol analyser. Used on a laptop in conjunction with a Riverbed AirPcap adapter [17], WiFi packets can be captured on-air and analysed to identify data, control and management packets, by source and destination.

2.5 Location

We performed these experiments at the Shirleys Bay campus; an aerial view of the campus showing the area used for our experiments is shown in Figure 1. A more detailed view of the zones used for the experiments is shown in Figure 2. For the preliminary physical layer experiments (Section 3), we used the grassy area marked as "zone 1" in Figure 2, and for the subsequent measurements we used the straight portion of a gravel road in the north-east area, marked "zone 2." These zones were selected because they are reasonably flat, and there are no buildings and little traffic; experiments were paused on the rare occasions that people or a vehicle passed through the experiment zone.



Figure 1: Aerial view of Shirleys Bay campus.

The nodes were placed on tripods approximately 1 m high, or held by an experimenter at the same height, for static experiments, while for mobile experiments, they were moved by hand at pedestrian speed or in a golf cart-type buggy for higher speeds.



Figure 2: Experiment zones.

3 Physical layer evaluations

3.1 Objectives

We would like to be able to geolocate WiFi transmitters using simple, commercially available sensor hardware which does not require precise time synchronisation between the measuring devices. One way is to use the WiFi chipsets themselves to measure received power, convert the received power into a distance, and then estimate the position of the transmitter by trilateration. To relate the received power to the separation distance from the transmitter, we can use an expression of the received power (P_{rx} in dBm) at some distance d from the transmitter [18]:

$$P_{rx} = P_0 - 10\gamma \log_{10} \frac{d}{d_0} \tag{1}$$

where P_0 is the power measured at a reference distance d_0 from the transmitter, and γ is the path loss exponent, which is a constant ranging between 2 and 4 (typically) depending on the local environment. The accuracy and consistency of the received power estimates reported by a measuring device will have a direct impact on the accuracy of the geolocation estimates.

The Received Signal Strength Indicator (RSSI) reported by a WiFi device is an approximation or an imperfect measure of the received power. The RSSI estimates the signal energy at the WiFi receiver during reception of the packet header; this data is available for successfully received packets only. As it is a measure of signal energy, the RSSI contains the energy from interference in addition to the received signal. The RSSI is reported differently by different chipset manufacturers, and is meant to measure signal quality within the chipset and its driver.

We performed a preliminary set of outdoor measurements for three purposes: 1) to investigate the effect of local motion on RSSI; 2) to determine the amount of variability between RSSI measurements on different devices having the same RF hardware; and 3) to examine the predictability of the relationship between distance and RSSI, as in Equation (1).

3.2 Plan and execution

To conduct the experiments, we used Nexus 5 Android phones, which had the roles of transmitting device, receiving device, and measuring device. For the measuring device, we used the custom IRN Sensor app. This app reads the WiFi radiotap header attached to the packet by the WiFi firmware; the radiotap header exfiltrates the RSSI to upper layer software. The app then logs the per-packet RSSI along with other data, such as the time, originating MAC address, and destination MAC address. The transmitting and receiving devices were loaded with the IRN Test app, which allows the user to initiate custom traffic flows. The transmitting device always sends packets at full power (approximately 10.5 dBm on the internal interface [19]).

The first experiment tested the variation in RSSI under local motion. We collected measurements in an open field environment (zone 1 in Figure 2). Two tripods were set up with a 50 m separation distance, where one tripod supported the transmitting Android device, and the other supported the measuring

device, both at roughly chest height. For each collection of RSSI data, we initiated a one minute flood of broadcast packets at a rate of 50 kB/s. To create local motion, with tripod holding the measuring device, the device was held away from the body and rotated in a horizontal circle roughly two wavelengths (25 cm) in diameter. Each rotation was executed over a period of a few seconds.

The second experiment tested the effect of replacing the transmitting device with an equivalent device; that is, the same make and model of Android phone, having the same WiFi chipset. We kept the experimental setup the same for these experiments as for those with local motion, however in this case, the devices were stationary. Five different equivalent transmitting devices were tested.

For the third experiment, the environment was a line-of-sight channel along a gravel road (zone 2 in Figure 2). The measuring device was placed in a tripod opposite the transmitting device, and the receiving device was nearby the measuring device. We varied the separation distance between the measuring device and the transmitting device from 25 m to 250 m in increments of 25 m. For a given separation distance, broadcasts were sustained for 90 seconds while RSSI data was collected. Naturally, as separation distance increased, the number of RSSI measurements collected within the 90 second period decreased due to a greater number of errored messages; for example, we collected 23,287 RSSI samples at 25 m, and 277 samples at 250 m. We note that as a consequence, the RSSI data collected does not represent the full range of received powers seen by the receiving devices, particularly at larger separation distances.

3.3 Preliminary results

3.3.1 Effects of local motion

We know that received signal strength can vary widely as the receiver is moved a few wavelengths, due to multipath propagation effects. To describe the effect of local motion in the first experiment, we report a simple mean and standard deviation in the RSSI values for each collection run. As shown in Table 1, for collection runs with stationary transmitting and measuring devices, the distribution of RSSIs has a smaller standard deviation¹ (i.e., 1.1–1.3 dBm) compared with collection runs with local motion (i.e., 3.6–3.7 dBm). However, referring to Figure 3 (run 1), even when the devices are stationary, the distribution of measured RSSIs can vary widely. Mean RSSIs varied by up to 6 dB in our experiment (Figure 3, run 1 vs. run 3), which could result in significant discrepancies in geolocation based on received signal strength [20], [21].

~3000 measured packets each:	Mean (dBm)	Std. Dev. (dBm)
Stationary, first run	-61	1.1
Stationary, second run	-65	1.3
Stationary, third run	-67	1.1
Local motion, first run	-64	3.6
Local motion, second run	-64	3.7

 Table 1: Mean and standard deviation of collected RSSI measurements, for stationary and moving receivers.

¹ Means and standard deviations were computed in log units.



Figure 3: Histograms of collected RSSI measurements for a static receiver and transmitter, runs 1–3.

Local motion appears not to be detrimental to the accuracy of the mean RSSI, as long as the RSSI measurements are averaged. In fact, our data suggests that local motion produces more stable mean RSSI by inducing channel variations causing a more evenly distributed set of RSSI measurements over the range of RSSI values. Figure 4 shows the histograms of measured RSSIs for the two collection runs with local motion. In these charts, we see a wide range of RSSI values spanning nearly 20 dB. However, the local motion produced a stable average RSSI (-64 dBm) over two runs and a very similar histogram of RSSI values between the two runs. By contrast, we note that the mean RSSI varied between -61 dBm and -67 dBm across collection runs when the measuring device was stationary.

