B. Table of Contents

List of Figures and Tables

C. Abstract

A Revolutionary Confederated Service Architecture for Future Telecommunications Systems

Principal Investigators:

Professor Randy H. Katz
Computer Science Division
EECS Department
University of California, Berkeley
Berkeley, CA 94720-1776

Professor Anthony D. Joseph
Computer Science Division
EECS Department
University of California, Berkeley
Berkeley, CA 94720-1776

Cooperating Companies:

Ericsson
HRL Laboratories
Nortel Networks
Sprint

Our goal is to understand how to create end-to-end telecommunications services with desirable and predictable properties, such as performance and reliability, when provisioned from multiple and independent service providers. We propose to develop a new architecture for future telecommunications services that supports the dynamic confederation of sometimes collaborating and sometimes competing service providers. Our first effort in this direction, the *Clearinghouse Architecture*, provides a resource management system based on predictive resource reservations, traffic-matrix admission control, and group policing for detecting malicious flows. The Clearinghouse is focused on dynamic trunking decisions within and across ISP clouds. It offers a starting point for our investigation into the generalization of the concept of *service level agreement* to multiple service providers and for properties other than bandwidth, latency, and packet loss rates. It illustrates the principles of improved scalability and predictability through aggregation, and the use of hierarchy and cooperation among service providers to make effective and agile resource allocation decisions. We propose to extend this work in the direction of more general application of economic mechanisms, such as dynamic auctions, for resource allocation problems in multi-provider telecommunications service architectures.

D. Introduction

1. Motivation: The Existing Operator Model is Failing

The expense of deploying Third Generation (3G) Telecommunications Systems will be huge. The European auctions for 3G spectra are likely to exceed $150 billion, with $45 billion already committed in Germany and $35 billion in the United Kingdom (U.K.). Equipment outlays are likely to match these spectrum expenses. And this is all before first revenue, without a clear understanding of the kinds of new services and applications enabled by 3G bandwidths for which subscribers will pay. Cheap (core) network bandwidth and the highly competitive environment brought about by widespread liberalization of the telecommunications sector is simultaneously driving bandwidth prices towards zero while yielding a financial crisis for the operators!

There is a growing recognition that highly integrated “all things to all people” telecommunications companies, like AT&T or British Telecomm (BT), on the one hand provide no effective economies of scale and on the other have encumbered very large debt in pursing their integrated visions[1]. A new, more functional specialization is being called for:

> “The new rules for success will be to provide one part of the puzzle and to cooperate with other suppliers to create the complete solutions that customers require. ... [V]ertical integration breaks down when innovation speeds up. The big telecoms firms that will win back investor confidence soonest will be those with the courage to rip apart their monolithic structure along functional layers, to swap size for speed and to embrace rather than fear disruptive technologies.” [Economist 2000]

2. The Need: A New Service Model for a New Business Model ... Beyond the Third Generation

We believe that a radically new service architecture is needed, and one that cannot be separated from an equally radical view of the telecommunications operator’s role in the new value chain of service provision. Simply stated, future telecommunications systems will be organized not as monolithic structures deployed by a single business entity, but rather as a dynamic confederation of multiple—sometimes cooperating and sometimes competing—service providers.

We propose to investigate just such a service architecture, and to deploy a prototype and evaluate it in a testbed environment, on a building, campus, and potentially regional-scale, in collaboration with our industrial partners. These partners include network equipment manufacturers (Ericsson, Nortel, HRL representing Boeing and Raytheon) as well as network operators (Sprint, HRL representing General Motors/DirecTV).

2.1. Overarching Themes

There are several assumptions that we make and themes that we expect to emerge from the work we propose:

- *Diversity rather than homogeneity*: Future telecommunications systems will be characterized by many kinds of end devices, access networks, services, applications, service providers and content providers. While common interfaces are highly desired, it is inevitable that not every participant will make use of even the most well developed standards. Therefore, the service architecture must embrace diversity and exploit software mediation to achieve interoperability.
- *No overarching access network*: A variety of access networks and data services will co-exist

1. Sprint appears to be an intriguing exception to this observation, based on a clear focus of enhanced service (speech-activated dialing) and aggressive and innovative rate plans (“Five cents a minute, anywhere, anytime”).

for some time to come. For example, in the wireless access network, we will have EDGE, HSCSD (high speed circuit switched data), GPRS (general packet radio service), 3G packet, and so on [Bi 2001]. The architecture must achieve application transparency by providing the necessary services to integrate these.

- *Enable emerging business models*: The service architecture must support technologies that enable new business models, such as *Mobile Virtual Network Operator* (MVNO) to achieve enhanced efficiency of resource usage (MVNO is described in Section E.3.1).
- *It's all about services*: The service architecture must enable business entities to provide enhanced services, as a primary means of differentiating one provider from another. Support for sophisticated capability negotiation and service level peering, and concepts like enhanced preferences management in support of the *Virtual Home Environment* (VHE), will be essential pieces of any such service architecture (VME is described in Section E.3.1).
- *It's all about confederated services*: The service architecture must support overlapping service provider regions, with subscribers able to roam among them for service provisioning, even without actually "moving." The architecture must support "co-opetition" among service providers, that is, sometimes cooperating and sometimes competiting relationships, forming dynamic syndicates for the purpose of service provision.
- *Network-application awareness*: To better support subscribers in their tasks, the service architecture must provide mechanisms that make the network more aware of applications (e.g., near-term future indications of needed bandwidth), while the applications become more aware of the availability of system parameters and resources (e.g., user location, proximity of system resources to the user's current location, etc.).

2.2. What is New: Dynamic and Adaptive Services, Collaboratively Provisioned By Multiple Providers, Exploiting Applications Awareness

Over the last few years our research group has explored a number of issues in multi-network service architectures. Our BARWAN project focused on overlay networks (e.g., vertical handoff [Stemm 1999, Wang 1999]) and interoperation of media services [Brewer 1998]. We are also completing an architecture for Internet-based service management in the ICEBERG Project (e.g., service mobility) [Raman 2000, Wang 2000]. Several of the concepts mentioned above, such as MVNO and VME, and new multi-provider business models, are well under way for incorporation into cellular industry's Third Generation (3G) architectural specifications in some form [3GPP]. These features have also been recognized as essential components of UMTS [UMTS XX].

Our proposed architecture builds on these prior efforts, but we will extend them in several significant directions. Our primary goal is aggressive support for the dynamic and adaptive deployment of services and network resources (bandwidth, processing, and storage) to users and their applications. This will be accomplished through architectural support for multiple providers at a variety of system levels, who confederate to provide services and resources, yet also compete for business. An essential feature of the architecture will be the support for radically new kinds of business models, admitting of a new more horizontal layering of service provision. Finally, it will be important for applications to interact with the underlying system, to better represent their needs and to become better aware of system capabilities, thus enabling both the applications and the system to be come more adaptive. The research we will perform will provide a firm foundation upon which to define the telecommunications service architecture of 2010 and beyond.

