A Cooperative Approach to User Mobility

Robin Kravets
Department of Computer Science
University of Illinois, Urbana-Champaign
rhk@uiuc.edu

Casey Carter
Department of Computer Science
University of Illinois, Urbana-Champaign
ccarter@uiuc.edu

Luiz Magalhães
Department of Computer Science
University of Illinois, Urbana-Champaign
magalhae@uiuc.edu

ABSTRACT

We propose a networking model that treats a user's set of personal devices as a MObile grouPEd Device, a MOPED, which appears as a single entity to the rest of the Internet. All communication for a user is directed to this point of presence. As the user moves through different environments, the devices cooperate to provide the user with access to all available communication resources. We present the basic networking functionality necessary to enable the operation of MOPEDs and their integration into the Internet. We introduce a middleware layer to extend IP routing to work with MOPEDs and a lightweight IP encapsulation protocol, Multipath Routing enCAPsulation (MRCAP), used to implement that middleware.

1. INTRODUCTION

Trends in mobile communications have resulted in two significant developments. First, advances in processor technology, both in increased processing power and decreased energy consumption, have led to the creation of a new breed of small intelligent devices, including palmtops, PDAs, smart phones, and other wearable devices. Many, if not all, of these devices have some form of wireless communication. As users collect multiple small computing devices, the amount of communication resources available to the user increases, as does the demand for coordination of these resources. Second, increasing demand for wireless connectivity has produced rapid deployment of many new wireless technologies, with overlapping coverage in some areas. Considering the set of devices supporting a user, at any point in time, some subset of these devices may have connectivity. The convergence of these two developments presents a new challenge to provide coordination of a user's devices to provide better connectivity, and potentially more communication resources, to the user.

The efficiency of a user's personal devices is limited by their isolation from each other. When the resources of a device are completely consumed (e.g., a dead mobile phone battery), the user is cut off from key services. Similarly, if a user leaves the coverage area of a device, the services accessed via that device become unreachable. As a user moves through different environments, the cooperation of devices brings the potential for increased bandwidth and better connectivity by exposing to all devices the aggregation of services available to each individual device. Current technology and communication support provide connectivity between devices, but do not enable cooperation. The goal of our research is to bridge this gap from communication to cooperation.

Our model for mobile communication is based on two basic assumptions. First, a user should be able to create a representative presence on the Internet. All communication to a user is directed through this presence. A user may even create multiple presences (e.g., business, personal). If the presence has a unique network name, a single IP address, the user is in essence built into the network infrastructure. All communication destined for that presence is addressed to a unique identifier. Second, correspondent hosts need not, and in fact should not, be aware of how the user realizes a presence. The mapping of this identifier to an actual end host is dependent on the devices and infrastructure used to support the user. A solution should provide flexibility in the coordination of the user's device or devices, while maintaining transparency to correspondent hosts.

Cooperation of the collection of devices to support a mobile individual demands the extension of the mobility paradigm from an individual device to a network of devices with multiple points of connectivity, a MObile grouPEd Device (MOPED). All communication traffic for a MOPED user is delivered to the MOPED, where the final disposition of traffic is determined. Since a MOPED is designed to support a single user, communication with any of the devices in the MOPED is considered successful communication with the user. This model enables a group of devices to be mapped into a user's point of presence on the Internet. To the outside world, this MOPED appears as a single device with a single interface and address. In reality, the group of devices cooperates to provide better services to the user.

The goal of the MOPED project is to provide service to a user through the cooperation of the MOPED devices that is better than the service provided by the devices working individually. Our solution provides three key benefits. First, a user can be connected via any of the services currently available to the individual devices. Second, if multiple devices have connectivity in a certain environment, the MOPED can take advantage of the additional bandwidth by routing through multiple connections. Finally, such connectivity enables smooth handoffs as individual devices gain and lose connectivity, maintaining external connectivity to all devices when at least one device in the component has external connectivity.

In addition to improved service, the design of the MOPED architecture provides three additional benefits that ease the integration and deployment of MOPEDs. First, our design supports the commonly accepted idea that non-mobile users should not have to be aware of the extra infrastructure needed to support mobile users. Our architecture supports communication with non-mobile-aware users as well as optimizations for mobile-aware users. This abstraction also provides the benefit of hiding the topology of the MOPED from external hosts, providing flexibility and anonymity. Second, any

new device acquired by a user can be integrated into the MOPED if it can become part of the PAN connecting the MOPED. This covers the easy inclusion of new technology as well as legacy devices. The level of cooperation of each individual device in the MOPED depends on whether or not the device is MOPED-enabled. Finally, the sharing of communication resources across devices allows for separation of concerns in the design of personal technology devices – a smart watch need not also be a phone.

The remainder of this paper presents an in-depth discussion of the design and implementation of the MOPED architecture. In Section 2, we present the motivation for MOPEDs, and discuss constraints on their design and in Section 3, we present the challenges involved in integrating MOPEDs into the routing structure of the Internet. In Section 4, we present our design in the context of related research in the area of mobile computing. Section 5 describes our solution, the MOPED Routing Architecture (MRA), including the lightweight IP encapsulation protocol, MRCAP. Finally, Section 7 presents our conclusions and directions for future research.

2. MOPEDS

Consider a group of devices connected in a personal area network (PAN) via a wireless technology such as Bluetooth, Infrared or wireless Ethernet. If any one of the devices is within its service area, cooperation between the devices can provide connectivity to all of the devices. In this context, a user's devices make up a mobile network, which may have multiple means of connectivity to the Internet at any point in time. In reality, we do not expect all of a user's devices to always be connected in one PAN. If a user leaves their laptop on a desk and walks away with their phone, the short-range connectivity between the two devices will disappear.

In order to provide support for such scenarios, we place no constraints on the topology of a MOPED. A MOPED may be composed of many devices, only some of which can communicate directly with each other and some of which have direct Internet connectivity. The MOPED architecture enables localized cooperation of devices through the concept of a MOPED *component*, a subset of the user's devices connected via a Personal Area Network (PAN), enabling the desired sharing of resources among those devices. As the user moves through different environments, the devices cooperate as a distributed virtual device. As part of a MOPED component, a device's resources are added to the pool of resources available to the component in that environment. Since the goal of the MOPED is to support a single user, management of the MOPED's resources can be solved based on the needs and preferences of the user. In contrast to traditional networks, a MOPED component can be considered as an ad-hoc network that represents a distributed virtual device. This model supports external connectivity to all devices when any one device has external connectivity.

Consider a MOPED component comprised of a PDA with a cellular modem, a mobile phone, and a laptop with no connectivity, all connected in a PAN (see Figure 1). While the user is talking on the phone, other connections to the laptop can be routed through the PDA. To support this, the proxy can route flows for different endpoints to the appropriate external interface. If the user is participating in a videoconference on their laptop, the audio could be routed through the phone, while the video is routed through the cellular modem, providing more bandwidth to the application than if only one of the interfaces is used. Traditional IP routing support does not allow the separate flows to the same endpoint to be routed along different paths. Finally, consider a single application whose

Figure 1: A simple MOPED

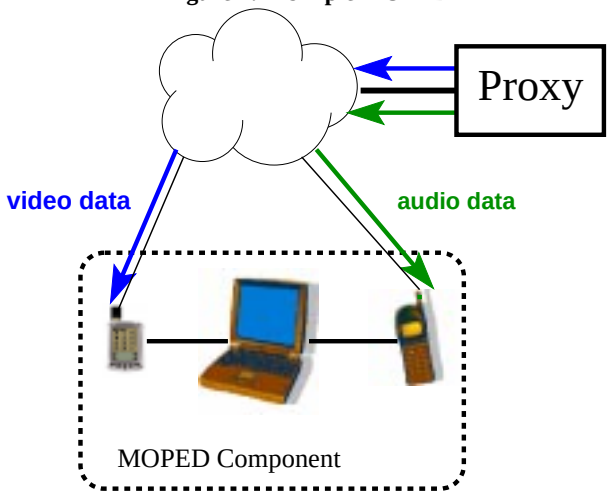

communication requirements for a single flow are more than the bandwidth provided by either of the external connections. In this case, specialized transport protocols can inverse multiplex a single flow's data across multiple paths [8]. In order to support these last two scenarios, the MRA provides flow-based and packet-based routing in order to realize the potential bandwidth improvements.