To explore the effect of scattering, we did a one-minute test in which people walked in and out of the line-of-sight between the transmitting device and the measuring device during the RSSI measurement collection (Figure 5). The separation between transmitter and measuring device was 50 m. The human scatterers moved both away from the receiver, and back and forth through the line-of-sight. In the chart, the x-axis orders the RSSI measurements as to when they were collected as time passed, but does not represent how much time passed between measurements—all three collections were taken over a one minute period.

As can be seen from the chart, there was a greater and more variable channel attenuation evident from the received RSSIs when the scatterers were present. As the scatterers moved towards the transmitter, the attenuation was reduced; when the scatterer is near the transmitter, most scattered signal components do not reach the receiver, leaving only the strongest specular components. In contrast, near the receiver, scattered signal components can combine constructively or destructively at the receiver, causing greater signal power variations. As seen in Figure 5, these local scatterers caused up to a 15 dB drop in RSSI. This partially obstructed line-of-sight channel would pose a challenge for a power-based geolocation algorithm.



Figure 4: Histograms of collected RSSI measurements for a static transmitter and a circularly moving receiver, runs 1–2.

3.3.2 Effect of device variation

Table 2 shows the mean and standard deviations for RSSIs collected using the five different stationary transmitting devices, and Figure 6 shows the histograms of the collected RSSIs. For the separation distance of 50 m, we observed a maximum difference in average RSSI of 3 dBm. A difference in received power of 3 dBm (where the received power is about -72 dBm) translates to a significant difference in estimated range using a power-distance relationship like that in Equation (1).

Furthermore, if we compare the mean RSSI results in Table 2 to those in Table 1, conducted the day before under the same conditions using transmitter 4 (identified as ".16" in Figure 6), we find mean RSSI values approximately 10 dB lower. Consequently, we are not confident in the repeatability of RSSI measurements at a given distance. Power-based geolocation, using local RSSI values, would therefore yield inconsistent estimates of location.



Figure 5: Collected RSSI measurements for two collection runs with unobstructed line-of-sight, and one run with scatterers (people) in the line-of-sight, moving from the transmitter toward the receiver.

Table 2: Mean and standard deviation of collected RSSI measurements for 5 different	
transmitters. Two runs were collected for each transmitter.	

1000-3000 measured packets each:	Mean (dBm) Run 1	Mean (dBm) Run 2	Std. Dev. (dBm) Run 1	Std. Dev. (dBm) Run 2
Transmitter 1 (.13)	-72	-71	2.0	1.8
Transmitter 2 (.14)	-72	-72	1.7	2.0
Transmitter 3 (.16)	-73	-72	2.3	1.8
Transmitter 4 (.17)	-71	-71	1.4	1.6
Transmitter 5 (.19)	-73	-74	1.8	1.9



Figure 6: Histograms of collected RSSI measurements for a static receiver and different transmitting devices. (Numbers .xx identify the devices' final IP address digits).

3.3.3 Relationship between distance and RSSI

The range experiment was conducted in an uncluttered environment, with stationary receiver and transmitter, and a line-of-sight, using many measurements per distance point. Yet, while our results confirm that RSSI is inversely correlated with distance, the RSSI data is a somewhat poor fit to a linear relationship with the log of distance, using a least-squares regression technique (see Figure 7). While other propagation models might fit parts of the data better, analytical or empirical models are not expected to fit any specific set of data well, as they are developed for averaged conditions.

Figure 8 shows the distance errors that would result from the solution of the fitted linear equation, compared to the actual distances, and using the mean RSSIs from the collected data. The distance error at 150 m is 77 m, more than half the separation distance itself. Such a large error casts doubt on the validity of a simple RSSI-distance relationship. Furthermore, given that the range of a WiFi transmitter is modest—even in outdoor settings—we would like to know our distance to the transmitter to within an error which is significantly less than the transmission range itself.

Our measurements do not support a recommendation to use RSSI as the sole basis for a geolocation result. Further, it was seen that local scatterers can cause wide variations in received power. Though we have not included it in our experiment, varying antenna orientation also has a blurring effect. These confounding factors obscure the distance-RSSI relationship further.



Figure 7: Collected RSSI measurements as distance is increased from 25 to 250 m.



Figure 8: Distance error, resulting from difference between propagation model and actual mean RSSI, vs. actual distance in metres.

4.1 Objective

The objective of this set of experiments was to gain an appreciation of the expected throughput of WiFi ad hoc network devices operating in a relatively benign environment free of major scatterers and with line-of-sight connectivity among neighbouring nodes. Note that our intent was not to arrive at exact throughput values for the networks in question, nor to identify precise ranges of expected throughput values, but was rather to make phenomenological observations of throughput in a number of different scenarios. The conditions we were interested in were: (1) the distance between nodes; (2) connectionless versus connection-oriented transport protocols; (3) direct communication between adjacent nodes versus communication through a single relay; and (4) mobile versus static nodes.

The remainder of this section details the experimental setup and procedure followed in conducting this set of experiments, and a summary and discussion of the observed results.

4.2 Experimental setup and procedure

This set of experiments was undertaken in zone 2: a more detailed aerial view is shown in Figure 9. Prior to running the experiments, we selected a location along the road as a "starting position," shown with a pink marker. We measured and marked intervals spaced 25 m apart from the starting position, running along the road, as indicated by the blue line; the yellow line in Figure 9 indicates a distance of 100 m, for reference.

As noted in Section 2.2, these experiments used Android phones; in each case, one phone acted as a transmitter (denoted "Tx") and another phone acted as a receiver (denoted "Rx"). Unless otherwise specified, the Tx phone was placed in a wooden holder on a tripod at the starting position for all experiments. Certain experiments also involved a third phone acting as a relay (denoted "H"). The phones operated as a MANET using IRN MC, and the Tx and Rx nodes used the IRN Test application to generate traffic for measuring network throughput. The internal WiFi interface was used for transmitting and receiving on all phones.

The following sub-sections summarise the procedure we followed for the experiments we conducted to evaluate MANET operation and link connectivity.



Figure 9: The experimental area for MANET operation experiments.

4.2.1 UDP and TCP measurement between two static nodes

The first experiment focused on direct communication between two ad hoc nodes with varying distances and two different transport protocols. Specifically, we examined the packet delivery rate as a function of distance for UDP transport, and examined the maximum throughput as a function of distance for TCP transport.

4.2.1.1 UDP transport—packet delivery rate

For a particular separation of Tx and Rx, the IRN Test app was configured to send UDP datagrams from Tx to Rx at a rate of 50 kB/s for a duration of 120 seconds using the "UDP Fire and Forget" broadcast feature in IRN Test. Packet data was recorded for the duration of the test on both Tx and Rx using the tcpdump function,² with the resulting packet capture data saved to a file.

² The exact tcpdump command used was "tcpdump -i wlan0 > filename", where this saves all packet data observed on the wlan0 (the internal WiFi interface) to a user-specified file.

The experiment was conducted once for each of the following Tx/Rx separations: 1 m, 50 m, 100 m, 150 m, 200 m, 250 m. In all cases, the Rx phone was held upright by the experimenters at waist height.

