Experiences in Deploying a Wireless Mesh Network Testbed for Traffic Control

Kun-chan Lan
National Cheng Kung University, Taiwan
klan@csie.ncku.edu.tw

Zhe Wang, Mahbub Hassan and Tim Moors
University of New South Wales
{zhewang, mahbub}@cse.unsw.edu.au, t.moors@unsw.edu.au

Rodney Berriman, Lavy Libman, Maximilian Ott, Bjorn Landfeldt and Zainab Zaidi
National ICT Australia
{Rodney.Berriman, Lavy.Libman, Maximilian.Ott, Bjorn.Landfeldt and Zainab.Zaidi}@nicta.com.au

ABSTRACT

Wireless mesh networks (WMN) have attracted considerable interest in recent years as a convenient, flexible and low-cost alternative to wired communication infrastructures in many contexts. However, the great majority of research on metropolitan-scale WMN has been centered around maximization of available bandwidth, suitable for non-real-time applications such as Internet access for the general public. On the other hand, the suitability of WMN for mission-critical infrastructure applications remains by and large unknown, as protocols typically employed in WMN are, for the most part, not designed for real-time communications. In this paper, we describe the Smart Transport and Roads Communications (STaRComm) project at National ICT Australia (NICTA), which sets a goal of designing a wireless mesh network architecture to solve the communication needs of the traffic control system in Sydney, Australia. This system, known as SCATS (Sydney Coordinated Adaptive Traffic System) and used in over 100 cities around the world, connects a hierarchy of several thousand devices — from individual traffic light controllers to regional computers and the central Traffic Management Centre (TMC) — and places stringent requirements on the reliability and latency of the data exchanges. We discuss our experience in the deployment of an initial testbed consisting of 7 mesh nodes placed at intersections with traffic lights, and share the results and insights learned from our measurements and initial trials in the process.

Categories and Subject Descriptors

C.4 [**PERFORMANCE OF SYSTEMS**]: Reliability, availability, and serviceability

General Terms

Experimentation

Keywords

Wireless Mesh Network, Traffic Control, Deployment

1. INTRODUCTION

Adaptive traffic control systems are employed in cities worldwide to improve the efficiency of traffic flows, reduce average travel times and benefit the environment via a reduction in fuel consumption. One of the main and most common functions of such systems lies in adaptive control of traffic lights. This ranges from simple lengthening or shortening of green and red light durations in an intersection according to the actual presence of cars in the respective lanes, to coordination of green light phases among neighboring intersections on main thoroughfares. This adaptivity is made possible with the use of sensors (typically in the form of magnetic loop detectors embedded under the road pavement) that feed data to roadside traffic light controllers, and a communications infrastructure that connects among the intersections and a traffic management centre, as well as, in some cases (typically in large cities), a hierarchy of regional computers (RC) that perform the control decisions for respective portions of the system.

Traditionally, the communications layer of traffic control systems has been based on wired connections, either private or leased from public telecommunications operators. While for many years such leased lines (operating at 300bps) have served their purpose well, they have several shortcomings, such as a significant operating cost, inflexibility, and difficulty of installation in new sites. In certain cases, alternative solutions, operating over public infrastructure, have been deployed for specific sites where private or leased lines were not a viable option; these ranged from ADSL, regular dialup, or cellular (GPRS). However, using public network for traffic control could suffer from inconsistent delay jitters and reliability issues. For example, previous experimental studies [1] have shown GRPS links could have very high RTTs (>1000ms), fluctuating bandwidths and occasional link outages.

In recent years, there has been considerable interest in wireless mesh networks and their deployment in metropolitan areas, from both a commercial and a research perspective [2]. Trials in several major cities in the US (e.g. Philadelphia, New Orleans, and others [3, 4]) and worldwide (e.g. Taiwan [5]) have shown mesh networks to be a viable tech-

nology that can compete well with alternative "last-mile" connectivity solutions to the public. Correspondingly, most of the research on metropolitan-area wireless mesh networks (MAWMN) has focused on maximising the throughput that can be extracted from them, in the anticipation that their major use will be public, for purposes such as accessing the Internet or conducting voice calls [6]. On the other hand, little attention has been directed to the aspects of reliability and latency, which are most important if MAWMN are to be considered for replacement of mission-critical infrastructure, such as traffic control system communications.

The Smart Transport and Roads Communications (STaR-Comm) project at National ICT Australia (NICTA), started in August 2005, sets out to develop protocols that enhance the reliability and reduce the latency of mesh networks, and thereby enable them to be used as the communications layer of traffic control systems. In this paper, we describe the testbed that has been built in the first stage of this project. Our initial testbed covers seven traffic lights in the suburban area of Sydney. These intersections are chosen because they represent a typical suburban area with lots of traffic, foliages, pedestrians and high-rise residential buildings. In addition, the inter-node distance (ranging from 200 to-500m) is representative of 90% of the distance between traffic controllers in the Sydney CBD (Central Business District) area. In the next phase, we plan to extend our testbed to 15-20 nodes. Our nodes have been custom-built to meet the need of our research.

The contribution of this paper are three-fold. First, to the best of our knowledge, our work is one of the first efforts to study the feasibility of using wireless mesh networking for traffic control. Second, we describe the details of our testbed implementation and some experiences we gained during the deployment of the testbed in an urban environment. Finally, we present some initial measurement study of link characteristics of different wireless and wired technologies used in our testbed (including the use of 900MHz, 2.4GHz and 3.5GHz radios and Ethernet-over-powerline). Although our results are still very preliminary, they are useful to serve as a reality check toward the goal of applying wireless mesh networking to traffic control applications.

The rest of this paper is structured as follows. In section 2, we describe the details of SCATS, the traffic control system used in Sydney and many other cities worldwide, and its communication requirements. We describe related work in Section 3. Section 4 presents a simple analysis of the topology of traffic lights in the Sydney metropolitan area, and in particular the dependence of the degree of connectivity of the mesh on the radio range. Section 5 describes the details of our testbed implementation. We present some initial measurement results of link characteristics of different radio technologies used in our testbed in section 6. Section 7 discuss the experiences we gained during the deployment of our testbed. We conclude the paper and discuss the future work in section 8.

2. SCATS OVER WIRELESS

In this section, we first describe the details of SCATS (Sydney Coordinated Adaptive Traffic System) and its communication requirements. We then discuss the benefits and research challenges when running SCATS on a wireless mesh network.