Figure 2 depicts an example MOPED, demonstrating the architectural components of a MOPED, including six devices configured into two MOPED components, and their associated interfaces. A device that can communicate with the Internet directly is termed a *perimeter node*, while a device that can communicate with the Internet only indirectly through other devices is an *internal node*. A set of devices that can reach each other using paths that pass only through internal interfaces forms a MOPED *component* (a connected component in a graph representation of a MOPED). It is possible, and even expected, for a MOPED to be partitioned into multiple components, and still continue to function normally, provided that each component has at least one external interface through which it can communicate with the other components. We expect partitioning to occur frequently in a MOPED, such as when the user carries some of the communicating devices out of their limited range. We also expect devices to enter and leave the MOPED with reasonable frequency.

The components of a MOPED can be considered as nodes in a star-like overlay network, where the Internet point of presence for the MOPED (accessed via the user's unique identifier) is at the center of the star. This point of presence is supported via a proxy that maps the single IP address onto the several nodes of the MOPED. Due to the potential of multiple external connections for each MOPED component, there may be multiple paths between each component and the proxy. This multiplicity of connections goes beyond simple connectivity and provides the possibility for increasing the resources available to the component, as well as to an individual device. As we mention above, only one external connection is necessary to support connectivity to a component.

Due to the potential of the partitioning of a MOPED into multiple components, MOPED communication can be represented by three distinct types of traffic, each of which requires routing support:

Figure 2: A MOPED with two components and six devices

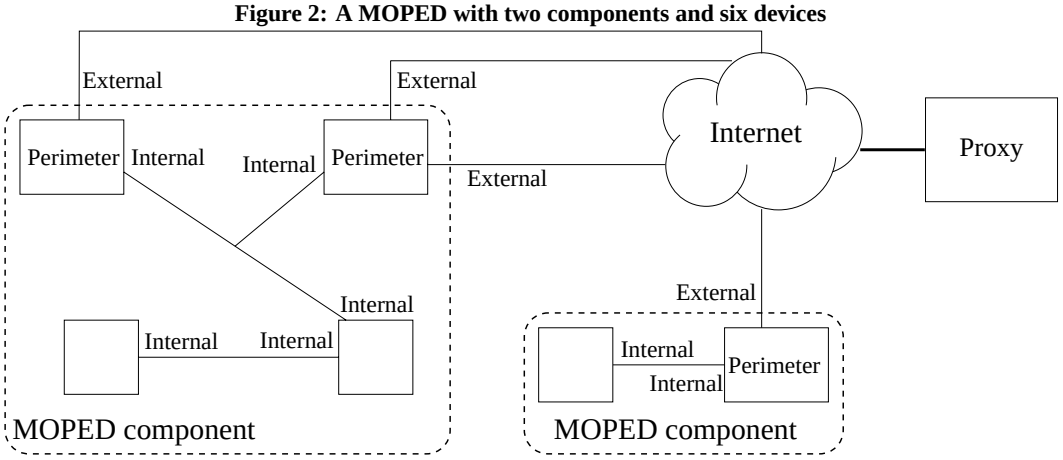

- Intra-component: between two devices in the same MOPED component.
- Inter-component: between two devices in different components of the same MOPED.
- Extra-MOPED: between a MOPED device and a host outside of the MOPED.

In this Section 5, we present the network model for MOPEDs, which we call the MOPED Routing Architecture (MRA), and describe the techniques for supporting all three types of traffic.

3. MOPED DESIGN RATIONALE

One consequence of the proliferation of personal communication devices is the complexity of individually addressable communication endpoints. Increasing the possible number of ways to communicate with a user makes it significantly more challenging to determine which method is the "best" in a given situation. The introduction of MOPEDs changes the paradigm for mobile communications, defining communication with any of the devices in the MOPED to be equivalent to communication with the user. Many current solutions are focused on a single host with one network interface [1, 12, 15], while other research seeks the ability to address and locate a person and the device they are currently using [10]. We believe that the appropriate next step is mobility management for a MOPED, the network of devices that is associated with one person.

In order to better support the user's needs, a MOPED may interact with networks and services in the surrounding environment to determine local connectivity and communication service availability. The design of a network routing architecture for MOPEDs must enable the exploitation of knowledge about the devices of the MOPED, the MOPED network topology, and available communication and routing.

3.1 Cooperating Devices and the Internet

Traditional networking and system models have several shortcomings when viewed in this cooperative context. Nodes on the Internet are identified by IP addresses, which statically specify where a packet should be sent to reach the identified node. When users carry devices with them, the location-specific nature of IP addresses becomes a significant burden. The realization of a user's presence can involve one or more devices. If a user has a single device with a single interface (e.g., a mobile phone, laptop), the device, and so the user, can be supported via existing techniques such as cellular telephony or MobileIP [15]. If the device has multiple interfaces, each interface can be used as available [24] or simultaneously [29]. The next logical step is to support a user via multiple devices.

We present a coherent network model for MOPEDs, enabling them to participate fully in the Internet. We adhere to the common philosophy that any modifications to support mobility should be localized to the mobile hosts themselves, and possibly some support systems associated with the particular MOPED. We do not require the replacement of Internet routers or any alterations to non-MOPED hosts. In fact, our solution preserves end-to-end semantics and is transparent to endhosts. Network applications running on a MOPED device require no modification, and Internet hosts communicating with a MOPED require no knowledge of the MOPED's structure and mobility.

The particular contribution of this work is a framework for integrating a MOPED into the Internet; we define an architecture through which data packets can be routed between the various devices in a MOPED and correspondent hosts at large on the Internet, which may not be MOPED-aware or even mobility-aware. The MRA handles the basic connectivity problem for a MOPED: directing traffic to and from the set of mobile devices.

3.2 MOPED Requirements

The goal of our research is to support a large range of devices in the architecture of a MOPED—from personal computers to mobile telephones to smart cards. In order to support such diversity, MOPED capability must place minimal requirements on the processing power, storage space, and bandwidth available to any given device. We expect that most personal technology devices will eventually include efficient short-range wireless communication interfaces (e.g., low-power 802.11, Bluetooth) for communication between the devices. In addition, we expect that many devices will have additional wireless connectivity to the Internet, such as wireless Ethernet or a cellular modem. In order to provide portability, devices in a MOPED can use any channel of communication that

can carry IP traffic[1]. We expect to support lighter-weight communication than IP in the future.

The design of Internet support for MOPEDs is modeled after MobileIP [15]. A MOPED has a single official IP address by which it may be reached, and must have a supporting proxy to direct traffic to the MOPED from its home network. We examine the role of MobileIP in a MOPED in greater depth in Section 4.2.

3.3 MOPED Routing Challenges

There are three key aspects of routing between a MOPED and a correspondent host that prohibit the use of traditional IP routing: addressability, the necessity of addressing each internal device individually although the MOPED itself has only one public IP address; mobility, managing the mobility of the MOPED devices relative to the proxy and each other; and path selection, the ability to selectively utilize multiple paths between a MOPED device and its proxy to enhance throughput and reliability.

3.3.1 Mobility

There are two types of mobility we consider in a MOPED: mobility of devices with respect to the proxy (i.e., mobility through the Internet proper) and mobility of MOPED devices within the same component. For mobility relative to the proxy, we take advantage of the existing machinery of MobileIP.

A static Home Address for use with MobileIP is assigned to each interface of a MOPED device. Devices with external connectivity then use MobileIP to establish a channel to the proxy, effectively forming a link between the device's MOPED component and the core of the overlay network.

Interestingly, the use of MobileIP is not exposed to upper layers in the architecture; it simply provides a "tunnel" into which a MOPED perimeter node or proxy can send packets, expecting them to arrive at the other (mobile) endpoint. In the future, this insulation from MobileIP allows it to be easily replaced with another mobility management protocol such as MobileIPv6 without affecting higher layers in the architecture.

