The Case for Wireless Overlay Networks

Randy H. Katz and Eric A. Brewer

Electrical Engineering and Computer Science Department, University of California, Berkeley, CA 94720-1776

ABSTRACT

Wireless data services, other than those for electronic mail or paging, have thus far been more promising than successful. We believe that future mobile information systems must be built upon heterogeneous *wireless overlay networks*, extending traditional wired and internetworked processing "islands" to hosts on the move over coverage areas ranging from in-room, in-building, campus, metropolitan, and wide-areas. Unfortunately, network planners continue to think in terms of homogeneous wireless communications systems and technologies. In this paper, we describe a new wireless data networking architecture that integrates diverse wireless technologies into a seamless internetwork. In addition, we describe the applications support services needed to make it possible for applications to continue to operate as mobile hosts roam across such networks. The architecture described herein is being implemented in a testbed at the University of California, Berkeley under joint government/industry sponsorship.

Key words: wireless communications, mobile computing, overlay networks, bandwidth adaptive applications;

1. INTRODUCTION

Given the recent award of the PCS licenses, combined with the rollout of digital direct broadcast satellite services, the digitalization of cellular services, and the deployments of cellular packet services, alternative wireless data services are poised to proliferate. Unfortunately, network planners continue to think in terms of homogeneous wireless communications systems and technologies, providing either low bandwidth connectivity over the wide-area or high bandwidth connectivity over a small area.

We believe that future mobile information systems are likely to be built upon heterogeneous *wireless overlay networks*, extending traditional wired and internetworked processing "islands" to hosts on the move over a wide area. For example, a high bandwidth in-room infrared (IR) network can be "overlaid" with a more moderate bandwidth radio frequency (RF) network to provide connectivity between rooms. Or a low-tier/low mobility) in-building PCS system could be overlayed with a high-tier/high mobility PCS system to provide cost effective connectivity bridging the local and wide areas. Figure 1 illustrates this concept.

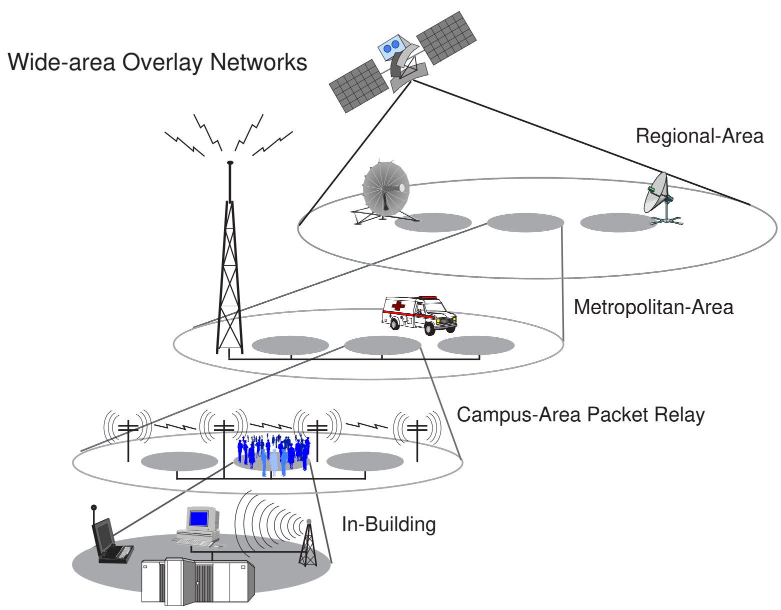

Figure 1 Wireless Overlay Networks

Because they are constructed from diverse wireless technologies with different engineering constraints, overlay networks vary widely in bandwidth, latency, and coverage. The key technical challenge is to develop *overlay internetwork management* and *applications support services* that allow mobile applications to operate across a wide range of network performance, roaming among potentially competing service providers, and choosing among alternative overlays for best performance given the current network state and application needs.

To verify this concept, we are creating a wireless overlay network testbed, spanning from in-building to direct broadcast satellite systems. We are developing pilot multimedia applications to drive the design and validation of the interfaces between applications and the network. Our overarching goal is to demonstrate a scalable architecture that can support wireless access across multiple overlay networks while delivering high levels of end-to-end performance to applications.

The rest of this paper is organized as follows. In the next section, we illustrate the overlay concept by giving an applications scenario enabled by such technologies. In Section 3, we describe existing and emerging wireless data technologies, and their strengths and weaknesses for supporting bandwidth and latency-sensitive applications. Our "gateway-centered" overlay architecture is presented in Section 4. Mechanisms for network management and handoff are described in Section 5, and our approach for application support services is given in Section 6. Section 7 describes related work. Section 8 presents our summary and conclusions.

2. APPLICATIONS ENABLED BY WIRELESS OVERLAYS

A police helicopter observes a serious auto accident on the Golden Gate Bridge; it immediately calls for an emergency medical team to be dispatched. The helicopter, using a wide-area overlay network, transmits low frame rate digitized video of the crash site to the medical team, to assist them in assessing the severity of the crash and the need for supporting units. They decide to request backup. The helicopter crew also establishes a two-way audio conversation with the ambulance driver to direct him to the scene via the least congested traffic route, perhaps forwarding the medical team a street map with a hand annotated route.

At the scene, the medical team and police units free the victims from the wreckage, stabilize their condition, and head them to Mt. Zion Hospital in San Francisco. One of the injured passengers has a serious head injury, requiring immediate emergency surgery. The team requests a suitable medical facility to be prepared. Critical patient information is downloaded from a regional archive at UCSF to the in-ambulance team via a wide-area wireless data network. Simultaneously, detailed patient records and archived neuro-images are forwarded to an image file server at Mt. Zion via a high performance wireline network. The in-hospital wireless network allows the emergency room physicians to retrieve and view the patient's images and records on their portable multimedia devices. They decide that a CAT scan should be performed on the patient upon arrival. The new scan, combined with existing patient information, will support their planning of the proper surgical suite and determination of a course of operating procedure.

After the ambulance arrives at Mt. Zion, the patient is routed to the CAT scanner. The archive system automatically appends the new neuro-images with the victim's existing patient information. The surgical team reassesses the injury's seriousness, and decides that a 3-D reconstruction of the CAT images is needed. Mt. Zion Hospital requests UCSF's high performance computers construct a 3-D visualization of the head trauma and provide a computer-assisted guide for the surgical procedure. These are instantaneously forwarded to Mt. Zion via the wireline network, where they are rerouted to the surgeon's portable display via the wireless in-hospital network while she is enroute to the OR. Its unusual nature prompts her to hold an impromptu video consultation with a distinguished neurosurgeon on his way to a lecture at the UCSF Medical School across town. Just as the two doctors agree on a course of surgery, the stretcher is wheeled into the operating room, and the neurosurgeon steps into his lecture hall.

