

Opportunistic Forwarding in Workplaces

Sabrina Gaito
Computer Science Dept.
Univ. degli Studi di Milano
v.Comelico 39, Milano, Italy
gaito@dsi.unimi.it

Elena Pagani
Computer Science Dept.
Univ. degli Studi di Milano
v.Comelico 39, Milano, Italy
pagani@dico.unimi.it

Gian Paolo Rossi
Computer Science Dept.
Univ. degli Studi di Milano
v.Comelico 39, Milano, Italy
rossi@dico.unimi.it

ABSTRACT

So far, the search for Opportunistic Network (ON) applications has focused on urban/rural scenarios where the combined use of mobility and the *store-carry-and-forward* paradigm helpfully recovers from network partitions and copes with node sparsity. This paper explores the chance of using ONs in workplaces, where the node distribution is denser, thus contributing to reduce the message delivery latency, and where we still find similar needs for informal and unplanned network platforms to support human social relationships and interactions. Both a survey and trace recording experiments have been used to support the analysis of this mobility setting. The ability of recording very short contact times (i.e. lasting few seconds) allowed to interestingly show the slightly different role the social relationships play in dense scenarios and how the large amount of contacts (both short and long), occurring in densely populated spaces, actually contribute to reduce the message-delivery latency and to increase the delivery probability.

Categories and Subject Descriptors

C.2.1 [Computer-Communication Networks]: Network Architecture and Design—*Wireless communication*

General Terms

Measurement, Experimentation

Keywords

Human mobility; message forwarding; opportunistic networks; mobility analysis.

1. INTRODUCTION

Opportunistic Networks (ONs) have recently received growing attention because of their ability to create unplanned and improvised urban/rural wireless connectivity among mobile nodes. Researchers in this area envision an urban sce-

nario where people carry radio devices that can be dynamically networked by exploiting human contact opportunities and, as a result, the term of pocket switched networks (PSN) has been used [8]. The growing interest in ONs is motivated by three main factors: 1. the pervasive multitude of portable devices has a huge amount of unused system and networking resources that may be profitably exploited to support a wide range of human interactions; 2. mobility can increase the capacity of wireless networks, as has been proved in [6]; 3. human social relationships provide quite a stable network of contact opportunities that can be profitably utilized to design forwarding algorithms [9].

So far, the search for ON applications envisions scenarios where the combined use of mobility and the *store-carry-and-forward* paradigm helpfully recovers from network partitions and copes with node sparsity. This paper focuses on a changed and enlarged application perspective. In fact, when we scale down from the metropolitan area to workplaces, buildings and small campuses, we obtain denser node distribution and observe similar needs for maintaining human social (or working) relationships that are likely to be profitably supported by an ON infrastructure, outside the institutional IT platform. The growing success of applications such as Twitter [13] – which keeps friends and co-workers frequently connected – and the constant growth of demand for mobile messaging (MM) applications¹ are showing this trend in human interaction and mandates verifying if ONs are suitable to support them. Apparently, this challenging idea is in contrast with the delay tolerant nature of ONs. There is the expectation, however, that the larger amount of contacts in a densely populated space can significantly contribute to reducing message-delivery latency. This motivates a research effort to understand the role of contacts and whether or not human mobility in workplaces shows behavior similar to mobility in more sparsely distributed settings.

When considering the application of PSNs in dense urban scenarios, people (and referees, as well) tend to ask why we should introduce a novel, delay-prone network when we have plenty of delay-free infrastructures. We too have wondered, without obtaining a definitive answer. The willing reader can find a more detailed analysis of this issue in [4], although we believe that we need not always have the applications before the technology. By contrast, we expect that, once deployed, the informal, easy, spontaneous access to this

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

WOSN'09, August 17, 2009, Barcelona, Spain.

Copyright 2009 ACM 978-1-60558-445-4/09/08 ...\$10.00.

¹A recent Report of Forrester Research claims that the use of MM has grown of the 9% in 2008 and that its demand is expected to grow up to 80 million users in 2013, when it is supposed to replace 13% of SMS traffic.

further form of connectivity will encourage the growth of a new family of applications centred on the innovative notion of ‘on-line social network’.

To improve our understanding of human social attitudes in a workplace scenario, we performed a survey involving nearly 300 computer science faculty members and students. The primary observations obtained are the following:

- (i) among faculty members the need for ubiquitous reception of notices about upcoming institutional meetings or events and for extemporaneous contacts with co-workers emerges. Students also need notifying and would like to be able to exchange messages with friends. Less than 30% of these communications are likely to have some attached file, so most of them could be fruitfully forwarded during brief contacts. In fact, the people surveyed seem to prefer downloading files over regular Internet when required (i.e. they want to be notified via thin device of attachments to download later);
- (ii) 80% of people surveyed usually see each other at least once a day;
- (iii) how much delivery latency people can tolerate varies according to service required, ranging from less than one to a few hours.