Although it may seem extravagant to have a distinct globally-valid IP address for each external interface, our goal is to support maximum flexibility of MOPED connectivity within the existing MobileIP infrastructure. One might imagine it possible for every external MOPED interface to use the public MOPED IP address as its MobileIP Home Address. This approach would require extensive modification to the MobileIP home agent to somehow differentiate these multiple registrations, but would still interoperate with existing MobileIP foreign agents. The catch, however, is that these unmodified foreign agents would only be able to support a single registration for an entire MOPED. Since we can envision a MOPED with multiple interfaces using the same link-layer technology possibly registering with the same foreign agent, our solution retains the flexibility of assigning unique home addresses to each external interface. We also note that this approach is no less expensive than a non-MOPED approach, which would also assign each interface a unique home address.

The second variety of MOPED mobility, mobility of devices within a component, has no such simple solution. A single MOPED component is an ad-hoc network. An ad-hoc link-state routing protocol is used to maintain (at each MOPED node) a topology graph of its component, including which external interfaces are active. The proxy need only know which MOPED devices are in which component, without having full internal knowledge of the components' topologies. This reachability information is relayed from the components to the proxy as part of the normal routing information dissemination. In the MOPED overlay network, the proxy effectively switches traffic between the components – when a node sends a packet to a correspondent host, or another node *not in its component*, it simply sends the packet to the proxy.

The ability to maintain mobility for communications with non-MOPED-aware Internet hosts is crucial to the success of our technology. Therefore, it is undesirable to enforce any requirements on the correspondent hosts to enable communication with a MOPED. Consequently, our approach to MOPED mobility more closely mirrors the proxy method of MobileIP [15] than the end-to-end approaches of other work [20]. Our future plans will integrate an end-to-end mobility method as an optimization for correspondent hosts that support it, but still retain the proxy for completeness. The proxy provides a fixed location via which a mobile host can be contacted and enables communication to continue even when both end-hosts move simultaneously.

3.3.2 Addressability

The MRA uses IP-based communication between devices in the MOPED, so that users can manage the MOPED with familiar applications. This makes it necessary for the MOPED to maintain a mapping between the single public IP address and the (possibly many) MOPED-internal IP addresses. The proxy must be able to use this mapping to deliver incoming packets to the proper end-device within the MOPED.

To realize the goal of MOPED-node addressability, the MRA assigns to each node a static, private IP address in the MOPED overlay network. Since hosts external to the MOPED will never use, or in fact be aware of, these addresses, these addresses are "private" in the sense that they need not have significance (or even be unique) outside of the MOPED proper.

3.3.3 Path Selection

The final task necessary for the MRA is path selection. We believe that the bottleneck for communication resources in MOPEDs will be the hop from the perimeter nodes to the infrastructure, and not in the MOPED component-internal links, or the Internet itself. Long-haul wireless technologies have lower bandwidths than the short-range wireless technologies used to link MOPED components together. Lower bandwidth technologies usually have larger coverage areas than high bandwidth technologies, so much of the time MOPED communication will be constrained by the available bandwidth through the perimeter of the components.

In order to provide cooperative resource utilization, in terms of bandwidth, power, or cost, the MRA supports flow- and packet-level load-balancing across the (possibly several) interfaces on the perimeter of a MOPED component. In Section 5.2, we present the Multipath Layer, which enables sources to choose specific exit interfaces on a MOPED-component. The MOPED's utilization of multiple communication channels is at a higher level than traditional cellular handoffs. Via the path selection mechanism, a MOPED uses multiple channels simultaneously to carry different traffic. This is distinctly unlike the simple failover mechanism of

[1] The authors have a special predisposition toward "Avian Carriers" [26].

vertical handoffs [24] or the MobileIP error-robustness technique of simultaneous mobility bindings [15].

3.4 Multiple Technologies, Multiple Interfaces

Coverage areas for different wireless communication technologies (and wired, for that matter) vary greatly, with some areas of overlap between multiple technologies. In order to maximize user connectivity, MOPED devices that are in their coverage areas must forward traffic for their peer devices that otherwise lack the ability to communicate. Further, a user with high bandwidth requirements should be able to utilize all of the bandwidth available to multiple MOPED devices, when several devices are all within their coverage areas. The MOPED therefore requires a routing mechanism capable of directing traffic flow through several simultaneously connected interfaces, and maintaining routes for this traffic in the face of user mobility or network failures.

As a basic mechanism without changes to the transport layer, the MRA supports bandwidth aggregation with flow-level granularity (i.e., all packets in a particular flow (identified by IP protocol, and source & destination IP addresses and transport-layer port numbers) should follow the same path. A finer level of granularity, directing packets from the same flow to follow different paths, could easily cause packet reordering, which traditional transport layers such as TCP may interpret as loss [18]. The splitting of packets from the same flow also makes it challenging for transport layer protocols to effectively collect channel quality statistics, such as TCP's estimates of round trip times, or path MTU discovery [11].

One of the goals of the MOPED project (although not of this paper) is to devise a family of transport protocols that use inverse multiplexing techniques to load-balance a single flow through multiple hardware interfaces in a wireless mobile host [7, 9, 8]. The key to these protocols is the measurement of end-to-end throughput available at each hardware interface that can be used for communication. This is achieved through the use of a rate-based transmission mechanism and the measurement of interarrival times at the receiver. Loss discrimination is used to compensate for the greater loss rate found in wireless communication. Much of this work has been completed in a non-MOPED context, but has yet to be integrated into the MOPED architecture.

Current results of these protocols for single hosts with multiple network interfaces promise that bandwidth aggregation is an achievable goal for MOPEDs. Experimental results show good applicability for both multimedia data, which requires timely delivery and for which the added bandwidth gained by the simultaneous use of multiple interfaces translates directly into greater data quality, and for bulk data, where the coupling of loss discrimination and higher bandwidth results in much higher throughput than TCP.

The same inverse multiplexing techniques can be extended from a single host to the multiple external interfaces of a MOPED. The same restrictions apply, namely that throughput only improves if the bottleneck link between the MOPED and the endhost is given by the last hop to the MOPED. The added resiliency of the aggregated link may justify the simultaneous use of multiple external interfaces even if bandwidth gains cannot be guaranteed.

4. RELATED WORK

The design of the MRA draws from several areas of research in mobile computing. In this section, we discuss the design in the context of such research in network technology, routing and user location management.

4.1 Infrastructure

A MOPED provides an infrastructure for several personal technology devices to connect to each other and communicate with the Internet. Integrating a set of personal devices is certainly not an idea original to MOPEDs. Technologies for personal area networks (PANs) such as Bluetooth and IEEE 802.11 have come into vogue in the networking research community, but they are simply mechanisms for physical connectivity among a set of devices—they do not address the question of *what* we should do with the PANs or *how* these tasks can be best accomplished. We see MOPEDs as a network-layer (or slightly above) entity that is complementary to the datalink-layer concept of a PAN. Although a PAN is useful, it is not necessary to our design; indeed, we enable separate components of connected devices to participate in the same MOPED using external channels.

4.2 Routing

Many projects address issues involved with mobility of and routing to groups of devices. A MOPED is a composite of many devices with many network interfaces: a mobile network with multiple points of attachment to the Internet. A single MOPED device might itself have multiple external interfaces, so our architecture is a mobility solution for one or more devices with zero or more Internet-mobile network interfaces each.

MobileIP handily solves the problem of mobility for a single device, but does not solve the problem of MOPED mobility without extension. The goal of MobileIP is to make it appear that a mobile host is not mobile, but is in fact at "home." The mobile node (MN) has a permanent home address at which other Internet hosts try to reach it. Some host on the MN's home network acts as a supporting home agent. When the MN wanders into a foreign network, it obtains an IP address on that network, a care-of-address. The MN registers this care-of-address with its home agent, which intercepts traffic sent to the MN's home address and redirects it to the care-of-address.