4.2.1.2 TCP transport—throughput

For a particular separation of Tx and Rx, the IRN Test app was configured to measure the maximum TCP throughput of the channel. The throughput test was conducted for a duration of 120 s, with the IRN Test app instance on the Rx node recording the "instantaneous" observed TCP data rate once per second.

The experiment was conducted once for each of the following Tx/Rx separations: 25 m, 50 m, 75 m, 100 m, 125 m, 150 m, 175 m, 200 m, 225 m. In all cases, the Rx phone was held upright by the experimenters at waist height.

4.2.2 UDP and TCP measurement between two mobile nodes

The second experiment focused on the observed effect on communications when one node moved at a constant velocity relative to the other, static node. Once again, we examined both UDP and TCP transport layer cases.

For all sub-experiments considered here, the Tx node was kept at the starting position. To create relative mobility between the Tx and Rx nodes, the Rx node was moved between the 50 m and 250 m markers at varying speeds.

4.2.2.1 UDP transport—slow mobility

The Rx phone was held at waist height by an experimenter such that their body did not obstruct the line-of-sight to the Tx phone, and the IRN Test app was configured to send UDP datagrams from Tx to Rx at a rate of 50 kB/s using the "UDP Fire and Forget" broadcast feature in IRN Test. As soon as IRN Test began sending data, the experimenter began walking from the 50 m mark to the 250 m mark at a slow pace, such that it took 175 s to cover the distance. This test was repeated, with the experimenter walking from the 250 m mark back to the 50 m mark. We ran this test twice (walking there and back twice), and used tcpdump to gather the packet capture data on both Tx and Rx nodes.

4.2.2.2 UDP transport—fast mobility

The Rx phone was held aloft by an experimenter in the back of a buggy located at the 50 m mark. The IRN Test app was configured to send UDP datagrams from Tx to Rx at a rate of 50 kB/s using the "UDP Fire and Forget" broadcast feature in IRN Test. As soon as IRN Test began sending data, the driver of the buggy accelerated to the maximum attainable speed (approximately 20 km/h) and continued driving to the 250 m mark. This test was repeated 4 times. The test always ran from the 50 m to 250 m mark for safety reasons (it was deemed unsafe to drive the buggy backwards from 250 m to 50 m, and we always wanted the rider in the back positioned such that they were facing the Tx phone so we could not turn the buggy around). Each buggy ride took roughly 35 s to cover the 200 m distance.

4.2.2.3 TCP transport—slow mobility

The Rx phone was held at waist height by an experimenter, and the IRN Test app was configured to measure TCP maximum throughput. As soon as IRN Test began sending data, the experimenter began walking from the 50 m mark to the 250 m mark at a slow pace, such that it took 175 seconds to cover the

distance. This was repeated twice. The reverse route (from 250 m to 50 m) was not measured, since the TCP session would not initiate at a separation of 250 m between Tx and Rx (the three-way handshake could not complete as the nodes had poor connectivity).

4.2.2.4 TCP transport—fast mobility

The Rx phone was held by an experimenter in the back of a buggy located at the 50 m mark. The IRN Test app was configured to measure TCP maximum throughput. As soon as IRN Test began sending data, the driver of the buggy accelerated to the maximum attainable speed and continued driving to the 250 m mark. This test was conducted only once.

4.2.3 TCP measurement between two nodes using a relay

The final experiment in this set focused on the communication between two ad hoc nodes via a relay. For this experiment we examined only TCP maximum throughput and did not look at UDP packet drop rate. For all sub-experiments conducted here, three phones were used: one phone was at the starting position, another was at the 250 m mark, and a third phone served as the relay (the "H" phone) between these two.

The IRN Test app was configured to measure the maximum TCP throughput. The throughput test was conducted for a duration of 120 s, with the IRN Test app instance on the Rx node recording the "instantaneous" observed TCP data rate once per second.

The experiment was conducted for each of the following positions of the relay node: 75 m, 100 m, 125 m, 150 m, 175 m. In all cases, the Rx phone and H phone were held upright by the experimenters at waist height. The first run of the experiment had the Tx node at the starting position and the Rx node at 250 m for each of the relay positions. A second run was conducted with the Tx and Rx nodes swapped such that the Rx node was at the starting position and the Tx node was at 250 m.

4.3 Results and discussion

All results discussed in this section correspond directly to the experiments described in Section 4.2. Section titles refer back in an identical fashion to the experimental setup and procedure titles, with the addition of the suffix "results" on each title.

4.3.1 UDP and TCP measurement between two static nodes—results

4.3.1.1 UDP transport—packet delivery rate

In this experiment, tcpdump packet capture logs were collected on the Tx and Rx nodes. We filtered the logs to include only the UDP broadcast packets, which were sent at a rate of 50 kB/s. For each Tx/Rx separation distance we computed the packet delivery rate by observing the number of filtered packets collected at the Rx node compared to the number sent by the Tx node over the 120 s test duration. The packet delivery rate as a function of distance is shown in Figure 10.

We observe that the packet delivery rate decreases as a function of Tx/Rx separation, as expected. At a separation of 1 m, we observe a packet delivery rate of 97.6%—indeed hardly any packets are dropped when the nodes are right next to each other. There is a precipitous drop observed between 100 m and 150 m. By 250 m, no packets are observed at Rx, meaning that the packet delivery rate during this measurement is 0%.

Figure 11 shows the total number of dropped packets observed at each Tx/Rx separation over time. Dropped packets were obtained from the tcpdump logs by taking the cumulative sum of the number of dropped packets for every 5 s of the test.



The total number of packets dropped appears to increase relatively linearly as time progresses when computed every 5 s. We suspect that this drop rate may exhibit more variability over shorter time durations.

Figure 10: Packet delivery rate for UDP broadcast from Tx to Rx.



Figure 11: Evolution of total number of dropped UDP packets.

4.3.1.2 TCP transport—throughput

In this experiment, logs were gathered on the Rx node to determine the TCP throughput observed for various Tx/Rx separations. These logs report the "instantaneous TCP throughput" once per second. In Figure 12, we show the average TCP throughput (the average of the instantaneous throughput values reported by the IRN Test app) as a function of Tx/Rx separation.



Figure 12: Average TCP throughput for a 120 s test.

We note that the general trend is a decrease in throughput as Tx/Rx separation increases, as we would expect. There is a notable exception, however, for the 25 m measurement. This is not easily explained, given the data available. We suspect that were we to re-run these tests multiple times and take the average, we would observe that 25 m would indeed have a higher throughput than the other (further) distances. However, it is also entirely possible that the geographical conditions in the test area (see Figure 9) are such that the 25 m separation results in particularly poor channel conditions—due to reflection, multi-path fading, or other factors. Without multiple repeated measurements to confirm our observation, we cannot explain the dip at 25 m.