E. Technical Discussion

1. Detailed Statement of the Problem

The existing approaches for building and operating wireless telecommunications systems are outdated, inefficient, and too expensive:

- Scarce spectrum resources are statically partitioned among license holders—independent of subscriber density or the nature of the users' workloads! One operator's spectrum resources are oversubscribed, while another's are underutilized. Why can't the latter resell capacity to the former, on a short term, as needed basis, perhaps informed by the near-term future needs of its current subscribers? Note that this implies more of a peer-to-peer model than the hierarchical model inferred by the relationship between Virgin Mobile and One2One described in more detail below.
- An operator must acquire cell sites for antenna deployment, often duplicating coverage already provided by other providers (and it is becoming more difficult for operators to acquire new sites due to local resistance to proliferating towers [New Scientist 2001]). Is there a faster way to achieve regional coverage than the current piecemeal and duplicative coverage approaches? Is there a role for the antenna "operator" to participate in the value chain, by accepting a percentage of the revenue stream generated by subscribers who make use of the resources—antenna access—it provides? Or will operators confederate to share the cost of building a common infrastructure, and recharge each other in order to recover the deployment costs?
- An operator must build its own backhaul network before interfacing to the PSTN and Internet. This is instead of sharing a common, better-provisioned backhaul capability.
- AT&T has followed a now somewhat discredited strategy to offer its subscribers a "one stop shop" with everything from wireless access to long distance services to cable TV/modem data access to portal services (i.e., Netscape). AOL's efforts to integrate Time Warners' cable TV networks and content with Internet Service provision may also be destined for failure. We believe that the future is horizontal disaggregation rather than vertical integration.

Our goal is to develop a revolutionary alternative architecture for future telecommunications systems, based on the following assumptions:

- Service providers will desire to form dynamic confederations to better share resources (spectrum, bandwidth, processing, storage, infrastructure like antenna sites), and deploy access and achieve regional coverage more rapidly. While this runs counter to the way that most operators behave today, there is an emerging realization that given the astronomical costs of building out the 3G networks, partnering to share the costs and the risks, or just to keep the network pipes filled, are driving operators in this direction. This is already happening in Europe, where some operators are forming partnerships to build a shared infrastructure for wireless access despite being competitors.
- Scarce resources will be more efficiently allocated in fine-grained units, on an as needed basis, using dynamic "market-driven" mechanisms, with awareness of applications demands to inform the allocation decisions. As resource limits are approached, prices for those resources will increase, thus providing current resource holders with incentives to relinquish them. This is an area that needs to develop new interfaces between applications and the underlying resource management infrastructure, so that applications become aware of the real costs of the

operations they are performing. Just how “fine-grained” such an allocation can be in the time domain will be an important aspect of this research; presumably the relationship between Virgin Mobile and One2One involves long lived business contracts.

- Trusted third partners (e.g., Clearinghouses, Resource Brokers, B2B Exchanges, etc.) will manage the resource marketplace in a fair, unbiased, audited and verifiable basis. While agent architectures and multi-agent coordination are well researched areas, the details of the potentially very rich web of relationships among intermediaries needs to be structured and better understood in the specific context of resource negotiation between subscribers and service providers (and between the various layers of service providers).
- The traditional vertical stovepipe of access providers (frequency, cell site, base station, backhaul) will be replaced by a radically different organization of horizontally organized “multi-providers,” open to increased competition and more efficient allocation of resources. The same is true of access network, core network, service provider, and content provider. A key contribution of this research will be a better understanding of the needed technological support for such horizontal service organization, through the design, prototype implementation, and system evaluation we will undertake.

The challenge is how to achieve “service level peering” and resource sharing in an environment of limited trust and cooperation. At this stage of our research, we can only identify some of the elements of a potential solution:

- An open service and resource allocation model, with independent service creation, establishment, and placement in overlapping domains of service providers. Service providers must be able to advertise available slack resources, husband some resources for near-term peak needs, and be recharged for usage of resources by third parties. There is an open question of how much information about the current availability of resources should be made—as this would allow a competitor to understand the weaknesses of the advertising service provider.
- Provider resources, capabilities, and current status need to be described, and exchanged between confederated service providers, whether it is for the purposes of “reselling resources” where they are collocated, or for implementing VHE functionality through local provisioning. The nature of capability negotiation needs to be understood in this more complex, more heterogeneous environment than that found in conventional telephony networks.
- Allocation mechanisms based on economic methods, such as electronic auctions, coupled with real-time accounting/billing/settlement systems for the resources used. This in part solves the question raised in the first bullet above—if the intermediary “market maker” can hide the details of the state of the participating providers’ resource pools. This raises another issue: the desirability of peer-to-peer systems versus a centralized market mechanism that might serve as a potential single point of failure.
- Mechanisms for managing trust relationships among clients and service providers, and between service providers, based on trusted third party monitors. Existing interrelationships among service providers are based on contracts between large and well-established commercial enterprises. Such assumptions may not be suitable as the number and scope of service providers increase dramatically. Trusted third parties might audit the behavior of service providers, a kind of “better business bureau,” providing a rating service that separates the reliable participants from those who are unreliable.

- General services for forming dynamic confederations, such as: discovering potential confederates, establishing trust relationships, managing transitive trust relationships (if A trusts B and B trusts C, what can you infer about the relationship between A and C?) and managing the levels of transparency (e.g., reveal aspects of internal information to peers or trusted third parties about performance, user information, etc.).
- Not all confederates need be potential competitors. Confederations can be formed of heterogeneous, collocated access networks to better support a given subscriber's application needs. For example, a WLAN operator (i.e., an enterprise) may confederate with a cellular network provider to provide two separate pipes for high bandwidth data and low bandwidth voice streams in support of a user's multimedia conferencing application. We assume that the user carries both a cell phone and some kind of screen-oriented computing device, and that these devices have now formed a confederation to transparently support an application, with the audio coming over the phone and the images over the WLAN. Similarly, the relationship between an enterprise level GSM on the Net (GoN) operator and the overlay cellular operator is an example of a non-competitive confederation.

2. Progress Report for Continuation Project

Not applicable, as this is a new proposal.

