

Performance of data dissemination among mobile devices

Maria Papadopouli

Henning Schulzrinne*

Abstract

This paper presents 7DS, a novel peer-to-peer data sharing system. 7DS is an architecture, a set of protocols and an implementation enabling the exchange of data among peers that are not necessarily connected to the Internet. Peers can be either mobile or stationary. It anticipates the information needs of users and fulfills them by searching for information among peers. We evaluate via extensive simulations the effectiveness of our system for data dissemination among mobile devices with a large number of user mobility scenarios. We model several general data dissemination approaches and investigate the effect of the wireless coverage range, 7DS host density, and cooperation strategy among the mobile hosts as a function of time. We also present a power conservation mechanism that is beneficial, since it increases the power savings, without degrading the data dissemination. Using theory from random walks, random environments and diffusion of controlled processes, we model one of these data dissemination schemes and show that the analysis confirms the simulation results for this scheme.

1 Introduction

Wireless devices are becoming smaller, more user friendly and more pervasive. They are not only car-

ried by humans, but are integrated into physical objects. This is simply the natural evolution of the historical trend in computing and communications technologies. The expansion of the Internet and wireless communications have amplified this trend by making information easier to share and by increasing the amount of information that is shared. Wireless information access will become as important as voice communications, since people are beginning to heavily depend on on-line information. People access local and general news, traffic or weather reports, sports, maps, guide books, music and video files, games. We classify mobile information access methods into three main categories. The first approach provides “continuous” wireless Internet access; examples include CDPD, 3G wireless, 802.11 and two-way pagers. Currently, this access either have sparse coverage, low cost and high speed (802.11) or have major-cities-only coverage and high cost (Metricom [1]) or have wider coverage, but extremely low rates and high costs (CDPD, RIM). The second approach provides information access via fixed (stationary) information servers or *infostations*. The *infostations* are “information kiosks”, for example, located at traffic lights, building entrances and airport lounges. Typically, an *infostation* is a server attached to a data repository and a wireless LAN. When a wireless device is in close proximity to an infostation, it can query the server and access the information. Both of these approaches need an infrastructure. If the wired infrastructure is low-bandwidth, they can be combined by having caches at the base stations.

*The authors are with Department of Computer Science, Columbia University in New York, USA. Email: {maria,hgs}@cs.columbia.edu

In that case, when the wireless device is within the wireless range of the servers, it can use the cache or the information server to access the data or directly the Internet via the base station. Given the exceedingly expensive license fees attained in recent government auctions of spectrum, the bandwidth expansion route is bound to be expensive. Similarly, the cost of tessellating a coverage area with a sufficient number of base stations or infostations coupled to the associated high speed wired infrastructure cost is prohibitive. Also, wireless devices are bound to become smaller and more pervasive and not only be carried by humans, but integrated into physical objects (such as cars, electrical appliances). It is unlikely that the density of base stations and infostations will keep pace, due to regulatory, environmental and cost barriers in deploying them. In addition, there are situations where the infrastructure is not available (e.g., emergency, disaster relief, tunnels) or it is overloaded. Thus, we focus on a third approach that does not need the support of any infrastructure (i.e., ad hoc), based on peer-to-peer data sharing among mobile, wireless devices. The devices are autonomous and not necessarily connected to the Internet. For the next few years, continuous connectivity to the Internet will not be available at low cost for mobile users roaming a metropolitan area.

In this paper we focus on the challenge of increasing the data availability to users roaming a metropolitan area that experience intermittent connectivity to the Internet. Two characteristics of the environment motivate our approach, namely the high spatial locality of information and the co-existence of a heterogeneous set of information providers and access methods. In an urban environment, such as part of Manhattan during rush hours, a platform of a train or an airport, a commercial center or a campus, we anticipate that the mobile user's access patterns will include high spatial locality of informa-

tion (such as local and general news, sports, schedules) and also popular information (such as music files, video games). The high spatial locality of information results from the type of services we expect a mobile user will run, namely location-dependent, service location and news services as well as collaborative applications. Mobile users are likely to be more flexible in their information tastes, (media) quality and information accuracy requirements. More important, we anticipate that there will be multiple wireless information access providers with servers not necessarily cooperating with each other or part of the same infrastructure. A system that is capable of transparently accessing different information providers based on the data and connection availability can increase the information availability of the users. We propose *7DS*¹ as a system that complements the three mobile information access approaches we describe in the previous paragraph. It is an architecture and set of protocols enabling the exchange of data among peers that are not necessarily connected to the Internet. *7DS* runs as an application and operates in two modes, namely prefetch and on-demand. In prefetch mode, it anticipates information needs of users. In on-demand mode, it searches for information among peers when the device failed to access the data via the Internet (e.g., the user tried to access a web site unsuccessfully). *7DS* uses transparently different mobile access methods depending on the network availability. It runs on heterogeneous devices (with different capabilities) that are mobile or stationary. For example, a *7DS*-enabled server can either be dual-homed device connected to the Internet or to a wired infras-

¹"*7DS*" stands for "Seven Degrees of Separation", a variation on the "Six Degrees of Separation" hypothesis, which states that any human knows any other by six acquaintances or relatives. There is an analogy with our system, particularly, with respect to data recipients and the device with the "original" copy. We have not explored if a similar hypothesis is true here.

tructure of other servers or an autonomous server attached to a cache with an access to a wireless LAN. When *7DS* runs on handheld devices (e.g., PDAs), it will use power conservation and collaboration methods different from the server's ones. We distinguish two principal interaction types, namely peer-to-peer (P-P) and server-to-client (S-C). In P-P, *7DS* hosts are cooperative with each other. S-C schemes operate in a more asymmetric fashion: there are some cooperative hosts (e.g., *7DS* servers) that respond to queries and non-cooperative, resource constrained clients (e.g., PDAs). *7DS* nodes can collaborate by data sharing, forwarding messages (i.e., "rebroadcasting" queries and responses) or by caching popular data objects. For example, an autonomous *7DS* server may monitor for frequently requested data, request them from other peers and cache the data locally to serve future queries. The information is typically web pages or any data elements of modest size.

In this new framework, we address some general questions related to the effect of wireless coverage range, density of devices, cooperation among the hosts and their power conservation strategy on data dissemination. For example, how fast does information spread in such setting if all nodes are cooperative with each other? How does it change when only a few nodes are cooperative (e.g., the *7DS*-enabled servers)? What is the percentage of the nodes that acquire a data item over time? What is the average delay that a node experiences till it receives the data? How does the server-to-client compare to a peer-to-peer approach? How does the wireless coverage range, power conservation, density of devices and servers affect the data dissemination? The investigation of these issues can also give insight for the design of an wireless information infrastructure in a metropolitan area. It appears to be not amenable to an analytical solution except for simplified settings with respect to the node layout, mobility pattern and user interaction pattern. In addition, it is

difficult to perform actual experiments using the current testbed (primarily lack of a large number of mobile devices with wireless interfaces, and difficulty of "approximating" the user's social behaviour). To investigate these issues and also assess the efficiency of information dissemination via *7DS*, we perform a simulation-based study. In addition to the simulations, we also present our initial analytical results using diffusion controlled processes theory. The simulations and analysis are not tied to *7DS*, and provide more general results on data dissemination.

