Secure Mobile Networking
First Quarterly Report

John McHugh and Jim Binkley

3 February 1966

Introduction

The Secure Mobile networking project at PSU got off to good start, but almost immediately encountered administrative difficulties with which we have yet to come to grips. After a rush on the part of ARPA to get us under contract by mid September, we were informed that the ARPA reorganization had resulted in our ARPA point of contact being moved from Teresa Lunt to Mike St. Johns. Major St. Johns then informed by email that it would be necessary to reduce the scope of our effort.

These issues have not been completely resolved and in the face of conflicting information, we decided to proceed with limited system development efforts in the absence of directions to the contrary. The remainder of this report discusses our activities in setting up an initial system and out participation in the The International Cryptography Experiment, a requirement added by ARPA during the contracting process.

The Initial Mobile Network

Current Project Status and Operating Assumptions

As of winter quarter, one of the PIs (Jim Binkley) is working nearly full-time on the project and is only teaching one course this quarter. We are assuming that we will be able to do at least a two year project. Under this assumption, the budget could be cut back by one third; i.e., we are dropping the final year from the original proposal. The first year will be as previously planned; i.e., we will implement mobile-ip and network layer security and integrate them with attention to tunnels. We will use NCR Wavelan as a link-layer medium and will construct at least a radio-based networking infrastructure in the CS/EE building. We have already mapped the building and have determined that 4-5 base stations are enough for coverage. We will also attempt to understand and analyze those protocols in a mobile LAN-based radio environment.

We would drop (or restructure as a third year option) the previous second-year effort aimed at distributed access protocols ("mobile firewalls"). We still think research in this area is interesting and it is certainly going to be necessary for widespread secure mobility, but based on the current state of mobile computing in general and secure mobile computing in particular, we think that the work on routing and agents, originally proposed for the third year is more critical. We will move the proposed third year research to the second year. This will focus on routing redundancy; i.e., multiple foreign agents, multiple home agents, and "ad hoc" networking; i.e., direct secure communication between mobile hosts.

Winter Quarter Results and Plans

It was hoped that we could put together a suitable development environment in fall quarter but that did not happen. One question we faced immediately was what operating system to use for a development environment. Two possibilities seemed reasonable. We could use linux, since wavelan drivers were available for it. Or we could use FreeBSD, because we had previous experience based on the BSD stack. Unfortunately drivers existed for operating systems that were close (BSDI) but not quite the same. We chose to not use BSDI because their stack source would be proprietary.

Due to the availability of NCR wavelan drivers for linux, we first decided to install linux and check it out. At this point, we have decided to not use linux, if we can use FreeBSD. Certain aspects of the current linux TCP/IP stack are immature. E.g., the routing/forwarding code in the kernel consists entirely of a linear search that does not even bother to try and find the longest matching prefix in subnet terms. The TCP protocol itself has some questionable areas. We foolishly managed to bring down our own local network for a while by blindly trusting the linux installation of "rip" (the switches are wrong by default as they assume that the box in question is meant to be the "internet gateway"!). Jim Binkley has considerable previous engineering experience with BSD stacks, including BSDI and vxWorks (a real time operating system) plus we are in contact with primary FreeBSD developers who have contributed cdroms to students at PSU. The BSD stack is probably a good choice for other reasons too since there is a widespread basis here for sharing with other efforts elsewhere. As an example, we are hoping to at least examine and possibly reuse work done by Perry Metzer (ipsec) as he implemented prototype ipsec protocols in NetBSD (a close ally of FreeBSD).

Work is underway to port the BSDI wavelan drivers (supplied by Harvard) to FreeBSD. We need drivers for both ISA bus cards (for wired infrastructure agents; i.e., HAs and FAs) and for the pcmcia card that goes in laptops. We have successfully ported the ISA driver and a port for the laptop pcmcia card is underway.