In the scenario outlined, this paper is a first attempt to answer questions such as the following: 1. Do social relationships in workplaces lead to the same structure of the contact topology as in the scenarios that have been considered so far? 2. has this new setting some different impact on the forwarding algorithms in ONs? 3. can the application needs, as they emerged from the survey, be satisfied by the changed mobility setting and social attitudes?

To answer these questions, we developed a test bed with on purpose designed devices, named Pocket Mobility Trace Recorders, or PMTRs. The main and original behaviour of PMTRs is the ability to observe fine-grained contacts, so that even contacts of few seconds can be recorded. Moreover, unlike other experiences reported in the literature, the trial involved people profiled in order to enable the verification of the results emerging from trace analysis. Both the above aspects distinguish our experiments from other experiments conducted in restricted areas, e.g. [2].

The trace analysis enables to characterize the considered setting in terms of contact times distribution, latency times, social attitudes and their impact on the forwarding algorithms. The traces we obtained in a dense setting are compared with the traces obtained in similar experiments, where, however, it was not possible to observe short contact times [11]. We can identify three main contributions of this paper: firstly, we show that the node density actually helps to reduce the delivery latency and, secondly, that short contact times (few seconds), produced by non-intentional mobility, give a great contribution to reduce the delivery latency and to increase the delivery probability. Finally, we show how the human social relationships can be profitably exploited for message forwarding in dense spaces.

2. POCKET TRACE RECORDER

The design of a specific device for trace recording is mainly motivated by the need of observing and recording very short contact periods, few seconds, that arise from random mobility in a dense area. As a consequence, PMTRs [5] have been designed to operate with beaconing times ranging from 1 sec. to some configurable value which depends on the mobility environment we wish to observe. Secondly, the devices have

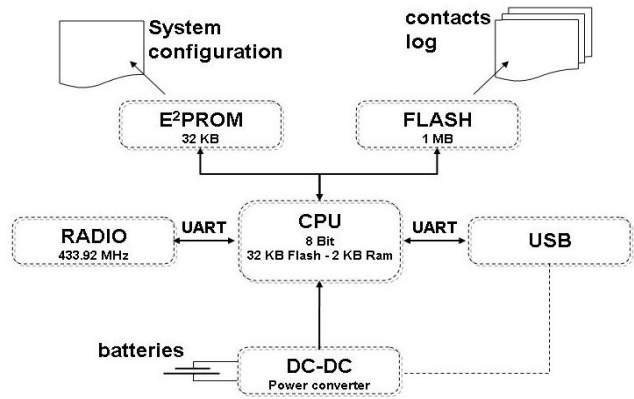


Figure 1: PMTR architecture.

to enable unmanned experiments lasting 3-4 weeks without batteries substitution. Layer 2 beacons are the unique frames a PMTR broadcasts to its neighbors. The first time a beacon is received from a given encounter, the current time is recorded for the contact start, together with the encounter ID; a contact ends when beacons from a certain encounter have been missing for more than t seconds, with $t = 60$ in our experiments. The local memory size should be dimensioned to store the contacts of the experiments. Our test beds have generated on average 2000 contacts per device with beaconing time set to 1 sec. No specific bandwidth and processing requirements have been envisaged for PMTRs.

The PMTR architecture is described in Fig.1. It uses the Cypress CY8C29566 micro-controller and the radio module AUREL, model RTX-RTLP. The radio range has been limited to 10 meters in order to reduce the power consumption and to maintain multi-hop paths between end-systems. This combination allows a very low power consumption that let the experiments last for the required time with common batteries NiMh, AA 1.2 V. Each PMTR has a 1 MB flash memory where more than 50K contacts can be stored. The PMTR implements a CSMA non-persistent MAC protocol. The local clock value is set at the configuration time. Each PMTR uses a USB interface to communicate with the Pocket Viewer PC, running the Desktop application software, which has been used to configure the devices, collect the recorded data at the end of the experiment and support data analysis and device monitoring.