Unfortunately, the multiple-interface, single IP address nature of MOPED mobility does not align well with MobileIP. MOPEDs must have a way to multiplex traffic destined for many devices onto a single IP address. The impedance-mismatch of MobileIP to MOPEDs is exacerbated by the fact that MOPEDs may contain devices that have no direct Internet connection, and thus cannot participate in MobileIP. Clearly, a different solution is necessary to support mobility of MOPEDs. There has been some work in the MobileIP community to address the mobility of a network of hosts with a *single* point of attachment, a "Mobile Router" [2, 15]. The work on mobile routers does not address the MOPED goals of user addressability or resource aggregation, or multiple points of attachment for a mobile network.

The MRA provides a mechanism to access exactly one of a set of several Internet hosts using a single IP address. In this way the MRA resembles anycasting [13]. The MRA, however, provides a much more structured environment for distributing traffic to specific devices in that set, additionally providing for mobility and resource aggregation. Explicit control over access paths into the MOPED makes it possible for the MRA to provide better resource utilization than that possible via blind anycast.

Our use of multiple interfaces and multiple paths for data transmission is greatly influenced by [29]. The prototype implementation of the Multipath Layer (described in Section 5.2) uses their mechanism for binding sockets to particular interfaces. We generalize their work to a multiple-device, network environment.

4.3 User Location

One of the main goals of MOPEDs is to bind communication mechanisms together and create a single point of presence for a user. Our work approaches the user location problem by defining a single Internet address (the MOPED address) to which all data for a user should be directed, replacing the user location problem with a more traditional network location problem. This conversion of person-location to network-location enables interoperability with unmodified, legacy network applications.

The user-location-management aspect of MOPEDs is similar to the goal of Stanford's Mobile People Project [10]. Mobile People is an architecture for allowing application-level mobility: it provides a name service to map from user names to application-specific addresses at which that user can be reached, a process Mobile People calls "person-level routing". Mobile People does not address the grouping of several devices into a single logical entity, and certainly does not support aggregation of device resources in a cooperative fashion. It provides an intermediary between communicating parties where they may record their current application-specific address and learn the addresses of others. Both Mobile People and the MRA provide location privacy; they make it possible to communicate with a given person through a proxy, hiding that person's actual location.

Our work is complementary to ICEBERG [27]: a comprehensive framework for communication and service adaptation, transforming communication datatypes to suit different devices. A MOPED would be an interesting basis for a communication network atop which to implement ICEBERG. ICEBERG does not address the issues of cooperative resource aggregation and network-layer connectivity upon which the MRA is focused. Although the MRA is an enabling technology for the goals to which ICEBERG aspires—namely, the idea of using a person as a communication endpoint—the two are in fact complementary, as they provide services at differing layers of the network hierarchy.

Hewlett Packard Lab's CoolTown project integrates people, places, and things into the web by augmenting each with a "web presence" [3]. CoolTown is an application layer solution for user-location, and therefore requires application modification.

4.4 Hierarchical Mobility and Paging

Many researchers have proposed hierarchical mobility solutions [19, 21] in which a node may be mobile within a local domain without updating its home agent, only updating its location with a local mobility agent. This hierarchy of mobility agents reduces the signaling load across the Internet between mobile nodes and home agents, and also ensures that the mobility agents' notion of the mobile node's location is, on average, more up-to-date.

The MRA is similar to such hierarchical mobility solutions in that the proxy is not informed when a MOPED node changes location within its component, but only when it changes components. A MOPED component performs as a hierarchical mobility domain with multiple points of Internet attachment.

Recent work in mobility has extended the notion of hierarchical mobility by adding paging: an idle mobile host only keeps the home agent vaguely informed of its location [25, 28]. A network is divided into a number of paging domains, and a mobile node (when idle) only informs its home agent of movement across paging domains. Since the home agent does not necessarily know the exact attachment point of the mobile node in the network but only its paging domain, there must be infrastructure in place to "page" the mobile node within the paging domain. Upon receiving the page, the mobile node registers its exact location.

The major potential benefit of a paging scheme is extension of battery life in mobile devices. If a device's link layer has support for a low-energy-consumption standby state in which the device can detect pages, then idle devices can consume less power than devices that actively update their mobility registrations. To keep the MRA free of link-specific concerns, we choose not to address issues relating to the paging of idle nodes in a MOPED component, although such functionality can be added to the architecture in the future. Paging of external interfaces by Internet infrastructure is essentially a feature of the mobility agent for that interface, and the presence or lack of support for paging on such an interface is orthogonal to the design of the MRA.

5. MOPED ROUTING ARCHITECTURE

The goal of the MRA is to provide a mapping from the single point of contact of the user, the MOPED IP address, to the destination node in the MOPED. First, the MRA provides addressing capabilities for each of the individual nodes, as well as each of the individual nodes' interfaces. Addressability is provided through the use of Network Address Translation (NAT), which is an approach commonly used in conjunction with connection tracking to compress address space. Second, the MRA determines and sets up an appropriate route to the destination node. Path Selection is supported in the Multipath Layer, which tracks connectivity and topology in order to make appropriate routing decisions. The Multipath Layer's sole responsibility is to maintain a partial graph of the MOPED and use its tracking information along with possible application input to choose external interfaces by which packets enter or leave the MOPED. Finally, the MRA handles the mobility of the destination. Although it seems contradictory to our earlier claim, mobility of individual nodes is supported through the use of MobileIP. In this section, we present our solutions for each of these functionalities and discuss our solution in the context of a concrete example.

5.1 Addressability

Network Address Translation (NAT) in the MRA solves the problem of addressing specific nodes and interfaces in a MOPED. NAT has traditionally been used as a solution to the problem of address space pressure in IPv4 [22]. Although NAT has accrued a certain negative connotation for its violation of the end-to-end argument [6], there are situations for which NAT is the ideal solution. When the exact topology of an internetwork should be kept hidden from the Internet (e.g., a firewalled corporate network), NAT is the ideal tool for isolation. Our goals of user privacy and isolating correspondent hosts from MOPED topology mesh perfectly with NAT.

It is common for a "secure" IP network to be assigned non-routable, private addresses and hidden behind a firewall. Such networks use NAT in conjunction with port translation to achieve what is commonly called "IP Masquerading:" enabling the hidden hosts to communicate with the Internet, while avoiding the problem of address space pressure by multiplexing the entire network of hosts

onto the single public IP address of the firewall [23]. We incorporate this approach in the MRA, using NAT and connection tracking to multiplex the MOPED onto a single public IP address. Address space pressure is not significant to the design of the MRA, but the ability to access an entire network of devices through a single address is a critical goal. NAT enables localized mapping from the public MOPED address to internal MOPED addresses in the proxy, reducing the complexity and state in the actual MOPED nodes.

To implement NAT in a MOPED, a unique node identifier (i.e., a private "internal" IP address) is assigned to each MOPED device. MOPED nodes use these addresses to communicate with each other internally; they are not visible outside the confines of the MOPED and its proxy. The NAT layer maintains a mapping from correspondent host IP and port number to internal IP.

For outgoing MOPED traffic, NAT recognizes packets that originate a flow and records a binding for that flow. The source address in the packet is then mangled so that the packet appears to come from the public MOPED address when it arrives at the destination. Any reply packets, or further packets sent from the MOPED in this flow, match the established binding so that NAT can determine to where they should be sent.

So that the MRA can handle incoming service connections to the MOPED from correspondent hosts, future work will develop protocols to control selective port forwarding at the NAT layer, enabling application programs to receive traffic destined for specific TCP/UDP ports on the public MOPED IP address. In the prototype implementation, any services exported from the MOPED require static port forwarding.

5.2 Path Selection

The Multipath Layer determines how data traffic is routed from the proxy to the internal devices addressed by this private space (and vice versa). The Multipath Layer maintains partial topology information for the MOPED, so that it can determine which MOPED devices compose each component and what external interfaces provide access to each component. This topology is used to determine paths for packets sent into the MOPED.