2.1 The SCATS traffic management system

Developed and maintained by the Roads and Traffic Authority (RTA, formerly Department of Main Roads) of the state of New South Wales, the Sydney Coordinated Adaptive Traffic System (SCATS) is one of the most popular traffic management systems used worldwide. Its main task is to adjust, in real time, signal timings in response to variations in traffic demand and system capacity. Real-time data from traffic controllers are collected and transported to a central traffic management centre (TMC) for analysis and optimum control of road traffic. The performance of SCATS, therefore, depends critically on the capabilities of the underlying communication system that transports roadside data to and from the TMC.

The existing communication system of SCATS relies strongly on third-party wired infrastructure (provided by Telstra, Australia's largest telco). The bulk of the communications to the intersections, namely the traffic light controllers and vehicle detectors, are predominantly made using serial point-to-point connections over standard voice-grade telephone lines, using 300bps modems. This is also the most common method of connecting between the TMC and other low-bandwidth devices, including variable message signs, variable speed limits, ramp meters, and over-height detectors.

At the core of the SCATS operation is a periodic exchange of messages between the controlling computer and each and every intersection (via the point-to-point links). This exchange happens every $1sec$, and is initiated by the computer which sends to the intersection's local controller a command message, instructing it about the next phase it should switch to and the timing of that switch. The controller, in turn, is required to reply with an acknowledgement, which includes information from the intersection's sensors. If an acknowledgement is not received within $1sec$ from the time the command message is sent, it is retried once; after the second time an acknowledgement fails to arrive, the communications link is declared failed, and SCATS instructs all controllers at the respective cluster of intersections to fall back into a 'default' self-controlling mode, where decisions about the timing of green light phases are made locally and independently. Likewise, a controller will fall back to this mode upon not receiving a command message. Once triggered, a controller will stay in the self-controlling mode for at least 15 minutes; if another communications failure happens during this time, the duration of this mode will be extended by another 15 minutes, and so on. Obviously, the self-controlling mode, where the decisions at intersections are uncoordinated, can lead to a severely suboptimal traffic control, particularly in a busy thoroughfare during rush hour. Accordingly, though the bandwidth required from the communication links is quite low (comfortably handled by 300bps modems), the $1sec$ latency is critical for an efficient operation of the system.

The currently used SCATS infrastructure, based on wired communications, suffers from the following problems:

- **Slow installation and inflexibility**. In most cases, installing a new line at a road site (especially a remote site) involves earth excavation, which is very slow and with adverse effects on existing infrastructure.
- **High capital and operating cost**. The installation of a wired connection at a new site, or repairs at an existing one, carries a high cost due to the material

and labour required. More importantly, the ongoing fees for leasing the wires from the telephone company run very high; currently, RTA pays nearly A$40 million annually to Telstra in leasing fees for connecting the traffic signals and other roadside devices to SCATS.

- **Low bandwidth**. Modem-based leased lines support bandwidth less than 32 Kbps. While these low-bandwidth telephone lines are adequate for connecting traffic signal sensors, they cannot provide adequate support for connecting high-bandwidth applications, e.g. high-resolution video cameras, that increasingly becoming necessary to effectively monitor traffic pattern on our roads.

2.2 Going wireless

With wireless solutions, there is no cabling involved. Wireless can therefore provide fast installation and exceptional flexibility. Cost can be reduced significantly by building a private wireless network, because there will be no monthly charges to be paid to telephone company (some small license fee may apply). Moreover, the installation cost will be low because there will be no cabling-related labour. The cost issue is, in fact, the major concern for most road authorities as well as the main factor that motivated RTA and us to start this work. Finally, it should be noted that recent advances in wireless technology provide bandwidth that is more than adequate for connecting many high-resolution roadside cameras to SCATS.

One possible option for going wireless is to build a dedicated RTA wireless network using widely available, standards-based, low-cost wireless technologies, e.g. IEEE 802.11x and 802.16x. 802.11x equipment is cheaper, less complex, and operates entirely in the unlicensed spectrum (no licensing fee). On the other hand, 802.16x is more reliable (has multiple carrier frequencies to avoid interference), has longer range, and better features to cater for a diverse range of communication needs of future roadside equipment. In addition, it is possible to operate 802.16x in both license and unlicensed spectrums.

Despite of its enormous benefits, there are several challenges when roadside ITS equipment are connected via wireless media:

- **Latency**. Wireless can potentially increase latency. For example, IEEE 802.11x, uses a common wireless channel (it is cheaper to share channel) among many contending devices causing potential conflict. To avoid such conflicts, some form of medium access control (MAC) is implemented by these technologies. MAC introduces some delay before data can be transmitted on the wireless channel.

- **Reliability**. Wireless signals are susceptible to interference from other signals in the vicinity operating in the same or adjacent spectrum. Given that ITS equipment is deployed in public area, such interference will be the norm rather than exception. Interference can corrupt messages transmitted over the wireless medium. Some frequencies do not work well (or at all) if there is no direct line-of-sight between the two communicating end points. In a dynamic context of public roads, roadside equipment may frequently face line-of-sight problems due to transient obstructions, e.g. a high vehicle carrying a tall crane etc. Also in vehicle-to-roadside communications, a car in the near-lane may obstruct communication between a far-lane car and the roadside equipment. Temporary outages, i.e., periods when no wireless signal is available, therefore, is a real issue to deal with.

- **Security**. What makes wireless so vulnerable is the fact that the attacker does not have to gain physical access to the channel from any predefined access point. Roadside wireless components are well within the wireless range of passing motorists and pedestrians, which make them vulnerable to intrusion, denial of service, and other forms of security threats.

- **Scalability**. As mentioned earlier, wireless systems are sensitive to interference from other communicating devices operating in the vicinity. Additionally, if a common wireless channel is shared among all devices within a given area (cell), the MAC delay increases rapidly as the number of competing devices increases. Another scalability issue arises from the processing overhead that is required at a central radio base-station. The more remote radios there are in communication with the central radio, the more processing that must take place. The radio controller at the base-station will simply not be able to process all incoming radio signals if there are too many of them.