3. Proposed Work

3.1. Related Work

UMTS/IMT-2000

The Universal Mobile Telecommunications Systems (UMTS) Forum has long focused on the issues of integrating multiple access networks into a single coherent system. More recently, the UMTS-Forum has been addressing the issues of multiple service providers and the means for integrating them into a coherent, multi-provider architecture:

> "This new environment is changing the role model for service delivery and the value chain for mobile service provision. Compared to a telephone service provider and network operator, more players are active (see Figure 1). The roles of the access network operators, core transport network operators, service providers, and content providers have to be distinguished. There are interrelationships between these different possible players. In addition, users and subscribers are not necessarily the same. Service and content providers play an increasing role in the value chain. The dominant part of the revenues moves from the network operator to the content provider. Both are using the services of the access network and core network operators. It is expected that value-added data services

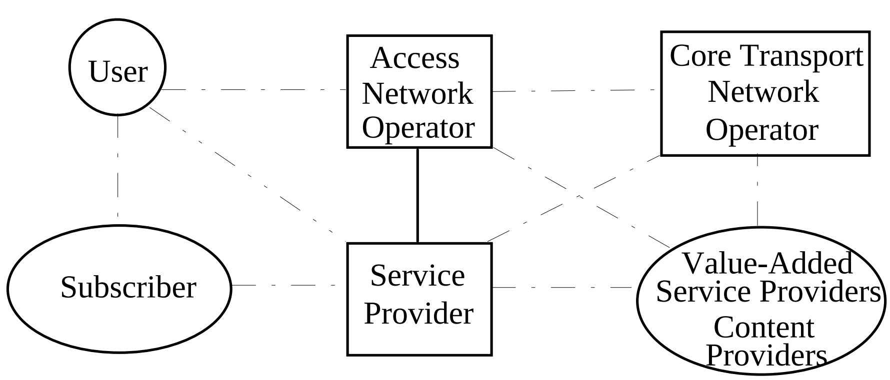

Figure 1. New Role Model and Service Delivery for UMTS

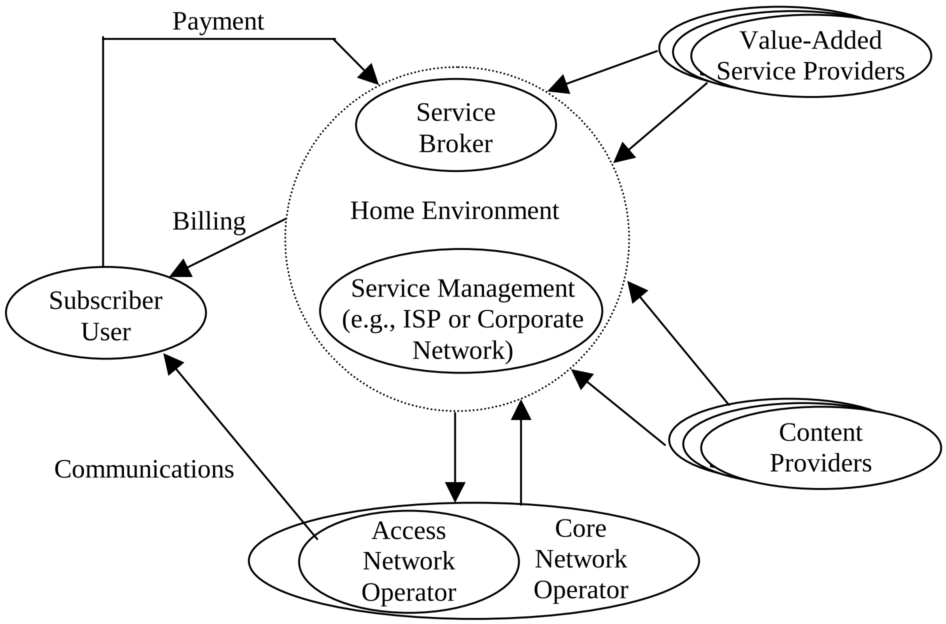

Figure 2. UMTS Service Model and Relationships

and content provisioning will create the main growth"[Mohr 2000].

In the emerging UMTS Service Model, there is an explicit separation of the Access Network and Core Transport Network operators, as well as Service, Value-Added Service, and Content Providers. Such a model implies that multiple providers can confederate and/or compete in delivering a service to the end user.

The UMTS Service Model also admits of a sophisticated accounting, billing, and settlements architecture to support third party brokering between subscriber needs for service and multiple service providers, as indicated in Figure 2. The italics are our own, for emphasis:

> "The available, emerging, and evolving access technologies have basically been designed in the classical vertical communication model that a system has to provide a limited set of services to users in an optimized manner. … *[S]ystems beyond third generation will mainly be characterized by a horizontal communications model, where different access technologies … will be combined into a common platform to complement each other in an optimum way for different service requirements and radio environments.* These access systems will be connected to a common, flexible, and seamless core network. … *sharing of frequencies for better use of resources between different systems and possibly different operators* … new network types and network management such as ad hoc and self-optimizing networks, automatic and dynamic network reconfiguration, and *dynamic frequency allocation to support a variety of access systems on a common platform.* … Vertical handover is combined with service negotiation to ensure seamless service, because in general different access systems support different user data rates and other bearer and service parameters. … Spectrum sharing between different systems and the investigation of coexistence conditions between different radio access systems" [Mohr 2000].

The UMTS Reference Architecture is based on five distinct layers, and it possible for different service providers to be active at each one:

- Distribution Layer, e.g., DAB/DVB: Full coverage, global access, full mobility, no individual links;
- Cellular Layer, e.g., GSM, IMT-2000, UMTS: Full coverage and hot spots, global roaming, full mobility, individual links;

- Hot Spot Layer, e.g., WLAN: Local coverage, hot spots, global roaming, local mobility, individual links;
- Personal Network Layer, e.g., BlueTooth: Short range, global roaming, individual links;
- Fixed/Wired Layer: No mobility, global roaming, individual links;

Note that there is no explicit discussion of where service provision and service mediation should exist in this architecture, other than in the core network that ties all access networks together. The architectural focus is on access networks, with insufficient emphasis, in our opinion, on the needed service architecture to bring such a system together.

3GPP MVNO/VHE

The 3GPP standardization group has been working on defining the concept of *Mobile Virtual Network Operator* (MVNO) [3GPP]. By definition, MVNOs do not have a licence to use radio spectrum, but rather have access to the radio networks of multiple current mobile operators. An MVNO can offer new services and new bundling and/or pricing of services to customers using that spectrum. The MVNOs pay the real network operators for the use of and access to the mobile network. This is sometimes called *retail-minus pricing*, wherein the charges are determined by deducting from the retail prices of the calls the costs of the elements of the calls handled by the MVNO rather than the mobile network operator.

In the U.K., Virgin Mobile is often referred to as an early generation MVNO, providing its own branded service on top of the network of the operator One2One. Virgin's marketing muscle helps it to sign-up subscribers who generate additional demand for One2One's network. Yet One2One views itself in competition with Virgin Mobile to sign-up as many subscribers as possible. Note that much of the differentiation is in the branding, pricing, and packaging of services, rather than in any new technology, capability, performance advantage, or network coverage. Virgin, for example, collects a fee of 10% of any pay-for-service invoked by its subscribers.