The choice of multipath routing algorithm is crucial to the proper operation of the Multipath Layer. This is an open problem, and one we will address in future work. The MRA enables the study of multipath policy algorithms by providing an infrastructure that enables path specification and facilitates communication between peer multipath policy agents on different devices. The current implementation of the Multipath Layer binds all packets of a particular flow—identified by a tuple (local IP, local port, correspondent IP, correspondent port, IP protocol)—to follow the same path. This binding is dynamic: when an alternative path with more suitable characteristics is discovered, the binding is easily altered.

The Multipath Layer is the key active entity in extra-MOPED traffic. It piggy-backs path information on the data packets, to be used by Multipath Layers on other devices in determining how to handle other packets from the same flow. This need to attach arbitrary data to packets encouraged the development of the light weight, extensible IP encapsulation protocol, Multipath Routing enCAPsulation (MRCAP), described in Section 5.5.2.

The Multipath Layer requires a mechanism to intelligently schedule packets through particular perimeter interfaces on their route to the proxy. At first, source routing seems the obvious solution to this problem; a MOPED node could specify the IP address of the chosen external interface in a Loose Source Route (LSR) IP option [17]. This ensures that the packet is routed through the correct perimeter node, but does not enforce routing through the desired external interface if that perimeter node has several active external interfaces. Even worse, if two interfaces on the same perimeter node are bound to the same MobileIP foreign agent, then the next hop on *both* interfaces is the foreign agent's IP address – there is no way for an LSR to distinguish the two paths. Clearly, some other mechanism that allows sources to choose a perimeter interface is needed. MRCAP provides exactly such a mechanism for the Multipath Layer to utilize.

5.3 Component Management

The dynamic nature of the MOPED's composition is a challenge to traditional IP routing that the MRA must address. In order to satisfy the cooperative goal of MOPEDs, the MRA supports the ability to 1) route traffic with differing QoS requirements, 2) balance load across external interfaces and 3) conserve transmission power so as to maximize MOPED lifetime. A MOPED serving the needs of a single user is in a unique position to perform global routing optimizations.

Supporting several routing types simultaneously requires the dissemination of multiple metrics. Link state protocols distribute complete global information about the network, allowing any MOPED node to make optimal decisions about the network state within its component. The high overheads often associated with proactive protocols, especially link state protocols, are not a problem in the MOPED environment due to the expected size of a MOPED component (on the order of 4-6 devices communicating on two internal links with 2-3 external links). To enable the global dissemination of extensive link, interface, and node state throughout the MOPED, an ad hoc hierarchical proactive link state routing protocol is the most viable choice.

The lowest level of the hierarchy is a single MOPED component. Each node in a component disseminates the state of its interfaces, and its set of neighbors on each interface, by flooding a link state packet throughout the component. Accumulation of link state packets enables every node to find routes to every other node in the same component. MOPED components self-assemble in this way.

One node is elected as a component leader and is responsible for representing the component in the second level of the hierarchical link state protocol. The leader synchronizes a portion of its link state database (LSDB) with the proxy. This abbreviated LSDB contains only state information for the external interfaces to the component and a list of the MOPED nodes that form the component. In this way, the proxy obtains link state information for external interfaces (for making routing decisions for MOPED-destined traffic) and learns the mapping of nodes to components. This partial component LSDB is also synchronized with other components to enable direct inter-component routing. The MRA's inter-component routing effectively stitches all the components together into a single virtual network; at the inter-component level, the entire MOPED is a single routing domain.

5.4 Mobility

Although not sufficient for supporting mobility of MOPEDs, we adopted MobileIP into the MRA to support mobility of individual nodes. Mobility of IP network interfaces in the Internet is a well-

studied problem. Instead of casting aside this body of work, we intend to leverage MobileIP as much as possible in handling MOPED mobility. Recall that mobility of a MOPED is unlike traditional MobileIP clients, in that a MOPED has many mobile interfaces to manage, and may be able to deal with mobility by routing traffic through another MOPED device.

In MobileIP, data from the Internet for the mobile node is delivered to its home agent, and then "tunneled" to the mobile node's care-of-address. MobileIP can also use reverse tunneling, in which all outbound traffic *from* the mobile node is tunneled to its home agent, and then sent on to the true destination. This is necessary to traverse firewalls that use reverse packet filtering—discarding packets that come from the "wrong" side of the firewall. The MRA always uses reverse tunneling, for that reason, as well as to ensure that the multipath Layer in the proxy has complete and timely information on MOPED topology.

5.5 Realizing the MRA

Now we assemble the pieces of the MRA, and give an operational description of its function. A short summary of the different address types used in the MRA illuminates a discussion of the forwarding path taken by packets as they traverse the MRA. We also describe the IP encapsulation protocol that enables the Multipath Layer to select paths, and briefly describe the status of our implementation of the MRA. Finally, a brief example of a web transaction over the MRA is presented to help the reader to comprehend the MRA as a whole.

5.5.1 Address Hierarchy

Recall that there are four distinct types of addresses used in the MRA; we summarize them here:

1. The MOPED IP address. This is the official, public IP address used to identify the MOPED, and therefore its owner.
2. Internal IP addresses. These addresses are used to identify particular MOPED devices; they are private in the sense that they have meaning only within the MOPED and its proxy.
3. Interface IP addresses. These are the MobileIP home addresses of the external interfaces on the MOPED devices. They are distinct from the Internal addresses, as there may be some devices that have only internal addresses.
4. Care-of-Addresses. These are the IP addresses to which the MOPED external interfaces are currently bound by MobileIP.

Each of these address types performs a distinct function in the overall operation of the MRA. Intuitively, when a packet arrives from a correspondent host addressed to the MOPED, the proxy determines:

1. To which MOPED device the packet is delivered—an internal IP address.
2. Through which external interface the packet is routed to reach the target device's component of the MOPED.
3. Exactly where in the Internet that external interface is.

Upon arrival at that external interface, the receiving MOPED perimeter device then:

1. Marks the packet as having passed through the external interface.
2. Uses the internal routing protocol to deliver the packet to the correct destination MOPED device.

Upon final packet delivery, the destination MOPED device may record the path taken by the packet, to ensure that return packets follow the same path.

Conversely, when a MOPED device needs to transmit a packet to some other host, it:

1. Determines if the target is in the source's component; if so, it uses the intra-component routing protocol to deliver it. Otherwise, the target is in some other component, or must be routed via the proxy.
2. Chooses an external interface from the device's component through which the packet is sent to the proxy (or peer component), and marks the packet appropriately.
3. Delivers the packet via the intra-component routing protocol to the device where that external interface is located.

Once the internal routing protocol delivers the packet to the desired perimeter MOPED device, that device simply transmits the packet through the chosen external interface to the proxy, or an external interface of the target component. At the proxy, the process used for delivery to the MOPED is reversed:

1. The proxy records the external interface through which the packet was directed out of the MOPED component, for use in later routing decisions.
2. The proxy modifies the source address in the packet, so that it appears to come from the official, public MOPED address.
3. Traditional IP routing delivers the packet to the target host.

A purpose-devised lightweight IP encapsulation protocol, Multipath Routing enCAPsulation (MRCAP), facilitates communication between peer Multipath Layers, and packet redirection.

5.5.2 Multipath Routing enCAPsulation

The design of the MRA, in conjunction with our goal of implementing the entire architecture in user space, necessitates the use of IP encapsulation. Since the Multipath Layer is responsible for routing pre-formed IP packets and may require the communication of some small amount of state to a peer Multipath Layer, we need a lightweight mechanism to:

- Encapsulate any IP packet.
- Dynamically alter the source and/or destination address.
- Track the packet's original source/destination addresses.
- Facilitate coordination between peer multipath policy agents.