3. RELATED WORK

Roofnet [7] is an experimental 802.11b/g mesh network built by MIT. Each node in Roofnet has an antenna installed on the roof of a building. Aguayo et al. [8] analyzed the link-layer behavior on the Roofnet testbed and described the impact of distance, SNR and transmission rate on the packet loss. While Roofnet's propagation environment is characterized by its strong Line-of-Sight component, our work differs from the prior work in that our links are generally heavily obstructed [1]. In addition, our planned deployment strategy is different from the unplanned topology in Roofnet.

[1]For example, our antenna is mounted at a height of about 4 meters from the ground. But the trees on the road are typically higher than 7 meters.

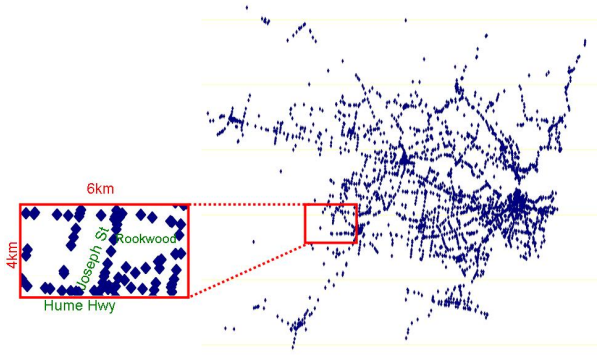

Figure 1: Location of traffic controllers

Similar to our work, The WAND project [9] has built a multi-hop wireless testbed in the centre of Dublin. They have 11 nodes mounted on traffic lights along a 2km route in urban area. However, their topology is simpler than ours (i.e. a chain topology) and the measurements they performed on their testbed were relatively limited.

TFA project [10] aimed to provide broadband access to low income community in Houston area via wireless mesh network technology. Their architecture consist of two wireless tiers: an access tier to connect homes, businesses, and mobile users to the infrastructure, and a back-haul tier to forward traffic to and from the wired entry point.

Jardosh et al. [11] discussed the correlation of link reliability with the frame retransmissions, frame sizes and data rate by collecting trace data from a structured 802.11b network during a international conference. They concluded that sending smaller frames and using higher data rates with a fewer number of frames improves the performance of congested network.

All the previous studies have been centered around maximization of available bandwidth for non-real-time applications such as broadband access for the general public. On the other hand, to the best of our knowledge, our work is the first to focus on using wireless mesh networking for traffic control. which places stringent requirements on the reliability and latency of the data exchanges.

4. A SIMPLE ANALYSIS OF THE SYDNEY TRAFFIC LIGHT TOPOLOGY

This analysis was based on data provided by RTA, indicating the latitude and longitude of traffic controllers. Figure 1 shows points at each traffic controller location. As shown in Figure 1, there are around 70 controllers in a 4km × 6km area.

To understand the effect of radio range on the degree of connectivity when the traffic controllers are forming a mesh network, we calculate the shortest distance (assuming that the radio has a circular radio range and have perfect coverage in that range) between every pair of traffic controllers, and the output was then sorted so that for each controller, its neighbours were listed (from nearest to furthest) with the distance to each neighbour. We then processed this data to determine how many neighbours of each traffic controller were within a specified range (from 100m to 1250m). The results of this analysis are shown in Figure 3, which shows on the y-axis how many traffic controllers had $0, 1, 2, 3, \ldots$ neighbours when a given radio range is assumed (x-axis).

The results shown in Figure 3 provide a rough indication of what radio range is needed if we are aiming to interconnect a certain number of nodes with each node having a certain degree (number of neighbours within range). For example, if we seek to interconnect 90% of the nodes (accepting that some alternative technology may be needed for the minority 10% of nodes) and require that each node has three neighbours (to provide redundancy and hence fault tolerance), then we require a radio technology with range of at least 1km. Note that, while city environments may have large densities of traffic controllers in both (lat/long) dimensions, in suburban environments controllers (particularly those that are important to synchronise with communication links) often lie linearly along main arterial roads, with few controllers in close range orthogonal to the main arteries. Neighbours that form a line would not provide the same level of fault tolerance as those that are better separated angularly around a controller, since the links are less likely to fail independently.

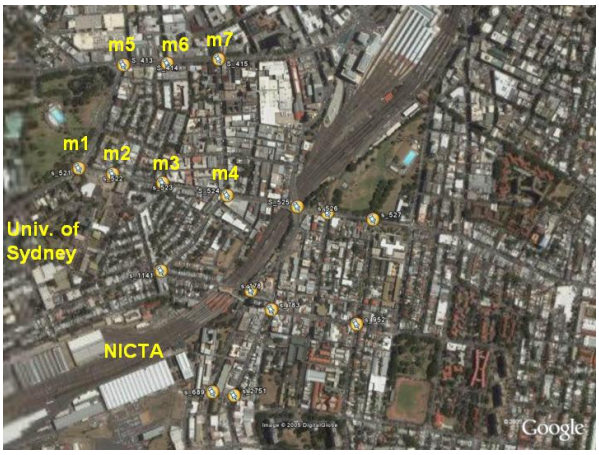

Figure 2: Map of Intersection locations

5. TESTBED

In this section, we provide the details of our testbed. We first describe the environment that the testbed is located. Next, the hardware used for the initial seven nodes and the software installed on each of these nodes are discussed.

5.1 Environment

The testbed is located in the Sydney CBD (Central Business District) area. We selected seven intersections initially to deploy the testbed, as shown in Figure 2 (specifically, intersection number m1 to m7). We plan to extend our testbed to 15-20 nodes in the next phase. We use a number of custom-build embedded PCs with multiple wireless interfaces. The nodes are mounted on the traffic lights at a height of about 2-3m from the ground, and distributed along the streets in the form of rectangle covering an area of 500 × 1000 square metres at a distance of 200-500m apart. None of the nodes is in a clear line of sights of its neighbor-

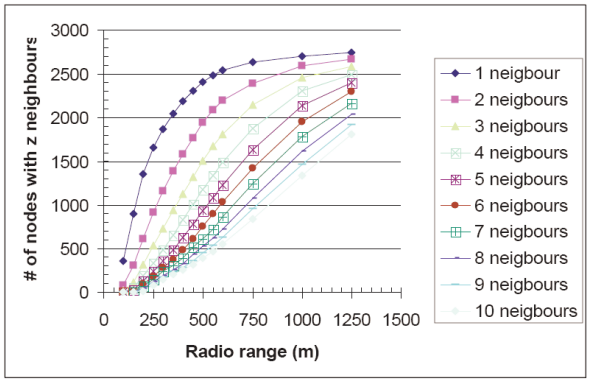

Figure 3: Numbers of neighbours within certain radio range of RTA controllers.

