Collaborative Research: Researching, Prototyping and Deploying 7DS

1 Introduction

For at least the next decade, we foresee two classes of packet communication networks existing side-by-side and dominating the data communications scene:

- The **broadband Internet**, consisting of abundant fiber bandwidth in the core coupled with broadband access technologies, including Ethernet, cable-modem, DSL and 802.11. The broadband Internet is largely a wired packet-switched network, although it includes spotty wireless broadband access at 802.11 access points.
- Second- and third-generation **cellular data networks**, and transitional packet networks such as the 2.5G data extensions (GPRS and EDGE). The cellular data networks provide nearly ubiquitous access to the Internet in major population centers and along interstate highways, although this access is costly and low speed.

The broadband Internet offers high-speed, unmetered access to general Internet services, including information access (web and media-on-demand), asynchronous messaging (e-mail), and synchronous services including voice-over-IP and instant messaging. Access to the broadband Internet will be readily available at home, in enterprises, and – via hotspots with access points (APs) – in high-traffic areas like airports, upscale hotels, cafes and convention centers. Some outdoor areas such as city parks may also be covered by APs. However, given the limited range of these APs, coverage will remain spotty, with users having to seek out hotspots. (For example, the island of Manhattan, with an area of 60 km^2, would require approximately 60,000 APs for full coverage, optimistically assuming that 802.11 coverage extends to more than 100 m in dense urban environments and discounting the need for indoor coverage beyond the first and second floor.)

The cellular data networks offer complementary strengths and weaknesses. They provide nearly ubiquitous access, at least in densely populated areas. However, their effective access speed is roughly equivalent to dial-up modems, with costs of currently several dollars per Megabyte. In addition, adding new cellular capacity to an area is a lengthy and difficult process, so that, for example, supporting ad-hoc gatherings of people with sufficient capacity is difficult.

In an ideal world, mobile users would be able to enjoy affordable broadband Internet access everywhere, including in public transportation vehicles (buses, trains, planes), in outdoor spaces not covered by 802.11 APs, and in disaster areas where the broadband wired infrastructure may be broken. But unfortunately, there does not seem to be such a deployable technology on the horizon. The two preponderant packet communication networks – broadband Internet and cellular data – even when taken together will fail to satisfy all our communication needs.

We propose to research, prototype and deploy **7DS**, *a new network modality that helps bridge the gap between the broadband Internet and the cellular data networks. 7DS bridges the gap by exploiting an increasing density of roaming devices that are endowed with both gigabyte storage and short-range, high-speed wireless communication capabilities.*

The basic idea of 7DS is as follows. Consider a device with an 802.11 or BlueTooth network interface that is currently *not* within range of an AP:

- If the device wants a data object (e.g., the front page of a specific newspaper, a map, or the current weather forecast), it may be able to obtain the object at high speed either immediately from another device within range, or in the near future from a device that roams within range. It can also use keyword searches to discover information of interest on other devices within range.

- Conversely, if the device wants to send an object into the broadband Internet (for example, e-mailing a picture taken with the device's digital camera), it can pass the object to multiple nearby roaming devices, with the expectation that at least one of the roaming devices will become within range of an AP in the near future.

Thus, 7DS[1] is a message-based application-layer service that propagates information objects among mobile nodes that are within high-speed wireless range of each other. 7DS scales with the user population in a given area and can be used instantaneously, without infrastructure planning or deployment. In many ways, 7DS is a pragmatic instantiation of the theoretical work on "mobility increases capacity" [GT02].

We envision 7DS as a component of a *multi-modal* network architecture, i.e., a network that uses many different modalities to transport data to and from users.[2] One mode is the broadband Internet, another mode is cellular data, and yet a third mode is 7DS. Just as trucking fills a niche in vehicle transport, 7DS fills a niche by combining low-cost high-peak bandwidth with an expansive coverage area. Table 1 shows that 7DS can leverage the low cost of traditional wired infrastructure with the ubiquity of much more expensive wireless communication modes. As we will discuss in greater detail in Section 3, combining cellular data networks with 7DS further enhances the performance of 7DS.

Transport	packet/circuit	speed	$/MB (= 1/3 MP3)
OC-3	P	155 Mb/s	$0.0013
Australian DSL	P	512/128 kb/s	$0.018
GSM voice	C	8 kb/s	$0.66-$1.70
HSCSD	C	20 kb/s	$2.06
GPRS	P	25 kb/s	$4-$10
Iridium	C	10 kb/s	$20
SMS (160 char/msg)	P	0.39 kb/s	$62.50
Motient (Blackberry)	P	8 kb/s	$133

Table 1: Bandwidth and data transport cost for different transport modalities

7DS also leverages the dramatic decrease in storage cost for mobile devices. CompactFlash memory is now available at about $250/GB, which is roughly the cost of disk storage in 1996. The decrease in storage cost has been markedly larger than the decrease in wide-area wireless communications costs: Average per-minute cost for cellular calls decreased from 56 to 11 cents from 1995 to 2002, while disk storage costs decreased by a factor of about 350. At current cellular data prices, a PDA with a 1 GB compact flash memory card costing $250 can hold between $1,000 and $4,000 worth of cellular data.

1.1 Applications of 7DS

We envisage many important applications for 7DS, including:

Vehicle transportation: Passengers in a bus, train or plane can share information, as well as "drop off" messages and data objects that are transmitted to the Internet as the vehicle passes by a wireless hotspot.

---

[1] The acronym stands for "seven degrees of separation," indicating its slightly more indirect connection to the Internet.

[2] "The term multimodal transport is often used loosely and interchangeably with the term intermodal transport. Both refer to the transport of goods through several modes of transport from origin to destination." (United Nations)

Wireless web caching: Before retrieving content from a costly, low-speed cellular data system, users can query nearby 7DS devices; similarly, before sending an e-mail photo into a low-speed cellular data system, users can rapidly pass the content to mobile 7DS nodes.

Emergency infrastructure: During natural and man-made disasters, Internet access may be unavailable or overloaded. Under those circumstances, 7DS can facilitate data exchange (e.g., photos and maps) between emergency workers or can be used to propagate "I'm alive" [TIK+00] messages. Beyond emergencies, large groups of people congregating in one place may strain available communication resources that cannot be efficiently designed for rare load peaks. 7DS can provide additional high-bandwidth channels for large objects that are delay-tolerant.

Wireless peer-to-peer object sharing: Content such as class notes in lectures, maps, public web pages, bus and train schedules, weather reports, tourist information and MP3s can be shared among users.

Mobile sensor networks: In sensor networks with mobile nodes that are too far apart to form an ad-hoc network, moving sensor nodes could serve as data carriers. For example, one could imagine attaching sensors to roaming wildlife, where data gathering stations are placed at watering holes. Vehicles can collect and propagate information about traffic conditions and road hazards to other vehicles passing in the opposite direction. Similarly, tourists with digital cameras can upload their pictures to the Internet before they reach a network access point.