Note however that the relationship between Virgin and One2One is static: One2One recognizes a SIM card assigned to a Virgin subscriber, and recharges Virgin for the use of its network by such a subscribers. A more general MVNO would have such a relationship with multiple mobile operators, manage its own billing/accounting system, and have an independent method of settling its accounts with the underlying MNOs [Oftel 1999]. With a more general model of selecting ("roaming") among co-located MNOs, one could imagine a dynamic auction managed by the MVNO whereby competing MNOs aggressively bid to carry its traffic when their networks are not busy, eschewing such efforts when their networks are approaching full utilization.

Within IMT-2000 systems, the *Virtual Home Environment* (VHE) permits users to retain their "home" settings on their handsets, regardless of their current geographic location and providing access network operator. It allows users to gain access to their own personalized services in a consistent way no matter where they are, e.g. at home, in the office, on the move. Users are able to roam into different technology networks and experience the same "look and feel" depending on the capabilities of the serving network and the terminal equipment in use. Wherever the user is, the attributes and configuration of the handset such as telephone number and options appear to be the same as if they were at home.

While the service could execute back within the subscriber's home network, it is highly advantageous if the services—if they exist locally—are executed there. This requires cooperation between the home and visited networks, as well as compatible technological infrastructures. In a sense this is a limited example of "service level peering": a negotiation among the home and visited access providers to determine how to best to create the image of the service, given local resources and capabilities.

3.2. SAHARA: a Service Architecture supporting Heterogeneous Access, Resources, and Applications

SAHARA is a new research activity, covered by the proposed project described herein. It builds on our Internet-based telecommunications service architecture developed in the ICEBERG Project. SAHARA extends that work towards the application of market-based economic mechanisms to resource allocation decision making, with support for multiple service provider confederation, to achieve end-to-end service provision with desirable performance and reliability. An alternative to conventional Diff-Serv and Int-Serv performance "guarantees," based on communicating agents which we call the Clearinghouse Architecture, is our first effort to understand the general issues of multiple provider service management in wide-area networks. It is described next.

A First Approach: The Clearinghouse

Our initial work on SAHARA involves the definition of what we call the *Clearinghouse architecture* [Chuah 2000]. The Clearinghouse is a distributed control architecture for scalable resource provisioning, based on prediction-informed resource reservations across multiple administrative domains, to drive specific network management algorithms such as admission control and traffic policing. These algorithms work at the edges of the service provider cloud. The architecture must achieve adequate performance without maintaining per-flow state at the edge router, while achieving robustness against traffic demand fluctuations and misbehaving flows. It is equally important that the mechanisms we develop introduce minimal overhead to edge routers to achieve the desired scalability and ease of deployment.

While the goal is the same as Int-Serv and Diff-Serv [Braden 1997, Blake 1998], the Clearinghouse differs from these in a number of ways. Unlike Int-Serv, which requires per-flow state maintenance in the network core, the Clearinghouse aggregates flows to achieve better scalability and avoid performance bottlenecks. Like Diff-Serv, the Clearinghouse is focused on soft quality of service provisioning, and does not seek to achieve hard performance "guarantees" as does Int-Serv. And unlike Diff-Serv, the Clearinghouse achieves performance provisioning in an end-to-end manner rather than pairwise between cooperating service providers.

The initial design has concentrated on the single service provider environment, allocating bandwidth resources between ingress and egress points of the ISP network cloud. This is essentially provisioning the trunking "pipes" from one edge of the cloud in one geographic region to another edge in another region, such as provisioning the aggregated bandwidth demands between San Francisco and New York in a given ISP's backbone. Our ultimate goal is to extend the work to the multi-service provider case, and to use the architecture to provision other resources beyond bandwidth, such as processing and storage.

The Clearinghouse (CH) architecture builds on five distinct functionalities: (1) it monitors network performance on ingress/egress links, (2) it estimates traffic demand distributions to drive near-time predictions of future resource needs, (3) it adapts trunk/aggregate reservations within and across domains based on traffic statistics, (4) it performs admission control based on estimated traffic matrix, and finally, (5) it coordinates traffic policing at ingress and egress points for detecting misbehaving flows.

We envision implementing the CH as a hybrid of hierarchical and flat structures. Within a

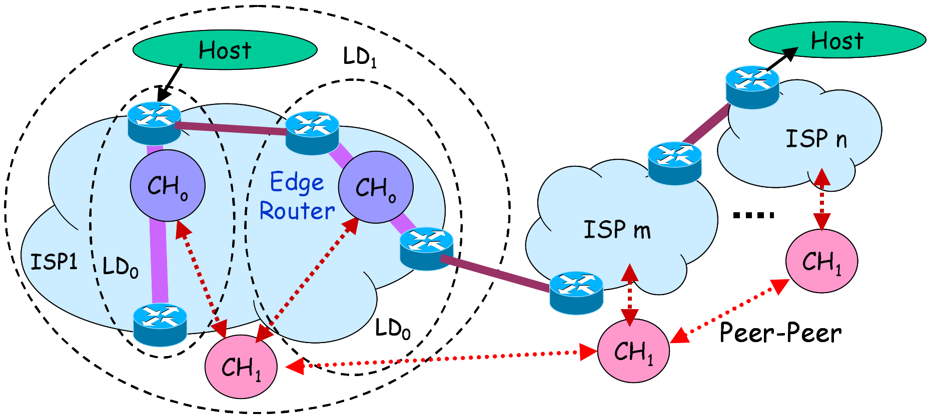

Figure 3. The Clearinghouse's Hierarchical and Flat Structure

single large-scale ISP, the CH is organized as a hierarchy, with traffic measurements within a local domain being aggregated and passed up to higher level CHs that collaboratively make resource allocation decisions for regions of the ISP cloud. This helps reduce the amount of state information that needs to be maintained, thus enhancing the scalability of the approach. A flat structure is used for peer-to-peer service agreements among the CHs of collaborating ISPs.

The conceptual architecture is shown in Figure 3. Within the ISP cloud, the Parent CH-node adapts trunk reservations across local domains (LDs) for aggregate traffic within an ISP. The CH appears to be flat at the top level, where the top level CH coordinates the peer-to-peer trunk reservations with the top-level CHs of other cooperating ISPs.

Resource Management with the Clearinghouse

Within a single administrative domain, the CH dynamically allocates bandwidth in the ISP's core network between the endpoints defined by ingress and egress routers. Reservations are established for aggregated flows on intra- and inter-domain links. They are adapted dynamically as the traffic between end points fluctuate over time. Such an aggregated approach allows the core routers to be stateless while the edge routers maintain only aggregate rather than per-flow state information. Traffic monitoring, admission control, and traffic policing for individual flows are performed only at the edge of the cloud. As a result, access routers have smaller routing tables and experience lower aggregation of traffic relative to backbone routers. It is our belief that most congestion, yielding packet loss and delay, occurs in access networks and peering points between ISPs.