To observe the trends in TCP throughput over time, as opposed to the "average throughput" as shown in Figure 12, we computed the 10 s moving average of the TCP throughput, shown in Figure 13. The reason for computing a 10 s moving average is to show general trends in the throughput, since the instantaneous throughput reported was quite "noisy" and it was difficult to see the trends when the data were plotted over a 120 s interval.

Of note is that the throughput over time is quite variable. Whereas the observations of dropped packets over time reported in Figure 11 were somewhat predictable, the TCP throughput is not. Without significantly more testing and data, we cannot definitively identify why we observe such variability here that was not as apparent for the UDP case. One possibility is that the TCP algorithm is always searching

for an "optimum" rate at which to operate; when data is lost, TCP uses various backoff algorithms to attempt to re-establish its operating rate. The cycling we see in Figure 13 may be evidence of TCP hunting for its optimum rate; however, there is too little data to draw firm conclusions and further testing is necessary.



Figure 13: 10 s moving average of TCP throughput.

4.3.2 UDP and TCP measurement between two mobile nodes—results

4.3.2.1 UDP transport—drop rate

In this experiment, tcpdump logs were recorded at the Tx and Rx nodes; we filtered for the broadcast UDP packets of interest, which were being sent at 50 kB/s. We computed the total number of UDP packets dropped over time, observing the difference between the packets sent and received during each 5 second window over the 175 s test duration. Figure 14 shows the total number of packets dropped as a function of time for each of the four test runs (during runs 1 and 2 the Rx node moved at a slow walk from marker 50 m to marker 250 m; during runs 1b and 2b the Rx node walked from 250 m to 50 m).

For runs 1 and 2, we observe that the total number of packets dropped is quite low for the first 50 seconds, and then the packet drop rate accelerates as time increases. Qualitatively, this is not surprising, since initially the Tx and Rx nodes are nearby (50 m apart) and, based on our results in Figure 10, we do not expect significant packet loss at this distance. As the Rx node moves further away, we expect the packet loss to increase, which is precisely what we observe here. We note that while the qualitative behaviours of runs 1 and 2 are similar, the actual values of packets dropped are quite different. Again, without many runs to compare we cannot definitively explain the reason for the difference—however, a likely source of this discrepancy is the non-constant velocity of the mobile Rx node; we simply walked at a slow pace while holding the Rx node and it was difficult to ensure a constant pace throughout. Perhaps we walked slower during the first 100 seconds of run 1 than run 2, for instance, leading to fewer dropped packets in this case.

In runs 1b and 2b, we observe a rapid increase in dropped packets at the beginning of the runs (when the nodes are far apart and we expect many packets to be dropped); the drop rate decreases and levels off as time progresses since the Rx node moves closer. As we observed for runs 1 and 2, the qualitative shape of runs 1b and 2b are very similar, but the quantitative values are not—again, this could be due to the non-constant velocity of the Rx node, but without more test runs and deeper investigation we cannot know for certain the source of this discrepancy.



Figure 14: Nodes in relative motion, slow walk (175 s to walk 200 m), stationary Tx and mobile Rx.

Of interest is that the curves in Figure 14 are not monotonically increasing—this is initially surprising since we would expect the total number of dropped packets not to decrease over time. This result is an artifact of the way in which we processed the data. Specifically, we looked at the number of sent packets and received packets in contiguous 5 s windows. In any particular 5 s window, it can (and does) occur that some of the transmitted packets, i.e., noted as transmitted by tcpdump, are not received until the subsequent 5 s window. Thus, certain windows indicate a certain number of dropped packets which are subsequently delivered in the next window, resulting in a decrease in "total packets dropped."

Figure 15 shows the total packets dropped when a vehicle was used to drive the Rx node at a constant speed from the 50 m mark to the 250 m mark. The curves in Figure 15 are qualitatively similar to runs 1 and 2 from Figure 14, but occur over a more rapid time period. We also note that all four curves in Figure 15 are qualitatively similar to each other, and exhibit less variation than we saw between runs for the "slow walk" experiment. We suspect that the reduced variation seen here is a result of the fact that it was easier to maintain a constant speed in the vehicle than it was when walking. Once again, the shape of the curves is as expected with a low packet drop rate initially (when Tx and Rx are nearby) and an increasing drop rate as the Rx node moves away.



Figure 15: Nodes in relative motion, 33–40 s to drive 200 m, stationary Tx and mobile Rx.

4.3.2.2 TCP transport—throughput

In this experiment, logs were gathered on the Rx node reporting the "instantaneous TCP throughput" once per second. We computed the 10 s moving average of the TCP throughput logs and plotted the result in Figure 16.

In each case, we observe a general trend for the TCP throughput to decrease as time goes on, reflecting the fact that the Rx node is moving further away from Tx. Note, however, that especially in the case of the slower moving Rx node, the decrease in rate is not smooth—again, we suspect this is due to the more complex interactions of the TCP transport layer when searching for an optimal rate.



Figure 16: TCP throughput 10 s moving average, nodes in relative motion, stationary Tx and mobile Rx.

4.3.3 TCP measurement between two nodes using a relay—results

We gathered the logs from the Rx node and computed the average TCP throughput from Tx to Rx via a relay. Figure 17 shows the measured throughput for the various relay positions.

Qualitatively, we note that we tend to observe higher throughput when the relay is closer to the midpoint between the two communicating nodes, i.e., closer to the 125 m mark. This is intuitively satisfying as we expect a relay in the middle of two nodes would have equally good connectivity to either neighbour. With a relay much closer to one node than the other, we would expect a degradation in performance due to a "weakest link" phenomenon for the nodes with the larger separation distance.

We were interested to see that the curves (for both runs) were not symmetrical about the mid-way point of 125 m. In both instances, we observe that throughput with a relay at 75 m has far worse performance than the symmetrical case of a relay at 175 m. After our first run of measurements (the blue curve) we wondered if this observation was related to the relay being closer or further from the Tx node; consequently for the second run we reversed the positions of the Tx and Rx nodes, but observed the same poor performance for the relay at 75 m. This suggests that there could be an environmental factor related to the setup and test location that resulted in a poorer performance for the 75 m relay. Without significantly more testing and controlling for variability, however, it is impossible for us to say for certain.



Figure 17: Average TCP throughput between nodes separated by 250 m, communicating through a relay.

Finally, we note that including a relay appears to significantly reduce TCP throughput compared to direct communications. For instance, comparing the throughput of the case with a relay (at 125 m) to the throughput observed in Figure 12 (for a single link of 125 m), we see a drop of more than half. Without a relay, however, TCP communication between nodes separated by 250 m would be impossible; the relay clearly helps to extend range, but we should not expect the same throughput at nodes reached via relay as opposed to directly-accessible nodes.

5 Gateway-inspired experiment

The experiments summarised in this section were designed to provide insight into the behaviour of a MANET when two of its nodes were chosen to be relays (or gateways) on a three-link route, as depicted in Figure 18. A relay is a MANET node that repeats the source's information en route towards the destination. A gateway is a special relay with two radios and network stacks that connects two MANETs.