We would like to emphasize the differences of our approach and the setting we consider from that in related works. Ad hoc and sensor networks typically assume a relative high density of devices that results in a connected network (a host can access other hosts via multi-hop routing) [2, 3, 4, 5]. In particular, sensor networks are typically modeled as an infrastructure of high density short-range wireless devices, stationary and connected with each other and/or with a base station/controller. On the other hand, a *7DS* network is rarely connected and the time taken for one *7DS* node to come in close proximity to another can be of the order of minutes. More important, ad hoc and sensor networks assume cooperative nodes with similar capabilities, part of the same infrastructure, that relay packets on behalf of other nodes. As we mentioned, in our setting, peers have different capabilities and cooperation strategy and they are not necessarily all cooperating with each other. Both in ad hoc and sensor networks the emphasis has been on routing protocols. In this study, we concentrate on single-hop multicast. *7DS* acquires the data from other peers within its wireless coverage. Due to the highly dynamic environment and the type of information, *7DS* does not try to establish more permanent caching or service discovery mechanisms. Instead we explore the transient aspect of information dissemination. Note that this setting is orthogonal to the service discovery in the wide area network.

In the latter, typically, there is an infrastructure of cooperative servers that create indices to locate data based on the queries and the content of the underlying data sources of their local domain [6]. Finally, an extensive amount of work has been done in the context of infostations. They use a single server/multiple clients model in which the server broadcasts data items based on received queries. They mostly address issues related to efficient scheduling algorithms for the server broadcast that minimize the response delay and power consumption of mobile devices, and utilize efficiently the bandwidth of the broadcasting channel [7, 8, 9]. Section 5 discusses related work in greater detail.

In our simulations, we consider variations of the P-P and S-C schemes as well as some hybrid ones. We consider a simple power conservation mechanism that periodically enables the network interface. During the *on* interval *7DS* hosts communicate with their peers. In its asynchronous mode, the *on* and *off* intervals are equal (but not synchronized). In synchronous mode, the *on* and *off* intervals are synchronized among hosts (and not necessarily equal). We also vary the wireless range of the network interfaces from 55 m to 230 m. We evaluate these approaches by measuring the percentage of hosts that acquire the data item as a function of time and their average delay. At the beginning of each experiment, only one *7DS* host has the data item of interest and the remaining hosts are interested in this data item. We found that the cooperation, the density of the cooperative hosts and their mobility, and the transmission power have great impact on data dissemination. For a region with the same density of hosts, P-P outperforms S-C with no cooperation among the mobile devices. The simulations indicate that, the probability that a host that queries for a data object will acquire it by time t follows the function $1 - e^{-a\sqrt{t}}$ when using FIS (S-C mode with fixed server and no cooperation among the mobile devices) and the function

$1 - e^{-at}$ when using P (P-P mode with data sharing and asynchronous power conservation). In case of high density of cooperative hosts, the data dissemination using P-P grows even faster. For example, in P-P, in a setting of 15 hosts with wireless range of 230 m, after 25 minutes, 99% of the users will acquire the data compared to just 42% of the users in the FIS. For the same average delay of 6 minutes, a host using FIS will get the data with a 42% probability, whereas using synchronous P, even in a setting of only five hosts per km^2 , this probability is double. For lower transmission power, P-P outperforms FIS by 20% to 70%. In the case of only five hosts, the two approaches differ by 3% to 43%, depending on the transmission power. We also discover two important scaling properties of data dissemination by expanding the area and varying the density of wireless coverage (i.e., average wireless coverage per space unit) of cooperative hosts and the density of cooperative hosts (i.e., average number of cooperative hosts per space unit). First, the performance remains the same when we scale the area but keep the density of the cooperative hosts and transmission power fixed. Second, for fixed wireless coverage density, the larger the density of cooperative hosts, the better the performance. In S-C, this implies that for the same wireless coverage density, it is more efficient to have a larger number of cooperative hosts with lower transmission power than fewer with higher transmission power. These results can assist in the design of wireless data infrastructures.

The contributions of this paper are

1. An overview of the design and implementation of *7DS*, a novel system that enables information dissemination and sharing among mobile hosts in a peer-to-peer fashion.
2. An evaluation via extensive simulations of *7DS* and the effects of the wireless coverage range,

7DS host density, cooperation strategy among the mobile hosts as a function of time.

3. The synchronous power conservation, a beneficial mechanism that increases the power savings substantially, without degrading the data dissemination.
4. An analytical model for FIS using theory from random walks and environments and the kinetics of diffusion-controlled processes. The analytical results on data dissemination are consistent with the simulation results for FIS.

In an earlier work [10], we presented a part of the performance analysis of *7DS* for a fixed simulation time. As we mentioned in the previous paragraph, here, we investigate its performance over time as well as the scaling results (in respect to the density of hosts and wireless coverage) and the synchronous power conservation mechanism. The remainder of this paper is organized as follows. Section 2 gives an overview of the main components of *7DS*. Section 3.1 describes in more detail the P-P and S-C models. Section 3.2 presents simulation results and Section 4 the modeling and analysis of FIS using kinetics of diffusion controlled processes. Finally, in Section 6, we summarize our conclusions and discuss directions for future work.

2 System architecture overview

We assume that the mobile host has a network connection to access the Internet, e.g., via a wireless modem or a base station, and is also capable of communicating with other hosts via a wireless LAN (e.g., IEEE 802.11). *7DS* runs as an application on mobile hosts and communicates with other *7DS* participants via a wireless LAN. We focus on information access from the Internet that takes place by retrieving data objects identified by URLs. When such access

fails (for example due to the loss of the Internet connection), *7DS* tries to acquire the data from other *7DS* peers. Figure 1 (a) illustrates how *7DS* operates. Mobile host A tries to access a data object (e.g., web page). The local *7DS* detects that the host has no connection to the Internet and tries to access the page from the peers in close proximity via the wireless LAN. Mobile host D has walked way and cannot listen the query. Both host B and C receive the query. Host C has a copy of the data in its cache and responds to A's query.