1.2 Relationship to Ad-Hoc Networks

Although we believe that current research in ad-hoc networking is of significant importance, we emphasize that our proposed 7DS research is not traditional ad hoc networking research. In traditional ad-hoc networking, when two nodes want to communicate, there needs to be a connected path between the nodes. Much of the research in ad-hoc networking is about network maintenance and end-to-end packet routing.

7DS shares some properties with ad-hoc networks. For example, it does not require network infrastructure or a planned topology. However, 7DS uses application-layer messaging rather than packet routing. Furthermore, rather than relying on existing end-to-end paths, 7DS exploits mobility and nodal storage capabilities to transport data objects. As will become evident in the body of the proposal, network maintenance and routing do not play important roles in 7DS research.

The applicability of ad-hoc networks are also limited to relatively dense configurations. For the Manhattan example cited above, at least 60,000 nodes would have to participate in the network to ensure full connectivity. 7DS operates in networks with any node density; performance degrades proportionally as the node density decreases. In [PS00], we showed that even with only 10 walking-speed mobile nodes per km^2, most users can obtain information in about 25 minutes. And while most ad-hoc networks are designed for networks of modest size of maybe a few hundred nodes, 7DS networks can scale to effectively any size since they do not need to maintain or propagate system state. (7DS networks are much less efficient than well-connected ad-hoc networks, so they are not replacements for these.)

1.3 Relationship to Peer-to-Peer Networks

One – but not the only – application of 7DS is wireless P2P object sharing, in which mobile nodes share class notes, maps, public web pages, bus and train schedules, MP3s, etc. A 7DS node indeed acts as both a client and a server; there are no "distinguished" participants – all nodes are equivalent. (As discussed below, some nodes may have special capabilities, such as a connection to a cellular data network, but the protocol behavior remains the same.)

However, the range of applicability of 7DS extends far beyond P2P object sharing, including asynchronous messaging, emergency networking, and mobile sensor networks. Furthermore, even within the context of peer-to-peer object sharing, the central research issues in 7DS are distinctly different from those in wired P2P object sharing [RR]. In wired P2P, all participating nodes are within range of each other, since the wired infrastructure can provide a path between any pair of nodes. However, in 7DS, the nodes move in and out of range of each other; because a desired object may not be immediately retrievable, a node may have to repeatedly query over a period of hours before locating another 7DS node with the object. Furthermore, the recent research in distributed hash tables (DHTs) [SMK$^+$01, RFH$^+$01, RD01, ZKJ01, RBF$^+$02, GEBF$^+$03] assume relatively static node relationships and is thus not particularly relevant to 7DS.

1.4 Summary of Proposed Research

Papadopouli and Schulzrinne developed the initial ideas for 7DS [PS00, PS01b, PS01a, PS02, Pap02]. The initial research has focused on dissemination of web pages that a mobile device downloads from the Internet and that other 7DS nodes obtain by querying. As part of this initial research, we have developed simulation and analytical models that predict the performance of such systems for random walk and other mobility patterns. We have also developed a preliminary prototype for data dissemination which runs in Linux machines. The initial results are extremely promising.

In parallel to the 7DS work, at Polytechnic we have developed an Infostation prototype, which is a node that wirelessly disseminates useful content as users walk by [FG03]. The initial 7DS and Infostation prototypes are highly complementary.

We now propose to leverage this initial work to (i) study new algorithms and protocols for significantly extending the range of applicability of 7DS, (ii) develop a comprehensive prototype for 7DS which builds on our initial 7DS and Infostation prototypes, (iii) actually deploy 7DS in a number of real-world environments. For the new algorithms and protocols, we propose the following:

- **Research algorithms and protocols for propagating content from an 7DS device through other mobile 7DS devices to the Internet.** The content is typically encrypted for privacy. The algorithm requirements for device-to-Internet transfer are very different from the dissemination protocols studied in earlier work. The sender-initiated content is often only of interest to the sender 7DS node. Thus, while dissemination relies on spreading popular information widely, upload tries to limit the number of replicas to the minimum number that is sufficient to reach an Internet-connected node.

- **Research incentive mechanisms to ensure that 7DS nodes actively participate in both caching and 7DS-to-Internet content transfer.** It is widely documented that the popular P2P file sharing systems are havens for “free riders”: a significant fraction of users do not contribute any resources, and a minute fraction of users contribute the majority of the resources [SGG02, AH00]. The need for incentives is even larger for 7DS, since users contribute precious battery power by responding to queries and propagating content.

- **Research pseudo-reliable multicast for content dissemination.** The current prototype of 7DS multicasts queries, but uses unicast to transfer content via HTTP. This has the advantage of simplicity, but non-querier “bystanders” can get the same content if we use a form of reliable multicast instead of unicast. For this research, we’ll build on our Infostation research, which has begun to explore reliable multicast for a wireless LAN content dissemination system [FG03].

- **Research in distributed cache management for data dissemination.** In earlier work, we used a simple scheme for managing the content in a 7DS community (e.g., the 7DS devices in a bus or in a campus building). Specifically, we assumed that only those devices interested in a particular data object participate in storing the object. In order to improve availability and acquisition time, we shall build on our work in distributed caching and develop new algorithms that aggressively replicate popular content [Ros97, KRT03].

- **Research multi-modal network architectures.** As noted in the introduction, we envision a future where nearly ubiquitous cellular data networks coexist with roaming "broadband data carriers" that occasionally connect to the Internet via hotspots. Given the huge disparity in access rates and costs, it makes sense to leverage both.

We describe the proposed research and prototype in greater detail in the subsequent sections of the proposal.

1.5 Research Team

The research team will consist of three PI's – Henning Schulzrinne of Columbia University and Phyllis Frankl and Keith Ross of Polytechnic University – a postdoc, one Ph.D., and several master's and undergraduate students. Professor Schulzrinne has made significant contributions to Internet research and protocols, including the lead author of RTP, RTCP, RTSP, and SIP, all of which are enjoying extensive implementation and deployment in industrial products. Professor Schulzrinne also supervised the initial work on 7DS algorithms and prototypes. Professor Frankl is the director of Wireless Internet Center for Advance Telecommunications (WICAT) at Polytechnic University. Over the past few years, she has been investigating wireless Infostations and the trade-offs in using unicast versus multicast data distribution. Professor Ross is co-author of the leading textbook on computer networking, and has been active researcher in distributed caching and P2P networking for several years. He also co-founded the VoIP messaging software start-up, Wimba, and led Wimba's research and development in its first two years.

The three PI's, all affiliated with universities in New York City, have complementary knowledge and skill sets. They will be working closely together, holding bi-weekly group meetings with the postdoc and students. Furthermore, New York City provides an ideal environment for deployment and experimentation of 7DS prototypes, due to its large pedestrian density, the presence of all kinds of public transportation, a large fleet of taxis cruising the avenues, and delivery trucks making their rounds.

2 Initial Work

2.1 Initial Work on 7DS

The proposal significantly extends the earlier work on 7DS that showed the usefulness of mobility-based data distribution for data dissemination. Our preliminary prototype consists of mobile nodes that respond to multicast queries by wireless peers. If a mobile node finds suitable content, it advertises its availability to the querying peer. The peer, in turn, uses a regular HTTP query to obtain the content.