The CH supports three crucial mechanisms upon which to support the reservation-based architecture. The first is the ability to predict near-term bandwidth demands using Gaussian traffic predictors. We call these *Predictive Reservations*. While it may be difficult to predict the traffic demands of a single flow, the aggregation of many flows yields better statistics and more predictable behavior. This has been observed independently by Roberts [Roberts 2001]. A given ISP's CH maintains a traffic-matrix that estimates the aggregate traffic demand between each pair of access routers in its domain. This enables what we call *Traffic-Matrix based Admission Control* (TMAC). The final mechanism is a method of detecting and penalizing misbehaving flows, which we call *Group Policing for Malicious Flow Detection* (GPMFD).

We assume the network is capable of supporting at least two quality of service (QoS) classes, such as real-time and best effort, with priority scheduling among these. High-priority flows are admission controlled and best-effort traffic cannot preempt high-priority traffic. The edge routers support passive monitoring of aggregated flows. Individual flows send explicit request and teardown messages, and specify a required peak rate. The network topology and all-pair shortest paths between IR and ER can be determined from the underlying routing protocol.

We treat ingress and egress routers as endpoints, and estimate the traffic demand for the *ingress-egress pipe* (IE-Pipes) between each pair. We dynamically compute an Upper-bound Matrix, indexed by ingress and egress points, splitting the bottleneck bandwidth between competing IE-Pipes in proportion to their estimated bandwidth demands. The matrix is updated at regular intervals, at a time scale that is smaller than that at which the traffic demand fluctuates. We outline our approach to TMAC and GPMFD in the following sections.

Traffic-Matrix Based Admission Control (TMAC)

The admission control policy admits a flow to an IE-Pipe if its requested bandwidth plus the estimated rates of the already admitted flows does not exceed the value in the Upper-bound matrix, with extra capacity factor capturing the effects of statistical multiplexing. The last factor is the essential control knob that allows the admission control mechanism to trade off between over commitment and underutilization of the available bandwidth resources.

Figure 4 illustrates the components of the TMAC admission control algorithm. It requires some modification to edge routers to include traffic monitors that passively measure the aggregate rate of existing flows between entry points (IR-s) and exit points (ER-d). The IR-s forwards the admission control messages (that is, *Request*, *Accept*, *Reject*) between the CH and the host or its Clearinghouse-aware proxy. It is also responsible for estimating traffic demand distributions, which it in turn reports to the CH. The CH leverages its knowledge of topology and the traffic matrix in making admission control decisions.

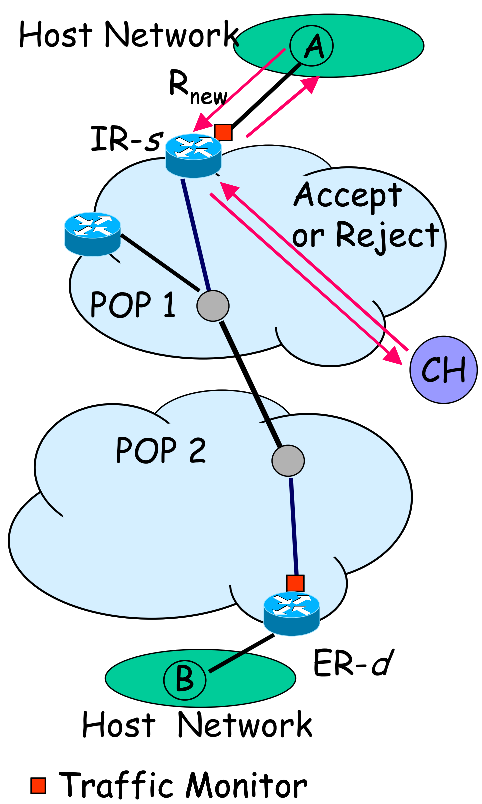

Figure 4. TMAC Algorithm

Initial performance simulations have demonstrated that each ingress-egress pair achieves 10-20% of its ideal bandwidth allocation. Given our focus on soft QoS, the inherent imprecision in traffic measurements and predictions, and our goal to develop a practical admission control scheme for aggregated flows, we find this result to be very encouraging.

TMAC-based admission control has several advantages over traditional approaches. It does not rely on hop-by-hop signaling protocols like RSVP. It allows the core routers to remain stateless. Only the ingress and egress routers need to keep aggregate state information. It can perform traffic monitoring without requiring end-to-end probing by end hosts. The traffic matrix provides a nice abstraction of the traffic demand distributions within an ISP. It is both topology sensitive and responsive to traffic fluctuations.

Group Policing for Malicious Flow Detection (GPMFD)
We assign a Flow ID (FID) based on a flow's ingress and egress points. Token bucket filters (TBFs) are used to police traffic at the ingress and egress routers. Admitted flows from the same source router and headed for the same destination router are policed as an aggregate. It is important to understand that traffic policers at the ingress and egress points only maintain information about total allocated bandwidth by group, not on a per flow basis.

If the ingress and egress routers agree that a given aggregated flow exceeds its bandwidth allocation, its token bucket can be reduced, thereby penalizing the flow. While this sounds straightforward, it is actually quite difficult to insure that well behaved flows within the aggregate are not penalized while a misbehaving flow manages to escape detection. Our approach is to enter a "micro-monitoring" mode whenever a token bucket overflows, thus allowing us to sample the rate of each flow within this group. The worst offenders within the aggregate are reported to the CH. The CH maintains a special *Policy Repository* (PR) that stores the bandwidth allocation for each admitted flow. It determines whether the estimated rate exceeds the allocated bandwidth plus a threshold, and if this condition is violated a threshold number of times, the flow is reported as malicious. That particular flow will then be penalized at the edge router. Initial analysis indicates that the algorithm we have outlined here can achieve good scalability by allowing edge routers to aggregate flows for policing without giving up the ability to identify misbehaving flows.

In terms of deployment, edge routers must embed traffic monitors to estimate the existing load. They also need the mechanisms for token bucket based traffic policing. Finally they need to be able to communicate with the CH, to identify bad flows as well as to participate in admission control decisions. To understand the overhead and complexity of such mechanisms, we are building a prototype using MIT's Clickrouter [Kohler 2000]. Initial performance comparisons with an unmodified Clickrouter indicate a small overhead penalty for implementing the enhanced functionality of GPFMD, in the range of 5% in lost throughput.

Summary and Directions
The Clearinghouse Architecture represents a first attempt at defining a distributed control architecture to coordinate intra- and inter-domain resource provisioning. It makes use of aggregate reservations for scalability and hierarchical control for better manageability and end-to-end support. As part of the CH, we have developed initial algorithms to accommodate scalability, performance, and robustness issues within a large ISP domain, namely a traffic-matrix based admission control mechanism that is topology sensitive and responsive to traffic fluctuations. We use aggregate flows for group policing and malicious flow detection, to reduce the amount of state maintained at the edge routers.