7DS uses three types of messages to communicate with other peers: queries, reports and advertisements. A query consists of a set of attributes with their values, such as the URL of the data object and the MAC address of the host that generated it. These two attributes, the URL of the data object and the MAC address, are also used as the query identifier. *7DS* forms queries based on the URL of the data object it tries to acquire. *7DS* maintains a query list, in which it also includes the URLs the system predicts the user will visit in the next few hours. *7DS* multicasts these queries periodically via the wireless LAN to a predefined multicast group. *7DS* uses different multicast groups for different queries/data items. It determines the appropriate group either by hashing the URL of the requested data item or using some application specific criteria. In order to conserve more power, a host may listen to a subset of these groups depending on the data objects it is willing to share. Both in the prototype and simulations, we consider single-hop multicast, using the "ad-hoc mode" of 802.11. When receiving a query, each *7DS* peer searches the cache. If a host finds a match, it forms and broadcasts a report. The report is describing the relevant data. After a defined interval, *7DS* selects from the received reports, the most relevant ones based on application-specific criteria and then it initiates an HTTP GET request. The advertisements are application-specific mes-

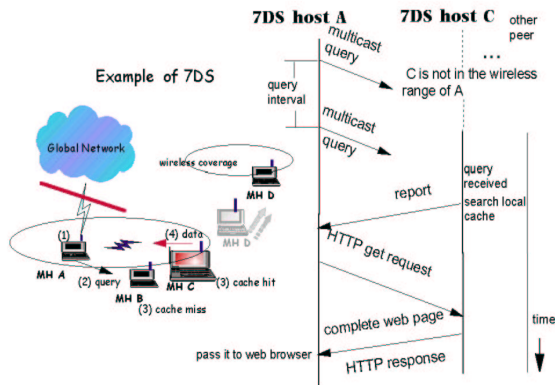


Figure 1: The arrows show the message exchange for the *7DS* communication. The ellipse denotes the wireless coverage of each host, the shaded signal the wireless LAN and the non-shaded one the (lost) connectivity to the Internet.

sages that announce the presence of *7DS*-enabled server. Power constrained devices use a “passive” mode for participating in the system. In particular, they participate only when the expectation for data availability is high, for example, when they receive an advertisement. A *7DS*-enabled server broadcasts periodically such advertisements with a description of the information or application it supports. A *7DS* with passive mode enabled host sends the query directly to the server, when it receives an advertisement. We call this “passive” querying, as opposed to active querying that takes place periodically until *7DS* receives the data.

We use XML to describe *7DS* messages. *7DS* extracts the metadata from the queries received from other peers and performs an attribute-matching search in its local cache. A report includes an identifier that matches the identifier of its corresponding query, and a data description field that contains the relevant information in the local cache of the peer

that responds. The report message also consists of some optional attributes with their values, such as the original URL, the time the object (of the data description field) was cached locally, time the original copy was created, its HTML title, size and format, the quality of the wireless transmission (using the signal-to-noise ratio value), the author, language, size and content type of the object. Some of this information is inherently provided by web objects, while others require adding additional (application-specific) meta information. *7DS* either passes the received reports for display or issues an HTTP GET request automatically (via the web client) using the local URL of the selected report to receive the complete object. A miniature web server is running as a part of *7DS* which services the HTTP GET requests. The primary information propagation is through the use of caching rather than reliable state maintenance. It is not a goal of the current prototype to resolve inconsistency among copies of a data object. *7DS* peers may have several objects matching a single query.

7DS organizes and indexes the cache. Through a GUI, it provides the capability to the user to view, browse and manage the cache. In the current prototype, the content of the cache is displayed in a tree-like structure. We are extending it to support grouping of the cache content to predefined categories and searching tools using the meta-data attributes of the stored objects. The user can set the access permissions for files and directories in the cache and specify the objects to share with other peers. To protect the user privacy, the system only transmits reports or pages that corresponds to public available objects. *7DS* can encrypt a private object before its transmission. Both the source and the destination are required to have PGP in order to encrypt (using the public key) and then decrypt (using the private key) the message, respectively. Alternatively, if the sender and receiver share a private key, they can use

that, too.

The user can also explicitly add or remove files from the cache and pending queries. Periodically, *7DS* removes expired objects, updates the index with these changes and also includes newly cached objects. Instead of deleting the expired objects, the system may try to prefetch them again transparently. Through the GUI, the user marks which pages need to be prefetched when they expire. When *7DS* searches the cache to remove the expired objects, for each marked page, it forms a request and adds it in the query list.

7DS trades power for data only when the battery level is above a threshold. Via a battery monitor and a power management protocol, *7DS* aims to adapt communication to reduce energy consumption during idle periods, when there is low expectation for data or collaboration and when the battery life is below a threshold. Generally, the prediction for data is a hard problem. In order to decide for the data availability, we currently use advertisements from the servers and the traffic in the network. When power conservation is enabled, the mobile host periodically turns off its wireless LAN interface. *7DS* periodically checks the battery level and adapts its collaboration accordingly. In particular, the system adapts its communication with other *7DS* peers by tuning several “thresholds” in the battery level. For example, it may set three values of the battery: when the battery level is above the highest value, the system can fully collaborate. Within the highest and second highest value, the system only applies a partial participation in the system. Below the third value, the system stops participating in the *7DS* network. Generally, the degree of participation depends on the querying (active or passive, frequency interval) and type of collaboration (data sharing and forwarding support). *7DS* is engaged entirely in the participation when it both actively and passively queries, and supports data sharing and forwarding. In a par-

tial participation, it disables forwarding and switches from active to passive querying. The default setting is as follows: for battery level above 75%, *7DS* uses data sharing and active querying, when it is within 50% to 75%, it switches to passive querying, and when the battery level falls below 50%, it stops participating in the system. The user can change this setting via a GUI. The power conservation method we consider alternates from the *on* state of the network interface to the *off* state. During the interval that the network interface is on, *7DS* communicates with the other hosts by sending queries, forwarding or receiving reports or data. The mobile host broadcasts a query at each *on* interval till it receives the data. It is clear that the smaller the *on* interval, the more the power savings. However, with smaller intervals, the degradation of data dissemination is larger. To prevent this degradation of data dissemination and still conserve power, we support the synchronous mode. When synchronous mode is enabled, *7DS* have synchronized the intervals their network interface is on. If there is a GPS receiver, it would get clocks synchronized to sub-microsecond level. Otherwise, the devices need to periodically synchronize their clocks, for example, querying other *7DS* clients with GPS. In Section 3, we discuss its effect on data dissemination and the power savings in more detail. In summary, *7DS* uses IP multicast for querying peers, HTTP for receiving the complete data and existing software for searching and indexing caches. An important feature of our architecture is its easy deployment. The system displays both the reports as well as the complete data information using any browser. It is transparent to wired and wireless networks as well as to different information providers that participate in the system. Also, *7DS* is flexible to support different applications and able to form queries and application-specific criteria for the selection of the appropriate cached copies as long these applications access the data by data object retrieval

using URLs. Users only need to install the software and *7DS* configures itself with minimal manual intervention; the system does not require any registration for data distribution. The system is resilient to failures and inconsistencies that occur in this dynamic environment. *7DS* is resource aware and tries to utilize the constrained resources efficiently.

The prototype is written in Java and uses the Glimpse search engine [11]. We are in the process of implementing it on Windows CE. Details on the implementation can be found at [12].