We have ordered and purchased a first round of equipment. The continuing drop in the overall prices of Intel/PC equipment is frankly delightful. We have purchased 3 IBM thinkpad laptops (2 cheaper 701c laptops and one pentium 755 thinkpad).

We have unix running on one thinkpad so far and are doing driver development on the machine. We purchased 2 pentium desktop stations for mobile agents. Both machines cost less than $2k apiece. We have purchased 6 wavelan cards initially, 3 pcmcia cards and 3 ISA cards. We decided to not purchase the Wavelan access points at this point, as we felt that getting the ISA cards going would be cheaper in the long term. This gives us the basis for an initial lab.

We have used the link-layer point to point diagnostics supplied by NCR to determine how best to deploy radio in our CS/EE building. Please see the attached report.

We should mention some possible assistance to our work. Jon Inouye is a Ph.D student at the Oregon Graduate Center who has chosen to work on network layer adaption to change. He is specifically interested in researching how a network stack can adapt to changes in adapters; e.g., pcmcia adapters and is interested in collaborating with us on FreeBSD/mobile-ip. Also one master's project student at PSU, Craig Hondo, who is an experienced industry engineer, has chosen to work on IP network security o.s. architecture for his master's project.

The plan for the remainder of winter quarter is to get the link-layer drivers established and then work on the design and implementation of mobile-ip.

Spring Quarter Plan

In the spring quarter, we will finish mobile-ip, design and implement network-layer security, and integrate them. Testing and analysis will probably continue into the summer. We also hope to have an implemented protocol deployed in the building to allow mobile roaming.

Research Notes

Ad Hoc routing protocols We have done a literature survey on the subject of ad hoc routing protocols. We will report on this in a little more detail in the next report and include references there. All we want to say for now is that there is a small body of highly theoretical work on the subject at the moment and that in general, you can divide it into two camps (barring flooding). Camp 1 wants to somehow reuse existing routing protocols (distance-vector and/or link-state) even though the need to minimize broadcast and involvement of every mobile node is not easy to achieve. Camp 2 consists of Dave Johnson at CMU who proposes that something might be learned from previous radio efforts and it is possible that a routing protocol based on demand source routing might minimize involvement of other hosts. A demand source routing system coupled with network authentication for security might be a very interesting research alternative indeed.

Link Layer Security Please see the section at the end of the attached report (Appendix ??) "Link-layer Connectivity."

International Cryptographic Experiment participation

During the contract negotiations, ARPA asked us to attend the meetings of the International Cryptography Experiment. This experiment, is best described by WWW entry at http://www.stc.nato.int/public/ice/ice.html from which we quote:

> ICE is an international programme to demonstrate the use of interfaces whereby computer based applications can invoke external cryptographic functions supplied either in hardware or software form. The overall aim is to provide and demonstrate Cryptographic Application Programming Interfaces and Cryptographic Algorithm Portability Interfaces (CAPI) which are available for international use and which are independent of particular applications, algorithms and vendors in order to achieve the following advantages:
>
> - application development without the need to become involved in cryptography
> - common secure applications working with a wide variety of different cryptographic modules
> - common cryptographic modules supporting many different applications
> - the ability to easily integrate cryptographic functions designed for civil, commercial and military needs
> - cryptographic modules can satisfy the regulations within each country without impact on applications

John McHugh attended the second ICE workshop at The Hague in September. The presentations from ranking NSA and NATO officials indicated that an approach like the one advocated by the ICE sponsors is likely to achieve official sanction because of the need by both national and multinational military forces for interoperability. The September workshop addressed the problem largely at the policy level, but a call for volunteers to develop pilot systems was made at the end of the workshop. Since our systems development effort is intending to use embedded cryptography, we have offered to be a participant in these developments.

The next ICE workshop will be held in the DC area in March.