Here, we present some preliminary results from the experiment run with relays, with a caveat that these results need more data points to support any reliable conclusions.



Figure 18: Experiment plan for a three-link route MANET, containing two relays.

5.1 Scenario and objectives

As depicted in Figure 18, the scenario of this experiment consisted of a three-link MANET, in which two relay nodes repeated the source's information towards the destination. The experiment was run with multiple use cases classified in three sets of sub-experiments for the following objectives:

- determine throughput under varying co-channel interference conditions (varied by distance);
- determine en route bottlenecks vs. distance.

The benchmark distance, i.e., links a, b or c in Figure 18, was selected as 125 m, based on the experiments in Section 4.3, which is the distance at which the UDP loss rate began to increase significantly. Two sets of sub-experiments were then designed, which varied either the relay-relay distance (link b), or one or both of the distances from source/destination to neighbouring nodes (link c only, or links a and c) to distances of 75 m and 100 m.

For throughput measurements, the IRN Test app was used to generate TCP and UDP data at the source, with a specified destination. Every test was run for 120 s and tcpdump data was collected at each node along the route.

Below, we present a summary of the results we collected from these experiments.

5.2 Varying the distance between relays

This set of sub-experiments focused on different distances between the relays, i.e., setting link b to distances of 125, 100 and 75 m. The throughput results associated with UDP and TCP traffic are presented in Figure 19 and Figure 20, respectively.

For UDP traffic, we varied the data rates, i.e., 10 kBps, 100 kBps, 500 kBps, 750 kBps, 1 MBps, 3 MBps, 5 MBps. We observed that at 125 m, the relays did not support transfer of data rates higher than 500 kBps; this is seen in **Figure 21**Figure 21, where two instances of the same sub-experiments are shown and the instantaneous throughput remains below the offered load at 500 kBps. The UDP results suggest that the MAC layer (see Section 1.1.2) becomes congested when the data rate pushed through the MANET is increased. This is observed, for example at 100 m, where throughput increases with higher rates, up to 3 MBps, and then falls at 5 MBps. The low rates, i.e., 10 kBps and 100 kBps, were not tried at 100 m because the 500 kBps run achieved good throughput. Similarly, at 75 m, rates below 3 MBps were not tested.



Figure 19: MANET UDP throughput via a route with two relays vs. distance between relays.



Figure 20: MANET TCP throughput via a route with two relays vs. distance between relays.

As discussed in Section 1.1.2, for the TCP traffic, both the MAC layer and the transport layer were connection-oriented for reliable transfer of traffic; this extra overhead leads to lower throughput compared to the results from UDP traffic. This is evident when the relays were 75 m and 100 m apart, as shown in Figure 20. However, when the relays were 125 m apart, the TCP throughput is higher than that of UDP, contrary to the expected effect of overhead from the connection-oriented transport layer. The higher TCP throughput at 125 m may be attributable to the impact of greater distances between the relays, which reduces the impact of their co-channel interference. This is also evident in Figure 22, where the averages of the TCP and UDP throughput results are displayed.



Figure 21: Two instances³ of use case MANET UDP instantaneous throughput via a route with two relays vs. time.



Figure 22: MANET TCP and UDP average throughput via a route with two relays vs. distance between relays.

³ Runs A and B refer to use case numbers 11 and 12 in the experiment data sets, respectively.

5.3 Varying the distance between destination and source to their respective neighbour relays

In this set of sub-experiments, we varied the distance to the destination from its neighbour relay, i.e., link c lengths 125, 100 and 75 m, and the distances from source and the destination to their respective neighbouring relay, i.e., links a and c lengths 125, 100 and 75 m. The throughput results associated with TCP and UDP traffic are presented in Figure 23 and Figure 24, respectively.

We expected that co-channel interference would be a contributor to low throughput when the source and/or destination was close to its respective relay, i.e., the sub-experiment where links *a*, *b* and *c* were lengths 75 m, 125 m and 75 m, respectively. The figures show that good connectivity at these distances kept TCP and UDP throughputs relatively comparable. However at distances 100 m, 125 m and 100 m, respectively, both the TCP and UDP throughputs decreased, which is unexplained and requires further investigation. One possibility is co-channel interference, especially for TCP where the traffic volume was exacerbated by re-transmissions at both MAC and transport layers.



Figure 23: MANET TCP throughput via a route with two relays vs. distance from destination to its neighbouring relay (link c) and vs. distance from source and destination to their respective neighbouring relay (links a and c).





We expected that UDP would consistently provide higher throughput than TCP because UDP lacks the features of TDP's transport layer re-transmission, i.e., interference and collisions. However, we observed that the UDP throughput is generally less than that of TCP, which indicates that the benefits of the transport layer retransmissions are more than compensating for the packet losses experienced by UDP. These results are preliminary, and more data points are needed to draw any conclusions about throughput and relay distance from destination and/or source.

6.1 Objectives

The objective of this field experiment was to collect data from several MANET applications (e.g., TCP, UDP, file transfer, voice, short message, and SA) reaching their destinations via 1-hop, 2-hop and 3-hop routes in a real environment, and to eavesdrop on the wireless communication traffic over-the-air. The data collected will be used to support traffic analysis and other future research activities. The eavesdropped wireless traffic should contain detailed information from the physical layer up to the application layer. To get diverse traffic for analysis, three testing scenarios were used for generating and collecting 1-hop, 2-hop and 3-hop traffic respectively.

6.2 Plan and execution

To conduct the experiment, we put four Android phones on tripods in a line beside the road in zone 2 in Figure 2, with 125 m spacing, as shown in Figure 25, such that each phone had a direct link with its neighbour phones only. The OLSR routing protocol was running on each phone, using the IRN MC app, to create a routing table for traffic forwarding.

We set the first phone as the source node to generate and fire application traffic and others as 1-hop, 2-hop and 3-hop destinations depending on their distance from the source node. Three laptops were located near the source node, the destination node, and the middle of the source and destination nodes respectively, where the destination node could be one, two or three hops away, depending on each testing scenario. Wireshark was used on each laptop to eavesdrop on the over-the-air wireless traffic. Note that the laptops were not part of the testing MANET.

To generate application traffic, the following apps were installed on each phone: IRN MC, IRN test, IRN blue force tracking and Serval Mesh.

For testing and collecting TCP, UDP, and file transfer application traffic, the IRN MC and IRN test app were launched in the source and destination nodes. The traffic generation command was fired through IRN MC at the destination node and sent to the source node to generate related application traffic.

For testing and collecting IRN blue force tracking traffic, the IRN blue force tracking app was launched on each phone, thereby broadcasting its SA messages to the others.

For testing and collecting voice and short message traffic, Serval Mesh app was launched in the source and destination nodes and sent related traffic to each other.



Figure 25: The experimental area and physical setting for traffic analysis data capture.