This illustrates several requirements that broadly characterize "tactical" applications: (1) Seamless interoperation of wide-area wireless overlay communications for operational units on the move, exploiting better communications inside buildings yet still able to provide connectivity over wider-area, lower bandwidth, higher latency wireless networks; (2) Access by mobile applications to the full computation and information resources of in-building infrastructures; (3) Rapid exchange of media-rich data, including video images, audio communications, annotated maps and graphics, handwriting and text, all in support of decision makers on the move, in real-time where possible, and in near real-time otherwise; and (4) Type-sensitive data exchange adapted to the available communications quality, exploiting mechanisms like bandwidth-adaptive compression, latency-adaptive mobile host frontend processing and wireline backend processing.

3. APPLICATIONS VIEWPOINT

Wireless data networks have yet to prove themselves. It is estimated that there are no more than 100,000 wireless data users in the U.S. today. This is in comparison with over 20 million cellular telephone users and 30 million paging subscribers! To yield ubiquitous connectivity, what is needed is a wireless internetwork formed from multiple wireless overlays interconnected by wired segments. For many emerging applications, (near) real-time audio/video and collaborative support for subscribers on the move is an absolute requirement. Table 1 identifies performance characteristics of representative wireless technologies. It also indicates approximate video and audio performance.

The marketing literature for wireless devices focuses on bandwidth, but there is little discussion of other important metrics, such as latency, expected packet loss rates, probability of packet retransmission, and available bandwidth per cubic foot. For example, a 10 mbps Ethernet typically supports 50 users; a 1 mbps in-room IR network designed to support 5 users collocated in the same room might be just as effective in terms of performance. Certain applications, such as web browsing, are somewhat tolerant of latency. Others, like decoupled mail access, can operate quite well even in a network environment characterized by very high latencies. Critical issues for wireless network technologies are scaling performance with increasing user densities, handling asymmetric bandwidth, over hybrid networks, with support for roaming and mobility, while constraining power consumption.

Table 1: Characteristics of Alternative Overlay Technologies

Type of Network (Coverage)	Bandwidth	Latency	Mobility	Typical Video Performance	Typical Audio Performance
Room/Building (Radio Frequency Infrared)	> 1 Mbps RF: 2-20 Mbps IR: 1-50 Mbps	<< 10 ms	Pedestrian	2-Way, Interactive Full Frame Rate (Compressed)	High Quality 16 bit samples 22 KHz rate
Campus-Area Packet Relay	≈ 64 kbps	≈ 100 ms	Pedestrian	Medium Quality Slow Scan	Medium Quality Reduced Rate
Wide-Area (Cellular, PCS)	19.2 kbps	> 100 ms Highly Variable	Pedestrian/ Vehicular	“Video Phone” or Freeze Frame	Asynchronous “Voice Mail”
Regional-Area (LEO/VSAT DBS)	Asymmetric Up/Down 100 bps to 4.8 kbps 12 Mbps	>> 100 ms	Pedestrian/ Vehicular Stationary	Seconds/Frame Async Video Playback	Asynchronous “Voice Mail”

Consider the following wireless technologies and their relative performance advantages and disadvantages:

- In-Room Infrared: Infrared technology is attractive for providing wireless connectivity for well defined physical spaces like offices and meeting rooms (typically 30 m x 30 m). Commercially available diffuse IR devices achieve 1 mbps. Observed latencies are comparable to wireline networks. Research devices have been demonstrated at 50 mbps [11]. This allows reasonable quality compressed video and high quality audio. Directed IR interfaces are expected to be inexpensive and ubiquitous in future PDAs and laptop computers, although they are primarily for interfacing to peripherals and docking stations.
- In-Building Radio Frequency: Numerous wireless local area networking products are commercially available using spread spectrum in the unlicensed bands. They yield bandwidth in the 1-2 mbps range, but whether this bandwidth is really reused among adjacent cells varies among the product implementations. Good quality compressed video and high quality audio can be achieved as long as user densities are modest. A typical cell size is a several hundred feet—enough to cover perhaps half the floor of a modern university building. Picocellular architectures with much smaller cell sizes are more interesting in terms of aggregate bandwidths, but today’s economics makes the approach somewhat expensive. The newest generation of WLANs support PCMCIA interface cards, making these useful for truly portable devices like PDAs and laptops. While typically implemented with a base station/mobile node architecture, the radios easily support packet radio techniques.
- Campus/Metropolitan Area Packet Relay Networks: Packet radio systems forward packets between mobile hosts, infrastructure radios, and wired access points. Radio link performance in the range of 100 kbps is possible, but this must be shared among all users within the microcell, typically 0.2 to 1 mile in diameter. Latency is a function of the number of radio hops between the mobile host and the wireline connection. 2-3 hops, at approximately 20-40 ms per hop, are typical.
- Wide-Area Packet Switched Data Networks: Cellular Digital Packet Data (CDPD) is a wide-area data overlay to the analog cellular system now being deployed. The system is designed for instantaneous transmission rates at 19.2 kbps, but this will be difficult to sustain. Pricing remains high, in the range of 10-18 cents per KByte, making the service most cost effective for short transactions. Latencies will be highly variable, depending on the cellular voice traffic. Other wide-area data systems currently deployed include ARDIS and RAM Mobile, providing two-way paging like services under 10 kbps.
- Regional-Area Satellite Data Networks: There are numerous low earth orbiting (LEO) satellite proposals planned for the latter part of this decade, offering low data rate services. Deployed Direct Broadcast Satellite (DBS) and Very Small Aperture Terminal (VSAT) services provide shared high data rate downlinks. However, these incur relatively high latencies due to geosynchronous orbit transit times. Furthermore, for DBS the uplink is over non-satellite links, adding some complexity to the system design.

It appears that the concept of multiple overlay networks will be impeded by the mobile host’s need for multiple transmitter/receiver systems. Yet even today it is possible to simultaneously configure a laptop with network adapters for in-room diffuse IR, in-building RF, campus-area packet radio (connected to the serial port) and wide-area CDPD (replacing the floppy drive). The rapid development of multimode radios described in [15] will yield radios that can integrate such alternatives into a more convenient package.