Using a random waypoint model, we simulated the data dissemination success ratio and the delay for successful data retrievals. A sample result is shown in Fig. 1, showing the probability that a host has acquired a data object after 25 minutes as a function of the host density , assuming that only one mobile host possessed the object at the beginning of the simulation[3]. The average delay is approximately 10 minutes. For the host densities studied, forwarding queries to other nodes that are hidden from the original querier did not improve performance substantially[4].

---

[3] Note that the host density only includes hosts that are interested in that particular data object, not all mobile hosts.

[4] We have not yet studied the impact of creating small ad-hoc networks for such queries where node density allows this.

Since mobile hosts are energy-constrained and since 802.11 devices consume substantial power even when in receive mode [FN01], we evaluated the impact of a number of simple energy conservation measures [PS01a], such as only enabling the wireless interface at roughly synchronized periodic intervals, and found that such techniques affected data dissemination performance only marginally.

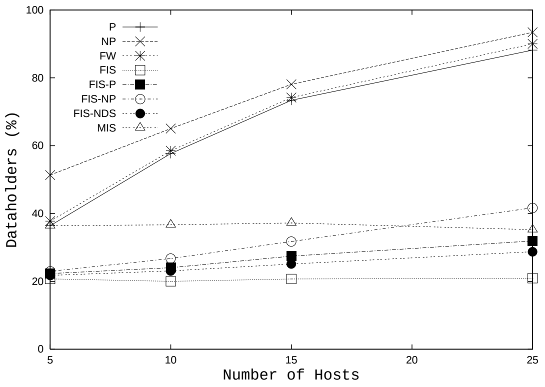

Figure 1: Number of data carriers after 25 minutes

We have analyzed simplified versions of data dissemination using epidemic models and diffusion-controlled processes.

We have built an early Java-based prototype for Linux that demonstrates the web data delivery application. Fig. 2 shows its architecture. It is designed so that it can distribute any content type, using standard web browsers for rendering. The 7DS system acts like a proxy cache for the mobile node, enhanced by a local search engine. The search engine indexes the cache, allowing not just queries identifying a URL, but also queries by content. These queries increase the chances that a peer will find suitable content without having to know details about its original URL. In addition, in many cases, peers care about particular content, say, a news story on baseball, not whether it was published by the *New York Times*, the *New York Post* or CNN.

2.2 Initial Work on Infostations

As part of the proposed research, we'll be integrating Infostations and pseudo-reliable multicast with 7DS, both in the basic research and in prototyping. The Infostation concept was originally proposed at WINLAB at Rutgers University. In its essence, an Infostation is a node that wirelessly disseminates useful content as users pass by, with little or no human interaction. While the Rutgers work focuses primarily on the data link layer and below, the Polytechnic work focuses on applications and transport layer mechanisms for Infostations. The Polytechnic prototype is a "walk-through" Infostation, where users can be expected to spend 20 seconds to a few minutes in the coverage area [FG03].

Polytechnic's work on Infostations blends nicely with Columbia's work on 7DS for two reasons. First, a "walk through" Infostation is a natural component in a comprehensive 7DS architecture. Indeed, a 7DS node absorbs fresh content as its passes by an Infostation; it can then disseminate the fresh content when it roams in the vicinity of stranded 7DS nodes. Second, as part of the Infostation research, researchers at Polytechnic have begun to investigate reliable multicast in a wireless LAN environment, both at the application and the transport layers. This wireless reliable multicast research has a direct bearing on 7DS research.

Polytechnic's walk-through Infostation prototype currently uses a data carousel to repeatedly broadcast data items. The Infostation software is written in C++ running on Windows 98 and Windows CE

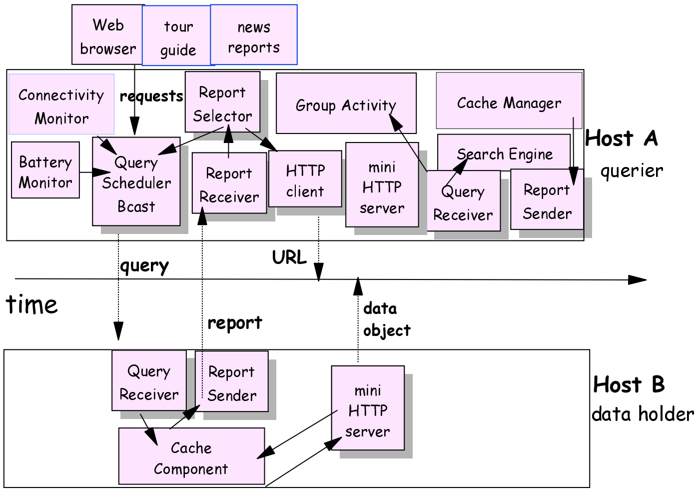

Figure 2: Architecture of initial 7DS prototype

PDA. The server and client software each consist of three main components: a manager, a collection of encoder/decoders, and application programs. We have developed both a course notes delivery/updating application and a music delivery application.

3 Fundamental 7DS Research

3.1 A Taxonomy of Storge-Enhanced Wireless Communications

This proposal explores a new class of storage-enhanced wireless devices. Fig. 3 illustrates the models that we are considering, filling in combinations of mobility, data direction and network backing.

Connected Infostation: A wireless hotspot, in a fixed location, is backed by an object cache that allows the use of low-speed access links to the wired network while still delivering large content objects. With no or little human interaction, mobile 7DS nodes absorb data objects as they pass through the Infostation's coverage area.

Disconnected Infostation: Here, the Infostation is "primed" with content via some non-network means, e.g., a flash memory module or an explicit "refill" operation by a mobile node.

Access inverse multiplexing: In access sharing, multiple nodes with low-speed wireless interfaces cooperate to retrieve large data objects faster.

7DS: In 7DS, mobile nodes exchange data with each other, occasionally connecting to the wired Internet and possibly using a low-speed wide-area wireless network for control functionality.

All modalities have in common that data transfer takes place without direct human intervention. Rather, it is based on profiles, cache management and distribution policies and earlier user access requests.

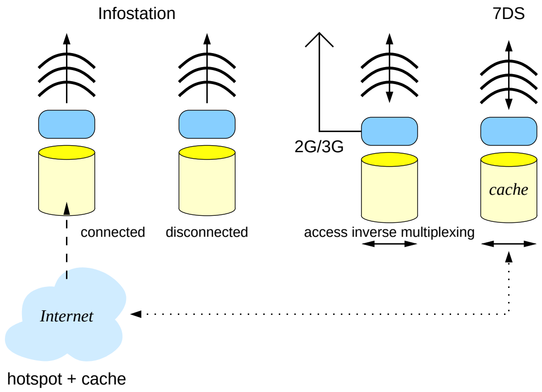

Figure 3: Network modalities considered in this proposal

3.2 7DS-to-Internet

One new mode of operation is the delivery of data from mobile nodes to a small set of Internet-connected nodes. (Applications include all kinds of sensor data, including images collected by still or motion-video cameras, weather, traffic, military and scientific data.) The mobile nodes generally do not know the identity and location of these connected nodes, so traditional or geographic [NI97, HBH00, JPS01, LJC^{+}00] routing approaches are not likely to be effective. One approach, explored in [VB00], is "leaky flooding", i.e., replicating data until it exceeds a certain hop count or is pushed out of the cache. We anticipate that we can significantly improve the performance by taking into account knowledge about the direction of travel of nodes, even if only approximate. Nodes may know their direction of travel via geolocation techniques, but a simple electronic compass is likely to suffice. Location beacons [BSSW03] offer additional directional hints. Given movement prediction, we can then distribute data to a maximally diverse set, favoring directions that have not yet been explored.