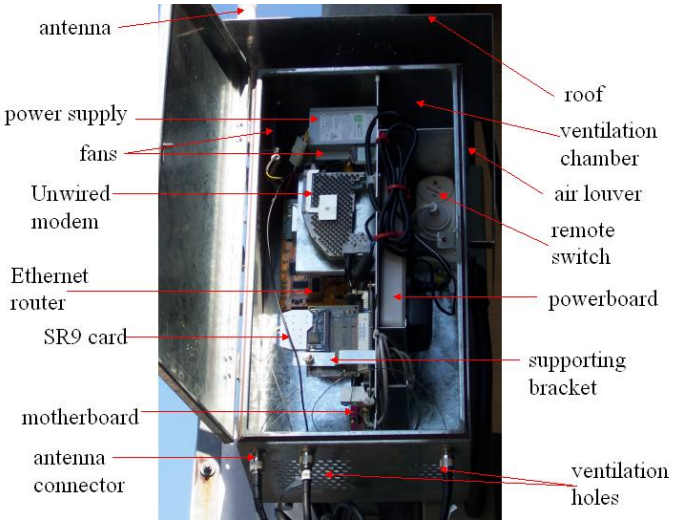

Figure 4: Hardware Component

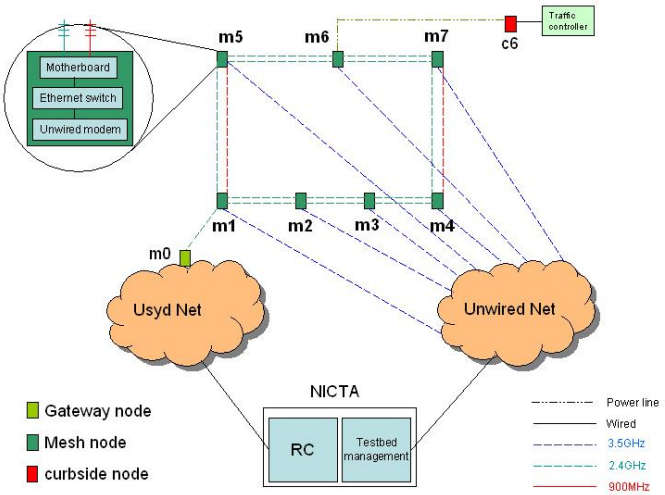

Figure 5: Testbed topology

ing nodes. The node at intersection m1 is connected to a gateway node in University of Sydney.

The streets where the network is deployed are about 10-20m wide and surrounded by building at least two stories high. The majority of these buildings are made of concrete and steel that block the propagation of radio signals into the neighboring streets. All these streets have busy public traffic during business hours. Most of the vehicles on the street have a height of less than 2.5m. But some double-decker buses (such as Sydney Explorer Bus) or truck can have a height of more than 5m.

5.2 Channel characteristics

Wireless channels can be coarsely characterized by its path loss exponent. Pathloss described the attenuation experienced by a wireless signal as a function of distance. The amount of variations in pathloss between similar propagation scenarios is called shadowing. In other words, shadowing represents the difference between the signal power at different points in the same environment with the same estimated pathloss. Prior study [12] showed that shadowing can be modeled as a zero-mean Gaussian random variable. Specifically, one can predict the received signal power at a given distance d with the following formula:

$P_{dBm}(d) = P_{dBm}(d_0) - 10\alpha\log_{10}(\frac{d}{d_0}) + \epsilon$

where α is the pathloss exponent, ϵ is the shadowing and d_0 is the reference distance.

The prior work [12] suggested that the pathloss can range from 2 to 5 for outdoor urban environment. To accurately estimate the range and reliability of mesh links, we performed extensive measurements at various locations and distances to find our environment's path loss exponent and shadowing. Such physical level measurements are important for an efficient deployment (i.e. overestimating pathloss can result in overprovioning network while underestimating pathloss can produce disconnected network). By using linear regression, we find our environment has a pathloss α =3.1 and shadowing ϵ = 7.2. Note that the observed pathloss in our environment is significantly lower than the suggested urban pathloss of 4 in the literature [12].

5.3 Hardware

The hardware components used for the nodes of our initial testbed are all off-the-shelf products including the following, as shown in Figure 4. All the components are mounted on two sides of a metal plate for easy maintenance (for example, we can simply swap an old plate with a new plate when we want to upgrade the node). A custom-built enclosure is made to house this component plate.

- **Motherboard**. A VIA MB720F Mini-ITX motherboard featuring an 1GHz processor and 1G memory is employed as the core in our system.

- **Storage**. The traffic pole sometimes vibrates a lot due to the passing traffic. Since that our node is mounted on a traffic pole, instead of using a hard-drive, we employ a 2G USB flash drive for storing OS and data. Unlike a hard-drive, a flash drive does not have a high-speed spinning platter and is less failure-prone.

- **Wireless interfaces**. Each node has two wireless interfaces to connect to its neighboring nodes, as shown in Figure 5. To allow the testbed users to experiment with different radio technologies, two different radio frequencies are currently used on our testbed: 2.4GHz (802.11b/g) and 900MHz radios. Specifically, the nodes at intersection m2, m3 and m6 are installed with two 2.4GHz mini-PCI wireless cards from Ubiquiti (SR2). The nodes at intersections m1 and m5 are equipped with one 2.4GHz Ubiquiti SR2 card (with a transmission power of 400mW) and one 900MHz Ubiquiti SR9 card (with a transmission power of 700mW). Finally, the nodes at intersections m4 and m7 are installed with two Ubiquiti SR2 cards. One of these two SR2 cards is connected to a 2.4GHz-to-900MHz converter (from RFlinx) to send 2.4GHz signal output by the wireless card on 900MHz band. Due to its better penetration rate for buildings and trees, theoretically the use of 900MHz radios could result in a better connectivity than 2.4GHz radios (i.e. 802.11x). Hence, we decided to install 900MHz radios on the nodes for

intersection pairs m1-m5 and m4-m7. These two intersection pairs have a longer distance (i.e. 400m and 500m respectively) than the other intersection pairs.