4. GATEWAY-CENTRIC NETWORK MANAGEMENT

4.1. Introduction

IP and TCP/UDP provide the basic network and transport layers upon which overlay and wireline networks will be integrated. Nevertheless, effective management of overlay networks requires protocol extensions that reach down to the network layer. The following challenges must be addressed:

- *Seamless Integration of Overlay Networks*: No general network management architecture exists for effectively integrating multiple overlay networks. Mobile applications roaming across overlays require network intelligence to determine that the mobile has moved from acceptable coverage in one network to better coverage in another. But a global network management algorithm is still needed to control handoffs across overlays based on current mobile connectivity. Link quality is only one metric that determines handover; priority of access, applications needs, and relative cost are equally important. Since overlays may not cooperate with one another, it is crucial to support *mobile assisted handoff*.
- *Support Services for Mobile Applications*: Handover across overlays will change an application's network bandwidth and latency. Needed is a *new applications interface to the network management* layers to allow them to initiate handovers, to determine changes in their current network capabilities, and to gracefully adapt their communications demands. Mobile applications and scalable wireline processing and storage capabilities could be better integrated through agent processing architectures that exploit data type specific transmissions to manage the communications demands over dynamically varying wireless links.
- *Managing Mobile Connections to Support Latency-Sensitive Applications*: Handoffs must be executed with lower latency than is now possible if (near) real-time multimedia applications are to be well supported. One strategy moves the routing and resource allocations to local subnets. For example, roaming authentications can be cached locally to avoid repeated remote, latency-intensive transactions. Needed are *algorithms that exploit information* about the location of mobile devices, the geographical adjacency of cells, and the likely routes devices might take, to improve handoff processing. Hierarchical approaches, which localize information collection to the region or the subnet containing the users, are needed.
- *Load Balancing for Scalable Mobile Processing*: Repositioning within wireless networks will become common. Traffic patterns will not be uniform, with high correlations between mobile host location, their repositionings, and their requests for service. We require network management *architectures that build on decentralized algorithms* to allocate network and processing resources on demand, avoiding the static and centralized schemes.

4.2. The Gateway-Centric Architecture

Conventional wireless networks place the mobile host at the center, with gateways to other networks around the boundaries. Figure 2 illustrates this concept. Such an architecture requires that the wireless network be homogeneous. Our approach, on the other hand, is *gateway-centric* (see Figure 3). In this view, the subnet-to-subnet gateways are placed at the center of the architecture. Diverse wireless (and wired) networks are integrated through software that mediates between the mobile host and the networks it could possibly connect to, supporting the mobile host as it roams among the multiple wireless networks. The figure shows specific instances of commercially available wireless networks that we are building into our gateway-centric architecture; those in black are networks we currently have access to, those in gray are other networks of interest.

The figure also introduces *cooperative* overlays. Most wireless networks, deployed by competing service providers, provide little controllability above the subnet for routing, network management, or applications support purposes. We call such networks *black pipes*, because packets transfer between the wireline gateway at one end and a wireless link to/from the mobile host at the other with no control over routing, priority, quality of service, etc. Cooperative networks, on the other hand, provide greater visibility of the

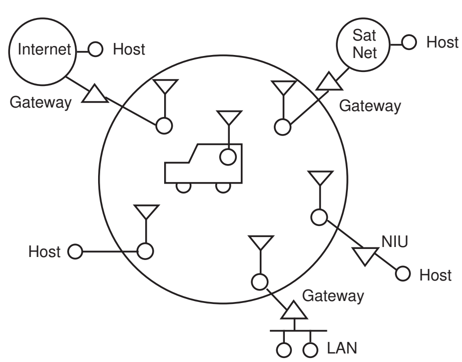

Figure 2 Network-Centric View

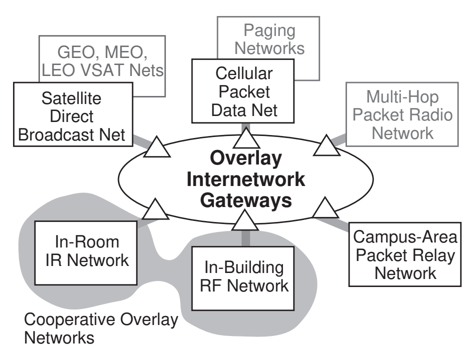

Figure 3 Gateway-Centric View

network control functions to the gateway as well as application-level software, enabling management software to balance the network load among the cooperating networks or to choose a particular subnetwork for priority applications, etc.

4.3. The Architecture in More Depth

Figure 4 provides a conceptual architecture for wireless overlay internetworking and mobile applications support services. More details of our technical approaches towards implementing these functions are presented in Sections 5 and 6. Working bottom up, the layers and the functionality they provide are described in the following subsections.

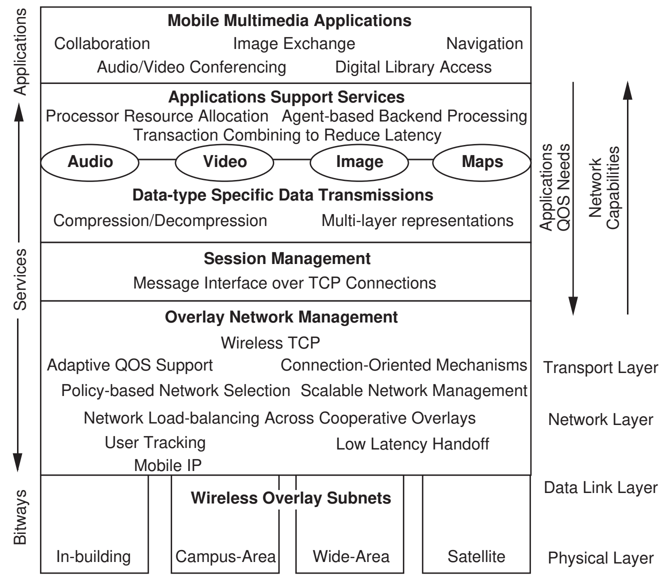

Figure 4 Layered Architecture

Wireless Overlay Subnetwork Layer

This corresponds to the physical, data link, and network layers in the OSI terminology. It represents the independent wireless networks we are integrating into our overlay architecture. These networks interface to the rest of the architecture through IP protocols. In general, the details of the underlying physical channels, media access, link-level, and routing protocols are not visible above this layer, because these are provided by a commercial network service provider, our so-called “black pipe” networks. There is no way for us to take advantage of them at the Overlay Network Management Layer. However, for in-building networks, we have the ability to fully expose the underlying protocols, including those that control handoff decisions.

For example, IR is particularly well suited for confined spaces, like offices and meeting rooms. However, handoff between IR cells is not easily accomplished, especially when crossing into large open spaces like corridors. An RF overlay advantageously provides connectivity for these gaps in IR coverage, as well as providing an alternative source of bandwidth for areas of heavy traffic. We can exploit this exposed information, such as the degree to which network resources are already allocated to mobile hosts, to demonstrate dynamic load sharing between overlay networks.

Overlay Network Management Layer

This layer corresponds to the network and transport layers in the OSI terminology, with extensions to manage the multiple overlay networks. It builds an internetwork on top of the independent wireless (and wired) subnetworks. It must handle the complex issues of routing packets across heterogeneous subnets, while also choosing the most appropriate wireless network to provide end-to-end connectivity to the mobile host given the application’s specifications for quality of service.