We plan to explore a prototype where city buses, postal vehicles or taxis are equipped with 7DS nodes. These "wireless mailboxes" pick up messages along their appointed rounds and then empty their mailbox into an access point at the bus depot or garage. Often, their arrival time at the access point can be predicted with reasonable certainty, so that mobile nodes can choose only carriers that arrive earlier than the ones they have already entrusted their data with.

3.3 Incentive Mechanisms

It is widely documented that the popular P2P file sharing systems are havens for "free riders": a significant fraction of users do not contribute any resources, and a minute fraction of users contribute the majority of the resources [SGG02, AH00]. The need for incentives is even larger for 7DS, since users contribute precious battery power by responding to queries and propagating content. 7DS nodes need incentives for:

- Caching content;
- Transmitting cached content when requested;
- Acting as couriers to move 7DS-generated messages to Internet-connected nodes;
- Absorbing fresh content when passing in front of Infostations.

At one extreme, 7DS could use a simple upload-scheme, similar to what is currently used by Kazaa [KaZ], the popular P2P file sharing system. In the simple upload-scheme, the 7DS software in each 7DS node counts the number of times it successfully transmitted an object upon request; it also counts the number of 7DS messages it successfully delivered to an Internet-connected node. In parallel, the counters are slowly decremented at periodic intervals. On the basis of these counters, the local 7DS software classifies the local node into levels of bad, medium, or good. When a 7DS node, say Alice, receives multiple requests (for different objects) from different 7DS nodes, she uses the levels of the requesting nodes to determine the order in which she is going to service the requests. Good nodes get serviced more quickly; bad nodes may have to wait.

Although the above scheme is attractive due to its simplicity, it is very easy to manipulate, even if the user cannot modify its 7DS software. For example, an owner of two 7DS nodes could have the two nodes repeatedly transfer files to each other at home at night. In this manner, both nodes start the day at the good level without having truly contributed to the 7DS community. As an alternative, a node could increment its counter for each successfully transmitted byte and decrement its counter for each received byte. Yet another alternative is to continue to simply count uploads and keep track of the diversity of recipients, to prevent collision as described above. (Devices can be identified by their MAC addresses, for example.) Thus, a good 7DS citizen is distinguished by having delivered lots of data to or from a wide variety of devices. Variations on these schemes may prove sufficient if users cannot easily hack the counters in their devices.

A related problem is that an 7DS node may not transmit authentic objects. For example, Alice may be malicious, and send other users pollution. Therefore, for 7DS to be successful, a reputation scheme may need to be put in place. Developing a 7DS reputation scheme is challenging since a 7DS node is often disconnected from the Internet. One possible scheme is as follows. We install an 7DS reputation broker in the Internet. Each receiver of data occasionally sends a summary to the reputation broker: “Node 184884 gave me n bytes of useful content”. This then adds to the “good karma” of user 184884. When this user 184884 connects to the Internet, it receives from the reputation broker a certificate that certifies the good karma of 184884. Later, when a 7DS node wants to download an object from 184884, before beginning the download, it obtains 184884’s certificate directly from 184884. Conversely, the claims of good deeds described in the previous paragraph can be verified off-line, leading to a list of “leeches” that a trusted authority propagates. There are some dangers of collusion, but reputation systems can mitigate the concern by requiring reports from a large number of users [KSGM03].

In other spheres, such as airline travel, the promise of prizes is used to encourage participation, even if the expected value per transaction is very small. It may also be feasible to extend these incentive models to the 7DS realm.

As part of the proposed research, we will study, classify and evaluate a broad range of incentive and reputation schemes. We will then choose the incentive and reputation schemes that strike the best compromise between simplicity and effectiveness, and implement those schemes in our prototype. Since our system offers a new way of providing connectivity, there is no data on how users will behave. By deploying the system (see Section 4) across large campus populations, we anticipate learning about how strongly users need to be motivated to participate.

3.4 Pseudo-Reliable Multicast for Content Dissemination

Although the current 7DS prototype multicasts queries, it uses unicast to transfer content via HTTP. This has the advantage of simplicity, but non-querier “bystanders” can get the same content if we use a form of reliable multicast instead of unicast.

As an example, consider the case when most of the nodes in a 7DS community have GByte storage capacity. When Alice requests a file from Bob, and Bob responds with the file, all of the bystander nodes within in range of Bob could also capture and cache the file. If Claire is one such bystander, she could read the file herself at a future time, or she could make the file available to other 7DS nodes as she roams about.

As a refinement, Claire could define filters (e.g., only English-language news) and only grab the content that passes through her filter. Similarly, different types of content could be sent over different multicast addresses, and Claire could subscribe only to a subset of multicast addresses.

One interesting twist on this reliable multicast problem is that the requester wants to receive the file in its entirety without error, whereas the bystander is prepared to abandon the file if it is difficult to get a complete copy. We refer to protocols that reliably deliver the file to the requester and attempt to simultaneously deliver files to bystanders as "pseudo-reliable multicast protocols".

Rather than relying on the availability of transport-layer reliable multicast protocols, we will instead focus on providing reliable data transfer at the application level. One natural scheme for pseudo-reliable multicast is as follows. Files are broken up into chunks, on the order of 100 KB per chunk. When Alice requests a file from Bob, Bob multicasts all the chunks over UDP. All 7DS nodes within range absorb the broadcasted chunks. If Alice doesn't receive all the chunks without error, she queries for the individual chunks, thereby discovering which nodes can provide the chunks at the highest signal strength. She then requests the missing chunks from those 7DS nodes. If she is still missing chunks after the chunks are multicast, she again queries for the remaining missing chunks. The bystanding nodes absorb all the chunks being sent. Once Alice has the complete file, some of the bystanding nodes may still be missing chunks. Because these bystanding nodes have not expressed explicit interest in the file, one simple rule is that these nodes abandon the file and simply flush their caches of the chunks for this file. Another rule might be that a bystanding node requests gap-filling chunks only if it lacks a small number of missing chunks.

The scheme described in the above paragraph is only one of a myriad of possible pseudo-reliable multicast schemes. For example, FEC and data carouseling can also be used to provide a form of pseudo-reliable multicast [FG03]. For the most promising schemes, we propose to develop performance and simulation models for comparison. We will also implement one or more of the more promising reliable multicast schemes in our prototype.