Figure 3: MRCAP Packet Format

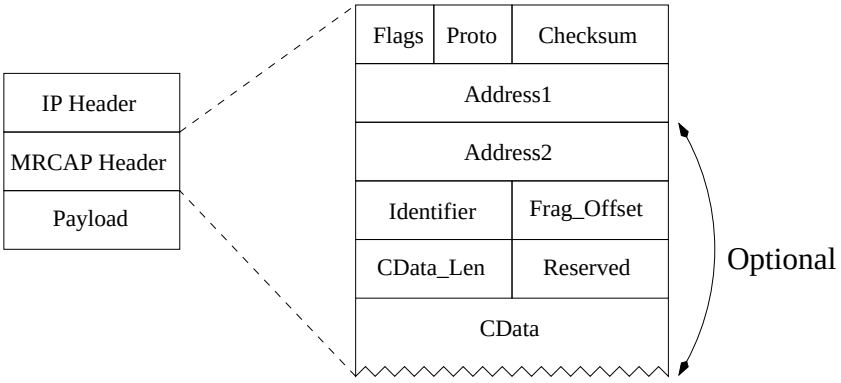

Since last-hop bandwidth is a primary bottleneck, per-packet overhead in excess of the costs of MobileIP must be minimized, a problem that is especially acute when MobileIP reverse-tunneling is employed. We evaluated existing encapsulation protocols, but were concerned with either their consumption of data space (IPIP [14], or GRE [5]), or lack of flexibility and general applicability. Minimal Encapsulation [16] is an optimization to IP in IP encapsulation: instead of adding an entire envelope IP header to the encapsulated packet, it stores a single extra IP address (the original destination address of the tunneled packet) and 4 bytes of accounting overhead. The very low overhead (8 bytes per packet) of Minimal Encapsulation is attractive, but sacrifices extensibility, and incurs an inability to encapsulate fragmented IP packets. Unwilling to sacrifice bandwidth, we developed a general, extensible encapsulation protocol tailored to the needs of MOPED routing.

Multipath Routing enCAPsulation, or MRCAP, has very low per-packet overhead, usually 8-12 bytes, comparable to Minimal Encapsulation. The MRCAP packet format (see Figure 3) includes a tiny fixed-length header inserted between the original IP header and the packet payload, as in Minimal Encapsulation. The presence of various extension headers is indicated by option flag bits. The fixed-length header occupies 8 bytes of payload space, while still retaining the flexibility to add optional extensions as necessary. (One of those extensions is a 4-byte Fragment header, so that MRCAP can encapsulate fragments.) All communication between Multipath Layers on the proxy or in the MOPED occurs in-band, in the control channel of MRCAP.

5.5.3 Implementation

The MRA has been partially implemented atop the Linux 2.4 kernel, running on a MOPED test bed of several laptops communicating over IEEE 802.11b wireless Ethernet. We use the netfilter NAT functionality built into Linux as the MRA's NAT layer, and Dynamics MobileIP to provide mobility. The modularity of the architecture enables the combination of these unmodified components with the Multipath Layer implementation in the Multipath Routing Daemon (MRD). The MRD is a user-space application that uses the Linux kernel's Universal Tun/Tap driver to intercept packets and inject packets directly into the kernel network stack. The MRD handles extra-MOPED (via the Multipath Layer) and intra-component (via the inner routing protocol) traffic only, and delegates the delivery of inter-component traffic to the proxy.

The MRD operates by carefully manipulating routing tables to direct externally destined packets into the tunnel device, so that they may be captured, MRCAP encapsulated, and have their path to the proxy selected appropriately. As with any packet tunneling implementation, great care is taken to ensure that packets are not multiply encapsulated, and that ICMP error messages are forwarded to the correct originators, as for IPIP encapsulation [14].

At perimeter nodes, the Linux `SO_BINDTODEVICE` socket option [29] is used to force MRCAP packets through raw sockets bound to the proper external interface, as directed in the MRCAP header.

5.5.4 A Concrete Example or A Day In the Life of a Packet

Consider a MOPED comprised of a PDA (with internal address I1) and mobile phone (with internal address I2, and external interface X, currently registered to care-of-address COA). We follow an example World Wide Web transaction between the PDA and a correspondent web server outside the MOPED, with reference Figure 4. The packets for the first half of the example are depicted in the figure.

1. The web browser on the PDA sends off a packet to the web server (WS).
2. The intra-component protocol fails to find a route to WS, and passes the packet to the Multipath Layer.
3. The multipath policy algorithm evaluates the packet and sends the packet through external interface X on the mobile phone. The Multipath Layer encapsulates the packet via MRCAP.
4. The internal routing protocol delivers the packet to the mobile phone.
5. The Multipath Layer receives the MRCAP encapsulated packet. It sees that this packet is to be sent to the proxy, through interface X. It mangles the MRCAP header, effectively storing the original source and destination addresses in the MRCAP header, while directing the packet to the proxy.
6. MobileIP on the phone intercepts the packet, and encapsulates it again, annotating the fact that the packet is actually traveling from the care-of-address.
7. The packet is shipped out on interface COA toward the proxy.

Figure 4: A packet's path through the MOPED Routing Architecture

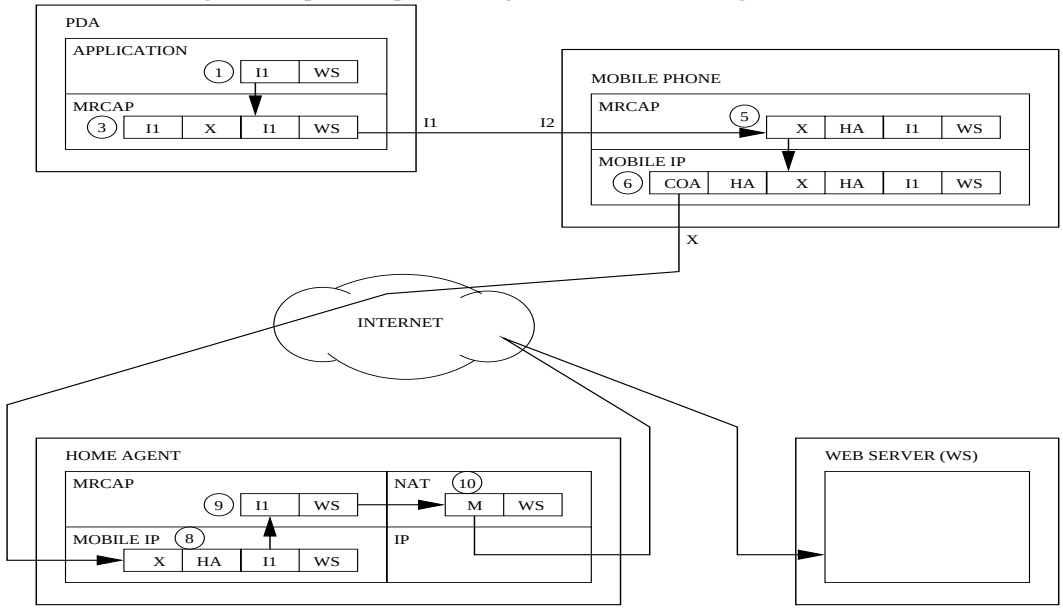

8. At the proxy, the packet is delivered to MobileIP, and the outermost header is stripped.

9. The new outermost header directs the packet to the Multipath Layer in the proxy. It records any info communicated to it by its peers in the PDA or the mobile phone, and decapsulates the MRCAP payload.

10. NAT on the proxy recognizes I1 as being MOPED-internal, and maps it to the public MOPED address M, after recording the binding (WS, I1), but before delivering the packet on to the web server.

The packet has been successfully routed out of the MOPED, and to the destination. We follow the return to the PDA of the response.

1. The web server initiates a response packet:

IP	
Src	Dst
WS	M

2. The proxy intercepts the packet, where it is recognized by the NAT layer as matching the earlier binding (WS, I1); the destination address of the packet is mangled appropriately, and the packet sent on to the Multipath Layer:

IP	
Src	Dst
WS	I1

3. The Multipath Layer determines the path into the MOPED for this packet. Our prototype, with its simple binding mechanism, uses the interface through which the first packet passed as it came to the proxy:

IP		MRCAP	
Src	Dst	Src	Dst
WS	X	—	I1

4. MobileIP tunnels the packet to the care-of-address COA:

MobileIP		IP		MRCAP	
Src	Dst	Src	Dst	Src	Dst
HA	COA	WS	X	—	I1

5. IP delivers the packet to the mobile phone, using interface COA. It is delivered to the local MobileIP layer, and the outer header is stripped:

IP		MRCAP	
Src	Dst	Src	Dst
WS	X	—	I1

6. The Multipath Layer receives the packet, and records its path through X before sending it on to the PDA I1:

IP		MRCAP	
Src	Dst	Src	Dst
WS	I1	X	—

7. At the PDA, The Multipath Layer once again has a chance to examine the packet and record any state attached by the multipath policy agents in the proxy or mobile phone. It then decapsulates the MRCAP payload to deliver to the web browser:

IP	
Src	Dst
WS	I1

The MRA directs traffic in its complicated dance, but in the end, the web server and browser are none the wiser; they have participated in the MRA transparently.