- **Back-haul connection**. In addition to the two Ubiquiti wireless cards, each node is equipped an "Unwired" modem [13] to establish a back-haul link back to NICTA for the purpose of remote management, as shown in Figure 5. Unwired is a Sydney-based metropolitan wireless ISP. The Unwired modem uses a proprietary protocol but claims to be a variant of WiMAX and operates in a licensed 3.5GHz band.

- **Ethernet router**. A Linux-based Ethernet router (Diamond Digital R100) is installed in each node. We employ this Ethernet router for several purposes. First, it is used as an Ethernet switch to connect the motherboard and the Unwired modem (and any IP-based devices such as a camera in the future). The Unwired modem is connected to the WAN port of the router, thus the router get a public Internet IP address. The motherboard has an Ethernet connection with the router's 4-port switch. Second, the Diamond Digital router has an USB port which allow the motherboard to have a serial connection with the router's USB port through an USB-to-serial adapter. Being able to establish a serial link to the motherboard allows the user to remotely login into the serial console for troubleshooting when the Ubiquiti wireless interfaces are not responding. Third, given that the router is a Linux box itself (runs on openWRT), we can run all the existing software (e.g. we are currently running DNS, NTP and VPN clients on it). Finally, the Diamond Digital router has an 802.11 wireless interface which can be used as an alternative link to remotely connect the mesh node in addition to Unwired and Ubiquiti links.

- **Power**. As shown in Figure 4, we use an off-the-shelf power-board (with surge protector and fuse) and a PC power-supply to provide the power to all the components in the node. The power-board takes a 240AC power from the traffic light.

- **Antenna**. Nodes on our testbed are all installed with omni-directional antennas due to the following

  - **Cost**. An omni-directional antenna is typically cheaper than a directional antenna. In addition, for a node which has n neighbors, n directional antennas are needed. On the other hand, one omni-directional antenna per intersection is sufficient to cover all the neighbors.
  - **Mounting**. The space on the traffic light for the mounting of antennas is quite limited. It is comparatively more difficult to mount a directional antenna on the traffic pole in practice.

  We use an 8dBi omni-directional antenna for the 2.4GHz wireless card and an 6dBi omni-directional antenna for the 900MHz wireless card.

- **Weatherproof**. The temperature in the summer can be above 40 Celsius degree in Sydney. The temperature inside the node can be even higher. As shown in Figure 4, to provide enough air circulation inside the node, we drilled many holes on the bottom of the enclosure and made some air louvres on the side. Two temperature-controlled fans are used in the node to dissipate the hot air out through the louvres. In addition, we mount a roof on top of the enclosure to shield the enclosure from direct sunlight and rain.

- **Remote recovery**. Due to the fact that the testbed is deployed in an outdoor environment, it is time consuming to visit the nodes when something goes wrong. In addition, given that our nodes are mounted on the traffic lights which is a public asset, visiting any node on the testbed required calling out the RTA maintenance crew to gain access to the node. Therefore, some means of remote recovery are necessary. Currently, we have one wireless remote switch installed on each node (runs in the unlicensed 433MHz band), which allows us to reboot the node on-site when accessing the node via the 2.4GHz or 3.5GHz links fails.

The ultimate goal of our project is to control traffic lights using wireless mesh networks. However, due to practical consideration, we do not connect the mesh node directly to the real traffic controller in the first phase of the project. A "dummy" traffic controller board is used instead. The main difference between the real traffic controller and the dummy traffic controller is that the data coming from the dummy traffic controller is fake data (and not the real sensor data coming from the road-side sensor). A pair of power-over-Ethernet adapters (Netcomm NP285) are used to connect the node to a dummy traffic traffic controller board in the curbside housing through the powerline. The dummy traffic controller board sends and receives data via a serial interface. Hence, a serial-to-IP conversion is required for the communication between the dummy traffic controller and the testbed (which runs IP). We mount the traffic controller board inside an embedded PC and connect the traffic controller board to the embedded PC's motherboard's (VIA MB770F) serial port. A serial-to-IP converter software is written and run on the PC to encasuplate the SCATS data from the serial port of the traffic controller board into an IP packet as well as to descapsulate the IP packet from the regional computer and send its payload to the serial interface.

In order to connect the testbed to the regional computer which is located at our facility, we deploy a gateway node at University of Sydney. The gateway node has a reasonable line-of-sight to intersection m1 and connects to the m1 node with a 802.11 link. Note that we do not use the Unwired links to connect the regional computer (RC) to the testbed due to the consideration of reliability, latency and cost issues. More details about the characteristics of Unwired links are described in Section 6. The RC is connected to the gateway node via AARNet [14]. Given that both NICTA and University of Sydney are members of AARNet, there is no cost to send traffic over AARNet. The round-trip delay between RC and the gateway is about 1.2ms, and the throughput is typically over 100Mbps.

5.4 Software

We use a custom-built Linux OS image that consists of the following components:

- The image size is small enough to be fit into an USB flash drive. and run completely in RAM (1GB). This

allows us to enable 'clean' reboots uncontaminated by previous experiments

- We add some kernel modifications for various driver support for USB, madwifi and PXE reboot.
- We modify Grub to activate the watchdog timer at the time of boot-loading so that the watchdog timer can be started before any program start. Watchdog timer is used to reboot the motherboard when the system fails.
- We include various tools including timesync, OpenVPN, some network support tools and software from Orbit project [15] in our image. The image is built to be Debian-based for the compatibility with Orbit software.

We build our OS image based on DSL-N [16]. DSL-N provides most of the software support we need out of the box. The default syslinux bootloader of DSL-N is replaced with grub. We use OML software [17] from Orbit project to build the measurement collection infrastructure for the testbed. Two security mechanism is currently implemented on our testbed. First, OpenVPN is used for the Unwired links from NICTA to each of the mesh nodes. Second, ssh and friends are used on all network interfaces. We plan to implement host-based and certificate-based access in the next phase. In addition, root access is disabled on all the machines.

6. LINK CHARACTERSITCS

In this section, we describe some preliminary results of measured link characteristics from the testbed. Specifically, we discuss some statistics of the wireless link performance in terms of round-trip delay, loss and throughput. We use ping to measure the round-trip delay and iperf for the throughput measurement.

6.1 Link latency

The round-trip delay increases as the number of hops increases on the 802.11 links. In addition, the variation also increases significantly when there are more hops. We do not observe such a strong correlation between distance and link latency though. As shown in Figure 6, the latency does not increase from 300m to 400m. However, the variation increase significantly as the distance increases. One possibility is that there are more retries at 300m than at 400m due to different line-of-sight conditions. We are currently investigating this issue.