A Report on Wavelan/Link-layer Connectivity in the CS/EE PCAT Building

Jim Binkley

Introduction

I took an IBM laptop configured with a Wavelan PCMCIA radio lan card and measured the connectivity between the laptop and a PC workstation located in my office, 128D. The goals were to:

1. establish that Wavelan would work with the PC hardware in question. There were two types of wavelan cards (PCMCIA for laptop) and ISA for a workstation box. Everything worked fine and this question will not be further discussed.
2. determine how many wavelan "base stations" (workstations equipped with ISA cards) would be needed on the wired infrastructure in order to establish building-wide connectivity.

PCAT was built as a showplace building for Blue Cross many years ago. Our offices are all on one floor in a quadrant with an empty courtyard containing a disfunctional fountain in the center of the building. My office is well situated at one angle of the quadrant, in the NE corner.

Equipment

Software: The laptops were both running PCDOS 6.3. I used the wavelan PCMCIA enabler device which configures the PCMCIA socket and simply ran the ptpdiag.exe utility on both boxes. The ISA card needed no setup or drivers. The diagnostic software implements a link-layer diagnostic between both pieces of hardware used that measures the signal quality (signal to noise ratio) in percent terms 0..100% and also gives a thumbnail rating in terms of "good", "acceptable", and "poor". "Poor" effectively means you are losing packets. With "Acceptable" you are not losing packets. 4 packets are sent and echoed per second.

Hardware:

- **Laptop:** IBM 701c laptop. This is a cheaper 486-based system that is quite compact and has two PCMCIA slots.
- **Workstation:** IBM PC clone pentium 90mhz with a number of PCI and ISA slots.

Wavelan cards: Wavelan is a 2mbit LAN technology; i.e., broadcast in the bus-based or ethernet sense is supported. Both cards use different Intel-based lan controllers, and function in some sense as close-to or equivalent Ethernet systems; i.e., they have IEEE MAC addresses. In general the card architecture consists of Intel lan controller and a NCR "radio-modem" components. Both are addressable from the ISA/PCMCIA buses in typical Intel fashion (i.e., ports/IRQs, etc.) Connectivity is said to be 800 feet in a radius from a base station; sans walls, metal cabinets, etc. It is limited by power restraints due both to physical reality and the FCC.

- 1 ISA card in the workstation equipped with an omni-directional antennae.
- 1 PCMCIA card in the laptop with a roughly equivalent antennae (the lan controller is in the PCMCIA card. The "antennae" actually includes the radio modem hardware.)

Raw Data

The raw data given in Table ?? shows signal strengths based on a single entry point located near the NE corner of the PCAT building. The actual locations can be seen in Figure ??.

Interpretation

Certainly we can get good to acceptable coverage with little or no packet loss in the CS/EE building. We could cover the building in an adequate fashion including lecture and conference rooms by putting one "base station" in each corner. 4 or 5 base stations would do it. We might be able to cover the building by putting one center base station outside in the fountain area, but we would probably certainly lose PCAT 28. Possibly two base stations might work here. (However the building seems to have extra walls in its corners and that might be one reason for putting base stations in the corners).

Later on, we should repeat this experiment with two laptops both equipped with PCMCIA cards. We can station one laptop in the center of the building and walk around the edges (both in the inner hall and out on the street).

As an additional equipment test, I also configured both systems with a PC TCP/IP stack and using the supplied ODI (Novell) drivers ran ping between the stations. This worked with no problems.

Observations

One interesting point is that there is an appreciable difference between Len Shapiro's office and the 113A area in the NW corner. Len's office is almost line of sight down