3 Performance Evaluation

3.1 System models and operation modes

7DS can operate in different modes that depend on the cooperation strategy among peers (data sharing, forwarding), power conservation and query mechanism (active, passive querying). To investigate its performance and in particular the effect of transmission power and the different modes of operation on data distribution, we evaluate P-P and S-C along with their variants. For simplicity, we refer to the *7DS* hosts in these schemes as nodes or peers and the *7DS* host that has the data originally in the S-C schemes as server. In the P-P schemes, all nodes are mobile with active querying enabled. We simulate three variations depending on the type of cooperation, namely data sharing (DS), forwarding (FW) and both data sharing and forwarding enabled (DS+FW). When forwarding is enabled, upon the receipt of a query or data, *7DS* peers rebroadcast it, if they have not rebroadcasted another message during the last 10 s. The last condition is a simple mechanism for preventing flooding in the network. We separate the S-C schemes into the “straight” S-C without any cooperation among clients (namely, FIS and MIS) and some hybrid ones with cooperative clients. In FIS (MIS) scheme, there is a fixed (mobile) host

with the data that acts as a server. The remaining nodes (clients) are mobile, non cooperative with active querying enabled and without any power conservation mechanism. They receive data only from the server. The hybrid schemes are with passive querying enabled and fixed server. As we mentioned in Section 2, in passive querying mode, the server sends an advertisement every 10 sec. Hosts send queries upon the receipt of an advertisement.

Model	Cooperation	Options	Querying
S-C	only server, server mobile/fixed (FIS/MIS)	DS (only server)	active
P-P	all hosts	DS, FW, DS+FW	active
Hybrid	fixed server, cooperative peers	DS, FW, DS+FW	passive

Table 1: Summary of the schemes with their querying mechanism.

Let us describe the main motivations for the comparisons we make in the remaining section. The P-P vs. straight S-C comparison is to understand the effect of the cooperation among mobile peers. The P-P and MIS vs. FIS shows how mobility affects data dissemination. In particular, the MIS vs. FIS comparison focus exactly on the effect of server mobility on data dissemination.

3.2 Model assumptions

Nodes move in a 1000 m x 1000 m area according to the random waypoint mobility model [2]. This random walk-based model is frequently used for individual (pedestrian) movement [2, 3, 4]. The random waypoint breaks the movement of a mobile host into alternating motion and rest periods. Each mobile host starts from a different position and moves to a new randomly chosen destination. For each node, the

initial and end point for each segment are distributed randomly across the area. Each node is moving to its destination with a constant speed uniformly selected from (0 m/s, 1.5 m/s). When a mobile host reaches its destination, it pauses for a fixed amount of time, then chooses a new destination and speed (as in the previous step) and continues moving.


The query interval consists of an *on* and *off* interval. The broadcast is scheduled at a random time selected from the *on* interval. The asynchronous mode is the default power conservation method. We explicitly denote the schemes with synchronous mode enabled with the word “sync”. In schemes with no power conservation, the *off* interval is equal to 0 and the *on* and query interval are the same. The exchange of queries, reports and advertisements takes place during the *on* interval. Generally, the transmission of the complete data object (for example, web page) is scheduled separately. For example, the dataholder may select a time (“rendezvous point” in which the HTTP transmission takes place) and include it in the report message. At that time, both the querier and dataholder set their network interface on and the querier initiates the HTTP get request (as described in Section 2). In the simulations, we concentrate only on the exchange of *7DS* queries, reports and advertisements. A cooperative dataholder responds to a query by sending the data item in the report. In this simulation study, we assume one data object, and all hosts in the area are interested in this data item. When a host receives a report for this data item, it becomes dataholder. This simplification is reasonable in order to investigate the dominant parameters on data dissemination.

A *scenario* (file) “defines” the topology and movement of each host that participates in an experiment. We consider different number of hosts in the area. Later in this section, we scale the area and vary the density of the hosts and their wireless coverage. We would like to emphasize that this host density does


not necessarily represent the total number of hosts in that area, but just indicates the popularity of the defined data object. By varying the density of hosts, we study how data items of different popularity disseminate in such environment. We speculate that in an urban environment such as Lower Manhattan, near the platform of the train or subway stop in a rush hour, there will be from four to 25 users that could be interested to get the local and general news using PDAs or other wireless devices. A density of 25 hosts per km² corresponds to very popular data whereas a density of five hosts per km² corresponds to a more typical data object [13]. We generate 300 different scenarios for different density values.

In each of these scenarios, the mobility pattern of each host is created using the mobility pattern we described, except in the FIS-based schemes, that the server is stationary. We run simulations using these scenarios, for the different schemes of Table 1. The wireless LAN is modeled as an 802.11 network interface. We use the ns-2 simulator [14] with the mobility and wireless extensions [15]. We consider transmission powers of 281.8 mW (high), $\frac{281.8}{24}$ mW (medium) and $\frac{281.8}{28}$ mW (low). Assuming the two-ray ground reflection model these transmission powers correspond to ranges of approximately 230 m, 115 m and 57.5 m, respectively. Considering the simulation results on these ranges, we can speculate for the performance of *7DS* with lower transmission power. We do not deal with very low power transmissions (such as of Bluetooth devices) that correspond to less than ten meters. In these simulation sets, we concentrate on pedestrian users in an open campus, platform, or metropolitan area. After all, data sharing among low power devices is more likely to occur in a different setting, such as an office environment, where the mobility pattern appear to be deterministic and all wireless devices tend to cluster together.

We evaluate the effectiveness of our approaches, by computing the percentage of nodes that acquire



(a)



(b)

Figure 3: Percentage of dataholders after 25 minutes for medium and low transmission power, respectively. The query interval is 15 sec.

Parameter	Value
Pause time	50 sec
Mobile user speed	(0,1.5) m/sec
server advertisement interval	10 sec
Forward message interval	10 sec

Table 2: Simulation constants.




Figure 2: Percentage of dataholders after 25 minutes for high transmission power. The query interval is 15 sec.

the information after a period of time. In the percentage we do not include the node that has the data at the beginning of the simulation. We also compute the average delay until a mobile host receives the information from the time it sends the first query. We run the 300 generated scenarios for each test and computed the average of the percentage of hosts that *become* dataholders by the end of each test. The default simulation time is 25 minutes. However, we also investigate data dissemination over time and vary the simulation time from 150 sec to 50 minutes. The hosts start querying in average after 90 sec from the beginning of the simulation. The 95% confidence interval for the average percentage of dataholders is

within 0-11% of the computed average, with the variance tending to be higher for low host density. Figures 2 and 3 show the percentage of dataholders as a function of the density of hosts for P-P and S-C schemes. In this set of simulations, the query interval is 15 sec. For high transmission power, as in Figure 2, *7DS* proves to be an effective data dissemination tool. Even when the network is sparse, 77% of the users will acquire the data during the 25 minutes of the experiment. For networks with ten or more hosts, more than 96% of the users will acquire the data during the 25 minutes. For host densities of 25 hosts per km^2 , the probability of acquiring the data is very close to 100%. The P-P vs. FIS comparison illustrates the effect of data sharing among mobile peers. In Figure 2, in a setting of 25 hosts, P-P schemes outperform FIS by 55%. In particular, in P-P, 99.9% of hosts will acquire the data after 25 minutes, compare to 42% of the users in the FIS. For lower transmission power, P-P outperforms FIS by 20% to 70% (Figures 3). The impact of data sharing among peers is also apparent in hybrid schemes. In particular, the hybrid vs. S-C schemes for density of ten or more hosts per km^2 and medium or high transmission power.

