Opportunistic Data Dissemination Using Real-World User Mobility Traces

Andreas Heinemann
Darmstadt University of Technology
Darmstadt, Germany
aheine@tk.informatik.tu-darmstadt.de

Jussi Kangasharju
University of Helsinki
Helsinki, Finland
Jussi.Kangasharju@cs.helsinki.fi

Max Mühlhäuser
Darmstadt University of Technology
Darmstadt, Germany
max@tk.informatik.tu-darmstadt.de

***Abstract*—**

Opportunistic communication allows humans equipped with mobile devices to exchange information via a wireless link whenever they are nearby. This work examines the performance of a profile-based data dissemination scheme. We use a *2-step* simulation approach which combines realistic user mobility traces with synthetic mobility models.

We consider different system configurations and evaluate their effects on the data dissemination scheme. In particular, we look at how different wireless networking technologies affect the system and how much can we benefit from the installation of fixed information sprinkler nodes. Furthermore, we investigate the effects of user behavior by comparing selfish and generous users.

Our results clearly demonstrate that realistic user mobility patterns result in a feasible opportunistic data dissemination scheme. In most of the studies scenarios, the effectiveness of the dissemination is very high, with over 90% of the users reached within a week of real time. Our results on user behavior show that although the effects of selfish behavior can be mitigated, they cannot be completely eliminated.

I. Introduction

Opportunistic networks are gaining momentum in the research community. They exploit the ubiquitous wireless communication capabilities of small, portable devices to build larger systems. Such networks are attractive counterparts to infrastructure-based networks which typically require large investments. As users with wireless-enabled devices come into communication range, their devices spontaneously discover each other. This lets the devices exchange information stored on the devices, according to user preferences or wishes. The natural user mobility patterns bring users in contact with others; thus the data dissemination in an opportunistic network mimics that of word-of-mouth information propagation between people.

In this work, we evaluate the effectiveness of such an opportunistic data dissemination scheme based on real-world user mobility traces. We use the data from the Reality Mining Project [1] as a baseline for determining how users in our system meet each other. We couple these macro-level mobility traces with synthetic mobility models in a *2-step simulation* model and determine the effectiveness of our opportunistic data dissemination scheme. We explore several different system configurations, and consider fixed information sprinkler nodes and evaluate their effect on system performance. We also consider the effects of user behavior (selfish or generous) on the effectiveness of the dissemination scheme.

Our results clearly show that realistic user mobility traces lead to an effective data dissemination system, thus confirming the potential of opportunistic networks. In addition, we show how different kinds of networking technologies affect the system performance. Finally, our evaluation of user behavior shows that opportunistic networks are vulnerable to selfish behavior. Although some of the effects can be mitigated with appropriately chosen network technologies, we cannot isolate the system from user behavior. A deeper discussion and additional figures can be found in [4, chap. 7]

The structure of the paper is as follows. We describe our opportunistic system in Section II. The simulation model and its various settings are presented in Section III. After that, we present and discuss the simulation results in Section IV. Section V gives a short overview on related research. We conclude this work in Section VI by summarizing our findings and sketch our next research steps.

II. System Model

We model the opportunistic network and its users as follows. Each user carries a mobile device with some wireless networking capabilities. This allows the devices to communicate in a spontaneous manner without user interaction. For data dissemination, each device holds a user profile. This profile consists of two parts, which we have called *iHave* and *iWish* lists in our previous work [2], [3]. The iHave list holds information the device is willing to share and the iWish lists holds filter expressions that match information the device wants to collect. When two devices come into communication range, they exchange their profiles to see if there is a matching between one user's iHave list and the other user's iWish list.

Our communication scheme is based on a *one-hop* communication paradigm. Only directly connected devices exchange information, under the assumption the profiles match. Since users carry their devices, and thus the information stored on the devices with them, the information is physically transported from one place to another. Regardless of time and place, information can be passed on to other interested parties. Together, the physical movement and one-hop information passing yield what is effectively a *multi-hop* data dissemination.