We next examine the efficiency of powerline communication. As suggested in Figure 7, given a distance of 100m, the link latency of powerline communication is excellent. The average round-delay is about 3.6ms and the variations are very small. In addition, the largest delay for such a distance is less than 8ms.

As described in Section 5, we use the Unwired network to carry out our back-haul traffic. To understand the expected latency of running management traffic over the Unwired network, we measured the round-trip delay from a machine at NICTA to the mesh node. As shown in Figure 8(a), the average delay of sending traffic over the Unwired network to the mesh node is about 400ms. However, there is a large variation (the delay can be as long as 3 seconds) and significant number of outages. We find that the delay and outages over the Unwired network are mostly contributed by the wireless link between the mesh node and the Unwired base station. As shown in Figure 8(b), the average delay of the Unwired wireless link is about 200ms. The large delay variations and significant number of outages suggest that a public-shared wireless network like Unwired is not suitable for operating SCATS traffic.

6.2 Losses

As shown in Figure 9 and Figure 10, the packet loss seems to be distributed uniformly over time. However, the loss becomes more bursty as the number of hops and distance increase. Note that Figure 9 and Figure 10 are based on the results from a low probing rate (i.e. one-packet-per-second ping). The loss pattern might change if we change the probing rate. In addition, we do not find there is a strong correlation between packet loss and the distance. The line-of-sight condition (which is location-dependent) plays a more important role on the packet loss. We find the use of 900MHz radio results in a much lower loss rate (0.5%) than 2.4GHz radio (20%), which is not surprising though since 900MHz radio have a better penetration rate than 2.4GHz radio. Finally, for a distance of 100m, the loss rate of powerline communication is almost negligible (less than 0.1%).

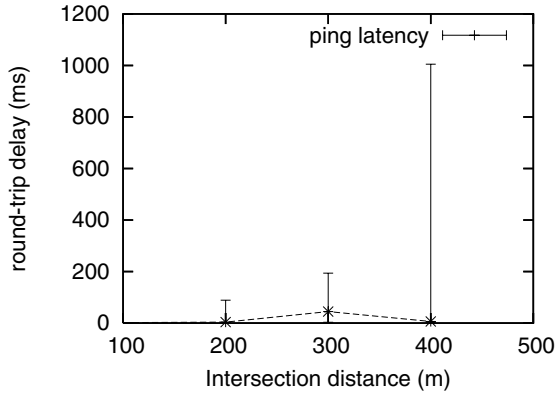

Figure 6: Effect of distance on round-trip delay

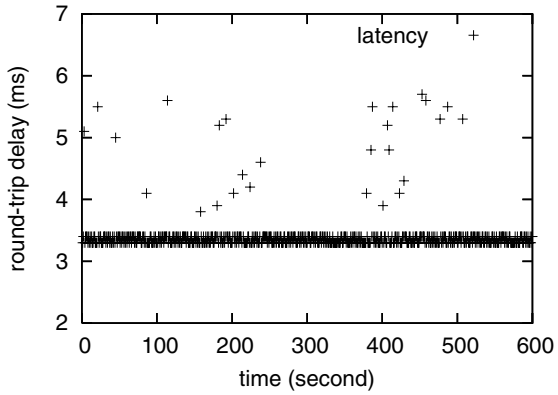

Figure 7: Latency of powerline communication

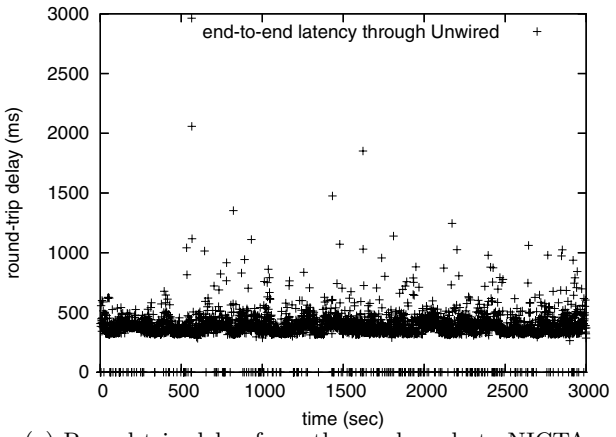

(a) Round-trip delay from the mesh node to NICTA

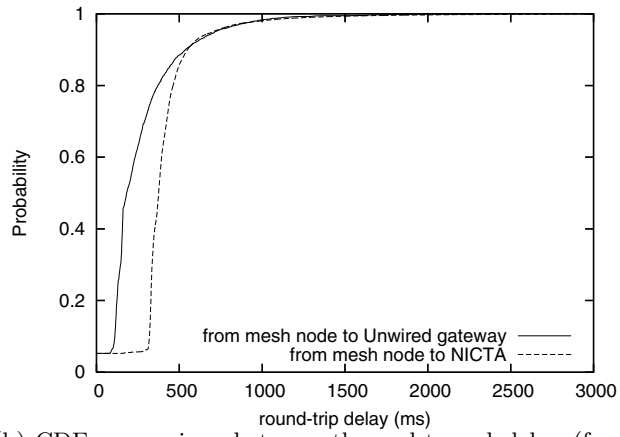

(b) CDF comparison between the end-to-end delay (from mesh node to NICTA) and the Unwired wireless link delay (from mesh node to Unwired gateway)