The obvious next steps are to generalize the CH to multiple service providers and to manage multiple resources beyond network bandwidth, such as processing and storage for executing services in the network. One critical direction is towards an Overlay Resource Management Infrastructure that can coordinate peering relationship across independent application/Internet/content/location/wireless service providers (xSPs) to maintain seamless connectivity when users cross network boundaries, and wide-area service mobility and transformation by coupling/decoupling functionalities offered by different xSPs.

3.3. Statement of Work

Based on this research plan outline, we propose to:

- Develop architecture and tools for marketplace-driven resource allocation, in particular, spectrum resources, backhaul bandwidth, on as fine a grain basis as is practical. A critical issue is convergence: in a highly dynamic environment, can an efficient allocation method be determined? How long should such “leases” be retained?
- Develop architecture and tools for the dynamic marshalling, observation/verification of participant behaviors, and ultimate dissolution of confederations. Enable “Mom and Pop” service operators. Management of trust relationships, both static (e.g., AT&T and BT trust each other) and dynamic (e.g., Operator X did not work in a trustworthy fashion last time, and I won’t form a confederation with them again).
- Develop new handoff concepts: hand over to alternative service provider in midstream of activity in order to balance load/achieve lower cost access.
- Develop mechanisms to “audit” third party resource allocation agents, to insure fairness and freedom from bias in their operation. It is important to determine if it is possible to detect a “misbehaving” service provider.

As part of the proposed work, we are interested in constructing several testbeds at a variety of scales to allow us to investigate whether a generalized architecture will emerge. In particular, we proposed to develop testbeds at the following levels:

- *Room-scale*: The room-scale testbed will make use of personal-area network (PAN) technologies, such as Bluetooth, to investigate how devices can work as ensembles, cooperatively sharing bandwidth within a microcell. In such environments, there is inherent trust, but we seek to understand how finer grained intelligent and active allocation of resources can be more effective, as opposed to the existing kinds of static etiquette rules that exist for these networks today. The essential question is just how lightweight such an architecture can be. Is it too heavyweight for Bluetooth style devices?
- *Building-scale*: On the Berkeley campus, we have deployed multiple alternative Wireless LAN technology within the buildings that house the Electrical Engineering and Computer Science Department, namely Soda and Cory Halls. This allows us to investigate a service provisioning environment of multiple independent wireless LAN “operators” (e.g., base station by base station, or floor by floor, with areas of overlapping access). Within this testbed, we can also experiment with “misbehaving operators” in order to develop third party audit mechanisms to determine offending participants. Ericsson’s *GSM on the Net* (GoN) technology will offer us a very attractive alternative for telephony services, in an environment that will allow for dynamic allocation of frequencies and/or time slots to competing and confederating service providers within a building.
- *Campus-scale*: The Berkeley campus is currently in the planning stages for a campus-wide wireless LAN infrastructure. Furthermore, the Ericsson Bay Area Research Center will have early access to the latest data over cellular capabilities of the GSM infrastructure. Using these alternative networks will allow us to explore cooperative provisioning of services, such as among Departmental WLAN service providers with overlapping coverage out of doors, e.g., Mechanical Engineering and Computer Science, as well as the possibility of forming ensembles among Internet-capable cell phones supported by a wide-area provider and in-building services such as WLAN or GoN.

- *Regional-scale*: Possibilities exist to work with wireless operators, such as Sprint, to obtain measurements, experiment with new services and their resource demands, and generally gain experience of the opportunities and constraints in the wide-area.

4. Relevance to MICRO

The United States currently leads the world in innovative information technologies. But the U.S. appears to be falling behind Europe and Japan in terms of advanced wireless telecommunications infrastructure, integration of wireless networks and the Internet, and identification of compelling new services that can be deployed within such integrated networks. We believe that the next generation of wireless systems is at risk because the existing architectures are ill suited to deploying a diverse collection of advanced services in a way that broadly supports a diversity of business models.

The research we have proposed focuses on a radical architectural approach for supporting services in future telecommunications systems. A key contribution will be dynamic mechanisms for service provision in the local and wide areas. If successful, this work could spawn whole new ways of organizing advanced telecommunications systems. This could help generate additional jobs in the telecommunications and software industries.

The Ericsson Bay Area Research Center (BARC), HRL Laboratories, Nortel Networks Architecture Laboratory, and the Sprint Advanced Technology Laboratories will work with us on the proposed research program. These are world leading companies in telecommunication equipment, services, and network operation.

Ericsson is a worldwide telecommunications equipment developer, with a market leading position in wireless networking infrastructure. They have recently established a research center in Berkeley, CA. We will work closely with Dr. Gunnar Nilsson on this project:

Technical Point of Contact: Dr. Gunnar Nilsson
Ericsson Bay Area Research Center
2100 Shattuck Avenue
Berkeley, CA 94704
Gunnar.C.Nilsson@era.ericsson.se
510-666-3900

HRL Laboratories is a major employer in the State of California and a long time supporter of the California MICRO program. HRL is a corporate research laboratory co-owned by Boeing, General Motors, and Raytheon. These firms retain major data communications and research facilities located throughout Southern California. We will work closely with Dr. Son Dao of HRL Laboratories.

Technical Point of Contact: Dr. Son Dao
HRL Laboratories
3011 Malibu Canyon Rd.
Malibu, CA 90265
310-317-5682
skdao@hrl.com

Nortel Networks is a leading developer of network equipment with a worldwide business. They are a major employer in the State of California, with major facilities in Santa Clara. We will work closely with Dr. Dan Pitt of Nortel Network’s Technology Laboratory.

Technical Point of Contact: Dr. Dan Pitt

Director, Technology Laboratory
Great America Parkway
Santa Clara, CA
408-495-7096
dpitt@nortelnetworks.com

Sprint is a major telecommunications network operator, spanning telephony, wireless Internet, and enhanced network services like voice recognition. They have a well established Advanced Technology Laboratory in Burlingame, CA. We will work closely with Dr. Bryan Lyles on this project:

Technical Point of Contact: Dr. Bryan Lyles
Sprint Advanced Technology Laboratory
One Adrian Court
Burlingame, CA 94010
lyles@sprintlabs.com
650-375-4200

References

[3GPP] http://www.3gpp.org/3G_Specs/3G_Specs.htm.

[Bi 2001] Q. Bi, G. Zysman, H. Menkes, "Wireless Mobile Communications at the Start of the 21st Century," IEEE Communication Magazine, V. 39, N. 1, (January 2001), pp. 110-117.

[Blake 1998] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, W. Weise, "An architecture for differentiated services," Internet RFC 2475, IETF Network Working Group, (December 1998).