Pseudo-reliable multicast has the potential to significantly improve content availability for 7DS users. However, it does so by expending additional battery energy. A bystanding node will expend energy when it grabs a file for which it does not have an immediate interest. We we will therefore also explore the trade-off of availability gains versus battery energy.

3.5 Distributed Content Management

In many 7DS scenarios, there will be a number of 7DS devices in proximity of each other, with new devices coming and leaving from time to time. Examples of such "7DS communities" include campus buildings and public transportation vehicles such as buses and subways. Whenever an object is transmitted, all the 7DS devices in the community and within range of the transmitting node have the option of caching the object. The cache management policy specifies when a device should store such a transmitted object (and evict other objects in storage) in order to minimize metrics such as hit probability and expected access time. Note that mobile devices have limited storage and objects can be large, so each mobile device cannot store an unlimited number of objects.

Intuitively, if there is typically a small number of devices in the community, then all of the devices should be replicating the most popular objects, as we want to maximize the probability that the a copies of the most popular objects remain in the community. One way to achieve this is each peer (1) listens and grabs everything that is transmitted; and (2) uses the LRU (least recently used cache replacement policy). In this manner, if the number of devices in the community drops to one, the remaining device will likely have some of the more popular content, which can be later be disseminated to a newly-arriving peer.

If, on the other hand, there is typically a large number of devices in the community (e.g., hundreds, such as in an academic building on a college campus), then LRU would be far from optimal. With LRU, the tendency is for all peers to replicate the same content, thereby wasting the aggregate storage capacity.

So in this case, for each object it is desirable that only a subset of peers to replicate it; the number of peers in the subset should depend on the popularity of the object among other things. One heuristic along these lines is as follows:

1. Each peer has the same hash function that maps arbitrary strings into $[1, 100]$.
2. When a peer grabs an object, it concatenates the name of the object (e.g. URL) with the MAC address of the peer, and runs the concatenation through the hash function to get h.
3. The peer also estimates the popularity x of the object, based on the recent request history for objects requested in the community.
4. The peer then calculates $y = h \cdot x$, and compares this y value to all the y values it has for the objects currently in the cache. If the y value for the new object is higher than some of the existing y values, it bumps out the object(s) with the lowest y value.

The heuristic should also account for how frequently peers come and go, and the density of peers in the community.

We propose to examine a variety of heuristics for distributed content management, including the simple heuristic outlined above. We'll use both discrete-event simulation as well as tools from stochastic ordering and optimization [KRT03].

3.6 Multi-Modal Architecture

As discussed in the Introduction, combining the strengths of 7DS and cellular networks appears particularly attractive. We have identified five areas that demonstrate the advantages of this symbiotic relationship:

- **Data sharing:** In data sharing, some nodes download material they consider popular, based on their own estimates or past queries. They then re-distribute this data to their 7DS peers, recouping the cost of the cellular data transmission. Second-generation 7DS peers may lower their cost in turn by distributing material to third-generation peers, and so on. This operation entails risk for the first downloader if it does not have an interest in the data itself, but may be viewed as an additional benefit if it would have downloaded the data anyway for its own use. Using economic models, we propose to investigate the optimal strategy for peers, based on running estimates of current data prices and popularity for particular items. (For example, it appears plausible that the value of objects does not increase linearly with their size or that certain media types are better bets.)
- **Data spreading:** If speed dominates cost, nodes may decide to spread large objects over multiple peer nodes, each using their cellular channel to upload or download chunks of the object in parallel[5]. This mode of operation may be particularly useful under emergency conditions. We had explored a simple version of this in [PS99], where connections were shared in real-time and synchronously, rather than spreading messages across multiple independent nodes. Upload spreading does not require that the peers remain in contact, while download spreading does. This is a version of swarming[6], but the mobility of nodes makes the analysis and algorithm much more challenging. We propose to investigate adaptive algorithms that recover gracefully from losing peers that have walked away.

---

[5] We are assuming a TDMA-style networks, as in GPRS, where per-user capacity is severely limited.

[6] Examples include SwarmCast, the Open Content Network, BitTorrent and Gnutella BearShare.

- **Content location:** The cellular data channel may also be used to point mobile users to promising data locations, e.g., fixed Infostations or the route of large mobile data carriers. (For example, we envision city buses traversing Broadway as high-volume mobile data carriers.) Thus, it offers a form of service location, requiring little bandwidth, while the large data objects are efficiently conveyed using 7DS. We plan to implement and test the utility of such location information in our field trials.

- **Cache clearing:** One of the core problems with epidemic data propagation is the difficulty of deciding when to remove the object from the node caches, since nodes have generally no clue when the data object has been delivered. We propose to use the low-bandwidth cellular data network as a cache invalidator. Nodes may query a registry periodically as to whether a particular data object, identified by its cryptographic hash, has been delivered. Alternatively, nodes that carry copies of a particular object can subscribe to the delivery event. Since we are only concerned about tracking large objects, not IP packets, the storage effort for such package tracking is likely to be modest. We propose to investigate appropriate invalidation algorithms and their performance.

- **Reputation checking:** We observed earlier that peer-to-peer systems are likely to need incentives to keep people from “freeloading”. We suspect that relatively simple approaches that separate the “contributors” from the “freeloaders” are sufficient, rather than having to build a complete micropayment system, for example. The cellular data network may be useful to check the reputation of a querier before contributing a large data object, particularly for peers with uncertain reputation. We plan to investigate low-bandwidth probabilistic reputation checking mechanisms that add minimal overhead to the overall system operation.

4 Prototype and Deployment

At the core of this proposed research is a comprehensive prototype for 7DS. The prototype development has three major axes: extension of current prototype to allow for pseudo-reliable multicast; a PDA version of the prototype, including 7DS-to-Internet messaging; and an embedded system version of 7DS for transport vehicles.

We plan to disseminate our results via open-source software, available through SourceForge or similar venues. Since the hardware design will be based on commercially available single-board computers, replicating the design by other research groups should not be difficult.

4.1 Pseudo-Reliable Multicast

As described in the previous section, as part of our fundamental research we will study different pseudo-reliable multicast and cache management schemes. Once we have identified the most promising schemes, we will implement them in the 7DS software. We will then deploy 7DS on Polytechnic students’ portable computers, which are all equipped with 802.11 cards. We will test 7DS for course content sharing and generic file sharing. As part of this effort, we will port existing Linux code to the Windows environment.

Also, as part of this axis, we’ll be integrating the walk-through Infostation, developed at Polytechnic, with 7DS, initially developed at Columbia. In order to integrate the two prototypes, we will need to ensure that data absorbed by a device from an Infostation is accessible to the 7DS software.

4.2 PDA Version with 7DS-to-Internet Messaging

The PDA version will be designed for common PDA operating system platforms (CE 4.0/PPC 2003, PalmOS, Linux) and be geared towards making 7DS operation as transparent to the user as possible.