Figure 8: Round-trip delay over the Unwired network

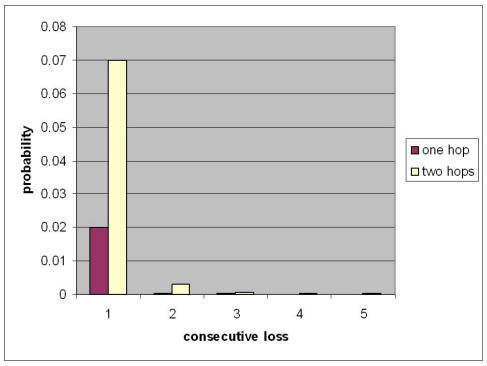

Figure 9: Effect of number of hops on consecutive packet loss

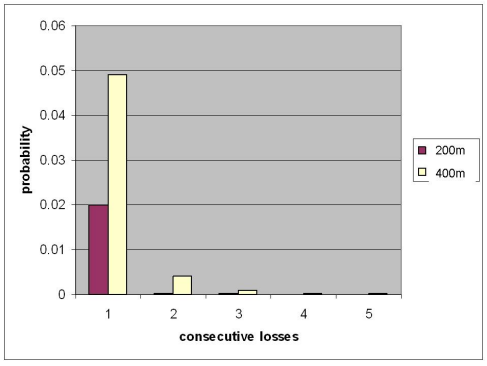

Figure 10: Effect of distance on consecutive packet loss

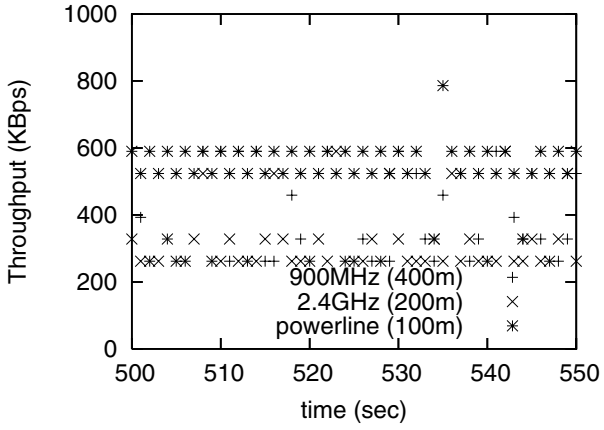

Figure 11: Comparison of throughput for different technologies

6.3 Throughput

As shown in Figure 11, the use of 900MHz radio results in a better throughput than when 2.4GHz radio is employed, even for a longer distance. However, we also observe there is a larger variation when using 900MHz radio, which might be due to MAC-layer retransmission. The throughput of Ethernet-over-powerline communication is very stable and typically maintained at about 600Kps.

6.4 Discussion

In this work, we built a testbed using off-the-shelf hardware within unlicensed bands. However, our initial results are somewhat discouraging since our system does not perform to the requirements of a traffic control system. Specifically, based on Figure 6, the dimension of the network will not be able to scale up to more than 20-hop assuming an average delay of 50ms. In addition, as shown in Figure 9, the loss is pretty significant (2% for 1-hop and 7% for 2-hop) as compared to the typical loss rate that one will see on a

wired network. Needless to say, there are many research challenges that need to be addressed before our mesh network can be used in a live system for traffic control. We are currently developing innovative multi-path routing and cross-layer techniques to address these issues.

7. EXPERIENCES

In this section, we discuss some experiences we gained in terms of the deployment and maintenance of our testbed in an urban environment.

7.1 Deployment

- **Hardware**. We observed that many antenna connectors were held on by weak glue or crimp. Gradual stress (e.g. vibration) could eventually loosen the plug and degrade the signal before it is transmitted into the air. Some protection of the antenna plug might be necessary be necessary for an operational network to ensure there is no signal leakage from the antenna connector. In addition, while the appearance of the hardware might look identical, it is safer to check if the hardware does comply to the specification before starting using it. For example, during our experiments we found some of our Senao wireless cards does not output a transmission power of 200mW as they should according to the specification.

- **Software**. Most of the wireless measurements are based on readings from the wireless cards. However, while the hardware can be identical, different firmwares and drivers could introduce inaccuracy in the measurement results. We strongly suggest, if possible, one should try to validate the readings from a wireless card against the results from a spectrum analyser.

- **Antenna locations**. As described in Section 5, each node is equipped with three antennas, including two 2.4GHz (or one 2.4GHz and one 900MHz) omni-directional antennas and one 3.5GHz directional antenna. To facilitate the ease of mounting, we first mount all three antennas on a pole and then mount this pole on the traffic light. Specifically, one omni-directional antenna is pointing upward and the other is pointing downward while the directional antenna is mounted in between. We found the location of antenna can have an effect on the link performance. Figure 12 shows the round-trip delays from node m2 to its neighboring node m3 using the omni-directional antennas. The use of the lower antenna results in a higher loss and a larger variation. One possible explanation is that the upper antenna might be less obstructed and hence have a better connectivity. At 2.4GHz, a quarter wavelength is approximately 30cm. Antenna position changes in the range 10-30 cm can cause dramatic changes in signal strength, presumably due to the presence of standing waves in the vicinity of the traffic light pole or more specifically in the vicinity of metal stop signs and the like! If multiple antennas are deployed, it is essential to have a means for independently adjusting their position.

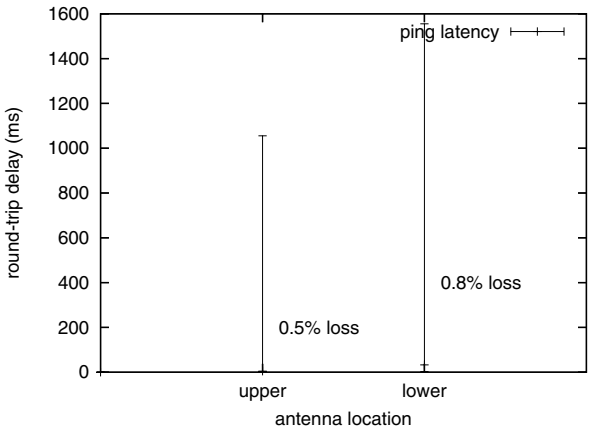

Figure 12: Effect of antenna locations on round-trip delay

7.2 Maintenance

- **Remote management** Remote management is challenging. Currently we provide the following ways to allow the user to access the nodes.

  – **To access the Linux-based Ethernet router** The router can be connected via the Unwired link over OpenVPN. In the case when the OpenVPN connection can not be established, given that a public IP address is obtained for each router from Unwired, one can connect to the router via its public IP address, although this could introduce a dependency on a dynamic DNS lookup. In addition, one can connect to the mesh node first and then connect to the router via the Ethernet or USB-to-serial link between the router and the motherboard.

  – **To access the mesh node** One can connect to the mesh node (i.e. the motherboard) by first connecting to the router and then connect to the motherboard via the Ethernet or USB-to-serial link. Being able to access the motherboard via its serial port is important since the Ethernet link might fail for various reasons. In addition, one can access the motherboard via its 802.11 interfaces from any reachable neighboring nodes.

