The Cambridge Wireless Broadband Trial

Raymond Liao[*,1] , Martin Brown[**], Glenford Mapp[**], and Ian Wassell[***]

[*] The Comet Group, Columbia University, USA, liao@comet.columbia.edu.
[**] AT&T Laboratories Cambridge, UK, {mgb, gem}@uk.research.att.com
[***] Lab for Communication Engineering, Dept. of Engineering, University of Cambridge, UK, ijw24@cam.ac.uk

***Abstract*-In this paper, we present a system overview of the Cambridge Wireless Broadband Trial which is being deployed in the Cambridge City area in the UK. The trial network is intended to serve as a testbed for research into new paradigms for wireless communications, and will provide an infrastructure for research on home networks. After a brief discussion of the network provisioning methods required to make the network operational, we present some preliminary measurement results. We then identify potential performance issues and propose some possible solutions as future work.**

I. Introduction

In recent years, there has been a significant and growing demand for wireless systems. These systems have the potential of seamlessly supporting mobile users and also have tremendous advantages over other access technologies in situations where a wired system, such as cable, would be either expensive or inconvenient.

In the past, these systems offered limited bandwidth, making it only suitable for audio applications, such as telephony. Recent developments have made it possible to achieve aggregate throughputs of 25 Mbps between subscriber units and a given base station - enough bandwidth to support data communications and even multimedia applications. The challenge now is to build networks using this equipment in order to understand the system issues involved in deployment on a large scale.

The Laboratory for Communications Engineering (LCE) located at the Department of Engineering at the University of Cambridge, in collaboration with the AT&T Laboratories Cambridge (ALC), and Adaptive Broadband Ltd (ABL) intend to explore these issues via the Cambridge Wireless Broadband Trial. The ALC has a long tradition in building networks to support novel and demanding applications. It pioneered the use of multimedia workstations in the Pandora project [1]. Medusa [2], the successor to Pandora, looked at developing a rich multimedia environment using network multimedia peripherals on a 100 Mbps ATM network. The Wireless ATM project [3] looked at building a wireless system based on ATM technology. This group later spun out as Adaptive Broadband Ltd (ABL) which was acquired by Adaptive Broadband Inc, formerly known as California Microwave.

II. The Trial

The trial will encompass the Cambridge City area and its immediate surroundings. Its first goal will be to provide access to the homes of the employees of the ALC as well as allow academic personnel access to the computing facilities of the University via the LCE. A key long term aim is to build up sophisticated management and monitoring tools which will enable us to understand and control these networks. Using these facilities, it will be possible to look at issues of network, protocol and application performance in this environment. It is hoped that this new infrastructure will be used as a testbed for research into new paradigms for wireless communications. The system will also provide an infrastructure useful for research on home networking.

The networking equipment for the trial has been acquired from ABL. These units presently offer broadband fixed wireless access. The deployed versions, known as Helium units, offer line-of-sight distance ranges of 3-5 km with 1W max. EIRP at 5.2 GHz and 8-10 km with 4W max. EIRP operating at 5.7 GHz. ABL has announced plans to develop a MMDS system in the near future. Base station units have a 60-degree wide cell sector with an aggregate data rate of 25 Mbps in each sector. Subscriber units can provide a 10 Mbps Ethernet or an ATM-25 connection to the end user. The first stage of the trial will use 20 units, it is hoped that this will be expanded to 50 units.

Deployment is taking place in phases. The first phase was the establishment of a connection over the air between the LCE and the ALC. This is shown in Figure 1 and was used to support a terrestrial TV application. This was also used as a testbed for initial performance tests including the round-trip delay of ping and the throughput of ftp applications. In the tests, an early prototype version of the ABL equipment, called the Hydrogen[2] unit, was used. The second phase of the deployment is shown in Figure 2. Using the LCE as the primary base station site, the network will have a number of video sources including a video repository (PC7), a Telemedia Systems Ltd. (TSL) network camera and cameras

[1] This work was done when the author was at AT&T Lab Cambridge for the summer of 1999.

[2] The Hydrogen units are not supported by ABL anymore.

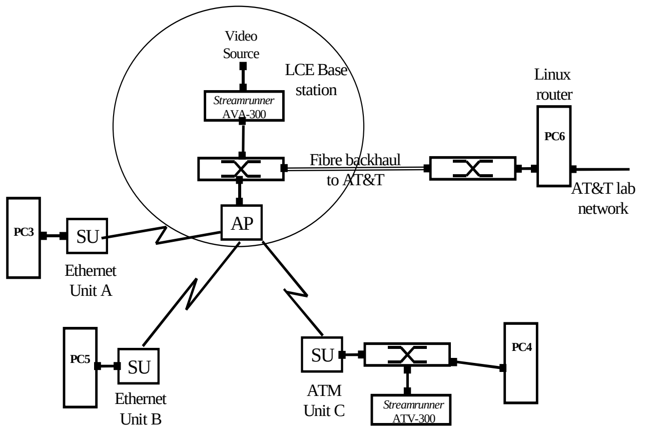

Fig. 1. Phase I Network Topology
(AP: Access Point, SU: Subscriber Unit)

connecting to StreamRunner AVA's from Fore Systems. A Linux router provides access to ALC's internal network for trusted clients.