978-0-7695-3096-3/08 $25.00 © 2008 IEEE
DOI 10.1109/WAINA.2008.91

In addition, we introduce fixed nodes, called *information sprinklers*. These nodes, mounted at dedicated places, may extend the usefulness of an application by acting as intermediates or by offering proximity based services. In addition to simply spreading out pre-programmed information, information sprinklers can also collect information from users they see and pass that on to other users. Consider a coffee shop with an information sprinkler. User Alice usually drops by that coffee shop in the morning, and user Bob typically visits the shop in the evening. Normally, Alice and Bob would never meet, thus they would not have any chance to pass information to each other. In this situation, an information sprinkler can help. The sprinkler is set up in the shop and collects all information of users visiting the bar. This allows Alice to leave her information at the sprinkler in the morning and Bob to learn about this information from Alice in the evening.

As an extension, sprinklers in different locations can be connected by a backbone network. Thus, information passed from a user to an information sprinkler is almost immediately available at all other connected sprinklers.

III. 2-Step Simulation Model

We evaluated the effectiveness of our one-hop based data dissemination scheme through simulations. In particular, we were interested in the following questions:

1) How good is the information dissemination coverage, i.e., how many individual information wishes could be fulfilled within a given time?
2) What is the dissemination benefit obtained by i) deploying information sprinklers in the network and ii) connecting them?
3) How does a user's sharing behavior, i.e., a selfish vs. generous attitude, affect the system effectiveness?

We developed our own 2-step simulator for the evaluation. The simulator combined realistic user mobility data with synthetic mobility models. In the *first step*, we use real-world data from the Reality Mining Project [1], to get a realistic *macro-level* mobility of users. In the *second step*, we use synthetic mobility models to model *micro-level* user movement. Our two-step approach attempts to remedy the short-comings of purely synthetic mobility models (as argued in [5], [6]). The combination of coarse grained user traces with synthetic mobility models on the micro-level allows to explore various communication ranges and user movements, while simulations based on pure user traces are bound to the used technology, for example, user traces collected from bluetooth devices, are not suitable to investigate WiFi like communication behavior.

In addition, our model does not make any assumptions about the underlying communication channel, i.e. bandwidth, link loss, etc.. This is not applicable, since opportunistic communication is a best effort communication approach by definition and does not assume any reliable communication guarantees.

A. Overall Simulator Operations

Our simulator operates as follows. The simulator core extracts a list of user whereabouts from the Reality Mining dataset. Each entry consists of a unique location ID and a list of unique user IDs, which determine which users were at that location at that time. This tells us which users were somewhere in the general vicinity of each other. This information determines the user's *macro mobility*.

When two or more users are co-located according to the macro-level model from step one, we create an environment for simulating the *micro-level* mobility. We create a square with certain dimensions, place the users in the square randomly, and let them move according to a micro-level mobility model for the time that the users were in that cell. We use BonnMotion [7] for the micro-level mobility generation.

If two users come within a pre-defined range of each other, this gets signaled to the simulator core so that a match and possible information pass can be performed. We use different ranges to simulate different kinds of wireless networking technologies When two users are within communication range, then the simulator core informs the profile matching component. We perform the matching of the iHave and iWish lists and pass information items as appropriate. Our simulation does not take any memory limitations on the mobile device into account, since state-of-the-art devices are easily shipped with several gigabyte of memory and the information items we consider are of much smaller size.

The number of users and locations are determined by the user traces. We get both values from the Reality Mining dataset, which we will discuss below.

B. Reality Mining Dataset Usage

The Reality Mining experiment conducted at MIT Media Lab captured communication, proximity, location, and activity information from 100 users at MIT between January 1st 2004 and May 5th 2005. Each user was given a mobile phone that runs a special software. This software logged communication behavior of the user as well as location information it learned from its surroundings. We use the cellular network tower ID to which a mobile phone was connected to at a certain time as location of the user.

This location information is used for the *micro* mobility step in the simulation process as follows. If two users are connected to the same tower ID, we put them in random locations in a virtual area of 1000×1000 meters and let them move around for 5 minutes according to the configured micro-level mobility model. If the users come close enough to each other to be in communication range, e.g., 10 meters for Bluetooth and 100 meters for WiFi, the simulator executes the profile matching algorithm between them. On a match, information between users is exchanged.