Tracking users and managing their handoffs in regions of high user density generates considerable additional processing load. The key to maintaining low latency in the network management algorithms is to design these algorithms to be able to exploit scalable processing techniques, such as harnessing network of workstations (NOWs). This layer takes full advantage of the distributed processing power of networks of workstations.

Existing handoff algorithms have been designed for homogeneous wireless subnetworks. Algorithms to support handoff between heterogeneous subnets are more complex. They base the handoff decision on application-specified policies as well as periodic measurements of the quality of the underlying network connectivity by the mobile host. These policies can constrain the feasible handoffs. For example, an application might specify that high priority traffic traverse the available connections with the lowest latency, regardless of cost. Less critical traffic might be constrained to travel over the lowest cost connections currently available. Issues like signal quality, bit error rates, and the resulting probability of packet loss and retransmission must be considered as part of the handoff decision process. Under certain conditions, the handoff algorithm might defer switching its connection to a higher latency channel, even though it has better signal quality due to stronger signal strength, from a low latency channel with a weak signal.

Maintaining multiple powered-on network interfaces consumes precious power. To extend battery life, we assume that only one network interface is normally active at any point in time. Other interfaces are in standby mode and powered up from time to time to determine the quality of connectivity to their respective overlays. A trade-off exists between the frequency of determining link quality, the power consumed by the operation, and how up-to-date the information is about the network state to drive handoff decision making. When the mobile host is low on power, a good strategy is to connect to the network with the lowest transmitter power demands, and the handoff algorithm in the mobile host may scale back the frequency of probing other overlays. The system context may also dictate the frequency of probes. For example, as the mobile host moves towards the fringe cells of its current network overlay (e.g., a cell that covers an exit from the building), a good strategy would be to increase the probes to alternative overlays.

Within the network management layer, we must extend existing protocols to improve reliable transport over wireless links. Existing real-time protocols achieve their performance guarantees through admission-control procedures. Once a link becomes fully subscribed, no new connections are allowed to use that link. New schemes for link sharing are being developed to better utilize link bandwidth [8]. Yet even in these schemes, once a connection is granted to an application, it does not need to tailor its demand. These approaches are of limited use in wireless overlay networks, where applications must be able to adapt dynamically to dramatic changes in the quality of their connectivity as they move between networks. Since entire subnets are "black pipes," we must depend on end host adaptation to changes in the network characteristics, rather than strong support from the network, such as resource reservations. And we must extend support for adaptation beyond the applications-independent protocol stack to the applications running on the mobile host.

Session Management Layer

This layer provides applications-independent session mechanisms. One example of a useful session level mechanism is a message-oriented interface built on top of a single TCP connection. Without such a capability, applications like the World Wide Web that use TCP but need a message-oriented interface are forced to use a separate TCP connection for each message. The inefficiencies of the standard approaches are considerable.

Data Type Specific Transmission Layer

This layer provides specialized functions for applications to manage the transport of data type specific objects. Multimedia applications process audio, video, image, and other semantically rich data types. Section 6 describes the mechanisms we are implementing to support alternative representations and compression/decompression schemes for commonly encountered data types for multimedia applications. This allows us to support flexible trade-offs between bandwidth and latency, based on the network selected for connectivity and the relative cost of using that network, such as network load and billing costs.

Applications Support Services Layer

This layer provides resource management and distributed processing services that allow applications to migrate computation between the mobile host's environment and the wireline infrastructure. The quality of network connectivity, in terms of bandwidth and latency, as well as the mobile host's processing capabilities and available battery life, determines how aggressively to partition the application. The management of backend processing, choice of data representation, and level of compression are determined by "agents" or proxies running on behalf of the mobile host in the wireline environment.

Certain application-specific mechanisms can also yield improved efficiency. For example, multiple TCP connections can be combined into a single connection to reduce round trip latency. All image link references within an HTML document could be packaged into a single TCP connection, thereby incurring a single round-trip network delay for a document spanning multiple images. This is particularly effective when the network exhibits highly asymmetric communications paths or particularly long latencies, such as for satellite channels.

Mobile Multimedia Applications Layer

This layer consists of the actual mobile applications exploiting the underlying services and network management functions. If the architecture is robust, we expect to be able to reuse major pieces of code across applications for conferencing, collaboration, and navigation. These code fragments will form the basis for a library of mobile computing modules.

Application/Network Interfaces

Figure 5 is a high-level architectural view of how the applications and the network layers interact. Applications interface with the applications support layer (combining application support services and data type specific transmission layers) through data type specific manipulation operations. The details of how data types, such as image, video, audio, and text, are represented for the purposes of network transport is chosen by the applications support layer. Options include lossless or lossy compression, multi-layer representation, etc. The applications support layer, in turn, uses information about network performance collected from the network monitor to choose a "best" representation. The (overlay) network management layer interprets network performance in terms of desired quality of service metrics (i.e., "service kinds"), such as effective bandwidth and latency, requested by the application layers. With this information, it chooses the best overlay network to support the current set of applications running on the mobile host. Note that elements of the application support layer and network management layer reside both on the mobile host as well as the server side of the client-server interaction.

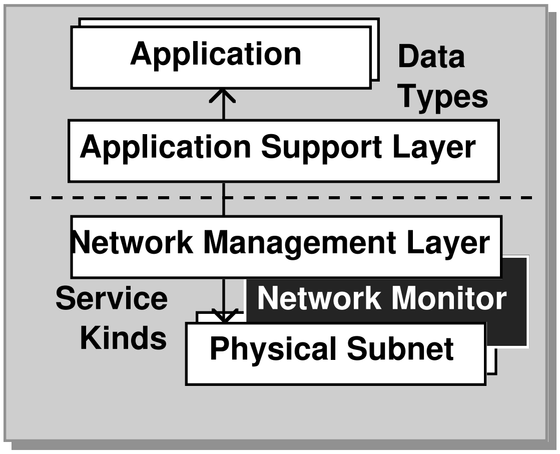

Figure 5 Application/Network Interface.