The third phase of the project involves the setting up of other base station sites to allow access to the network from several places in the Cambridge area. The New Addenbrooke's Hospital, located about 2 km south of the city, has been designated as the next base station site and will allow many homes in the South of Cambridge City to use the system. Another base station is planned to be deployed at the University Library in West Cambridge, providing access to various university departments.

III. NETWORKING ISSUES

The ABL equipment uses ATM technology in the backbone. IP support is provided where required by using

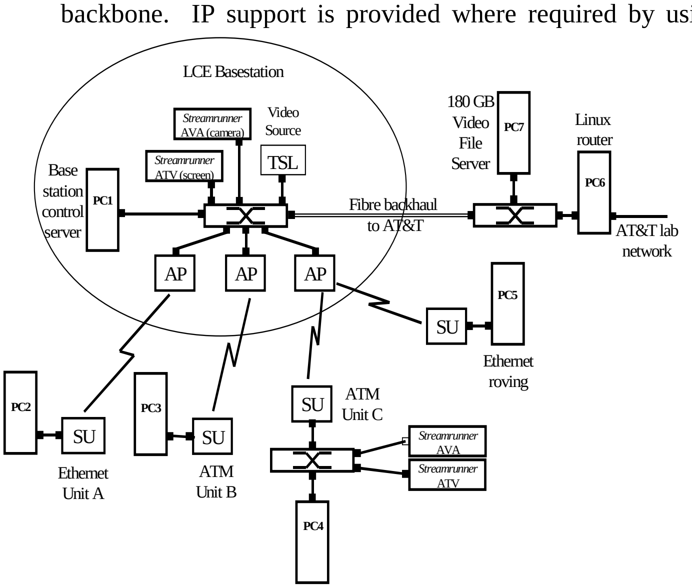

Fig.2. Phase II Network Topology

Classical IP over ATM or CLIP. There is no intention of running LAN Emulation (LANE) anywhere in the system. We have devised an address allocation scheme in which each home can have a subnet of 32 or 64 IP addresses. The Hydrogen units only provided support for the establishment of PVCs between the subscriber units. In order to establish a connection between two entities, two PVCs are required, one in either direction. Since CLIP does not support full ATM signalling, the use of ATMARP and other facilities are required but introduce end-to-end performance issues.

More specifically, for a host to communicate with another node in the ATM network, either a direct bi-directional VC has to be created between them, or both of them should have a bi-directional VC to a common router respectively. The first case is not scalable since a mesh of VCs are required for a fully connected network, while the latter case involves inefficient triangle routing through the router. A complete solution to this problem requires a complicated suite of ATM signalling and routing protocols, including SVC signalling, PNNI (Private Network-Node Interface) routing, and NHRP (Next Hop Resolution Protocol) for routing across different IP subnets.

Since the ATM equipment in our network does not support all of the above protocols, we have opted for a simpler but limited approach: besides creating PVCs from each network node to the router of the ATM subnet, we also create PVCs between dedicated multimedia application server (e.g. video file server PC7 in Figure 2) and service subscriber nodes. Since some application servers will be multi-homed with both an ATM interface in the trial network and an Ethernet interface in the lab network, a separate DNS (Domain Name Service) subdomain has been created for the trial network to enable automatic translation between the application server's hostname and the corresponding IP address of different interface in the lab or trial networks.

IV. PRELIMINARY PERFORMANCE

Using the phase I setup of the project as depicted in Figure 1, we measured the delay and throughput characteristics of the Hydrogen systems. The test setup is shown in Figure 3. The background traffic load used in the tests was generated from the terrestrial TV application. A TV channel was piped into a StreamRunner AVA-300 and then over the network as a pair of Motion-JPEG video and PCM audio streams. This was captured by a StreamRunner ATV-300 delivering analog video to a monitor and speaker. With the rate-controllable encoder of AVA-300, we varied the peak and average rates of the video stream and effectively changed the background traffic load.