6.3 **Preliminary results**

We collected most data as planned, and it will be analysed in future work. The field testing went smoothly and very well under most scenarios. Several software bugs were found during testing that were fixed later. For example, we found that we could not launch and run file transfer testing after TCP testing during the first day of field testing. Some bugs were caused by software updating issues, for example, we found that we could not receive blue force tracking message over two hops away. In addition, we did not collect much complex application traffic due to limited apps on our testing devices. In the future, we hope we can generate more complex application traffic from our traffic generator and replay them back to the network.

7 Lessons learned and future work

These experiments have enhanced our understanding of the effects of operational environments, and have given us the opportunity to stretch the capabilities of our current suite of network and cyber tools. In particular, these experiments proved valuable in providing us with a set of qualitative observations that give us a better "gut feel" for the operation of WiFi ad hoc links under various protocols and conditions.

We have learned a number of lessons for both operation and experimentation, as well as advanced requirements for software capabilities and directions for future experimental work.

7.1 Operational lessons learned

The dynamic nature of the WiFi receiver and the wireless environment make it difficult to predict meaningful things about the system in a reliable way. For example, to model and predict the probability of success of packet delivery in a wireless network, researchers have looked beyond RF propagation models and have turned towards experimentation with real radios operating in particular networks [22],[23].

Our findings follow and support what those in the wireless networking community have found, that simple abstract models of propagation are a poor match to empirical data [24]. For example, our measured RSSI values varied under local motion and differed between WiFi devices. In particular, non-line-of-sight channels caused significant inaccuracy in the measured RSSI. We find that in practical scenarios, RF measurements of the RSSI alone will provide poor predictions of distance.

These findings have an impact on the networked performance of the radios, as we found that throughput did not always match our expectations, which might be attributable either to varying channel conditions or to the complexities in the radio nodes themselves. For example, we observed that the "instantaneous throughput" of TCP was extremely variable even under conditions where the endpoints and environment were held static; the UDP packet delivery rate was much more stable under the same conditions.

We found that the limit on the usable range of a link was approximately 125 m in relatively clear line-of-sight conditions (these radios emit at approximately 10.5 dBm at 2.4 GHz), and beyond this the packet delivery rate of UDP dropped precipitously, and the TCP throughput degraded significantly as well. At 250 m, there was virtually no connectivity at all.

When communicating via a relay, higher throughput was achieved when the relay was closer to the centre of the nodes as opposed to closer to the edges (closer to the nodes themselves). This was attributed in part to co-channel interference and collisions between the two nearest nodes. It is not clear from our experiments where the bottleneck resides on a MANET route with two relays, given the complexity of factors at play.

The two main types of application data exchange used in wireless communication (UDP and TCP) experience different impacts of relays in network operation. Unicast UDP data transmission appears to experience congestion through re-transmission and throttling messages via RTS, CTS and ACK (more investigation is required), while TCP data transmission suffers congestion due to re-transmission and throttling messages, e.g., ACK. Both TCP and UDP traffic data in MANETs reveal the realities of the

protocols; experiments with one traffic type and lessons from its behaviour cannot be substituted directly for the other.

7.2 Experimental lessons learned

In addition to qualitative observations, we learned a number of lessons and best practices that will help inform future experimental work.

Pre-testing in the lab to make sure every app functions as expected, including being run multiple times back-to-back, is time well-spent. Preparing a detailed testing plan and documentation helps the field testing go smoothly. Preparation of the testing field before testing also saves time. Standardised methods for note taking to ensure accurate and thorough recording of conditions, rationale and results should be used, as plans need to be adjusted on-the-fly. It was also found that limiting testing time to half-a-day at a time is a more efficient use of the team's time, as this allows for data validation and preliminary analysis between measurement collections.

Even though we conducted throughput measurements over time periods of up to 120 s with the hopes of averaging out variability due to fading, it is clear that multiple, longer measurements are required to draw meaningful quantitative conclusions. This is especially true of TCP, which includes backoff intervals and timers on the scale of minutes (in some cases), meaning that single measurements cannot be relied on if the TCP session misses enough packets to cause a re-transmit and backoff cycle.

Interpreting and fully explaining TCP throughput measurements and coming to firm conclusions would require monitoring more than just the throughput—it would be valuable to have logs at various levels of the protocol stack to know when/if TCP is retransmitting packets and/or entering a backoff cycle.

Simple unicast UDP is not an appropriate choice to measure the raw packet delivery rate; although UDP is a connectionless protocol that will not "re-try" if packets are missed, the underlying MAC protocol (WiFi in this case) will nevertheless send re-transmissions if packets are missed by the receiver at layer 2. We learned this in early testing (results not reported here) and thus, for some of the testing reported here, we used broadcast UDP, in which the MAC will not re-transmit since no MAC-layer acknowledgements are expected in the case of a broadcast. One difficulty with this strategy, however, is that the broadcast packets are not typically forwarded by relay nodes, meaning that to measure the packet drop rate across a relay requires additional effort.

We noted catastrophic connection-oriented traffic failures during some of the experiments, and sometime even short-range UDP tests failed. Analysis of the collected data during failures leads us to suspect they were due to the inconsistent alignment of WiFi radio modes among the nodes along the route, i.e., nodes using 802.11b/g/n. If the source and its neighbouring relay attempted to use a mode that did not match that of the destination and its relay, even for a short duration, then the connection would fail. This inconsistent alignment of en route radio modes is suspected to be exacerbated by increasing the number of relays in a route. This could be detected using Wireshark, however our implementation of that software is limited to b/g messages only. Unless the nodes were forced to be in 802.11g mode, the connection tended to be established on the 802.11n mode, by default, and we would not capture them. The data suggests that gateways may help manage consistent alignment of en-route radio modes, as well as co-channel interference, because the gateways, unlike relays, would be equipped with two radios partitioning a route to three point-to-point connections.

For experiments involving mobility, we note that to quantitatively measure its effects requires a rigorous test environment where the speed and path traversed can be carefully controlled to ensure repeatability.

In many cases, it would be useful to perform preliminary tests (as we did here) to identify areas for further investigation and then follow up with more detailed testing to zero in on areas of interest.

Overall, it is difficult to derive specific conclusions about the factors related to the more complex MANET operations, such as the bottlenecks we observed. To understand these operations more clearly, we must invest in designing specific experiments, each with multiple data points, to provide more rigour.

7.3 Software development requirements

We experienced difficulties in the operation of the IRN Test app, which relies on iPerf for certain measurements [25]. We were able to confirm that the tool provides the correct average data rate by cross-referencing with the tcpdump logs. However, the IRN Test app was found to be unreliable for repeated, sequential TCP tests in the field at long distances; to ensure repeatability, the tool needs to be reset and reconfigured after every test. There were also oddities observed in the test app logs when sequential tests were performed, in which the timestamps recorded went beyond the test duration specified in the test condition. A more robust test app to support future experimental work would be useful to alleviate this shortfall for data collection.