[Braden 1997] R. Braden, L. Zhang, S. Berson, S. Herzog and S. Jamin, "ReSerVation Protocol (RSVP) version 1 functional specification" Internet RFC 2205, IETF Network Working Group, (September 1997).

[Brewer 1998] E. Brewer, R. H. Katz, E. Amir, H. Balakrishnan, Y. Chawathe, A. Fox, S. Gribble, T. Hodes, G. Nguyen, V. Padmanabhan, M. Stemm, S. Seshan, T. Henderson, "A Network Architecture for Heterogeneous Mobile Computing," *IEEE Personal Communications Magazine*, (October 1998), pp. 8-24.

[Chuah 2000] C. Chuah, L. Subramanian, A. D. Joseph, R. H. Katz, "QoS Provisioning Using A Clearing House Architecture," 8th International Workshop on Quality of Service (IWQOS 2000), Pittsburgh, PA, (June 2000).

[Economist 2000] "Telecoms in Trouble," *The Economist Magazine*, 16 December 2000, pp. 77-79.

[Kohler 2000] E. Kohler, R. Morris, B. Chen, J. Jannotti and M. F. Kaashoek, "The Click Modular Router," *ACM Transactions on Computer Systems*, vol. 18, no. 4, (November 2000).

[Mohr 2000] W. Mohr, W. Konhauser, "Access Network Evolution Beyond Third Generation Mobile Communications," *IEEE Communications Magazine*, (December 2000), pp. 122-133.

[New Scientist 2001] "Not in my backyard: Without tens of thousands of new masts 3G cell phones won't happen," New Scientist Magazine, (10 February 2001), p. 6.

[Oftel 1999] http://www.oftel.gov.uk/competition/mvno0699.htm.

[Raman 2000] B. Raman, R. H. Katz, A. Joseph, "Providing Extensible Personal Mobility and Service Mobility in an Integrated Communication Network," 3rd IEEE Workshop on Mobile Computing Systems and Applications (WMCSA2000), Monterey, CA, (December 2000).

[Roberts 2001] J. W. Roberts, "Traffic Theory and the Internet," *IEEE Communications Magazine*, V. 39, N. 1, (January 2001), pp. 94-99.

[Stemm 1999] M. Stemm, R. H. Katz, "Vertical Handoffs in Wireless Overlay Networks," *ACM/Baltzer Mobile Networking and Applications (MONET)*, Special Issue on "Mobile Networking in the Internet," V 3, N 4, (January 1999), pp. 319-334.

[UMTS XX] UMTS Forum, "Report #2: The Path Towards UMTS: Technologies for the Information Society," available at http://www.umts-forum.org/reports.html.

[Wang 1999] H. J. Wang, R. H. Katz, J. Giese, "Policy-Enabled Handoffs Across Heterogenous Wireless Networks," 2nd IEEE Workshop on Mobile Computing Systems and Applications (WMCSA '99), New Orleans, LA, (Feb. 1999), pp. 51-60.ùwÿ

[Wang 2000] H. Wang, B. Raman, C. Chuah, R. Biswas, R, Gummadi, B. Hohlt, X. Hong, E. Kiciman, Z. Mao, J. Shih, L. Subramanian, B. Zhao, A. Joseph, R. H. Katz, "ICEBERG: An Internet Core Network Architecture for Integrated Communications," *IEEE Personal Communications Magazine*, V. 7, N. 4, (August 2000), pp. 10-19.

F. Personnel

Randy H. Katz, Principal Investigator

Professor Randy H. Katz received an A.B. degree in computer science and mathematics from Cornell University in 1976. He received his M.S. and Ph.D. degrees in computer science from Berkeley in 1978 and 1980 respectively. After positions in industry and at the University of Wisconsin-Madison, he joined the Berkeley faculty in 1983, where he is now the United Microelectronics Corporation Distinguished Professor.

Professor Katz was an NSF Presidential Young Investigator from 1984-1989, received an IBM Faculty Development Award in 1984, won the Berkeley Distinguished Teaching Award in 1992, received the CRA Distinguished Service Award in 1995, shared the IEEE Reynolds Johnson Information Storage Award in 1999, won the ASEE Frederick Emmon Terman Award in 1999, and was given the ACM Karl V. Karlstrom Outstanding Educator Award in 2000. He is a Fellow of the IEEE and the ACM, and a member of the National Academy of Engineering. He has written over 220 technical publications on CAD, database management, multiprocessor architectures, high performance storage systems, video server architectures, mobile wireless networking, and computer network-based telecommunications service environments. He has supervised 34 masters theses and 19 Ph.D. dissertations. He led the implementation of the SPUR multiprocessor memory system, developed the RAID concept, organized a major wireless overlay testbed in the San Francisco Bay Area, and has developed a new service architecture for computer-telephony integration. His recent research has focused on Internet services. From January 1993 through December 1994, he was a program manager and deputy director of the Computing Systems Technology Office of ARPA. His relevant publications include:

C. Chuah, L. Subramanian, A. D. Joseph, R. H. Katz, "QoS Provisioning Using A Clearing House Architecture," 8th International Workshop on Quality of Service (IWQOS 2000), Pittsburgh, PA, (June 2000).

B. Raman, H. Wang, J. Shih, R. H. Katz, A. Joseph, "The ICEBERG Project: Defining the IP and Telecomm Intersection," *IT Professional*, V 1, N 2, (November/December 1999), pp. 22-29.

B. Raman, R. H. Katz, A. Joseph, "Providing Extensible Personal Mobility and Service Mobility in an Integrated Communication Network," 3rd IEEE Workshop on Mobile Computing Systems and Applications (WMCSA2000), Monterey, CA, (December 2000).

H. Wang, R. H. Katz, A. D. Joseph, "A Signaling System Using Light-weight Call Sessions," IEEE Infocomm 2000 Conference, Tel Aviv, Israel, (March 2000).

H. Wang, B. Raman, C. Chuah, R. Biswas, R, Gummadi, B. Hohlt, X. Hong, E. Kiciman, Z. Mao, J. Shih, L. Subramanian, B. Zhao, A. Joseph, R. H. Katz, "ICEBERG: An Internet Core Network Architecture for Integrated Communications," *IEEE Personal Communications Magazine*, V. 7, N. 4, (August 2000), pp. 10-19.

H. Wang, R. H. Katz, A. Joseph, "A Personal Communication Service Creation Model for Integrated Communication Systems on the Internet," Proc. IEEE International Conference on Communications (ICC2001), Helsinki, Finland, (June 2001).

Anthony D. Joseph Co-Principal Investigator.