We used several calendar weeks as data. This paper shows results based on calendar week 45 in 2004, with the highest amount of active users (78 users). The results based on other weeks are similar (see [4, chap. 7]).

C. User Behavior

We distinguish two types of users. One user-type, *free-rider*, acts purely selfishly and does not share any information he collects. The other user-type, *generous*, shares all information he collects with others. Because our data dissemination scheme relies on the users passing the information along, it is important to know how much user participation is needed for efficient dissemination. As our results show, information sprinklers can relieve the effects of free-riders, but they cannot eliminate them completely.

D. Micro-Level Mobility Models

We use three different mobility models for the *micro* mobility: Random Waypoint [8], Manhattan Grid [9], and Gauss-Markov model [8]. The three models were used to compare the data dissemination results, since the micro-level mobility model affects the chance of users coming into communication range and the chance of information being exchanged.

Random Waypoint is a simple and widely used mobility model. A mobile node chooses a random destination and speed and moves towards that destination. Upon arrival, the node pauses a random time and proceeds by choosing a new destination.

In the Manhattan Grid model, mobile nodes move only on predefined horizontally and vertically arranged paths. This model mimics a typical street network in an urban area.

The Gauss-Markov model eliminates sudden stops and sharp turns encountered in the Random Waypoint mobility model by allowing past velocities and directions to influence future velocities and directions.

IV. Simulation Results

We now present the simulation results for calender week 45 using 10 meters ("Bluetooth") and 100 meters ("WiFi") as a user device's communication range. For each week, the simulation was run 100 times (the figures display the averaged values) and with three different setups.

In setup 1, there are no information sprinklers; in setup 2 there are information sprinklers in every cell; and finally in setup 3, there is an information sprinkler in every cell and they are all connected by a backbone network. In this last case, as soon as an information sprinkler receives an information item, the item is available at all other information sprinklers. When an information sprinkler is present, it is in the middle of the 1000×1000 square and has the same communication range as users. Thus, even if the sprinklers are on, the user has to be near the sprinkler to pick up or leave items.

In all runs, one user, chosen randomly, owns an information item at start and all other user are interested in it. This is a simplified assumption due to the fact, that uninterested users neither contribute nor harm the information spreading. In real world, there will several different information items and only a subset of the overall population will show interest in one particular item.

In addition, we assumed that all users are generous. In Section IV-B we consider the effects of free-riders.

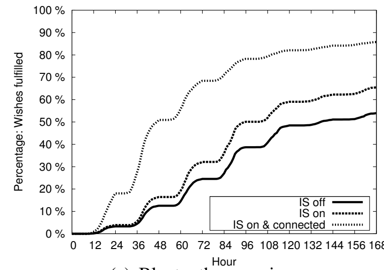

(a) Bluetooth scenario

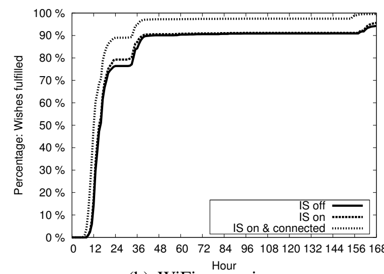

(b) WiFi scenario

Fig. 1. CW 45 - Hour/fulfilled wishes ratio

The metric we used was the percentage of users who have received the information item as a function of time. We first present the results showing the overall effectiveness of our data dissemination system. Then we vary the number of generous and free-riding users, and finally, we consider the effects of different micro-level mobility models.

A. Dissemination Effectiveness

Figure 1(a) shows the results for the Bluetooth scenario. It shows which percentage of users received the information item up to the time on the y-axis. The x-axis in all the plots starts at midnight between Sunday and Monday. In the figure, we show three lines, corresponding to the cases of no information sprinklers (IS off), information sprinklers turned on (IS on), and information sprinklers turned on and connected via a backbone (IS on & connected). The micro-level mobility model was Random Waypoint.