To support data collection for future traffic analysis and other research activities, and to be able to demonstrate the impact of data latency or loss, a traffic generation app would be useful. This should provide a variety of application data from stored files, such as voice via UDP, video via UDP, data file via TCP, and data file FTP via TCP.

WiFi incorporates an automated mode selection (802.11b/g/n), whereby the mode and data rate changes according to the instantaneous conditions. The ability to constrain modes (802.11b/g/n) and data rates within the test app would eliminate some of the complexity in evaluating network performance.

7.4 Future experimental work

Further experimentation is needed to support the Tactical Network Operations project, and subsequent Electromagnetic (EM) Cyber projects. In particular, we should develop an increasingly sophisticated set of experiments designed to collect rigorous data sets to increase our understanding of network operation under multiple relays, gateways and/or mobility. This will help develop new techniques to make networks more resilient, as well as identify opportunities for exploitation.

We noted several areas where we should repeat the experiments to obtain additional data sets for analysing consistency and variation within operational environments. In particular, we will repeat link data rate tests with fixed data rates so we can eliminate one of the variables (Section 4.3.1), and we will repeat this with better-controlled mobility (Section 4.3.2).

When a relay is introduced to the MANET, the testing becomes significantly more complex, and the test conditions need to be controlled carefully. We noted the need for more testing, in particular to verify the impact of different distances between the source and relay, and relay and destination. These experiments should be performed with fixed data rates and using Wireshark as a diagnostic tool to assess the impact of

MAC and TCP re-transmissions. This should also aid with determining when co-channel interference and collisions cause degradation in performance.

Beyond the experiments reported here, future experiments will replace the relay nodes of Section 5 with gateways with the objective of understanding how the performance changes—this will provide insights into when and where relays and gateways can be used most effectively.

As our research in traffic analysis techniques progresses, we will have a requirement for more complex over-the-air data captures, using more sophisticated network topologies and a larger variety of mixed traffic data.

In our research work, we use the simulation tool EXata [26] to evaluate MANET concepts. It would be instructive to understand how well this tool emulates real operating conditions. From our results in Section 3, we do not expect very close alignment with path loss measurements, but we hope that the impacts of MAC and transport layer protocols will be well represented. Designing experiments for this purpose will be challenging; like the ones reported here, they will have to start simply, and build up to more complex scenarios.

These experiments were performed in the simplest outdoor environment: unobstructed line-of-sight. Future experimental work should be planned in more challenging environments, such as urban and outdoor-indoor, where the propagation conditions and connectivity between nodes is even less predictable, and the effects of mobility will be more pronounced.

8 Conclusions

This report describes the field experiments completed by the Tactical Network Operations team in the summer of 2017, along with some preliminary analysis of the data. These experiments were informative to our research, as they have shed light on challenges experienced in fielded networks that are not seen in a laboratory setting. The experiments were built upon each other, starting at the physical layer and developing to static and mobile networks with relays.

The first key finding is that there is a significant variation in measured power levels (RSSI) among devices, and over a local region. These effects make geolocation of these low power devices to any useful resolution unrealistic using only RSSI measurements. These results also highlighted the challenges in relying on standard propagation models, even in fairly ideal conditions.

Measuring the performance of different types of data is complicated by many dynamic components, including the radio interfaces themselves. While future tactical radios will probably use some degree of automated data rate selection, the implications of that added level of dynamism in an already dynamic network are complex. Mobility and relays further complicate the characterisation of network performance.

The complexities of dynamic networks must be better understood, so that robustness can be built into the network operation, and the dynamism may be advantageous for increased resilience and defence. It is also important to recognise these dynamics in a target network, so that it can be correctly characterised for exploitation. To this end, further experiments are recommended, both for understanding the network operation and for data collection to characterise target networks.

References

- [1] Project Charter: Tactical Network Operations (05ab), last updated October, 2017.
- [2] "Tactical wearables and the self-organizing network," video available at <u>www.drdc-</u> <u>rddc.gc.ca/en/dynamic-article.page?doc=tactical-wearables-and-the-self-organizingnetwork/hyyl17cg</u> Accessed March 19, 2018.
- [3] T. Clausen and P. Jacquet, "Optimized link state routing protocol (OLSR)," RFC3626, Internet Engineering Task Force (IETF), <u>www.ietf.org/rfc/rfc3626.txt</u>, October 2003. Accessed March 19, 2017.
- [4] OLSRd, <u>www.olsr.org</u>, Accessed March 19, 2018.
- [5] Y.C. Huang, S. Bhatti, and D. Parker, "Tuning OLSR," *Int. Symp. Pers.*, Indoor and Mob. Radio Commun. (PIMRC), 2006.
- [6] C. Perkins, "Ad hoc on-demand distance vector (AODV) routing," RFC3561, Internet Engineering Task Force (IETF), <u>www.ietf.org/rfc/rfc3561.txt</u>, July 2003. Accessed March 19, 2018.
- [7] M. Salmanian, W. Pase, J. D. Brown, C. McKenzie. "A Gateway Prototype for Coalition Tactical MANETs," In: Zhou Y., Kunz T. (eds) Ad Hoc Networks. Lecture Notes of the Institute for Computer Sciences, *Social Informatics and Telecommunications Engineering*, vol 184. Springer, Cham, 2017.
- [8] LG Nexus 5 smartphone, <u>www.lg.com/ca_en/cell-phones/lg-D820-White-nexus5</u>. Accessed March 19, 2018.
- [9] Panda Wireless 300Mbps Wireless N USB adapter, www.pandawireless.com/Products%20%7C%20Panda%20Wireless.html. Accessed March 19, 2018.
- [10] TP-Link 150Mbps high gain wireless USB adapter, TL-WN722N, <u>www.tp-link.com/us/products/details/cat-5520_TL-WN722N.html</u>. Accessed March 19, 2018.
- [11] C. McKenzie, "TEC3 REALTIME blue force tracking: development and demonstration report," Internal DRDC report, April 2014.
- [12] J.D. Brown, M. Salmanian, D. Simmelink, H. Tang, and R. Song, "Tactical edge cyber command and control (TEC3) concept: a vision for network situational awareness and network command and control at the tactical edge," Defence Research and Development Canada, Scientific Report DRDC-RDDC-2014-R155, December 2014.
- [13] The Serval Project, <u>www.servalproject.org</u>. Accessed March 19, 2018.
- [14] Kali Linux NetHunter, <u>www.kali.org/kali-linux-nethunter</u>. Accessed March 19, 2018.
- [15] Tcpdump, <u>www.tcpdump.org</u>. Accessed March 19, 2018.