Professor Anthony D. Joseph received his B.S., S.M., and Ph.D. degree in computer science from Massachusetts Institute of Technology in 1988, 1988, and 1998 respectively. At MIT, he was supported by IBM and Intel Graduate Fellowships, and participated as a research assistant in the Parallel and Distributed Operating Systems Group. He joined the UC Berkeley faculty in 1998 as an assistant professor. His graduate work focused on application support for intermittent connectivity in a mobile computing environment. At MIT, he was one of the principal designers of the Rover mobile toolkit for mobile computing and he helped develop the Proteus parallel simulator and Prelude language. Professor Joseph was awarded a National Science Foundation CAREER Award in 2000, received an Okawa Foundation Research Grant in Telecommunications and Information Processing in 1999, and an IBM Faculty Development Award in 1999. While at MIT, he was awarded an Intel Foundation Graduate Fellowship in 1995 and an IBM Corporation Graduate Fellowship from 1987 to 1990. At Berkeley, Professor Joseph's research focuses on the integration of heterogeneous, multi-provider mobile telephony systems and IP-based networks, the development of a scalable, reliable, distributed computation environment, the creation of new models for analyzing and simulating wireless network channels, and the continuation of research on the tools and techniques embodied in the Rover toolkit, including service discovery, location-based services, and computation and data migration. He has supervised 5 master theses and is currently supervising 5 Ph.D. candidates. He led the design and implementation of the first GSM basestation to be directly connected to the Internet.

His relevant publications include:

Gribble, S., Welsh, M., von Behren, R., Brewer, E., Culler, D., Borisov, N., Czerwinski, S., Gummadi, R., Hill, J., Joseph, A., Katz, R, Mao, Z. M., Ross, S., Zhao, B., “The Ninja Architecture for Robust Internet-Scale Systems and Services”, to appear in a Special Issue of *Computer Networks* on Pervasive Computing. Invited paper.

Raman, B., Katz, R., Joseph, “A. Universal Inbox: Providing Extensible Personal Mobility and Service Mobility in an Integrated Communication Network”, In the Proceedings of the 3rd IEEE Workshop on Mobile Computing Systems and Applications, (December 2000).

Ross, S., Hill, J., Chen, M., Joseph, A., Culler, D., Brewer, E., “A Composable Framework for Secure Multi-Modal Access to Internet Services from Post-PC Devices,”in the Proceedings of the 3rd IEEE Workshop on Mobile Computing Systems and Applications, (December 2000).

Wang, H., Raman, B., Chuah, C-N., Biswas, R., Gummadi, R., Hohlt, B., Hong, X., Kiciman, E., Mao, Z., Shih, J., Subramanian, L., Zhao, B., Joseph, A., Katz, R., “ICEBERG:An Internet-core Network Architecture for Integrated Communications,” *IEEE Personal Communications* (2000): Special Issue on IP-based Mobile Telecommunication Networks, August 2000. Invited paper.

von Behren, J., Czerwinski, S., Joseph, A., Brewer, E., Kubiatowicz, J., “NinjaMail: The Design of a High-Performance Clustered, Distributed E-Mail System,” in the Proceedings of the International Workshops on Parallel Processing 2000. August 21-24, 2000, Toronto, Canada. (Edited by P. Sadayappan). pp. 151-158.

Chuah, C-N., Subramanian, L., Katz, R., Joseph, A., “QoS Provisioning Using A Clearing House Architecture,” International Workshop on Quality of Service (IWQoS), Pittsburgh, PA, pp. 115-124, (June 5-7, 2000).

Wang, H., Joseph, A., Katz, R., “A Signaling System Using Lightweight Call Sessions,” Proceedings of IEEE INFOCOM 2000, Tel-Aviv, Israel. (March 2000).

Raman, B., Wang, H., Shih, J., Joseph, A., Katz, R., “The Iceberg Project: Defining the IP and Telecom Intersection,” *IT Professional,* (Nov/Dec 1999). Invited article.

Czerwinski, S., Zhao, B., Hodes, T., Joseph, A., Katz, R., “An Architecture for a Secure Service Discovery Service,” Fifth Annual International Conference on Mobile Computing and Networks (MobiCom ‘99), Seattle, Washington, (August 1999), pp. 24-35. To appear in *ACM Wireless Networks* (WINET).

Ludwig, R., Konrad, A., Joseph, A., “Optimizing the End-to-End Performance of Reliable Flows over Wireless Links,” Fifth Annual International Conference on Mobile Computing and Networks (MobiCom '99), (August 1999), Seattle Washington. To Appear in *ACM Wireless Networks* (WINET).

Ludwig R., Rathonyi B., Konrad A., Oden K., Joseph A., “Multi-layer Tracing of TCP over a Reliable Wireless Link,” Proceedings of ACM SIGMETRICS 1999, (May 1999).

Joseph, A., Katz, R., Hohlt, B., Kiciman, E., Wang, H., Ludwig, R., “System Support for Multi-Modal Information Access and Device Control, Presented as a Works-In-Progress talk at the Second IEEE Workshop on Mobile Computing Systems and Applications (WMCSA’99), February 1999.

Joseph, A., Badrinath, B., Katz, R., “A Case for Services over Cascaded Networks,” Proceedings of First ACM/IEEE International Conference on Wireless and Mobile Multimedia (WoWMoM’98), (October 1998).

G. Current and Pending Support

H. Letters of Intent

I. Budget

J. Explanation and Justification of Budget Items

K. List of Reviewers

Name: Professor Mario Gerla
Address: Department of Computer Science
3732F Boelter Hall
UCLA
Los Angeles, CA 90095-1596
Reviewer's fields of specialization: Networking, Internet Technology
Tel: 310-825-4367 FAX: 310-825-2273 E-mail: gerla@cs.ucla.edu

Name: Professor J. J. Garcia-Luna-Aceves
Address: Computer Engineering
University of California, Santa Cruz
Santa Cruz, CA 95064
Reviewer's fields of specialization: Networking, Ad Hoc and Wireless Routing
Tel: 408-459-4153 FAX: 408-459-4829 E-mail: jj@cse.ucsc.edu

Name: Professor Mary Baker
Address: Computer Science Department
Gates Hall
Stanford University
Stanford, CA 94305-2140
Reviewer's fields of specialization: Networking, Operating Systems, Mobile Computing
Tel: 415-725-3711 FAX: 415-725-7411 E-mail: mgbaker@cs.stanford.edu

Name: Dr. Hamid Ahmadi
Address: AT&T Labs Research
180 Park Avenue, P.O. Box 971
Florham Park, NJ 07932-0971
Reviewer's fields of specialization: Networking, Mobile Computing
Tel: 973-360-8200 FAX: 973-360-8853 E-mail: ahmadi@research.att.com

Name: Dr. Jennifer Rexford
Address: AT&T Labs Research
180 Park Avenue, P.O. Box 971
Florham Park, NJ 07932-0000
Reviewer's fields of specialization: Internet Technology, Networking, Distributed Systems
Tel: 973-360-8728 FAX: 973-360-7011 E-mail: jrex@research.att.com