Notice that forwarding in addition to data sharing does not result in any further performance improvements. This is due to the low probability that a case as the following occurs: There is a querier A and a dataholder C that cannot listen to each other, and a third host B that can communicate with both and forward data. Moreover, A will not acquire the data directly from a dataholder till the end of the test. We would like to emphasize that this is true also for smaller simulation times, starting from 150 sec (just a few seconds after the hosts start querying). Due to the lack of space, we do not include the graphs that illustrate it; the reader is referred to [13] for more information. In a more dense setting of mobile hosts that forward messages independent on their data in-

terests, we expect forwarding to have a higher impact. Forwarding without data sharing results in a performance improvement. For example, in Figures 2 and 3 hybrid schemes with forwarding enabled outperforms FIS by 4%-40% depending on transmission power.

As we expect for both FIS and MIS, their performance remains constant as the number of hosts increases, since a data exchange takes place only when a querier is in proximity to the server. In addition, notice that in Figures 2, Figure 3 (a) and (b) MIS outperforms FIS by approximately 22%, 16%, 6%, respectively. An intuitive explanation is based on the fact that, in MIS schemes, the relative speed of the server from the clients is larger than in FIS schemes (where the server is fixed). Therefore, the mobile information server will meet with more hosts and disseminate the data faster. On the other hand, as we expect, the density of hosts affects the schemes that are based on peer-to-peer cooperation. As the number of hosts increases from five to 25 hosts, in P-P schemes with medium transmission power, the performance improves substantially.

Measurements of the power consumption of the wireless network interfaces have shown that they consume substantial power even when they are idle (powered on but not sending or receiving packets). Moreover, receiving packets costs marginally more energy than being idle [16]. Using the asynchronous power conservation mechanism we described with the *on* and *off* intervals to be equal there is 50% power savings, since the network interface is on only half the time. As Figures 2 and 3 illustrate there is some degradation in data dissemination. This is due to the decrease of the time interval the hosts can communicate. If we keep the query interval constant and reduce the *on* interval, the smaller the *on* interval, the higher power savings. However, with smaller intervals, the degradation of data dissemination is larger. To prevent this degradation, we enable the




Figure 4: Figures (a) and (b) illustrate the impact of synchronous mode on data dissemination in a P-P with data sharing and a Hybrid with data sharing and forwarding scheme, respectively. Query interval is 15 sec and the on period in “sync” schemes is 1.5 sec. The simulation time is 25 minutes.

synchronous mode. As we show in Figure 4, when the synchronous mode is enabled, even with a small *on* interval, the power conservation does not cause any degradation of the data dissemination. More specifically, Figure 4 (a) illustrates P-P schemes with data sharing and Figure 4 (b) hybrid schemes with data sharing and forwarding. The query interval is 15 sec, in which, during the first 1.5 sec the network interface is *on* and at the remaining time (13.5 sec) switches *off*. In an ideal setting without packet losses and need for retransmission, the number of messages exchanged in the P-P schemes without power conservation and the ones with sync power conservation are the same. Therefore, the power spendings due to message receiving/sending is the same, whereas the period in the idle state is reduced (the network interface is on only during 10% of the time). The synchronous mode may result in a 90% reduction in power dissipation. In general, hosts may query for different data items. In such dense setting of hosts retransmissions and packet losses may result in further power spendings. It is part of our future work to

investigate further the synchronous power conservation mode and the impact of retransmissions, packet loss and on interval in such environment.

In [10], we investigated the performance of the system as a function of the query interval using the asynchronous power conservation method (i.e., the on interval is half the query interval and not synchronized). The degradation in the FIS performance is relative small compared to P-P schemes as the query interval increases. This is due to the high probability that a mobile host that gets in close proximity to a server acquires the data (i.e., there is sufficient time to broadcast a query and receive the data). In P-P schemes the impact of the query interval can be more apparent. In a setting of 25 hosts with medium transmission power, data sharing enabled and no power conservation enabled, when the query interval increases from 15 sec to 3 minutes the degradation is approximately 30% and for five hosts, it reaches 50%. However, using the synchronous power conservation, even when we maintain a low ratio of the on-interval (e.g., 5% with on interval to be 6 sec and




Figure 5: S-C: Fixed information server schemes. Figure (a) illustrates the probability of not acquiring the data ($1-p$) as a function of time. Figure (b) depicts the average delay to receive the data as a function of the probability to acquire it in FIS.




Figure 6: P-P: data sharing without power conservation schemes. The Figures (a) and (b) illustrate the average delay to receive the data as a function of the probability to acquire it.

query interval 2 minutes), we expect the degradation to be much weaker. We need to investigate further what is the optimal on and query interval and when a mobile host need to switch to passive querying to utilize its battery more efficiently.

Let us now discuss the scaling effect in our simulation. We focus on expanding the area but keep the movement pattern the same (i.e., average speed and distance that the user travels each time, before pausing). Both in P-P and FIS schemes, when we

expand the area, but keep the density of hosts and their transmission power fixed, the performance of data dissemination remains the same. That is, our simulation results are robust. Let $p(t)$ denote the probability a host will acquire the data by time t . Figure 5 (a) shows the probability that a host will *not* acquire the data by time t , i.e., $1 - p(t)$ (or survival probability) on a logarithmic scale. This figure shows the percentage of data holders at time t using the transformation $(\log(1 - p(t)))^2$. Their shape indicates that $p(t)$ in FIS follows the $1 - e^{-\sqrt{at}}$. As we fix the density of servers to one server per square kilometer the scaling of the area (to 2km x 2km or 3km x 3km) does not affect data dissemination (for the same transmission power). In P-P settings (e.g., P-P with data sharing and power conservation) $p(t)$ grows faster than in FIS and our simulation results indicate that the P-P with data sharing and power conservation can be approximated by the function $1 - e^{-at}$, especially, for not very dense settings (e.g., with less than 20 hosts per km² transmitting with high or medium power). For very dense settings, this probability grows even faster.