6. EXTENSIONS TO MRA

The design of the original MOPED architecture assumed that all MOPED nodes would support the MRA, limiting the inclusion of legacy devices. The design also assumed that all traffic would be routed through the home proxy, causing triangle routing in situations where the MOPED component is topologically close to the correspondent host. This section describes two extensions to the architecture and implementation described in the previous section that first, enable inclusion of all user devices, even non-MOPED-enabled devices, and second, support route optimizations to avoid triangle routing.

6.1 Legacy Devices: Freeloaders

Our goal of enabling legacy devices to participate in a MOPED is challenged by the use of MRCAP for routing to nodes outside the component and the use of the ad hoc routing protocol for routing inside the component. We have developed a simple extension to the normal MOPED Routing Protocol to provide unmodified legacy devices, or *freeloaders*, limited participation in the MOPED. With support from one of the MOPED-enabled devices, called a *designated relay*, the freeloader is able to take advantage of all MOPED connectivity, though it does not participate in the routing of packets from other nodes in the MOPED. The freeloader connects to the MOPED by acquiring an internal address via DHCP [4], which is supported by the relays and the proxy.

When the freeloader device boots, it initializes its network interface and broadcasts DHCP requests (DHCPDISCOVER). Any overhearing relay can send the DHCP request on to the proxy, which allocates an unused MOPED-internal IP address for the freeloader. After this address is determined, the relay configures the freeloader (via DHCP's DHCPOFFER - DHCPREQUEST - DHCPACK sequence) to use this address, with the reserved relay-routing IP address as its default router. The relay then advertises its route to the freeloader through the MOPED routing protocol, so that other MOPED nodes—and the proxy—can reach it. The relay can intercept the freeloader's outgoing traffic and encapsulate it properly for MOPED routing. In effect, the relay is enabling the freeloader to enjoy the MOPED's advantages without carrying its share of the costs of MOPED maintenance.

Freeloader generated packets with destinations inside the freeloader's component require no special handling. When the freeloader transmits such a packet to its relay, the relay uses normal intra-component routing to forward the packet.

Inter-component and extra-moped traffic from/to a freeloader requires some special handling. Since the MRD intercepts all IP packets with extra-MOPED destinations in the course of its normal operation, the MRD on a relay also captures packets directed to it by a freeloader for handling. Simple inspection of the IP packet reveals that it was originated by a host other than the relay, but is destined outside the MOPED, and not MRCAP encapsulated – so that the originating host must be a freeloader. The MRD encapsulates the packet, handling it just as it would a locally generated IP packet, but additionally sets a flag in the MRCAP header (the *Freeloader* flag) indicating that the source of this packet is a freeloader. The flag makes it clear to any MOPED routing agent that the origin of the packet *does not* understand MRCAP. The MRD then directs the encapsulated packet through the normal routing process, choosing a path and directing it toward the proxy, or another MOPED component, as usual. In its binding cache entry for this packet's association, the proxy records that the packet's source is a freeloader, as indicated by the flag in the MRCAP header.

The return path for packets transmitted from correspondent hosts to freeloaders also requires extra handling. Upon receipt of such a packet, the proxy observes the annotation on its binding cache entry for the association and sets the freeloader bit in the MRCAP header of the encapsulated packet, which is then forwarded normally to the freeloader's MOPED component. Upon arrival at the MOPED perimeter device, the packet would normally be routed to its final destination after being marked with its ingress path. However, the MRD on the perimeter node observes that the *Freeloader* flag is set in the MRCAP header. Since freeloaders do not, by definition, process MRCAP-encapsulated packets, the perimeter MRD unencapsulates the packet and delivers it to the freeloader using normal intra-component routing procedures. (Recall that the freeloader's designated relay is advertising a route to the freeloader, which all nodes in the component have heard.)

ICMP error delivery is a critical facet of IP operation that needs special handling for freeloaders. Outside the domain of MOPED Routing, ICMP errors can be generated normally and are delivered to the freeloader by the MRA like any other packet. An ICMP error generated by an intermediate MRD in response to a MRCAP packet with the *Freeloader* flag set requires special handling. Since the IP layer of the freeloader would be confused by ICMP errors generated for MRCAP packets, the MRD generating such an ICMP error unencapsulates the packet before ICMP error generation.

Since the freeloader may be mobile relative to its designated relay, the relay-freeloader relationship must be maintained. First, the proxy offers the freeloader a short lease to force frequent renewals. This renewal is a "catch-all" in the event that other maintenance techniques fail: when the lease expires, the freeloader begins the relay registration process anew, hopefully reusing the same address. The designated relay monitors the presence of the freeloader with periodic ICMP probes. When/if the freeloader moves out of range of its designated relay, the designated relay initiates freeloader following.

Freeloader following entails a message flooded throughout the component by the designated relay, informing MOPED nodes that it has lost contact with a freeloader. Each relay in the MOPED tries to "follow" the freeloader by using ARP to resolve the freeloader's IP address. A relay that can successfully contact the freeloader volunteers to become the freeloader's new designated relay by contacting the original designated relay. If no volunteer asks the original relay to handoff responsibility for the freeloader, the designated relay deems that the freeloader has disappeared and deallocates the freeloader's IP address. After a freeloader handoff, the new designated relay becomes responsible for tracking the freeloader (with ICMP probes) and advertising a route to the freeloader.

6.2 Circumventing the MOPED Proxy: MOPED Route Optimization

In some cases, it may be desirable for a MOPED node to communicate directly with a correspondent host, avoiding the necessity to route through the proxy. Zhao, Casteluccia and Baker describe a system that performs flexible routing for a mobile host using MobileIP, allowing that host to selectively send some packets using regular IP [29]. For services/connections that do not require mobility (e.g., name resolution via a local name server), avoiding MobileIP is a useful optimization. The MRA should also support the ability for certain traffic to circumvent mobility and communicate

as directly as possible with a correspondent host. "As directly as possible" may not truly mean directly, as it does in the case of MobileIP, since the source node may need another MOPED device to route its traffic to the desired correspondent host. Nevertheless, this optimization avoids the triangle routing incurred by directing all traffic through the proxy, and can reduce network latency, as well as reducing overall demand on the network if the correspondent host is topologically close to the MOPED.

Essentially, the MOPED device that wants to circumvent normal MOPED operation in this way requires the perimeter node between it and the correspondent host to carryout the proxy's usual function in the MOPED communication path. That is, the perimeter node needs to use NAT and present the traffic from the internal node as if it comes from the perimeter node's external interface. To circumvent mobility as well, the perimeter node also uses an additional mechanism like that of Zhao, Casteluccia, and Baker to allow this traffic to bypass MobileIP.

Extending the MRA to perform this optimization is simple, although determining what scenarios are well-suited to its application is a topic of future research. When the MRD decides to route a packet directly to a correspondent host, it MRCAP encapsulates the packet as usual, and additionally sets the *Masquerade* flag in the MRCAP header. The MRD chooses a perimeter node through which to direct the packets of this association (if one is not already recorded in the binding cache) and sends the packet to that perimeter node via intra-component routing. Upon receipt, the perimeter MRD notes the *Masquerade* flag in the packet's MRCAP header, and, if this is the first packet in this association, sets up a NAT rule to translate the internal address of packets on this association to the care-of-address of the external interface through which the destination host can be reached. This packet, and others sent from the internal node on this association, is decapsulated and handed to IP for delivery.