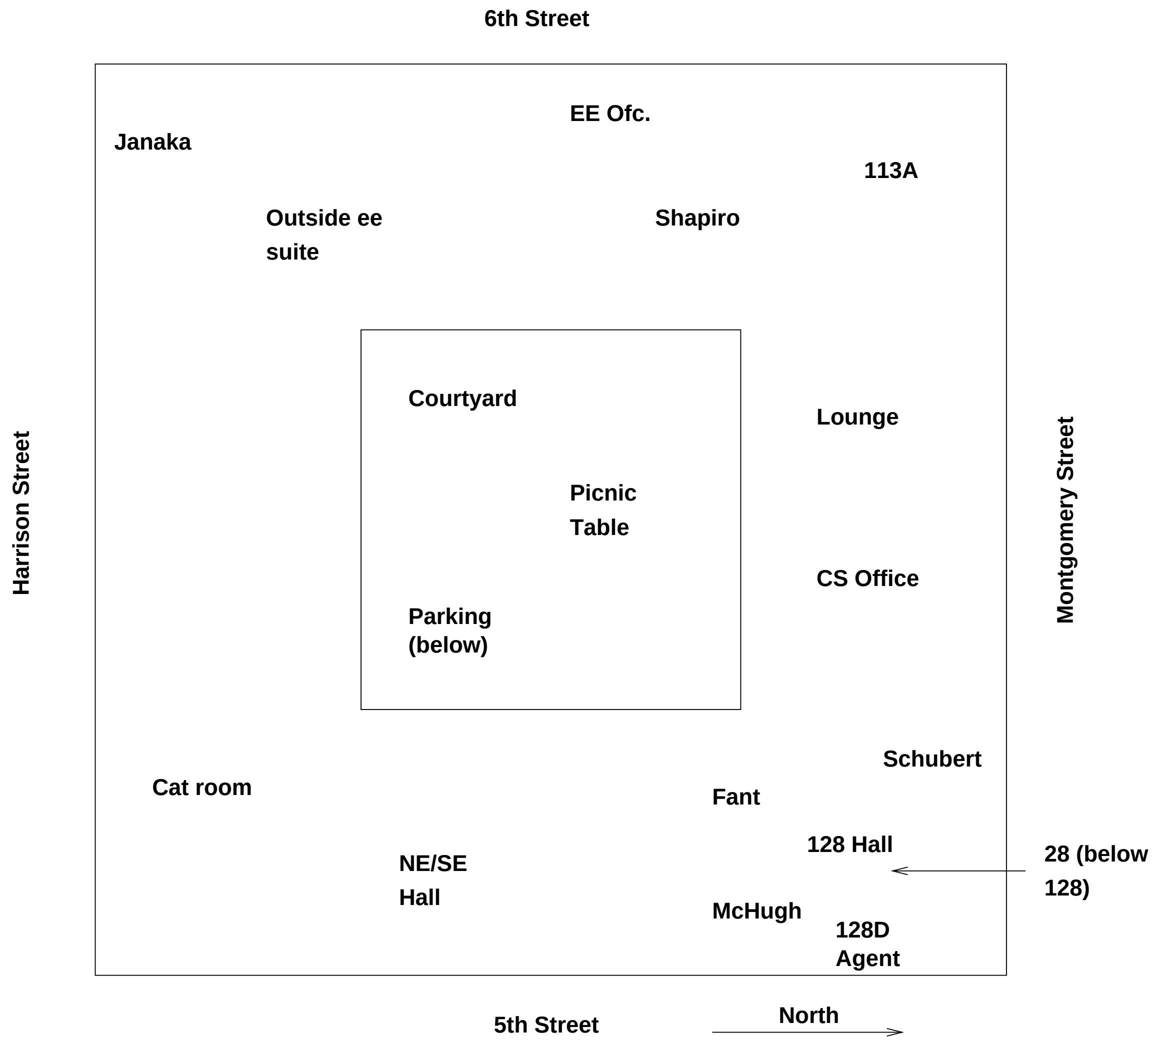

Figure 1: Floor plan of PCAT

	Location	Signal %	Signal Grade
1.	Jim's office (128D) comment: Still some dropoff here probably due to metal surfaces in office. My office is in the NE corner of the building.	90	good
2.	outside in 128 hall	70	good
3.	Schubert's office	52	good
4.	PCAT 28 comment: Downstairs lecture room. No loss on acceptable. signal quality varied acc. to how far from my office you were. Not as good in the front. of the room near the blackboard.	40-60	acceptable-good
5.	downstairs parking lot	< 20	poor
6.	John McHugh's office	70	good
7.	Karla Fant's office	50	good
8.	outside McHugh's office	65	good
9.	c.s. office (mid North quadrant)	39	acceptable
10.	student lounge	32	acceptable
11.	113A area, across from Joe Maybee's office	≤ 20	acceptable/poor
12.	Len Shapiro's office (end of hall)	40	acceptable
13.	EE office	14	poor
14.	outside EE office suite (where Janaka's office is) diagonally opposite from my office in quadrant (SW corner)	23	acceptable
14.	outside Janaka's office in far SW corner	10	no connection
15.	Picnic table in fountain square at center of building	37	acceptable
16.	cat room	varies 15-34	poor to acceptable
17.	halfway in hall on NE/SE side	34	acceptable

Table 1: Wavelan Signal Strength Measurements

the hall back to the base station in my office in the NE corner. 113A has to go through a number of walls, the cs office, etc to get there. John McHugh points out that the Wavelan radio frequency at 900 mhz is on the order of having a wavelength around the size of a metal cabinet. Walls and metal cabinets do indeed seem to cause signal loss and limit the range.

I have tested the wavelan setup at home with the same equipment. I have a 2 story wood house with about 3000 sq. feet. There was signal loss but the results were always rated "good" from one end to the other even with the base station at the far end downstairs.

Link-layer Routing and Security Observations

There are two possible physical networking configurations that might be employed for constructing a Wavelan-based LAN (not a WAN). The constant here is that you would use PCMCIA cards in laptops for mobility. What varies is how you might choose to setup connectivity to your wired infrastructure. You could choose to 1: employ ISA cards in PCs that would function as either bridges or routers, or 2: you could buy the Access Points from ATT that are effectively a packaged version of #1. The Access Point (AP) is effectively a 386-based computer with an ethernet card and a wavelan ISA card in it. In point of fact, you can buy the ISA card ($600) and construct your own AP or you can buy it. The AP itself is around $2k list. The interesting difference here though is that the AP comes with so-called "roaming" software.