- [16] WiresharkTM, <u>www.wireshark.org</u>. Accessed March 19, 2018.
- [17] Riverbed AirPcapTM, <u>www.riverbed.com/products/steelcentral/steelcentral-riverbed-airpcap.html</u>. Accessed March 19, 2018.
- [18] T.S. Rappaport, Wireless Communications: Principles and Practice, 2nd edn. Prentice-Hall, Englewood Cliffs (2001).
- [19] Certification Test Report for Model LG-D820, Report Number 13U15420-2A, July, 2013 Online at the Federal Communications Commission, Office of Engineering and Technology website, FCC ID ZNFD820. Accessed January 18, 2018.
- [20] A. Roxin, J. Gaber, M. Wack, and A. Nait-Sidi-Moh, 2007, November. Survey of wireless geolocation techniques. In Proc. IEEE Globecom Workshops, 2007.
- [21] S. Wang, R. Inkol, and B.R. Jackson. May 2012. Relationship between the maximum likelihood emitter location estimators based on received signal strength (RSS) and received signal strength difference (RSSD). In *Proc. Biennial Symposium on Communications (QBSC '12)*, pp. 64–69, 2012.
- [22] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris, Link-level measurements from an 802.11b mesh network. In Proc. Conf. on Applications, Technologies, Architectures, and Protocols for Computer Communications (SIGCOMM '04), pp. 121–132, 2004.
- [23] K. Jamieson, B. Hull, L. Miu, and H. Balakrishnan, Understanding the real-world performance of carrier sense. In *Proc. ACM SIGCOMM workshop on Experimental approaches to wireless network design and analysis* (E-WIND '05), pp. 52–57, 2005.
- [24] D. Kotz, C. Newport, R.S. Gray, J. Liu, Y. Yuan, and C. Elliott, 2004. Experimental evaluation of wireless simulation assumptions. In Proc. 7th ACM Int. Symp. on Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWiM '04), pp. 78–82, 2004.
- [25] iPerf The network bandwidth measurement tool, <u>www.iperf.fr</u>. Accessed March 19, 2018.
- [26] EXata network emulator, <u>www.scalable-networks.com/exata</u>. Accessed March 19, 2018.

	DOCUMENT CONTROL DATA									
	*Security markings for the title, authors, abstract and keywords must be entered when the document is sensitive									
1.	ORIGINATOR (Name and address of the organization preparing the doc DRDC Centre sponsoring a contractor's report, or tasking agency, is enter 8.)	ument. A red in Section	2a.	SECURITY MAR (Overall security n special supplement	KING narking of the document including tal markings if applicable.)					
	DRDC – Ottawa Research Centre Defence Research and Development Canada	CAN UNCLASSIFIED								
	Ottawa, Ontario K1A 0Z4	Ing Avenue Intario K1A 0Z4								
	Canada		NON-CONTROLLED GOODS DMC A							
3.	. TITLE (The document title and sub-title as indicated on the title page.)									
	Mobile network field experiments: Summer 207	17								
4.	AUTHORS (last name, followed by initials - ranks, titles, etc., not to be	used)								
	Willink, T.; Watson, S.; Brown, D.; Salmanian,	M.; Song,	R.;	Li, M.; Droui	n, PL.					
5.	DATE OF PUBLICATION (Month and year of publication of document.)	6a. NO. OF (Total pa excludin verso pa	PAGI ages, ir g DCE ges.)	ES neluding Annexes,), covering and	6b. NO. OF REFS (Total references cited.)					
	April 2018		48	3	26					
7.	DOCUMENT CATEGORY (e.g., Scientific Report, Contract Report, Scientific Report, Contract Report, Science Repor	ientific Letter.)								
	Scientific Report									
8.	SPONSORING CENTRE (The name and address of the department pro-	ject office or la	iborato	ory sponsoring the r	esearch and development.)					
	DRDC – Ottawa Research Centre Defence Research and Development Canada 3701 Carling Avenue Ottawa, Ontario K1A 0Z4 Canada									
9a.	 a. 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.) b. CONTRACT NO. (If appropriate, the applicable number under which the document was written.) 									
10a	0a. DRDC PUBLICATION 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 NO(s). (Any other numbers which may be assigned this document either by the originator or by the sponsor.)									
	DRDC-RDDC-2018-R032									
11a	. FUTURE DISTRIBUTION WITHIN CANADA (Approval for further di	ssemination of	the do	cument. Security cl	assification must also be considered.)					
	Public release									
11b	11b. FUTURE DISTRIBUTION OUTSIDE CANADA (Any limitations on further dissemination of the document, other than those imposed by security classification.)									

12. 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.)

This report describes field experiments undertaken in the summer of 2017 using network and data applications developed under the Tactical Network Operations project. The aims of the activity were to develop a field testing capability and to obtain insight into the performance of wireless networks in static and mobile operation. We completed a number of experiments, building on each other to understand the impact of the physical environment, the network topology and the applications. Preliminary data analysis has shown that path loss, even in relatively ideal unobstructed line-of-sight links, is highly variable over very small distances and across different devices: this impacts the potential for power-based emitter geolocation. Additionally, these experiments have revealed some of the complexities of application data transfer in mobile ad hoc networks, which involve data being relayed through other nodes en route from source to destination. In particular, we found that the traffic and MAC layer protocol messaging appeared to cause interference and collisions when routed via relays, resulting in unexpected degradations in throughput and packet loss performance. Future work to investigate this further will require more detailed over-the-air observations in addition to logging on the radio nodes. The data collected during these experiments will continue to support the project's R&D activities, and further experiments will be planned to gain deeper insights and to investigate performance in more challenging environments.

Le rapport décrit des expériences menées sur le terrain pendant l'été 2017 à l'aide d'applications réseau et de données développées dans le cadre du Projet d'opérations de réseau tactique. Elles visaient à développer les capacités d'essai sur le terrain et de mieux comprendre le rendement des réseaux sans fil en fonctionnement statique et mobile. Nous avons mené à bien plusieurs expériences, en mettant à profit les résultats des premières pour les expériences ultérieures, afin de comprendre les répercussions sur le rendement de l'environnement, de la topologie du réseau et des applications. Une analyse préliminaire des données démontre que l'affaiblissement de propagation, même pour les liens en visibilité directe, passablement idéaux car libres d'obstructions, varie considérablement sur de très petites distances et d'un dispositif à l'autre; cela ne peut que jouer sur la géolocalisation des émetteurs en fonction de la puissance. Ces expériences ont aussi mis en lumière certains des aspects complexes des communications de données interapplications dans les réseaux mobiles spéciaux, car ces données sont relayées par des nœuds intermédiaires entre la source et la destination. Plus précisément, nous avons observé que les messages des protocoles de couche MAC semblent entraîner interférences et collisions s'ils sont transmis par relais; il en résulte la perte de paquets et une dégradation inattendue du rendement. Pour étudier cela plus avant, il nous faudra observer les radiocommunications plus en détail et journaliser l'activité des nœuds radio. Les données recueillies pendant ces expériences appuieront les activités de R et D du projet, et nous prévoyons d'autres expériences afin d'approfondir nos connaissances et d'étudier le rendement des réseaux en environnements plus difficiles.

13. 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.)

Wireless networks; EM cyber