Three subscriber units were used during the test to communicate with the same access point. Two of them were the same with Ethernet wireline interface, the other had ATM

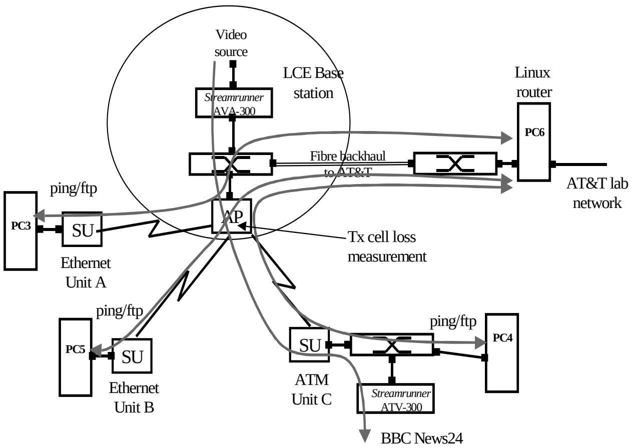

Fig. 3. Test Setup

wireline interface. While the two Ethernet subscriber units experienced the same traffic load, the ATM subscriber unit had an additional load of background video/audio traffic.

The purposes of the initial tests were to investigate system delay and throughput under different load conditions, and to compare the treatment of the two Ethernet subscriber units by the wireless scheduler at the access point. There were two group of tests: the delay tests comprised measuring the round trip delay of ping from the three subscriber units to the Linux router (PC 6 in Figure 3). The throughput tests comprised of measuring the throughput of ftp from the three subscriber units to a host in AT&T lab network.

The delay results of the Hydrogen system exhibit good performance as illustrated in Figure 4. The curves depict mean round trip delay of ping vs. varying background traffic load ranged from 1 to 15 Mbps. The mean delay is averaged over the round trip delay samples of 70 consecutive pings. From Figure 4, it is clear that the delay is not sensitive to whether each subscriber unit pings alone or all three ping together. The result is also not sensitive to the varying traffic load up to 12 Mbps. In addition, the delay results of the two Ethernet subscriber units are very close, even though unit A has shorter delay that unit B. The much lower delay experienced by the ATM subscriber unit C is due to the low protocol overhead as there is no protocol conversion from ATM AAL5 to Ethernet at the ATM subscriber unit.

Figure 5 illustrates the relative throughput results in correlation with the cell loss measured at the access point air interface. The figure comprises two plots, one for the relative throughput and the other for the cell loss, all with respect to the same varying background traffic loads. Each throughput and loss result is the average from 6 independent tests. Since the loss is the result of transmission buffer overflow and only counts for part of the total cell loss, it is used only as an indicator of traffic load condition in the system. The throughput results are shown relative to the maximum throughput achieved by any unit. The throughput does not peak under zero background traffic load because in that case the 10Base-T Ethernet interface of the ftp file server became the bottleneck.

There are two main observations from Figure 5. First, there is a distinct difference between the throughput achieved by units A and B. Unit A gets 15-40% more throughput than unit B. The better performance achieved by unit A is also observed in delay results shown in Figure 4, but with a smaller difference. These results indicate that the same type of subscriber units can get quite difference performance when sharing the same access point. The reason could come from their different locations that can cause difference in the quality of line-of-sight physical signals.

Secondly, we can observe the impact of ftp data traffic on the multimedia application when there is no service differentiation between them. As shown in Figure 5, cell loss occurs at quite an early stage when the video background