All the curves look alike, namely they rise during day times and remain constant from midnight to 6am. The leveling off during night was artificially introduced in the simulation, by turning off the micro-level mobility. The participants in the Reality Mining dataset did *not* turn off their devices during night, thus running a normal simulation would have meant that two users who sleep near the same celltower might have "sleepwalked" and exchanged information, even though they were sleeping too far from each other.

Likewise, the overall behavior of the three follows roughly the same pattern. In the beginning, only a few people possess the item, thus only a few exchanges are possible. In the middle of the week, more people possess the item, more exchanges are possible, and the item spreads rapidly. Towards the end of the week, the curves flatten out. The main reason for the flattening out is the relatively short communication range of 10 meters in this scenario.

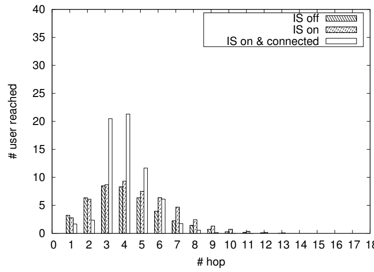

Fig. 2. Hop counts - Bluetooth scenario

In summary, without information sprinklers, we are able to reach on average 54% of the users. With information sprinklers turned on at every cell, we reach on average 65% of the users. Finally, connecting the information sprinklers via a backbone gives a reach of 86%.

If we extend the communication range to 100 meters (WiFi scenario), the results look as in Figure 1(b). Even without any information sprinklers, we are able to reach 94% of the users on average. Deploying information sprinklers brings only marginal gains with 96% of users reached. Finally, connecting the information sprinklers gives an average reach of 99.6%.

Figure 2 shows how many hops the information item took to reach a user in the Bluetooth scenario (results for the WiFi scenario look similar, see [4, chap. 7]). The x-axis is the hop count and the y-axis is the percentage of users reached with that many hops.

Most of the users receive the item with relatively small number of hops (about 3–5), although there are users who receive the item after over 10 hops. The peak is around 3 to 4 hops. For the case of information sprinklers turned on and connected, the peak is very pronounced. This is because of the way we calculate the hops. One hop is the item passing from a user to a user (where either user could also be an information sprinkler). In the "turned on and connected" scenario, one hop is the original user passing the item onto a sprinkler, one hop is the sprinkler passing it to all other sprinklers, and one hop is the item reaching another user. Four hops in this scenario means an additional pass from a real user to another real user either before or after the item got passed through the sprinklers.

Comparing the "IS off" and "IS on" cases, we can clearly see the effect of the additional hop caused by the information sprinklers.

B. Different User Behavior

We now evaluate effects of user behavior. We use the same simulation setup as in Section IV-A, with the difference that we considered only the case of information sprinklers being off in order to make the system rely purely on user behavior.

As above, the information starts from one user and all users are interested in this item. We randomly selected a fraction of the users who would act as free-riders. The free-riders would collect the information item from others, but would not pass it on to other users. All the other users were generous, i.e., they pass the item on to any other user they encounter.

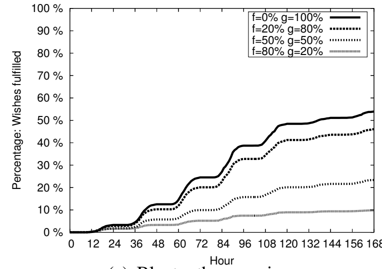

(a) Bluetooth scenario

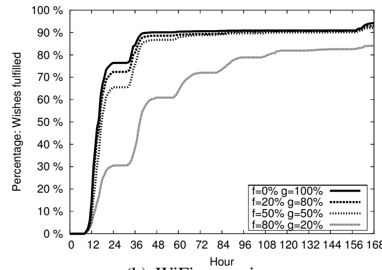

(b) WiFi scenario

Fig. 3. Comparing various free-rider/generous ratios (CW 45, no IS deployed)

Figure 3(a) shows Bluetooth results and Figure 3(b) WiFi results. We simulated three different free-rider/generous ratios: (20% free-riders / 80% generous), (50% free-riders / 50% generous), and (80% free-riders / 20% generous). The figures also show the case with all users acting generously (Figure 1).