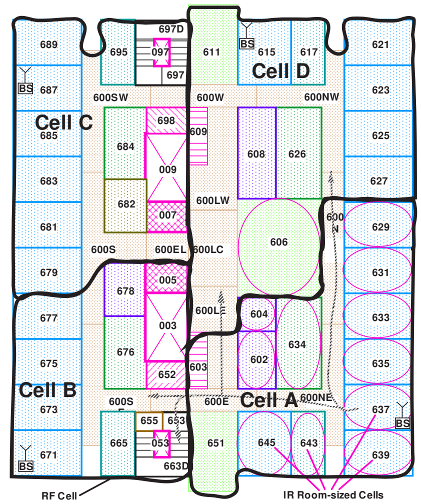

Figure 6 Exploiting Geographical Constraints to Assist Hand-off Processing

5. OVERLAY NETWORK MANAGEMENT

5.1. User Tracking, History, Geography

Network management exploits the location of users and their physical environment to yield low latency handoffs and better allocation of resources to high traffic areas. Particularly inside buildings, it is possible to use the layout to localize the collection and analysis of tracking information, and to drive the handoff algorithms. For example, consider a mobile host receiving a continuous media packet stream while moving through a building. Its path is constrained by the building's physical layout. Figure 6 shows a floor of a typical building, covered by four RF cells. It is not possible to traverse directly from Cell A to C.

Existing mobile routing algorithms, like Mobile IP, typically require high latency operations like communicating location updates back to the mobile host's home network. We have implemented a handoff strategy based on IP multicast to reduce these expensive location update operations [18]. By tracking the speed and trajectory of the mobile host through the network, and with knowledge of the geographic layout of base stations, we form multicast groups containing the mobile host's current base station as well as a small number of base stations that the mobile host can reach in the near future. In Figure 6, the base station in cell A could be formed into a multicast group with those in cells B and D. Packets for a mobile host in cell A are simultaneously delivered to cells B and D. As the mobile host comes within range of one of these, the mobile host uses the relative strength of base station beaconing signals to choose when to register with a new primary base station. The multicast groups are then reconstituted as the mobile host moves through the network. The new base station already has the tail of the packet stream destined for the mobile device. At the modest expense of additional buffering, backbone network bandwidth, and inter-base station control communications, handoff can be made sufficiently low latency to support real time streams.

Figure 7 shows the performance of a reliable file transfer as a function of frequent handoff operations. In the figure, the X axis is time in seconds, with a handoff operation forced every 10 seconds. Along the Y axis is packet sequence number. The relatively constant slope of the line indicates that packets are delivered without any significant delays across handoffs.

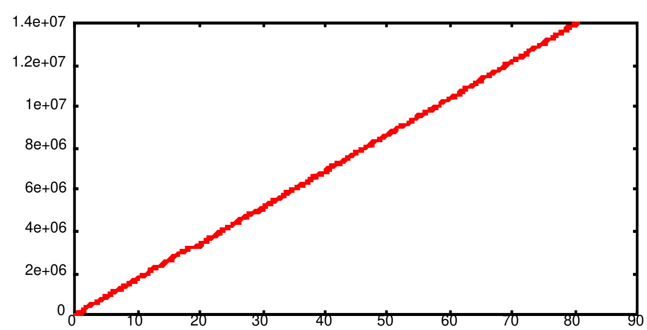

Figure 7 Packet Sequence Numbers as a Function of Time

Support for "Vertical" Cooperative Handoffs

Conventional wireless handoffs are *horizontal*, i.e., within a homogeneous wireless subnet. The mobile host or associated base station detects a degraded signal as it reaches the fringe area of its cell. In *mobile-assisted handoff*, the mobile host listens for beacon signals from base stations in adjacent cells, choosing to register with the cell with the strongest signal.

Vertical handoffs allow mobile hosts to roam between overlays (see Figure 8). The mobile host, or higher level network management, determines when to switch the connections to an alternative overlay network, driven by signal quality, network load, or the

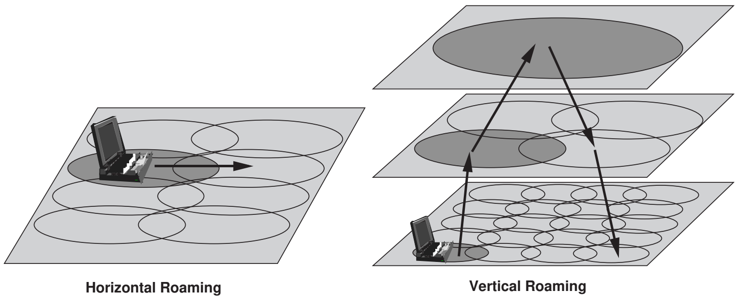

Figure 8 Roaming in Wireless Overlay Networks

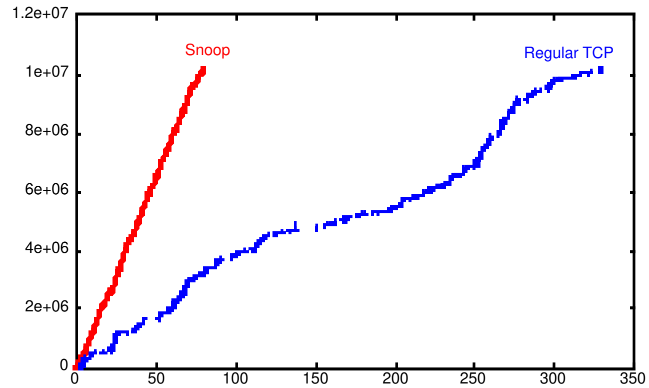

Figure 9 Snoop TCP versus Conventional TCP

costs of using one overlay versus an alternative.

For in-building networks, we selectively violate strict network layering to expose information from below the network layer to higher level overlay management software. The allows wireless subnets to cooperate to reduce vertical handoff latency as well as to better share connectivity and network loading between alternative overlays. Our cooperative overlays are formed from independent in-building IR and RF physical networks. They are treated as logically independent networks that can only be integrated through overlay management. Basically, the multicast method described above can span base stations for different wireless subnets.

Consider a group of mobile users in a meeting room. The in-room IR network provides the preferred connectivity. But if its bandwidth becomes oversubscribed, resulting in poor service to the mobile hosts in the room, some hosts can be forced to handover to the RF network. Furthermore, since the overlays are cooperating, new connections can be set up in the RF overlay in advance of actual handoff using the same multicast strategy as in homogeneous wireless networks. This ensures that the handoff is executed with low latency. Several management mechanisms are possible. The IR network could refuse new connections to mobiles in the high density area, generating mobile-initiated handover to the RF network. Or the IR network management algorithms could “request” assistance from higher level management layers, with a more global view of the condition of the subnets under its control. Certain mobiles could then be selected for overlay handoff based on their current network performance needs.

Connection-Oriented Services for Mobile Hosts

Connection-oriented TCP performance can also be improved for wireless links. The standard protocol interprets lost packets as congestion, employing mechanisms that unnecessarily decrease the rate of transmission. Our approach caches the unacknowledged TCP packets in the base station while performing aggressive local retransmissions over the wireless link and managing duplicate acknowledgements from the mobile host [3]. It requires no changes to TCP’s semantics nor to any implementations of TCP in the wireline portion of the network, except at base stations under our control.