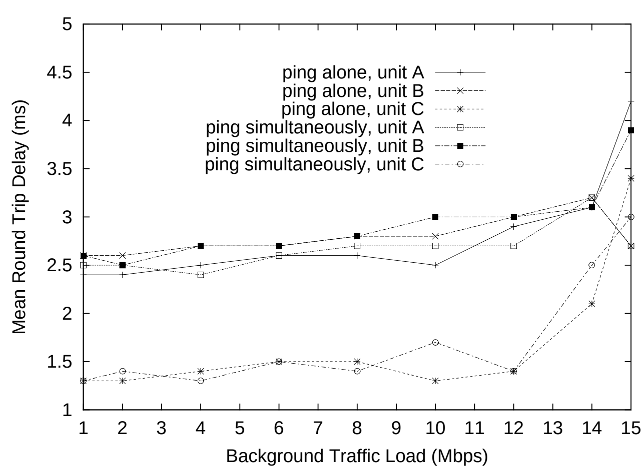

Fig. 4. Delay Results

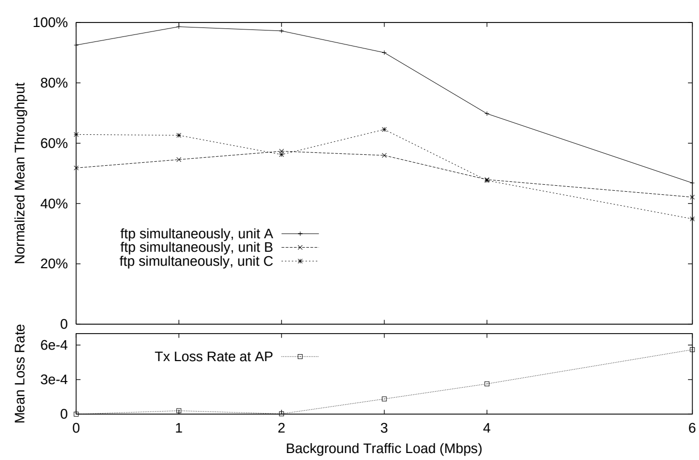

Fig. 5. Normalized Throughput vs. Loss Results

traffic load is as low as 1 Mbps. This is because the ftp applications from the three subscriber units use TCP's adaptive flow control algorithm that relies on the loss of packets to throttle bandwidth consumption which would otherwise continue to increase. In this sense, packet loss is inevitable and is helpful to TCP applications. However, multimedia applications will suffer when sharing the same service with data applications. In the experiments, the loss at 1 Mbps background traffic load caused noticeable interruptions to the audio stream, which manifests the need for service differentiation between multimedia applications and data applications.

V. Proposed Solutions

We are studying various ways of introducing quality of service differentiation into the wireless network, especially the IETF differentiated service approach [4] that avoids significant system complexity.

Two service priorities could be provisioned in the subnet. The high priority service would be used by rate-controlled multimedia sessions between the wireless subscriber units and the multimedia content servers within the wireless subnet. All the other traffic would use the low priority service, including the traffic from the ATT Lab network and the global Internet that may contain multimedia sessions whose rate is not controllable from the wireless subnet.

A bandwidth broker module could be created in the wireless subnet. With the knowledge of the subnet topology and the quality of different fixed wireless links (measured by the throughput/delay tests in Section IV), the bandwidth broker could perform admission control for sessions between the wireless subscriber units and the subnet multimedia content servers (e.g., the terrestrial TV service from the TSL box, and the video file server at PC7 in Figure 2). By adopting the utility function model [5], the bandwidth broker could apply application-aware and wireless-channel-aware utility-maximization or policies that compensate for different wireless link quality and application needs.

The enforcement of bandwidth allocation for multimedia sessions would be performed by the rate-controlled video encoders (e.g., the FORE AVA-300 or TSL box which regulate the peak rate of multimedia sessions at the terrestrial TV receiver). The border policing/shaping for traffic to/from the external network would be performed at the Linux router (PC6 in Figure 2) with the Linux traffic control module [6]. The sessions could be classified based on the wireless terminals' IP addresses. The policing rates would then be specified by the bandwidth broker so that an access point will not be overloaded with traffic to/from all of its wireless subscriber units.

VI. Conclusion and Future Work

The preliminary results presented in this paper are encouraging given that they are measured from the now superseded Hydrogen system -- ABL's early prototype of broadband wireless access system. With the deployment of the newer Helium system, the system performance, reliability and maintainability will significantly improve.

Our future work will be on two fronts. We will continue our deployment plan of the wireless broadband network to a larger user population. In the process we will investigate the issues of managing and controlling these networks. In addition, we will study various ways of introducing quality of service differentiation into the wireless network.

Acknowledgment

We would like to thank Steve Greaves, Damian Gilmurray, and David Pegler from ABL for their technical support during the trial. In addition, Raymond Liao would like to thank Andrew T. Campbell for directing him into this project and encouraging him to write this paper.

References

[1] A. Hopper, "Pandora – an Experimental System for Multimedia Applications", *ACM Operating System Review*, April 1990.

[2] S Wray, T. Glauert, and A. Hopper, "The Medusa Applications Environment", *Proc. of International Conference on Multimedia Computing and Systems*, May 1994.

[3] J. Porter, A. Hopper, D. Gilmurray, O. Mason, J. Naylon, and A. Jones, "The ORL Radio ATM System, Architecture and Implementation", *Technical Report 96-5, Olivetti Research Ltd.*, 1996, http://www.uk.research.att.com/abstracts.html#68

[4] S. Blake, et al., "An Architecture for Differentiated Services", *IETF RFC 2475.*, Dec. 1998.

[5] R. R-F. Liao and A. T. Campbell, "On Programmable Universal Mobile Channels in a Cellular Internet", *Proc. of ACM/IEEE Mobicom'98*, Austin, TX, Oct. 1998.

[6] W. Almesberger, "Linux Traffic Control -- Implementation Overview", *Technical Report SSC/1998/037, EPFL*, Nov. 1998. ftp://lrcftp.epfl.ch/pub/people/almesber/pub/tcio-current.ps.gz