In the Bluetooth case, increasing the number of free-riders significantly decreases the percentage of users reached. For the 20% free-riders case, the reach drops from 54% to 46%; in the 50% free-rider case it drops down to 23%, and in the 80% free-riders case down to 10%. In the WiFi case, the results are very different. With 20% or 50% or free-riders, the results do not significantly deviate from the case of no free-riders. If the number of free-riders increases to 80%, then we observe a drop of reachability to 84%.

These two cases show that if the communication range is short, then it is important for users to be willing to cooperate. This is because a short communication range implies that a user might not meet many other users, and thus the willingness to share becomes important. With a longer communication range, users do not have to come so close to each other and are not so reliant on the cooperation from others.

C. Different Mobility Models

The Reality Mining dataset tells us when two users are near each other, but the actual meeting of the users happens through the micro-level mobility model. Only users that come into communication range are able to exchange information. Until now, all results used the Random Waypoint mobility model. The chosen mobility model affects the likelihood of users coming into communication range and thus, affects the data dissemination process. We evaluated two other mobility models, namely Gauss-Markov and Manhattan Grid, to determine the effect of the micro-level mobility model on the data dissemination process.

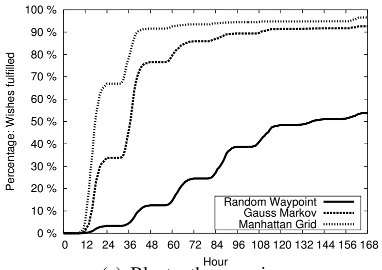

(a) Bluetooth scenario

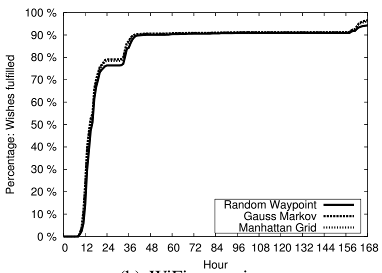

(b) WiFi scenario

Fig. 4. Comparing mobility models (CW 45, no IS deployed)

Figure 4(a) shows that Gauss-Markov and Manhattan Grid models reach a higher number of users than the Random Waypoint model. With Gauss-Markov, we reach 93% of the users, whereas with Manhattan Grid we reach 97% of the users. In contrast, Random Waypoint reaches only 54% of the users. In the WiFi scenario in Figure 4(b), there is basically no difference between the different mobility models. For reference, Gauss-Markov reaches 97%, Manhattan Grid reaches 96%, and Random Waypoint reaches 94% of the users.

As with the user behavior results above, we can conclude that for short range communication, user movement patterns are significant, since they determine when information can pass from one user to another. In the WiFi case, the communication range is relatively large (100 m) compared to the simulated cell area (1000 $\times$ 1000 m) and the details of the movement pattern become less important.

D. Discussion

We now discuss the implications of our results and what lessons can be learned from them. In particular, we will answer the three questions from the beginning of Section III in light of the results we obtained.

The Reality Mining dataset has only a small number of users and all users have some affiliation with MIT and there might be relationships between users which bias them to meeting more often than usual. In spite of these possible limitations, we believe the Reality Mining dataset provides a good trace of real-world user mobility and to be a valuable source for researchers. Other traces are available [10], and we plan to validate our results with those traces later.

For the micro-level mobility, we used three different synthetic mobility models. As our evaluation in Section IV-C shows, the effects of the mobility model depend greatly on the assumed communication range at which information exchanges can take place. For short range communications, it is vital that users come close to each other for the exchanges to happen. On the other hand, with the longer communication range, the mobility becomes less important. This is an important result in general, because it shows that the selection of a mobility model can *significantly* affect the results, and that the effects can depend on choices of other parameters. When choosing a mobility model, it is therefore vital to explore different parts of the parameter space, in order to evaluate the sensitivity of the results to the choice of the mobility model.

We now answer the questions from Section III.

Question 1: How good is the dissemination coverage?