Figure 9 shows the performance of this “snoop” approach to reliable transmission over wireless. Once again, the X axis is time in seconds while the Y axis is packet sequence number. We injected a single packet error every 256 Kbits transmitted. While errors in wireless channels tend to be bursty, the controlled errors we introduced represent a relatively low error rate for real wireless links. The figure shows the speed up and the predictable latencies that can be achieved when local methods are used to fixed local transmission problems, rather than relying on TCP’s end-to-end mechanisms.

6. APPLICATIONS SUPPORT SERVICES

6.1. Network Servers

Mobile applications will demand access to the same remote computing power now available from desktops. An attractive, cost-effective model for scalable compute resources are NOW-based compute servers. How best to use them varies with the performance and reliability of the wireless link. Nevertheless, we believe that shared computing power can enable very powerful mobile services.

For cases where the network performance is poor, we will use the network servers to execute agents on behalf of the mobile units, and to perform application- and data-specific compression before transmission over the wireless link. Agents mitigate poor latency by reducing the number of link crossings, while compression increases the link’s effective bandwidth. Both of these strategies can require substantial cycles from the network servers.

6.2. Uniform Architecture for Applications Support

Our development toolkit supports "data type aware" mobile applications, characterized by access to multiple databases, media-rich documents, and global network access. The toolkit must dynamically adjust the application's bandwidth and latency expectations to the capabilities of the network. Each mobile device has a *proxy*, which is a process running on the NOW, responsible for managing the wireless connection, deciding on its appropriate level of compression and encryption, and performing computation on behalf of the mobile client both interactively and in the background. The representative has the same access rights and security privileges as the mobile end user.

The architecture relies on strongly typed transmissions. A dynamically extendable type system enables type-specific compression levels and abstraction mechanisms for progressive object transmission, thus husbanding link bandwidth. For example, depending on link quality, we send a raw, compressed, or lossy/highly compressed bitmap to a mobile client. As an extreme example, we can perform edge detection and OCR on the bitmap of a topographic map, only transmitting the text and edges. This *distillation* process produces a summary version that can be used to evaluate the value of the full bitmap; in the case of the topographic map, the distilled version is usually sufficient for making decisions. These techniques exploit the backbone compute servers to trade software latency to compress for increased effective bandwidth, thus reducing overall latency. The compressed forms can also be cached to further increase their benefit.

Figure 10 shows an example of the image proxy at work. A multicolored, large scale image of the cartoon character Stempy is reduced to a four level gray scale image thumbnail in about two seconds of processing on a SparcStation20 workstation.The original image would require approximately 25 seconds to transmit over a 19.2 kbps modem. The translated image requires only a third of a second of transmission time at the same data rate. If the application is a web browser, the user can use the thumbnail to choose whether s/he wishes the full resolution image to be downloaded.

Figure 11 illustrates another real-time proxy we have constructed, this time for video transcoding [2]. The transcoder accepts video in several different formats, including motion JPEG, H.261, NV (network video) formats, and maps these into alternative formats. Motion JPEG may arrive over a high speed wired network at 3 mbps and be transcoded into an H.261 format suitable for transmission over a wireless local area network at 256 Kbps. The transcoding is accomplished by executing the inverse encoding (E^{-1}), quantization (Q^{-1}), and discrete cosine transform (DCT^{-1}) steps to map the video frame into the pixel domain, and then to encode forward from this representation for the appropriate target format. For example, the mapping from pixel domain to H.261 introduces a motion compensation step before the DCT, Q, and E steps. Such an approach can be improved by transcoding from the DCT rather than the pixel domain, and to introduce a conditional replenishment step that passes through blocks that do not change between the two representations. For low motion video streams, we have achieved full 30 fps transcoding rates on a SparcStation 20; for higher motion streams, the transcoder rate drops to a still acceptable 21 fps. Without the optimizations, the transcoding rates are 13 fps for low motion streams and 7 fps for high motion streams. For scalability, we are porting this transcoder to NOWs.

These early prototype proxies are first steps in developing a more comprehensive approach to building data type specific proxy services. We are developing a range of bandwidth/computation tradeoffs for each data type, distinguishing between *versions*, which are independent, from *layers*, which are part of one version and form a sequence of increasing detail useful for progressive transmission of objects. For example, a PostScript object has at least two versions: ASCII text and full PostScript description. Images in the doc-

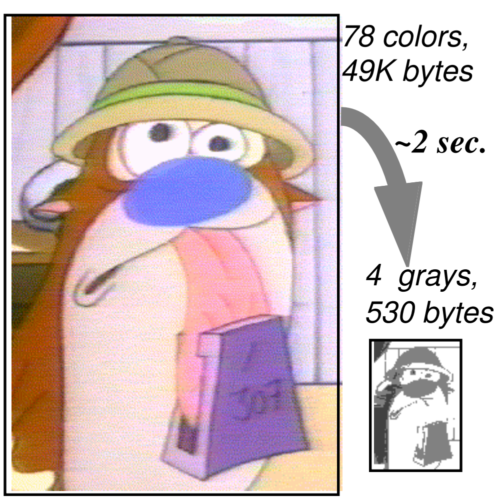

Figure 10 Real-Time Translation By Image Proxy

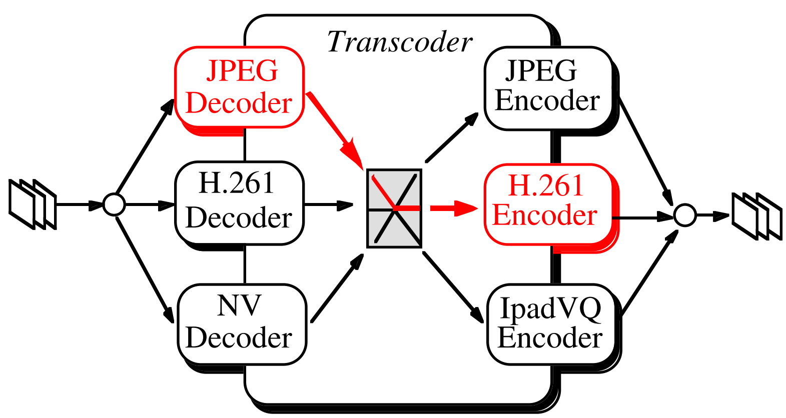

Figure 11 The Video Gateway Transcoder Proxy

ument, such as a map, use a corresponding type known to the application support subsystem. A map data type could exploit layering for progressive transmission, first sending the major roads and landmarks, and adding layers of increasing detail with each subsequent transmission. Layers/versions can have meaningful names; an application can ask for specific layers, such as a map's "elevation lines" or "county boundaries."