It is interesting to compare the effect of density of cooperative hosts vs. their wireless coverage density. This is related to a design decision for deploying servers to cover a certain area with wireless data access. Assuming the same total area of wireless coverage, would be more efficient to deploy a number of servers transmitting with high power or increase their number but reduce their transmission power lower. In respect of total power spendings and wireless throughput utilization the setting with the larger number of servers is more attractive. However, it is not clear which setting performs better in terms of data dissemination and average delay. We found that when the wireless coverage of the cooperative hosts are the same, the schemes with higher host density of cooperative hosts perform better. This is true

both with FIS as well as with P-P schemes. In Figure 5 (b) we compare two FIS settings. The first has one server in a 2kmx2km area with high transmission power and the latter four servers in a 2kmx2km area with medium transmission power. The case with higher density of servers perform better. For example, a 20% probability of acquiring the data, the FIS scheme with higher density of servers corresponds to an average delay of 500 s. In the case of same wireless coverage, but lower density of servers, the same data access probability corresponds to a double average delay. Figure 6 illustrates similar results in P-P schemes for different density of hosts. For a 40% probability of acquiring the data, the system provides an 600 sec average delay in the higher density of hosts setting (5 peers/sq.km) whereas in a lower density setting the average delay doubles.

Finally, for each test, we compute the average delay of the nodes that acquired the data by the end of simulation. Then, we take the average over all 300 sets, excluding the ones without new dataholders. First, let us fix the simulation time to 25 minutes and compare P-P and FIS schemes in terms of average delay for the same probability of acquiring the data (Figures 2, 5 (b) and 7 (a)). In P-P with data sharing and no power conservation, for high transmission power, the average delay is as high as 6 minutes for sparse networks and reaches 77 sec for dense networks (Figure 7 (a)). In the case of low transmission power, it reaches 13 minutes. Using FIS, the average delay is constant (over the number of hosts in the area) and for high transmission power is 6 minutes (Figure 5 (b)), whereas for low transmission power it reaches 9 minutes. Let us now compare FIS and P-P with data sharing and no power conservation (Figure 7 (b)). For a simulation time of 470 sec, the number of hosts that acquire the data in P-P with high transmission power reaches 40% with an average delay of 135 sec. With the same delay and using FIS, 30% of hosts will acquire the data. With FIS a 40% probab-

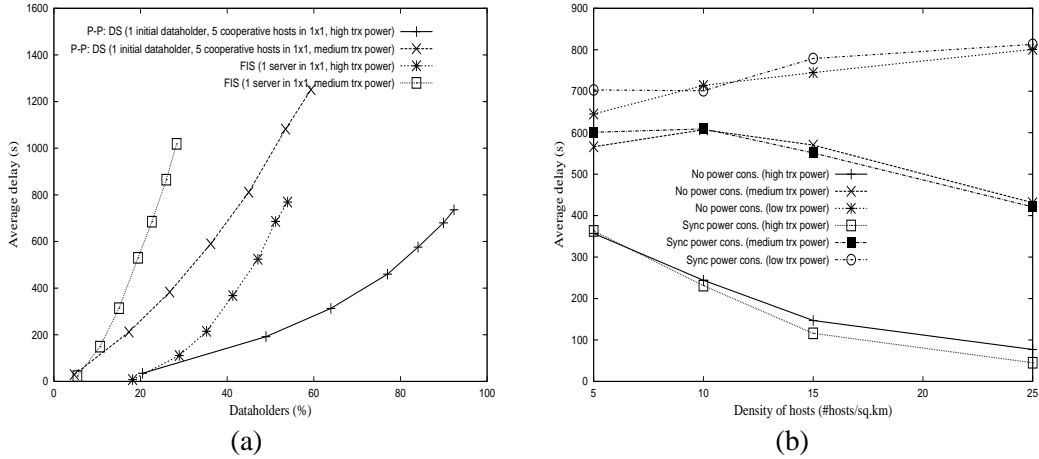


Figure 7: Figure (a): Average delay as a function of the probability to acquire the data. Figure (b): Average delay as a function of the host density for P-P with data sharing schemes and a simulation time of 25 minutes.

ity of acquiring data corresponds to an average delay of 6 minutes. For a higher average delay of 10 minutes, 85% hosts will acquire the data using P-P and 50% using FIS. In the case of medium transmission power, with an average delay of 315 sec, a host will get the data with a probability of 15% and 22% using FIS and P-P, respectively. For P-P, the 15% of dataholders corresponds to an average delay of 3 minutes. Increasing the average delay to 10 minutes, this probability becomes 21% and 37% for FIS and P-P, respectively. Using synchronous P-P with data sharing, the difference from FIS is more prominent. A host using FIS will get the data with a 40% probability and an average delay of 6 minutes whereas using (sync) P-P even in the case of low density, the probability is double (Figure 7).

To summarize the simulation results,

- P-P schemes outperform S-C schemes. The results indicate that, the probability that a host that queries for a data object will acquire it by time t using FIS and P-P, follows the $1 - e^{-a\sqrt{t}}$ and $1 - e^{-at}$, respectively. In case of high den-

sity of cooperative hosts, the data dissemination using P-P grows even faster. Generally, their difference becomes more prominent in cases of lower transmission power (medium and low) with more than ten hosts. In our setting with ten or more hosts per km^2 , P-P provides 60% or higher probability for acquiring the data item to hosts that move in the area for 25 minutes and transmit with medium or high power. This probability is three times higher than in FIS. In some of the cases, their difference in average delay is negligible, in other cases FIS has lower average delay (with a maximum difference of 100 sec).

- Forwarding (i.e., rebroadcasting $7DS$ messages upon their receipt) in addition to data sharing does not result in any performance improvements.
- The query interval has negligible effect on S-C schemes.
- The synchronous power conservation method is

beneficial. It increases the power savings without degrading the data dissemination.

- The performance remains the same when we scale up the area, but keep the density of cooperative dataholders and transmission range fixed.
- Dominant parameters are the density of cooperative hosts and their wireless coverage density. Also, their mobility can contribute to higher data dissemination. More specifically, for a given wireless coverage density, the larger the density of cooperative hosts, the better the performance. For example, in both the P-P and FIS schemes, for the same wireless coverage, it is more efficient to have a larger number of servers with lower transmission power than fewer with higher transmission power.

4 Data dissemination as a diffusion-controlled process

This section discusses our initial efforts to model data dissemination as a diffusion controlled process using theory from random walks and environment [17] and the kinetics of diffusion-controlled chemical processes [18]. In particular, we use the diffusion in a medium with randomly distributed static traps to model the FIS approach of data dissemination. Let us first define the static trapping model. Particles of type C perform diffusive motion in d -dimension space. Particles of type S (“sinks” or traps) are static and randomly distributed in space. Particles C are absorbed on particles S when they step onto them. The basic trapping model assumes traps of infinite capacity. For Rosenstock’s trapping model in d dimensions (with a genuinely d -dimensional, unbiased walk of finite mean-square displacement per step), it is showed that the large- n behaviour of the survival

probability

$$\log(\phi_n) \approx -\alpha[\log(\frac{1}{1-q})]^{2/(d+2)}n^{d/(d+2)} \quad (1)$$

In Eq. 1 α is a lattice-dependent constant and q denotes the concentration of the independently distributed, irreversible traps. For the modeling, we also assume that when a *7DS* querier comes in close proximity to the server, it always acquires the data. That is, its duration within the server’s coverage is more than the query interval. As we mentioned in Section 3.1, in FIS the information sharing takes place among the server and the querier. We model the stationary information servers as traps and the mobile peers as particles C. When a host acquires the data, it stops participating in the system, and with respect to the model is considered “trapped”. Figure 8 illustrates the analytical and simulation results on data dissemination as a function of time. The analytical results on Trap model are derived from Eq. 1 (i.e., Rosenstock’s trapping model) for high and medium wireless coverage.