Similar to the freeloader mechanism of Section 6.1, the return path of the association operates differently from the outgoing path. When packets from the correspondent host arrive at the perimeter node—which is the destination, as far as the correspondent host knows—the NAT layer translates the packet's destination address to that of the internal node, reversing the NAT rule. Intra-component routing can then forward the packet to the proper internal node.

If the MOPED node using this optimization detects that the assisting perimeter node has lost connectivity with the target (i.e., the perimeter node shutdown its external interfaces or crashed), it can select another perimeter node to be its proxy for this connection. In this fashion, such connections are insulated from some MOPED mobility events.

7. CONCLUSIONS AND FUTURE WORK

The MOPED Routing Architecture is a coherent network model for MObile grouPEd Devices. It addresses the three major challenges of MOPED routing—addressability, mobility, and route selection—allowing MOPEDs full Internet integration. The MRA enables a set of mobile communicating devices to cooperatively maintain Internet connectivity through multiple simultaneous points of access. All of this functionality is realized in a manner that is transparent to the remainder of the Internet, requiring no modification of infrastructure or changes to legacy network applications. An important factor contributing to the effectiveness of the MRA is the light-weight, extensible encapsulation protocol MRCAP.

The most important goal of our future work is the complete implementation of the MRA—no architectural design is truly complete until any errors and inconsistencies in the design have been exposed by implementation. Completing the implementation would entail the following:

- Examine alternatives for the MOPED-internal routing method (e.g., various static or ad-hoc routing protocols).
- Design a protocol for remote NAT configuration, enabling a MOPED device to declare itself the correct endpoint for a type of data.
- Implement and study alternative policy algorithms for multipath route selection.

After completing the basic structure of the MRA, we will scrutinize other aspects of MOPED networking that will enhance the overall MOPED design. We intend to replace the transport layer IP protocols with our family of multipath, bandwidth-aggregating protocols. We must develop the control structure and user interface to manage interface connectivity over the MOPED—an agent to set up and tear down external interfaces as appropriate for optimal resource utilization. We believe that collapsing the layered structure of the MRA, although complicating the implementation, may enable space optimization in the network packets by combining the MRCAP and MobileIP headers.

Security is one important issue of any set of personal technology devices that we do *not* explicitly address here. Existing solutions for network level security in the context of MobileIP apply perfectly well to our extended, MOPED-mobility environment.

We have shown how to extend the paradigm for communication from a mobile device to a mobile person, via the representative Internet presence embodied in a MOPED. The MRA enables efficient utilization of MOPED resources through cooperative communication. We make this all possible without necessitating any changes to Internet infrastructure or network software.

8. ACKNOWLEDGEMENTS

This work was sponsored in part by NSF ITR grant ANI-0081308. We would also like to thank everyone in the Mobius Group for their input.

9. REFERENCES

[1] M. Baker, Z. Zhao, S. Cheshire, and J. Stone. Supporting Mobility in Mosquitonet. In *USENIX Winter Conference*, 1996.

[2] L. Bellier and C. Casteluccia. Mobile Networks Support in Mobile IPv6. Internet Draft draft-ernst-mobileip-v6-network-01.txt, IETF, November 2000.

[3] P. Debaty and D. Caswell. Uniform Web Presence Architecture for People, Places, and Things. *IEEE Personal Communications*, 8(4), August 2001.

[4] R. Droms. Dynamic Host Configuration Protocol. Request For Comments (Draft Standard) RFC 2131, IETF, March 1997.

[5] D. Farinacci, T. Li, S. Hanks, S. Meyer, and P. Traina. Minimal Encapsulation within IP. Request for Comments (Proposed Standard) RFC 2784, IETF, March 2000.

[6] T. Hain. Architectural Implications of NAT. Request For Comments (Informational) RFC 2993, IETF, November 2000.

[7] L. Magalhães and R. Kravets. End-to-end Inverse Multiplexing for Mobile Hosts. In *19th Brazilian Symposium on Computer Networks (SBRC'01)*, 2001.

[8] L. Magalhães and R. Kravets. Transport Level Mechanisms for Bandwidth Aggregation on Mobile Hosts . In *9th International Conference on Network Protocols ICNP 2001*, 2001.

[9] L. Magalhães and R. Kravets. MMTP: Multimedia Multiplexing Transport Protocol. In *The First Workshop on Data Communications in Latin America and the Caribbean (SIGCOMM-LA 2001)*, 2001.

[10] P. Maniatis, M. Roussopoulos, E. Swierk, K. Lai, G. Appenzeller, X. Zhao, and M. Baker. The Mobile People Architecture. *ACM Mobile Computing and Communications Review*, 3(3), July 1999.

[11] J. Mogul and S. Deering. Path MTU Discovery. Request For Comments (Draft Standard) RFC 1191, IETF, November 1990.

[12] A. Myles, D. Johnson, and C. Perkins. A mobile host protocol supporting route optimization and authentication. *IEEE Journal on Selected Areas in Communications*, 13(5):839–849, 1995.

[13] C. Partidge, T. Mendez, and W. Milliken. Host Anycasting Service. Request For Comments (Informational) RFC 1546, IETF, November 1993.

[14] C. Perkins. IP Encapsulation within IP. Request for Comments (Proposed Standard) RFC 2003, IETF, October 1996.

[15] C. Perkins. IP Mobility Support. Request for Comments (Proposed Standard) RFC 2002, IETF, October 1996.

[16] C. Perkins. Minimal Encapsulation within IP. Request for Comments (Proposed Standard) RFC 2004, IETF, October 1996.

[17] J. Postel. Internet Protocol. Request For Comments (Standard) RFC 791, IETF, September 1981.

[18] J. Postel. Transmission control protocol. Request for Comments (Standard) RFC 793, Internet Engineering Task Force, September 1981.

[19] R. Ramjee, T. L. Porta, S. Thuel, K. Varadhan, and S.-Y. Wang. Hawaii: A domain-based approach for supporting mobility in wide-area wireless networks. In *Proceedings of the International Conference on Network Protocols (ICNP) '99*, 1999.

[20] A. Snoeren and H. Balakrishnan. An End-to-End Approach to Host Mobility. In *ACM Mobicom '99*, 2000.

[21] H. Soliman, C. Casteluccia, K. El-Malki, and L. Bellier. Hierarchical MIPv6 mobility management. Internet Draft draft-ietf-mobileip-hmipv6-04.txt, IETF, July 2001.

[22] P. Srisuresh and K. Egevang. Traditional IP Network Address Translator (Traditional NAT). Request For Comments (Informational) RFC 3022, IETF, January 2001.

[23] P. Srisuresh and M. Holdrege. IP Network Address Translator (NAT) Terminology and Considerations. Request For Comments (Informational) RFC 2663, IETF, August 1999.

[24] M. Stemm and R. H. Katz. Vertical Handoffs in Wireless Overlay Networks. *ACM Mobile Networking (MONET), Special Issue on Mobile Networking in the Internet*, 1998.

[25] A. G. Valko. Cellular ip - a new approach to internet host mobility. *ACM Computer Communications Review*, 1999.

[26] D. Waitzman. A Standard for the Transmission of IP Datagrams on Avian Carriers. Request for Comments (Proposed Standard) RFC 1149, IETF, April 1990.

[27] H. J. Wang, B. Raman, C.-n. Chuah, R. Biswas, R. Gummadi, B. Hohlt, X. Hong, E. Kiciman, Z. Mao, J. S. Shih, L. Sunramanian, B. Y. Zhao, A. D. Joseph, and R. H. Katz. ICEBERG: An Internet-core Network Architecture for Integrated Communications. *IEEE Personal Communications*, August 2000.

[28] X. Zhang, J. G. Castellanos, and A. T. Campbell. P-MIP: Paging Extensions for Mobile IP. *ACM Journal on Mobile Networks and Applications*, 2001 (to appear).

[29] X. Zhao, C. Castelluccia, and M. Baker. Flexible Network Support for Mobility. In *Fourth ACM International Conference on Mobile Computing and Networking (MOBICOM'98)*, 1998.