This is mainly determined by the communication range, which determines the probability that two nearby users actually can exchange items. In the Bluetooth case, the reachability ranges from 27% to 91%, whereas in the WiFi case, we reach on average 94% to 99% of the users. In the Bluetooth case, the reachability is also greatly affected by how the information sprinklers are configured.

The simulation time was 1 week. We also ran longer simulations and observed, that eventually the information item typically reaches all users (except isolated cases caused by "unlucky" micro-level mobility). Hence, we can conclude that realistic human movement patterns lead to a system where it is reasonable to assume that information distributed in an opportunistic manner reach (almost) all of the interested users.

Question 2: What is the benefit obtained by i) deploying information sprinklers in the network and ii) connecting them?

In the Bluetooth scenario, turning the information sprinklers on typically increased the percentage of reached users by 10%. Connecting the sprinklers via a backbone network would typically boost the percentage of reached users to over 90%. In the WiFi scenario, information sprinklers do help, but their effect is minimal.

As expected, the information sprinklers help the most in the cases where the communication ranges are short. In effect, they remove the strict requirement of two users being in communication range for an exchange to take place, and allow for time-shifted information passes. Our results from the WiFi case show that with longer communication ranges, information sprinklers are not essential and that it is reasonable to build an opportunistic information dissemination system with little or no infrastructure support.

Question 3: How does a user's sharing behavior, i.e., a selfish vs. generous attitude, affect the system effectiveness?

In the Bluetooth scenario, 80% free-riding users drops the reachability of the system down to 10%, whereas in the WiFi scenario, a similar amount of free-riders only has a small effect on the reachability of the information item. Currently there are no estimates on the number of free-riders in opportunistic networks, but 80% free-riders have been observed in Internet-based peer-to-peer networks [11], [12].

As with the information sprinklers, a longer communication range helps alleviate the dependency of the system on individual user behavior. However, the system remains vulnerable to free-riding, which implies that some kinds of incentive schemes might be necessary in opportunistic networks to encourage user participation.

Lessons Learned: To sum up the main points from our results, we conclude the following.

- Realistic user movement traces exhibit user behavior which enables building of opportunistic data dissemination systems. This is the main result of this paper and is clearly demonstrated by all our results.
- Sufficiently long communication range is enough to make the system relatively immune to free-riding and it also eliminates the need for any kind of fixed infrastructure support (i.e., information sprinklers).
- For shorter range communications, information sprinklers are efficient at helping with the data dissemination. It is also highly beneficial to connect all information sprinklers via a backbone network.
- Although the effects of free-riding can be alleviated, they cannot be removed. This implies a need for more research in incentives for opportunistic networks.

V. RELATED WORK

User traces have in general been used for two, slightly different kinds of opportunistic network applications: Opportunistic message routing and forwarding, and opportunistic data dissemination. The former assumes an end-to-end communication need between two or more communication partners but without a direct path between them. Communicating wirelessly and exploiting node mobility (and sometimes some intermediate infrastructure, e.g., an info station) eventually a delay tolerant duplex link is established [13].

The latter—opportunistic data dissemination—is similar to word-of-mouth communication among humans. This is also our focus of our work. Data is spread between nodes (devices/users) that expose an interest in a certain type of data.

Becker et al. [14] simulated the performance of epidemic-like diffusion algorithms. Although their results on the number of users reached is similar to ours, they only consider a very limited set of parameters, namely only one type of communication, no information sprinklers, and purely synthetic mobility models. This is also true for the work for Khelil et al. [15], who use random waypoint in their simulation.

Similar to our current work, Leguay et al. [16] evaluate user traces gathered by an experiment carried out in the city of Cambridge, UK. They placed 20 fixed Bluetooth devices at popular places and issed 40 Bluetooth devices among students. They consider the dissemination of a digital newspaper issued by the 20 fixed devices every morning. Similar to our Bluetooth results (Figure 3(a)), a selfish behavior of nodes results in a poor dissemination ratio and node collaboration increases the dissemination ratio significantly. Our work extends this by considering different types of communication technologies, and we evaluate the benefits gained from installing information sprinklers in different configurations.

VI. CONCLUSION AND OUTLOOK

In this paper, we have evaluated the effectiveness of opportunistic data dissemination using real-world user mobility traces. We have considered many different system configurations, including varying communication ranges, information sprinklers, and user behavior.