Extending the proxy model in the current 7DS prototype, we plan to provide instant messaging and SMTP (email) proxies (acting as Message Transfer Agent (MTA)) for uploading messages, so that users can continue to use their existing mail user agents (MUAs) and IM tools[7]. We believe that mechanisms that are transparent to users and do not require specialized user applications will find large-scale use. Fig. 4 shows a simplified block diagram of the software architecture.

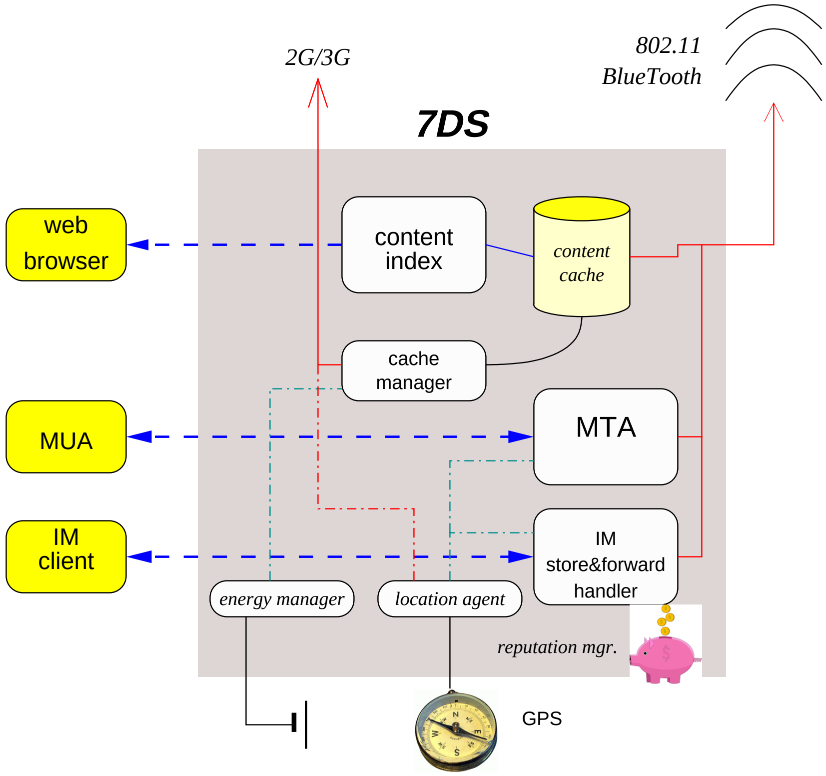

Figure 4: Architecture for 7DS node on PDAs

4.3 Embedded Version for Transport Vehicles

For a number of applications, the 7DS node is carried by a vehicle, not a person. We plan to develop a self-contained, rugged node, consisting of a low-power single-board computer contained in a watertight enclosure. The box can be attached to utility poles, taxi cabs, city buses, garbage trucks, delivery trucks and other vehicles that cover large areas. The node will be equipped with a large CompactFlash memory module, a USB wireless LAN interface, a GPS module and a wide-area cellular interface. The GPS module facilitates the use of location-based and heading-based selection of next-hop devices, but can also offer location information to PDA devices passing by.

We are currently considering the InhandElectronics Elf3 module, running Linux, as our candidate hardware, as it consumes only 350 mW and thus can be powered parasitically from a 12V automotive power supply or a small solar panel. This node can then serve as an outdoor Infostation as well as a mobile 7DS node, for both data delivery and pickup.

---

[7] Although there are many proprietary IM protocols, we will focus on the IETF-standardized SIMPLE architecture.

5 Broader Impacts

One of the features of this proposal is that it is rich in broader impacts:

- **Technology transfer:** It is our ambition that the architectures, protocols, and prototypes developed in this proposal will be a catalyst for a new industry, which fills an important niche between the broadband Internet industry and the cellular data industry. After having demonstrated proof of concept, we believe many existing and emerging industrial players (both equipment manufacturers and service providers) should take interest. Also, we will approach operators of vehicle fleets, such as courier services, and public transportation systems such as the MTA to explore large-scale deployments. Throughout the proposed research, we will be actively seeking industrial collaboration. This research fits well with the pending industry/university cooperative research center (I/UCRC) of Polytechnic and Columbia.

- **Dissemination:** We will strongly encourage open source development of 7DS software. We will be creating and managing a Web site for a clearinghouse of ideas and specifications for implementors. We will seek standardization of protocols at the IETF.

- **Diversity and outreach:** The principal investigators of the proposal are diverse in gender. Furthermore, the research involves close collaboration between researchers at Columbia University and Polytechnic University, which are diverse in their style. While Columbia University is an Ivy League institution, drawing outstanding undergraduate students from throughout the world, Polytechnic university, with its motto "realizing the American dream," has traditionally served children of families that have immigrated to New York City in the past 20 years or so.

- **Education:** (1) Keith Ross is co-author of the top-selling networking textbook, *Computer Networking: A Top-Down Approach Featuring the Internet,* [KR02b] which is currently read by over 30,000 professors and students every year. Much of the knowledge gained through the proposed research will be disseminated to students worldwide, with significantly expanded coverage of wireless LANs and cellular data in the third edition, which will likely be published by 2005. (2) Undergraduates will be involved at both Columbia and Polytechnic in undergraduate projects related to this research. (3) Through Polytechnic's YES program, we expect to involve summer students from Stuyvesant High School and Bronx School of Science (two of the best public high schools in the US) involved in the research.

6 Management Plan and Schedule

This proposal has three parallel components. The first component is the fundamental work on protocols and performance modeling of 7DS, as described in the proposal. This 7DS research will be carried out by a Polytechnic Ph.D. student, jointly supervised by Phyllis Frankl and Keith Ross, with significant input from Henning Schulzrinne. The second component is the comprehensive prototype. The prototype will be developed by the postdoc (resident at Columbia) along with master's and undergraduate students at both Polytechnic and Columbia. Undergraduates will also be employed to test the prototypes. The entire 7DS team will meet bi-weekly, alternating between the Columbia and Polytechnic campuses. The final component is the real-world deployment of the prototypes, preferably with industrial partners. Since all three PIs have enjoyed significant interaction with industry and startups, and since 7DS technology and services could lead to important sources of revenue, it should not be difficult to find enthusiastic deployment partners.

In the first year, our focus will be on modeling and implementing pseudo-reliable multicast, 7DS-to-Internet multicast, and Infostation integration. In the second year, our focus will be on researching and implementing incentive systems, 7DS cache management and location-aware upload optimization; in this second year we will also be implementing 7DS in PDAs. The third year of the proposed research we will study multi-modal architectures and we will complete the implementation of the embedded 7DS system for transport systems.

7 Results from Prior NSF Support

Phyllis Frankl is currently PI on NSF Grant CCR-9988354, "Techniques for Testing Database Applications". Although database application programs play a critical role in most modern organizations, little prior research had addressed the question of how to test whether they perform as intended. We have developed a set of tools, AGENDA, for testing relational database applications. AGENDA parses the database schema and application queries, extracts relevent information from them, and uses this information, in interaction with the user, to populate the database, generate test cases, execute them and partially check test output and resulting DB state. Extensions to testing for concurrency faults and empirical evaluation are in progress. Previous NSF grants have funded development of several other software testing tools, as well as resesearch evaluating the effectiveness of software testing techniques (most recently, CCR-9870270).