The vendors of such products as Wavelan are limited to what they can do to enable mobility and can only act at the link or hw layers (the domain of device drivers). They are limited to "same network" in terms of the network layer; e.g., IP routing will fail if you moved from one base station to another and the base station was on a different IP network segment (subnet change). Mobile LAN vendors by definition have to have a bridging solution. It is also desireable that the number of packets on the radio link be minimized. One partial solution is called "roaming" which effectively means that a link-layer registration protocol is used between PCMCIA-based clients and base stations (APs). This is a software solution. APs send "beacons" which are periodic special link-layer packets that have a beacon header on top of the normal link-layer header. The mobile node (client) uses a link-layer sign-on protocol in response to any received beacon that has the higher signal level. As a result, packet duplication is eliminated. APs that receive packets destined for a given mobile node can discard them if the node is not registered with them. This may be a useful optimization, although it does not address network layer mobility. There may also be room for criticism on a couple of points in terms of secure networking.

Link-layer beaconing may have a flaw in that an outside attacker might pull up in a van and run an Access Point that could have a better signal strength and thus could suck up all packets. This is in some sense a general problem with any routing protocol with a central system that says "send the packets to me". For example, this is also true for the Mobile-IP Home Agent since tunnels to and from it might be manipulated by attackers or the HA itself might be spoofed. It may be better to simply beacon at the network layer and thus take advantage of IPSEC mechanisms like authentication. In general, routing control (or all packets) would be much more secure with network-layer authentication mechanisms. For example, ARP could stand to be authenticated as well, although possibly the mechanism here should be kicked upstairs and become an ICMP registration protocol for ad hoc systems.

From a redundancy point of view, it may be that signing-on with one system and only using it for access to the wired infrastructure is not a good thing either. If the link-layer is under stress (packet loss and packet corruption), which is certainly a likely attribute of radio by definition of, multiple paths may be useful.