  In addition, the following mechanisms are currently implemented to recover the system from failure. First, and the default way to recover from a failure is to login to the offending router or motherboard using one of the above methods, analyse the problem and/or reboot the node. Second, the watchdog timer support on the MB720 is used. In addition, Grub is setup to fall back to a stable backup image which is installed in a separate partition in case when the default image fails. Finally, the BIOS is configured to give top priority to PXE network boot but we configure DHCP server in a way that it does not provide PXE boot information in the default case. Therefore the node defaults to its second priority, which is to boot from the USB flash

drive. However, in the event of a node failure (for example, due to a bad image) the DHCP server can be quickly reconfigured to support the PXE boot. Having rebooted the node using PXE, a new working disk image can be distributed to the node via frisbee [] or FTP. In practice, we use FTP instead of frisbee since that frisbee introduce more control traffic overhead on the Unwired links, where we are charged for every bits send to the nodes.

- **Security** Security is a major concern especially when our wireless mesh testbed is sharing public spectrum with an average of 50+ external APs at each intersection. Furthermore, our testbed is effectively connected to the Internet via Unwired network, and exposed to various password attacks. In a live deployment for traffic control, the mesh security should be integrated with the traffic control system security model, which may include e.g. segmentation to contain the damage of a denial of service or break-in attack, combined with multiple levels of fallback to local control.

8. CONCLUSION

In this paper, we discuss our experiences in deploying a testbed as a first step towards creating a fully functional wireless mesh network-based traffic control system. In addition, we describe some initial results of link characteristics of different technologies used on our testbed. While wireless mesh networks have been used in public safety and residential broadband for years, to the best of our knowledge, our work is one of the first attempts to use mesh network for traffic management. However, there are several research challenges such as latency, reliability, security and scalability that need to be addressed. We are currently developing innovative multi-path routing and fast anomaly and fault detection schemes to address these issues. In addition, in the next phase, we plan to extend our testbed to 15-20 nodes as well as to cover a larger area. Finally, the measurement results presented in this paper are very preliminary. We are in the process of performing more in-depth experiments to understand the link behaviour of the testbed.

9. REFERENCES

[1] R. Chakravorty and I. Pratt, "Performance issues with general packet radio service," *Journal of Communications and Networks (JCN), Special Issue on Evolving from 3G deployment to 4G definition*, vol. 4, no. 2, Dec. 2002. [Online]. Available: http://www.cl.cam.ac.uk/Research/SRG/netos/papers/comob-web/2002-jcn.pdf

[2] H. Lundgren, K. Ramachandran, E. Belding-Royer, K. Almeroth, M. Benny, A. Hewatt, A. Touma, and A. Jardosh, "Experiences from building and using the ucsb meshnet testbed," *IEEE Wireless Network*, 2006. [Online]. Available: http://user.it.uu.se/~henrikl/publications.html

[3] http://www.tropos.com/, "Tropos networks," http://www.tropos.com.

[4] http://www.locustworld.com/, "Locust world," http://www.locustworld.com.

[5] http://www.pwlan.org.tw/mp.asp?mp=3, "Mobile taiwan applications promotion project (m-taiwan)," http://www.pwlan.org.tw/mp.asp?mp=3.

[6] S. Ganguly, V. Navda, K. Kim, A. Kashyap, D. Niculescu, R. Izmailov, S. Hong, and S. Das., "Performance optimizations for deploying voip services in mesh networks," *Performance Optimizations for Deploying VoIP Services in Mesh Networks*, 2006. [Online]. Available: http://www.wings.cs.sunysb.edu/%7Eanand/papers/jsac06.pdf

[7] J. Bicket, D. Aguayo, S. Biswas, , and R. Morris, "Architecture and evaluation of an unplanned 802.11b mesh network," in *Proceedings of the 11th annual international conference on Mobile computing and networking (MOBICOM)*, ologne, Germany, Sept. 2005. [Online]. Available: http://www.pdos.lcs.mit.edu/papers/roofnet:mobicom05/roofnet-mobicom05.pdf

[8] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris, "Link-level measurements from an 802.11b mesh network," in *Proceedings of the 10th annual international conference on Mobile computing and networking (MOBICOM)*, Philadelphia, PA, USA, Sept. 2004. [Online]. Available: http://research.microsoft.com/mesh/papers/multiradio.pdf

[9] S. Weber, V. Cahill, S. Clarke, and M. Haahr, "Wireless ad hoc network for dublin: A large-scale ad hoc network test-bed," *ERCIM News*, vol. 54, 2003. [Online]. Available: http://www.ercim.org/publication/Ercim_News/enw54/weber.html

[10] J. Camp, J. Robinson, C. Steger, and E. Knightly, "Measurement driven deployment of a two-tier urban mesh access network," in *Proceedings of ACM MobiSys 2006*, Uppsala, Sweden, June 2006. [Online]. Available: http://networks.rice.edu/papers/sys7122-camp.pdf

[11] A. Jardosh, K. Ramachandran, K. C. Almeroth, and E. M. Belding-Royer, "Understanding congestion in ieee 802.11b wireless networks," in *Proceeding of ACM SIGCOMM Internet Measurement Conference*, Berkeley, CA, Oct. 2005. [Online]. Available: http://moment.cs.ucsb.edu/~amitj/jardosh-imc2005.pdf

[12] T. S. Rapport, "Wireless communications principles and practice," *Prentice Hall Prt. NJ*, 1996.

[13] http://www.unwired.com.au/, "Unwired wireless," http://www.unwired.com.au/.

[14] http://www.aarnet.edu.au/, "Aarnet - australia's research and education network," http://www.aarnet.edu.au/.

[15] D. Raychaudhuri, I. Seskar, M. Ott, S. Ganu, K. Ramachandran, H. Kremo, R. Siracusa, H. Liu, and M. Singh, "Overview of the orbit radio grid testbed for evaluation of next-generation wireless network protocols," in *Proceedings of the IEEE Wireless Communications and Networking Conference*, New Orleans, LA, USA, Mar. 2005.

[16] http://www.damnsmalllinux.org/dsl n/, "Damn small linux not (dsl-n)," http://www.damnsmalllinux.org/dsl-n/.

[17] M. Singh, M. Ott, I. Seskar, and P. Kamat, "Orbit measurements framework and library (oml): Motivations, design,implementation, and features," in *Proceedings of IEEE Tridentcom 2005*, Trento, Italy, Feb. 2005.