Our results clearly demonstrate that realistic user mobility leads to users encountering in a manner which makes an opportunistic data dissemination system feasible. Information sprinklers are useful at helping with the dissemination in cases where the communication ranges are short (e.g., Bluetooth-based systems). Furthermore, connecting the information sprinklers over a backbone yields significant gains.

User behavior has a large effect on the dissemination performance, although long enough communication range alleviates the problems. Regardless, our results show a need for incentives to encourage user participation in opportunistic data dissemination systems.

As a next step, we plan to verify our simulation results using other available user traces from the CRAWDAD archive [10] and take more communication ranges into account.

REFERENCES

[1] Massachusetts Institute of Technology, “Reality Mining,” http://reality.media.mit.edu/ (seen 10/2007), 2005.

[2] A. Heinemann, J. Kangasharju, F. Lyardet, and M. Mühlhäuser, “Ad Hoc Collaboration and Information Services Using Information Clouds,” in *Workshop on Applications and Services in Wireless Networks*. 2003.

[3] ——, “iClouds – Peer-to-Peer Information Sharing in Mobile Environments,” in *Euro-Par Conference*, 2003.

[4] ——, “Collaboration in Opportunistic Networks,” Ph.D. dissertation, Technische Universität Darmstadt, Germany, 2007.

[5] A. Jardosh et al., “Towards Realistic Mobility Models For Mobile Ad Hoc Networks,” in *MobiCom* 2003.

[6] F. Bai and A. Helmy, *Wireless Ad-Hoc and Sensor Networks*. Springer, 2006, ch. A SURVEY OF MOBILITY MODELS.

[7] University of Bonn, “BonnMotion - A Mobility Scenario Generation and Analysis Tool,” http://web.informatik.uni-bonn.de/IV/Mitarbeiter/dewaal/BonnMotion/index.html (seen 10/2007), 2002-2005.

[8] T. Camp, J. Boleng, and V. Davies, “A Survey of Mobility Models for Ad Hoc Network Research,” *Wireless Comm. and Mobile Computing*, vol. 2, no. 5, pp. 483–502, 2002.

[9] European Telecommunications Standards Institute (ETSI), “UMTS: Selection Procedures for the Choice of Radio Transmission Technologies of the UMTS,” ETSI, Tech. Rep. TR 101 112 V3.2.0, 1998.

[10] Dartmouth College, “CRAWDAD Project: A Community Resource for Archiving Wireless Data At Dartmouth,” http://crawdad.cs.dartmouth.edu/ (seen 10/2007), 2005.

[11] E. Adar and B. A. Huberman, “Free Riding on Gnutella,” *First Monday*, vol. 5, no. 10, 2000. http://firstmonday.org/issues/issue5_10/adar/index.html (seen 10/2007)

[12] S. Saroiu, P. K. Gummadi, and S. D. Gribble, “Measuring and Analyzing the Characteristics of Napster and Gnutella Hosts,” *Multimedia Systems Journal*, vol. 9, no. 2, pp. 170–184, 2003.

[13] L. Pelusi, A. Passarella, and M. Conti, “Opportunistic Networking: Data Forwarding in Disconnected Mobile Ad Hoc Networks,” *IEEE Comm.* , vol. 44, no. 11, pp. 134–141, 2006.

[14] C. Becker, M. Bauer, and J. Hähner, “Usenet-on-the-fly – Supporting Locality of Information in Spontaneous Networking Environments,” in *Ad Hoc Comm. and Collaboration in Ubiq. Comp. Environments*, 2002.

[15] A. Khelil, C. Becker, J. Tian, and K. Rothermel, “An Epidemic Model for Information Diffusion in MANETs,” in *WS on Modeling, Analysis, and Simulation of Wireless and Mobile Systems*, 2002, pp. 54–60.

[16] J. Leguay, A. Lindgren, J. Scott, T. Friedman, and J. Crowcroft, “Opportunistic Content Distribution in an Urban Setting,” in *CHANTS*, 2006.