Frankl is co-PI on NSF Awards 0210526 and 0210514, "Scholarship for service in information assurance", and "Capacity Building Project in Information Assurance Education", under which new educational programs in computer security are being developed and twenty students are receiving full scholarships and stipends. Previous support under DUE-9981035, "Undergraduate Information Systems Security Laboratory", led to development of a course sequence and laboratory for Computer Security Education.

Keith Ross, who was employed outside of the US from 1997 through 2002, has had no NSF support within the last four years.

Henning Schulzrinne has been co-PI or PI on previous NSF grants [NSC^{+}01, FAS01, Sch00, SKK^{+}02], designing a web hot spot rescue service [ZS03], system architectures for VoIP [RSC^{+}02, RS02a, RS02b, JLSS01, JLN^{+}02, KSW02, SJL^{+}02], Internet-standard programming languages for multimedia services [LS00a, LSR01], feature interaction of IP telephony services [WS00, LS00b], application-layer mobility [SW00, DVC^{+}01], measurement techniques for packet audio quality and reliability [JS00, JS02a, JS02a, JS03, KJS02, JS02b, JS02c], improved service location techniques [ZSG03, ZS01, ZSG00, Zha02], infrastructure-less wireless data replication [PS01a] and congestion-based pricing [WS01].

References

[AH00] E. Adar and B. Huberman. Free Riding on Gnutella, October 2000. First Monday, http://www.firstmonday.dk/issues/issue5_10.

[BSSW03] Stefan Berger, Henning Schulzrinne, Stylianos Sidiroglou, and Xiaotao Wu. Ubiquitous computing using SIP. In *ACM NOSSDAV 2003*, June 2003.

[DVC$^+$01] Ashutosh Dutta, Faramak Vakil, Jinghu Chen, Miriam Tauil, Shinichi Baba, and Henning Schulzrinne. Application layer mobility management scheme for wireless Internet. In *3Gwireless*, page 7, San Francisco, May 2001.

[FAS01] Steve Feiner, Peter Allen, and Henning Schulzrinne. Infrastructure for context-aware wireless network applications, September 2001. NSF grant ANI-0099184.

[FG03] P.G. Frankl and D.J. Goodman. Delivering Information "To Go" via Infostations. Technical report, Polytechnic University, 2003. http://www.poly.edu/wicat/doc/02-008.pdf.

[FN01] Laura Feeney and Martin Nilsson. Investigating the energy consumption of a wireless network interface in an ad hoc networking environment. In *Proceedings of the Conference on Computer Communications (IEEE Infocom)*, Anchorage, Alaska, April 2001.

[GEBF$^+$03] L. Garces-Erce, E. Biersack, P. Felber, K.W. Ross, and G. Urvoy-Keller. Hierarchical P2P Systems. In *Proceedings of Euro-Par 2003*, 2003.

[GT02] M. Grossglauser and D. Tse. Mobility Increases the Capacity of Adhoc Wireless Networks. *IEEE/ACM Transactions on Networking*, 10:77–486, August, 2002.

[HBH00] Larry Hughes, Omid Banyasad, and Evan Hughes. Cartesian routing. *Computer Networks*, 34(3):455–466, September 2000.

[JLN$^+$02] Wenyu Jiang, Jonathan Lennox, Sankaran Narayanan, Henning Schulzrinne, Kundan Singh, and Xiaotao Wu. Integrating Internet telephony services. *IEEE Internet Computing*, 6(3):64–72, May 2002.

[JLSS01] Wenyu Jiang, Jonathan Lennox, Henning Schulzrinne, and Kundan Singh. Towards junking the PBX: deploying IP telephony. In *Proc. International Workshop on Network and Operating System Support for Digital Audio and Video (NOSSDAV)*, Port Jefferson, New York, June 2001.

[JPS01] Rahul Jain, Anuj Puri, and Raja Sengupta. Geographical routing using partial information for wireless ad hoc networks. *IEEE Personal Communications Magazine*, 8(1), February 2001.

[JS00] Wenyu Jiang and Henning Schulzrinne. Analysis of on-off patterns in VoIP and their effect on voice traffic aggregation. In *International Conference on Computer Communication and Network*, Las Vegas, Nevada, October 2000.

[JS02a] Wenyu Jiang and Henning Schulzrinne. Comparison and optimization of packet loss repair methods on VoIP perceived quality under bursty loss. In *Proc. International Workshop on Network and Operating System Support for Digital Audio and Video (NOSSDAV)*, Miami Beach, Florida, May 2002.

[JS02b] Wenyu Jiang and Henning Schulzrinne. Comparisons of FEC and codec robustness on VoIP quality and bandwidth efficiency. In *ICN*, Atlanta, Georgia, August 2002.

[JS02c] Wenyu Jiang and Henning Schulzrinne. Speech recognition performance as an effective perceived quality predictor. In *IWQoS*, Miami Beach, May 2002.

[JS03] Wenyu Jiang and Henning Schulzrinne. Assessment of VoIP service availability in the current Internet. In *Passive & Active Measurement Workshop*, San Diego, CA, April 2003.

[KaZ] KaZaA Web Site. http://www.kazaa.com.

[KJS02] Kazuumi Koguchi, Wenyu Jiang, and Henning Schulzrinne. QoS measurement of VoIP endpoints. In *IEICE Group meeting on Network Systems*, Matsuyama, Japan, December 2002.

[KRT03] J. Kangasharju, K.W. Ross, and D. Turner. Adaptive Content Management in Structured P2P Communities. Technical report, Polytechnic University, 2003. http://cis.poly.edu/ ross/papers/AdaptiveContentManagement.pdf.

[KSGM03] S.D. Kamvar, M.T. Schlosser, and H. Garcia-Molina. The eigentrust algorithm for reputation management in p2p networks. In *Proceedings of the Twelfth International World Wide Web Conference*, 2003.

[KSW02] Petri Koskelainen, Henning Schulzrinne, and Xiaotao Wu. A SIP-based conference control framework. In *Proc. International Workshop on Network and Operating System Support for Digital Audio and Video (NOSSDAV)*, Miami Beach, Florida, May 2002.

[LJC^{+}00] Jinyang Li, John Jannotti, Douglas S. J. De Couto, David R. Karger, and Robert Morris. A scalable location service for geographic ad hoc routing. In *ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom)*, pages 120–130, Boston, Massachusetts, USA, August 2000.

[LS00a] Jonathan Lennox and Henning Schulzrinne. The call processing language: User control of internet telephony services. In *Lucent Technologies XML Day*, Murray Hill, NJ, February 2000.

[LS00b] Jonathan Lennox and Henning Schulzrinne. Feature interaction in Internet telephony. In *Feature Interaction in Telecommunications and Software Systems VI*, Glasgow, United Kingdom, May 2000.