We define q as $\pi R^2 N_{servers}/A$ and use an R equal to 230 m and 115 m for high and medium wireless coverage, respectively. For the simulation results on FIS in Figure 8, we use the FIS simulations we described in Section 3. Note that, using Eq. 1 the $1-\phi_n$




Figure 8: Simulation and analytical Trap model results on FIS.

expresses the fraction of hosts that acquire the data at time n . As Figure 8 illustrates, our simulations are consistent with Eq. 1 for α equal to 0.021. That is, using the Eq. 1, the $(1 - \phi_n) * 100\%$ match our simulation results for the percentage of dataholders at time n for the FIS scheme we described.

We need to emphasize that this is an initial investigation for the analytical problem and the analysis assumes a finite mean-square displacement per step, a mobility pattern that differs from the randway model we use in our simulations and consider more realistic.

5 Related work

In earlier work [19], we investigated a different facet of cooperation, namely network connection sharing. Using network connection sharing, mobile devices with multiple wireless interfaces can serve as temporary gateways to wide-area wireless networks.

Napster [20] and Gnutella [21] are two systems that explore the cooperation among hosts and enable data sharing among users in a fixed wired network. The first is focus in sharing music files, whereas the latter for any type of files. In contrast to Gnutella, *7DS* does not need to discover its neighbors or maintain connections with them, but only multicast its queries to a well known group. Unlike Napster *7DS* operates in a distributed fashion without the need of any central indexing server. Moreover, Napster requires user intervention and effort for uploading files, whereas *7DS* does this automatically.

Infostations have first been mentioned by Imielinski in the DataMan project [8]. Badrinath was among the first to propose an infrastructure for supplying information services, such as e-mail, fax and web access by placing infostations at traffic lights and airport entrances. Imielinski *et al* [7] investigate methods for accessing broadcast data in such a way that running time (which affects battery life) and access

time (waiting time for data) are minimized. They demonstrate that providing index or hashing based access to the data transmitted over the wireless can result in significant improvement in battery utilization. Barbara *et al* [22] propose and study a taxonomy of difference cache invalidation strategies and study the impact of client's disconnection times on their performance.

Caching and prefetching have been successfully employed to alleviate user perceived latencies and there has been extensive research. In the context of mobile users, hoarding is a similar technique to prefetching to improve the data availability (for users that experience intermittent connectivity) [23, 24]. In a similar context to ours, prefetching targeted for mobile users in a wide-area wireless network has been used in [25]. Tao Ye, *et al* [25] assume an infostation deployment. Their prefetching algorithm uses location, route and speed information to predict future data access. Their emphasis is on devising and evaluating techniques for building network-aware applications. They describe an intelligent prefetching algorithm for a map-on-the-move application that delivers maps, at the appropriate level of detail, on demand for instantaneous route planning. When a mobile user enters an infostation coverage it prefetches a fixed amount of k bytes that corresponds to a map with a certain level of detail, where k depends on user speed. They investigate the effectiveness of infostations as compared to a traditional wide-area wireless network. There are two main differences of their setting with our FIS based schemes. First, in their environment, mobile clients are constantly connected to a wireless network. Devices are using a high bandwidth link when they are within infostation coverage. Outside these regions, their requests are passed to the server via a conventional cellular base-station. In our case, the mobile hosts have no wide-area network access. Second, they investigate the effectiveness of (fixed) infostations as compared

to a traditional wide-area wireless network. For that, they vary the infostation density and its coverage. In our case, we consider a fixed infostation (i.e., FIS) in the region of 1 km x 1 km (that corresponds to low infostation density). As we explained in the Section 1, the focus of this paper is to investigate a different data access method, namely, peer-to-peer data sharing among mobile users. For its evaluation, we compare it to the access via an infostation. Also, we vary several parameters (like various mobility patterns, power conservation methods and querying schemes) that have not been investigated in [25]. We should mention that their qualitative result, that having many infostations that cover small ranges is a more optimal topology than having few infostations that cover large ranges is consistent with ours.

Kravets, *et al* [26] present an innovative transport level protocol that achieves power savings by selectively choosing short periods of time to suspend communication and shut down the communication device. It queues data for future delivery during periods of communication suspension, and decides predicting when to restart communication. This work motivated us to consider schemes for predicting high data availability in our setting to power on the communication device and start *7DS*. In Section 3, we discuss this in more detail.

There is substantial peer-to-peer work in the file system and OS literature that is relevant, including the Ficus [27], JetFile [28], Bayou [29] projects. All of them are replicated storage systems based on the peer-to-peer architecture. Ficus is a distributed file system aiming to a wide-scale, Internet-based use. It supports replication using a single-copy availability, optimistic update policy. Its main focus is on the consistency among the different copies and reconciliation algorithms to reliably detect concurrent updates and automatically restore consistency. Like Ficus, Bayou provides support for application-dependent resolution of conflicts, but unlike Ficus, it

does not attempt to provide transparent conflict detection. JetFile requires file managers to join a multicast group for each file they actively use or serve. Our system is targeted in a different environment and addresses different research issues. The primary concern of our work is the effect of the wireless coverage, collaboration strategy and power conservation method in the data dissemination across mobile hosts, rather than consistent replication.

Flooding and gossiping (a variant on flooding, that sends messages only to some neighboring nodes instead of all) protocols have been also studied extensively. For example, [30] presents a protocol for information dissemination in sensor networks. In their setting, the sensors are fixed and the network fully connected. They measure both the amount of data these protocols disseminate over time and the amount of energy the dissipate. It features meta-data negotiation prior to data exchange to ensure that the latter is necessary and desired, eliminating duplicate data transmissions, and with power resource awareness. They compare their work with more conventional gossiping and flooding approaches. A more theoretical work [31] assumes a system where the nodes are placed on a line. They present an optimal algorithm for broadcasting and compute the expected number of time steps required for it to complete. More theoretical studies on information dissemination have used percolation theory [32] or epidemic models.

6 Conclusions and future work

In this paper, we presented *7DS*, a new peer-to-peer data sharing system. *7DS* is an architecture enabling the exchange of data among peers that are not necessarily connected to the Internet. It anticipates the information needs of users and fulfills them by searching for information among peers. To as-

sess the efficiency of information dissemination via *7DS*, and also investigate the effect of the wireless coverage range, network size, query mechanism, cooperation strategy among the mobile hosts and power conservation over time, we performed an extensive simulation-based study. This study involves two main data dissemination approaches, namely P-P and S-C. We measured the percentage of data holders and also the average delay until a querier gets the data item since it sent the first query. We found that the density of cooperative hosts and their mobility are dominant parameters for data dissemination. We studied the scaling effect in our simulations and found that the performance remains the same when scale the area, but keep the density of cooperative dataholders and power transmission fixed. Also, the simulation results indicate that for fixed wireless coverage density, the larger the density of cooperative hosts, the better the performance. For example, both in P and FIS schemes, for the same wireless coverage, it is more efficient to have a larger number of servers with lower transmission power than fewer with higher transmission power. The combination of power conservation and synchronization of the network interface *on* periods is beneficial for *7DS* users, since it increases substantially the power savings without degrading data dissemination. For example, for an average delay of 6 minutes, a host using FIS gets the data with a 40% probability. Whereas with sync P, even in the case of low density (five hosts per km² interested in the same data), this probability is double.