To drive the process by which representations are chosen, we are also developing cost models based on link instrumentation and statistical model fitting. These techniques, first developed for mapping parallel programs onto a variety of platforms [4], can accurately predict the overall latency for each transmission option for a given data type. The models include the link latency and the compression/decompression time, and exploit knowledge of the cache contents. We can decide the optimal format by simply evaluating the models for each level and picking the one with the lowest predicted latency. The models typically pick the first version to transmit; subsequent versions are controlled by the user. We are also experimenting with predictive models to minimize access charges, by selecting the cheapest network and data representation.

The application architecture also provides for speculative downloads and client-side caching. This trades latency for bandwidth, which is useful in asymmetric communications environments. Client caching is advantageous for high-latency or low-bandwidth links avoiding round trips at the expense of more client memory. In fact, the user code on the mobile device will be cached using these mechanisms, so that the clients can dynamically download missing decompression or decryption functions as needed.

We can reduce required bandwidth and round trips by allowing applications to move computation across the wireless link via agent programs. An agent is just an instance of a small application, so it has access to the same strongly typed connections as any other, and can use the same mechanisms to adjust its behavior depending on dynamic link quality. In fact, user code is just another transmission type and can be sent in either direction. The recently announced Java system from Sun Microsystems provides an ideal substrate for these agent programs.

Consider an application that combines information from multiple heterogeneous databases, using a full-text web indexer and on-line library card catalog. Typically the end user must do the combining. Alternatively, an agent running on the representative can access, combine, and filter the requested information. The agent uses the card catalog to find relevant documents and the web-indexer to determine which are on line. The end user then receives an annotated reference list with hypertext links for the documents. The wireless link is only traversed to start the search and receive this list; alternatively the results could returned along an alternative path such as e-mail or FAX. The agent prefetches abstracts or whole documents should the end user decide to view an on-line document.

The proxy architecture allows uniform handling of disconnection, since it maintains state across sessions. The network management layer will notify the application support services layer of a disconnection, whether intentional or not. This allows the services layer to store results until the user reconnects without any help from the application. This is particularly useful if a connection is expensive or unreliable, since the user can start an agent (on the proxy) and reconnect later to get the results. This also gives us more freedom in implementing fast handoffs; if we temporarily lose the connection during a handoff, there are no adverse effects and applications will in fact not even know (unless they asked to be notified) that they were disconnected.

7. RELATED WORK

7.1. Packet Radio Networks

Driven by the military's need for communications in regions with poor infrastructures, ARPA initiated its pioneering packet radio program in the early 1970s [12, 13]. The original program focused on algorithms for self-configuring, self-managing packet switch networks with support for mobile nodes. The ARPA Survivable Adaptive Networks (SURAN) program of the 1980s focused on eliminating vulnerabilities, such as network partitioning, in packet radio networks. [14] reviews the critical packet radio network design issues at the physical, data link, and network layers, while [10] describes the implementation of the routing and network management protocols. Packet radio algorithms for link determination, discovery of network topology, packet routing, and route dissemination assume a homogeneous wireless technology. Wireless overlay networks introduce the additional flexibility (and complexity!) of alternative physical links, with varying performance, along which to establish network connectivity and to route packets.

7.2. Wireless Data Technologies

Metricom is deploying a wireless IP packet relay network, based on wireless modems, poletop radios, and wired access points to the Internet. Cellular Digital Packet Data (CDPD) is a commercially available cellular data overlay delivering a medium rate, "moderate" latency IP packet service [17]. It is offered by GTE Mobilenet in the San Francisco Bay Area.

Pacific Telesis Corporation has recently acquired PCS spectrum licenses in the San Francisco Bay Area, as well as in Southern California. The evaluation of alternative technologies is currently underway, with initial network services to be deployed in 1996. A likely candidate technology is based is the European digital cellular standard, GSM. Packet data services for GSM, called GPRS, are still in the early stages of definition.

The International Telecommunications Union (ITU) is moving forward with its effort to define Future Public Land Mobile Telecommunications Services (FPLMTS), which are meant to seamlessly integrate cordless and cellular telephone overlays with emerging satellite communications services. A similar Universal Mobile Telecommunications System (UMTS) is being defined by researchers in the European Community countries [5]. Both efforts are focused on voice rather than data services, but eventually packet data will no doubt be included in these evolving standards.

The Hughes Direct Broadcast Satellites are currently in operation, providing coverage to most of North America. It is *highly* asymmetric, with a shared multi-mbps satellite downlink, a dial-up Internet gateway uplink, and latencies greater than 500 ms. Individual users may be able to achieve up to 800 kbps on the downlink. DBS is attractive for wide-area data distribution or other applications in which more information can be usefully broadcast on the downlink than on the uplink.

7.3. Mobile IP

Mobile IP addresses the mobile routing problem, providing connectionless packet service as nodes move through internetworks. Existing implementations introduce intolerable latencies during handoff, making it difficult to support continuous media streams across such repositionings. The protocols take no advantage of local information to reduce handoff latency or to support media streams across handoffs. Variations of Mobile IP have been proposed to improve its performance, using a variety of short cut and tunneling schemes (e.g., [9, 16, 18]), but these have not as yet been included in the IETF Mobile IP protocol specifications.

7.4. Scalable Network Processing and Handoff Algorithms

The cellular telephone system, with the network intelligence embedded in large-scale switching offices, represents the kind of centralized network model that is inherently difficult to scale. The Intelligent Network (IN) architecture attempts to package services like handoff management as processes, decoupling them from the network switching nodes. We believe that a more general distributed networking architecture, based on NOWs, will be inherently more scalable.

Algorithms for cellular handoff have recently gained in attention by the research community, motivated by the need to scale up handoff processing for the emerging PCS networks characterized by higher user densities, smaller cells, and increased frequency of cell crossings. For example, [1] has proposed a “regionalized” scheme for managing handoff in cellular networks. When a connection is established for a given mobile host, virtual circuits are created and associated with all nearby base stations within its roaming region. The mobile host chooses among these to establish its connectivity, without requiring the intervention of any call processor. The latter only becomes involved when the mobile crosses between regions. Acampora’s scheme implements regions as a static concept, as a layer just above the base station leaves. Load balancing is achieved only through coarse-grained admission control at the regional boundaries. We believe that a fully hierarchical approach, allowing arbitrary nesting of regions that could be defined dynamically based on user densities, would be much more attractive. Furthermore, load balancing at the leaf level of base stations will be crucial for reducing latency, since in our model, mobile hosts make heavy use of wireline processing resources.