[LSR01] J. Lennox, Henning Schulzrinne, and J. Rosenberg. Common gateway interface for SIP. RFC 3050, Internet Engineering Task Force, January 2001.

[NI97] J. C. Navas and T. Imielinski. Geographic addressing and routing. In *ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom)*, Budapest, Hungary, September 1997.

[NSC^{+}01] Jason Nieh, Henning Schulzrinne, Edward Coffman, Predrag Jelenkovic, and Dan Rubenstein. The Columbia hotspot rescue service, September 2001. NSF ANI-0117738.

[Pap02] Maria Papadopouli. *Resource Sharing in Mobile Wireless Networks*. PhD thesis, Department of Computer Science, Columbia University, New York, May 2002.

[PS99] Maria Papadopouli and Henning Schulzrinne. Connection sharing in an ad hoc wireless network among collaborating hosts. In *Proc. International Workshop on Network and Operating System Support for Digital Audio and Video (NOSSDAV)*, pages 169–185, Basking Ridge, New Jersey, June 1999.

[PS00] Maria Papadopouli and Henning Schulzrinne. Seven degrees of separation in mobile ad hoc networks. In *Proceedings of the IEEE Conference on Global Communications (GLOBECOM)*, San Francisco, California, November 2000.

[PS01a] Maria Papadopouli and Henning Schulzrinne. Effects of power conservation, wireless coverage and cooperation on data dissemination among mobile devices. In *Mobihoc*, Long Beach, California, October 2001.

[PS01b] Maria Papadopouli and Henning Schulzrinne. Performance of data dissemination among mobile devices. technical report CUCS-005-01, Dept. of Computer Science, Columbia University, New York, New York, July 2001.

[PS02] Maria Papadopouli and Henning Schulzrinne. Performance of data dissemination and message relaying in mobile ad hoc networks. Technical Report CUCS-004-02, Dept. of Computer Science, Columbia University, New York, New York, February 2002.

[RBF$^+$02] K.W. Ross, E. Biersack, P. Felber, L. Garces-Erce, and G. Urvoy-Keller. Topology Centric Lookup Service. Technical report, Polytechnic University, 2002. http://cis.poly.edu/ ross/papers/P2Pemail.pdf.

[RD01] A. Rowstron and P. Druschel. Pastry: Scalable, distributed object location and routing for large-scale peer-to-peer systems. In *Proceedings of IFIP/ACM International Conference on Distributed Systems Platforms (Middleware)*, pages 329–350, November 2001.

[RFH$^+$01] S. Ratnasamy, P. Francis, M. Handley, R. Karp, and S. Shenker. A scalable content-addressable network. In *Proceedings of the 2001 ACM SIGCOMM Conference*, 2001.

[Ros97] K.W. Ross. Hash-Routing for Collections of Shared Web Caches. *IEEE Network Magazine*, November 1997.

[RR] K.W. Ross and D. Rubenstein. P2P Systems: Infocom 2003 Tutorial. http://cis.poly.edu/ ross/papers.

[RS02a] J. Rosenberg and Henning Schulzrinne. Reliability of provisional responses in session initiation protocol (SIP). RFC 3262, Internet Engineering Task Force, June 2002.

[RS02b] J. Rosenberg and Henning Schulzrinne. Session initiation protocol (SIP): locating SIP servers. RFC 3263, Internet Engineering Task Force, June 2002.

[RSC$^+$02] J. Rosenberg, Henning Schulzrinne, G. Camarillo, A. R. Johnston, J. Peterson, R. Sparks, M. Handley, and E. Schooler. SIP: session initiation protocol. RFC 3261, Internet Engineering Task Force, June 2002.

[Sch00] Henning Schulzrinne. Internet telephony services and mobile multimedia services, September 2000. NSF CAREER grant ANI-9958325.

[SGG02] S. Saroiu, P. Gummadi, and S.D. Gribble. A Measurement Study of Peer-to-Peer File Sharing Systems. In *Proceedings of International Conference on Distributed Computing Systems*, 2002.

[SJL+02] Kundan Singh, Wenyu Jiang, Jonathan Lennox, Sankaran Narayanan, and Henning Schulzrinne. CINEMA: columbia internet extensible multimedia architecture. technical report CUCS-011-02, Department of Computer Science, Columbia University, New York, New York, May 2002.

[SKK+02] Henning Schulzrinne, Gail E. Kaiser, John R. Kender, Steven K. Feiner, and Kathleen R. McKeown. CISE research infrastructure: Pervasive pixels, September 2002. NSF grant EIA-0202063.

[SMK+01] I. Stoica, R. Morris, D. Karger, M. F. Kaashoek, and H. Balakrishnan. Chord: A scalable peer-to-peer lookup service for internet applications. In *Proceedings of the 2001 ACM SIGCOMM Conference*, 2001.

[SW00] Henning Schulzrinne and Elin Wedlund. Application-layer mobility using SIP. *Mobile Computing and Communications Review (MC2R)*, 4(3):47–57, July 2000.

[TIK+00] Nobuhiko Tada, Yukimitsu Izawa, Masahiko Kimoto, Taro Maruyama, Hiroyuki Ohno, and Masaya Nakayama. IAA system (“I am alive”): The experiences of the Internet disaster drills. In *INET*, Yokohama, Japan, July 2000.

[VB00] Amin Vahdat and David Becker. Epidemic routing for partially-connected ad hoc networks. Technical report, Department of Computer Science, Duke University, Durham, North Carolina, April 2000.

[WS00] Xiaotao Wu and Henning Schulzrinne. Where should services reside in Internet telephony systems? In *IP Telecom Services Workshop*, Atlanta, Georgia, September 2000.

[WS01] Xin Wang and Henning Schulzrinne. Pricing network resources for adaptive applications in a differentiated services network. In *Proceedings of the Conference on Computer Communications (IEEE Infocom)*, Anchorage, Alaska, April 2001.

[Zha02] Weibin Zhao. Selective anti-entropy. In *PODC (Principles of Distributed Computing)*, Monterey, California, July 2002.

[ZKJ01] B. Y. Zhao, J. Kubiatowicz, and A. D. Joseph. Tapestry: An infrastructure for fault-tolerant wide-area location and routing. Technical Report UCB/CSD-01-1141, UC Berkeley, 2001.

[ZS01] Weibin Zhao and Henning Schulzrinne. ARMS - automated replication for meshed servers. In *3rd NMADS (New York Metra Area Distributed Systems) Day*, New York, NY, April 2001.

[ZS03] Weibin Zhao and Henning Schulzrinne. Predicting the upper bound of web traffic volume using a multiple time scale approach. In *International World Wide Web Conference*, page 243, Budapest, Hungary, May 2003.

[ZSG00] Weibin Zhao, Henning Schulzrinne, and Erik Guttman. mSLP - mesh-enhanced service location protocol. In *International Conference on Computer Communication and Network*, Las Vegas, Nevada, October 2000.

[ZSG03] Weibin Zhao, Henning Schulzrinne, and Erik Guttman. Mesh-enhanced service location protocol. Internet draft, Internet Engineering Task Force, January 2003. Work in progress.