An attractive feature for *7DS* is a mechanism that would indicate the appropriate interaction (P-P, S-C with active or passive querying) based on several parameters, such as data availability prediction, cooperative users in close proximity and battery level. Advertisement messages from servers or other cooperative hosts and traffic measurement tools can provide hints. It is part of future work to investigate how

they can be used to improve the power utilization and provide cost effective solutions. Also, it would be interesting to investigate the percentage dataholders that actually read/use the data. For that, a better understanding of the environment and the applications (average delay requirement, accuracy of prefetching prediction, notification mechanism for newly arrived data) is required. Our current research direction includes the integration of *7DS* with a tour guide and an academic news notification system, and cache servers and its deployment on the Columbia campus. Design issues related to the security and user privacy is a part of the on-going effort and the reader can find a relevant discussion in [13]. We also plan to continue the theoretical study for modeling the data dissemination for both the P-P and S-C schemes.

References

- [1] Metricom, “The Ricochet wireless network overview,” 1999.
- [2] J. Broch, D. Maltz, D. Johnson, Y.-C. Hu, and J. Jetcheva, “A performance comparison of multi-hop wireless ad hoc network routing protocols,” in *ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom)*, (Dallas, Texas), Oct. 1998.
- [3] E. M. Royer and C. E. Perkins, “Multicast operation of the ad-hoc on-demand distance vector routing protocol,” in *ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom)* (ACM, ed.), (Seattle, Washington), pp. 207–218, Aug. 1999.
- [4] Y. Xu, J. Heidemann, and D. Estrin, “Geography-informed energy conservation for ad-hoc routing,” in *ACM/IEEE International Conference on Mobile Computing and*

- Networking (MobiCom)*, (Rome, Italy), Aug. 2001.
- [5] D. Estrin, R. Govindan, J. Heidemann, and S. Kumar, "Next century challenges: Scalable coordination in sensor networks," in *ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom)* (ACM, ed.), (Seattle, Washington), pp. 263–270, Aug. 1999.
- [6] P. Castro, B. Greenstein, R. Muntz, and M. Papadopouli, "Locating application data across service discovery domain," in *ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom)*, (Rome, Italy), Aug. 2001.
- [7] T. Imielinski, S. Viswanathan, and B. R. Badrinath, "Energy efficient indexing on air," in *International conference on Management of Data, ACM SIGMOD*, (Minneapolis), 1994.
- [8] DATAMAN. <http://www.cs.rutgers.edu/dataman/>.
- [9] D. Barbara and T. Imielinski, "Sleepers and workaholics: Caching strategies in mobile environments," 1994.
- [10] M. Papadopouli and H. Schulzrinne, "Effects of power conservation, wireless coverage and cooperation on data dissemination among mobile devices," in *ACM Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)*, (Long Beach, California), Oct. 2001.
- [11] Glimpse. <http://www.webglimpse.org>.
- [12] 7DS. <http://www.cs.columbia.edu/~maria/7ds>.
- [13] M. Papadopouli and H. Schulzrinne, "Effects of power conservation, wireless coverage and cooperation on data dissemination among mobile devices," Tech. Rep. CUCS-005-01, Columbia University, 2001.
- [14] K. Fall and K. Varadhan, "ns: Notes and documentation," tech. rep., Berkeley University, 1998. Technical Report.
- [15] "Wireless and mobility extensions to ns-2." <http://www.monarch.cs.cmu.edu/cmu-ns.html>.
- [16] M. Stemm and R. H. Katz, "Measuring and reducing energy consumption of network interfaces in hand-held devices," *IEICE Transactions on Communications*, vol. E80-B, no. 8, p. 1125–31, vol. E80-B, no. 8, pp. 1125–31, 1997.
- [17] B. D. Hughes, *Random Walks and Random Environments*. Oxford Science Publications, 1995.
- [18] A. Ovchinnikov, S. Timashev, and A. Belyy, *Kinetics of Diffusion controlled chemical processes*. Nova Science Publishers, 1989.
- [19] M. Papadopouli and H. 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)*, (Basking Ridge, New Jersey), pp. 169–185, June 1999.
- [20] Napster. <http://www.napster.com>.
- [21] Gnutella. <http://gnutella.wego.com>.
- [22] D. Barbara and T. Imielinski, "Sleepers and workaholics: Caching strategies in mobile environments," 1994.
- [23] J. J. Kistler and M. Satyanarayanan, "Disconnected operation in the coda file system," *Thirteenth ACM Symposium on Operating Systems Principles*, vol. 25, no. 5, pp. 213–225, 1991.

- [24] H. Lei and D. Duchamp, “An analytical approach to file prefetching,” in *USENIX Annual Technical Conference*, (Anaheim, CA), Jan. 1997.
- [25] T. Ye, H.-A. Jacobsen, and R. Katz, “Mobile awareness in a wide area wireless network of info-stations,” in *ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom)*, (Dallas, Texas), Oct. 1998.
- [26] R. Kravets and P. Krishnan, “Power management techniques for mobile communication,” in *ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom)*, (Dallas, Texas), pp. 157–168, Oct. 1998.
- [27] T. W. Page, R. G. Guy, J. S. Heidemann, D. Ratner, P. Reiher, A. Goel, G. H. Kuenning, and G. J. Popek, “Perspectives on optimistically replicated peer-to-peer filing,” *Software—Practice and Experience*, vol. 28, pp. 155–180, February 1998.
- [28] B. Gronvall, A. Westerlund, and S. Pink, “The design of a multicast-based distributed file system,” in *Operating Systems Design and Implementation*, pp. 251–264, 1999.
- [29] M. Spreitzer, M. Theimer, K. Petersen, A. Demers, and D. Terry, “Dealing with server corruption in weakly consistent, replicated data systems,” in *ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom)*, (Budapest, Hungary), pp. 234–240, Sept. 1997.
- [30] J. Kulik, W. Rabiner, and H. Balakrishnan, “Adaptive protocols for information dissemination in wireless sensor networks,” in *ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom)* (ACM, ed.), (Seattle, Washington), pp. 24–35, Aug. 1999.
- [31] K. Ravishankar and S. Singh, “Broadcasting on $[0,L]$,” 1994.
- [32] R. Durrett, *Lecture notes on particle systems and percolation*. Pacific Grove, CA, 1988.