Initial work specifically focused on latency reduction for handoff has also begun to emerge. Ghai and Singh [8] describe a scheme similar in concept to what we have proposed: dynamic formation of multicast groups based on adjacent cells through which a mobile is likely to move, with continuous media packet streams buffered at all base station members of the group. Knowledge about the physical structure of the space is used to define the groups, i.e., only physically reachable cells can be combined into a group. Although the proposed buffer management schemes consider mobile velocity, coarse trajectories, and time within cell, the protocols do not attempt to exploit likely routes through the space of cells.

7.5. Adaptive Applications and Network Performance Guarantees

Applications built on top of overlay networks must be able to adapt to dramatic changes in performance as the mobile client moves among alternative networks. Existing network support is not flexible enough. Early work on performance guarantees [6] used admission control algorithms to deny service when applications requirements could not be met. Recent research on packet scheduling and controlled link sharing [7] provides mechanisms for better utilization of links, but a dynamic renegotiation of quality of service, either down as more applications demand access to link or up as contention for the link is relieved, is still not supported. Because the bandwidth- and latency-adaptive applications we will develop for our overlay networks, we are in a strong position to develop more dynamic mechanisms for specifying and negotiating quality of service guarantees in the underlying networks. Upcall mechanisms from the network layer are clearly needed, to alert applications to changes in their network connectivity.

8. IMPLEMENTATION STATUS, SUMMARY, AND CONCLUSIONS

In this paper, we have described the wireless overlay network concept as a way to combine the advantages of wide-area coverage while simultaneously achieving the best possible bandwidth and latency for mobile devices at any point in time. If we can demonstrate effective, low latency handover between overlays, we believe that whole new classes of compelling wireless applications will be enabled. Furthermore, we are developing support mechanisms that make it possible for applications to adapt to the changes in the quality of their network connectivity, either being changing data type representations or by moving the place where computations are performed. In our model, agents moderate these changes in network connectivity on behalf of applications.

Our architecture is being designed, deployed, and evaluated in a wireless overlay testbed being fashioned in the San Francisco Bay Area, with wireless technology and network services from AT&T, DEC, GTE Mobilenet, Hughes Aircraft Corporation, IBM, Metricom, and Pacific Telesis. We call this testbed BARWAN—Bay Area Research Wireless Access Network. Soda Hall, the home of the Computer Science Division at U. C. Berkeley, is being equipped with a building-wide wireless local area network based on AT&T WaveLAN technology. In addition, Metricom has deployed their wireless packet relay technology on our campus (their first toehold in the East Bay), and they have generously provided our research group with an independent Metricom network for experimentation. We also have access to CDPD service and the Hughes DBS system. We have implemented mechanisms for low latency handoff [18], reliable wireless transport [2], real time video transcoding for wireless transmission [19], and a prototype image proxy for wireless web browsing. A rudimentary form of user-initiated vertical handoff between multiple wireless local area networks and the Metricom packet relay network is now operational. Over the next several months, we plan to complete the detailed design of the wireless overlay and application support architecture described herein. A prototype system is expected for Summer 1996.

We plan to make this prototype network and associated technologies available as widely as possible to other members of the research community, to demonstrate their evolving mobile and wireless technologies. If successful, it will provide an ideal environment for experimenting with technology interoperability, especially between operational and experimental networks.

9. REFERENCES

[1] A. S. Acampora, M. Naghshineh, “An Architecture and Methodology for Mobile-Executed Handoff in Cellular ATM Networks,” *IEEE J. on Selected Areas in Communications*, V 12, N 8, (October 1994), pp. 1365-1375.

[2] E. Amir, S. McCanne, H. Zhang, “An Application Level Video Gateway,” ACM Multimedia ‘95, S.F., CA, (October 1995).

[3] H. Balakrishnan, S. Seshan, E. Amir, R. H. Katz, “Improving TCP/IP Performance over Wireless Networks,” ACM Mobile Computing and Networking Conference, Oakland, CA, (November 1995).

[4] E. A. Brewer, “Portable High-Performance Supercomputing: High-Level Platform-Specific Optimization,” MIT Ph.D. Dissertation, September 1994.

[5] J. S. DaSilva, B. E. Fernandes, “The European Research Program for Advanced Mobile Systems,” *IEEE Personal Communications Magazine*, V 2, N 1, (February 1995), pp. 14-19.

[6] D. Ferrari, “Client Requirements for Real-Time Communication Services,” *IEEE Communications Magazine*, V 28, N 11, (November 1990), pp. 65-72.

[7] S. Floyd, V. Jacobson, “Link Sharing and Resource Management Models for Packet Networks,” IEEE/ACM Trans. on Networking, V 3, N 4, (August 1995), pp. 365-386.

[8] R. Ghai, S. Singh, “An Architecture and Communications Protocol for Picocellular Networks,” *IEEE Personal Communications Magazine*, Third Quarter 1994, pp. 36-46.

[9] J. Ioannidis, D. Duchamp, G. Q. Maguire, Jr., “IP-based Protocols for Mobile Internetworking,” Proceedings 1991 ACM SIGCOMM Conference, pp. 235-245.

[10] J. Jubin, J. D. Tornow, “The DARPA Packet Radio Network Protocols,” *Proc. IEEE*, V 75, N 1, (January 1987), pp. 21-32.

[11] J. M. Kahn, *et al.*, “Non-Directed Infrared Links for High-Capacity Wireless LANs,” *IEEE Personal Communications*, V 1, N 2, (Second Quarter 1994), pp. 12-25.

[12] R. E. Kahn, “The Organization of Computer Resources into a Packet Radio Network,” *IEEE Transactions on Communications*, V COM-25, N 1, (Jan 1977), pp. 169-178.

[13] R. E. Kahn, *et al.*, “Advances in Packet Radio Technology,” *Proceeding IEEE*, V 66, N 11, (Nov 1978), pp. 1468-1496.

[14] B. M. Leiner, D. L. Nielson, F. A. Tobagi, “Issues in Packet Radio Network Design,” *Proc. IEEE*, V 75, N 1, (January 1987), pp. 6 - 20.

[15] J. Mitola, “The Software Radio Architecture,” *IEEE Communications Magazine*, V. 33, N. 5, (May 1995), pp. 26-38.

[16] A. Myles, D. Skellern, “Comparing Four IP-based Mobile Host Protocols,” *Computer Networks and ISDN Systems*, V 26, (November 1993), pp. 349-355.

[17] K. Pahlavan, A. Levesque, “Wireless Data Communications,” *Proceedings IEEE*, V 82, N 9, (September 1994), pp. 1398-1430.

[18] C. Perkins, “Providing Continuous Network Access to Mobile Hosts using TCP/IP,” *Computer Networks and ISDN Systems*, V 26, (November 1993), pp. 357-369.

[19] S. Seshan, “Providing Network Services to Mobile Hosts,” Ph.D. Dissertation, University of California, Berkeley, CS Division Report, (December 1995).