3. EXPERIMENTAL RESULTS

The experiment has been run for 19 days in November 2008; 49 PMTRs were involved, distributed to faculty members, PhD students, and technical staff. People work in offices and laboratories located in a three-floors building, large roughly 200×100 m., and they take lunches or coffee breaks in a nearby cafeteria. Some of the lessons take place in a different building 3.5 Km far, where faculty members and students may temporarily displace. These locations are equipped with fixed PMTRs. At the end of the experiment, 5 PMTRs (with IDs 5, 8, 15, 30 and 35) showed misbehaviors in data recording. Excluding weekends, we were able to trace contacts among 44 PMTRs for 15 working days with no gaps in the traces. A collection of 11895 contacts remained for analysis.




Figure 2: Inter-contact time ccdf along the whole experiment.

Table 1: Comparison between PMTR and MIT traces.

	PMTR	MIT
power law exponent	0.21	0.15
% contacts within a day	80%	47%
mean inter-contact time	11.81 hours	4.12 days
median inter-contact time	41.49 minutes	16 hours
mean intra-contact time	8.69 min.	57 min
median intra-contact time	0.8 min.	32 min.
mean # neighbors	2.18	2.79
median # neighbors	2	2

Characterization of the Environment.

We consider the following statistics to usefully characterize the environment we consider: the aggregated inter-contact times, defined as the time elapsed between the end of a contact and the beginning of the successive contact for every fixed pair of nodes, the duration of the contacts called intra-contact times, and the number of neighbors seen by every node at the beginning of a contact. A coarse analysis of the traces shows the well known power law distribution of inter-contact times [1, 8, 12, 2, 10] whose complementary cumulative distribution function (ccdf) is shown in fig.2.

In Table 1, we compare these statistics estimated on our traces with those yield by the MIT Reality Mining experiment [11]. The MIT traces are considered for comparison as they characterize a scenario different from that of our experiment. Differently from MIT, PMTRs are distributed over a reduced area. As a consequence, bearers are more likely to be in range. This is confirmed by the first six parameters. As far as applications are concerned, fine-grained beaconing allows to detect a high number of short contacts, as evidenced by the median intra-contact time lower than 1 min. These contacts are not present in other traces publicly available, as the usual beaconing time adopted by other experiments [3] is no less than 120 sec. However, short contacts can effectively contribute to message forwarding, according to the observation that the most part of sent data has a small size.

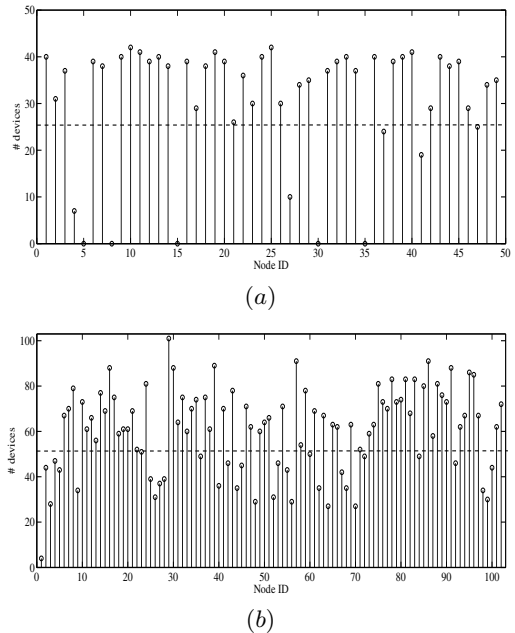


Figure 3: Number of persons encountered throughout the experiment in (a) PMTR and (b) MIT traces.

The difference between the two scenarios also reflects in a different exponent of the power law distribution of the inter-contact times. More in detail, 80% of inter-contact times in PMTRs are within a day since people, sharing the same workplace, repeatedly meet. But 50% of contacts lasts less than one minute because people meet by chance in corridor. In fig.3, the number of devices encountered by each device along the experiment is shown, for both MIT and PMTR. Only a 10% of PMTRs encounters less than 50% of the other devices, thus confirming a lower sparsity of the environment with respect to MIT, where the experiment lasted 9 months and 30% of the devices fall below the 50% line. This measure apparently contrasts with the mean number of neighbors when a contact occurs, reported in Table 1. In fact, the merge of the two results could indicate that the PMTR environment is less sparse: it is thus easier for a node to frequently detect other devices passing by, but these contacts have short duration showing a prevalence of opportunistic rather than planned encounters. Opportunistic encounters are likely to happen in presence of a lower number of neighbors with respect to planned meetings. On the other hand, the PMTR capability of detecting short encounters might more accurately capture the real connectivity graph.

Detailed Trace Analysis.

As a consequence of the latency requirements emerged from the survey, trace analysis focused on one-day behavior. In fig.4, the inter-contact and intra-contact time cdfs are plotted for each day of experimentation. All days show comparable behavior; in the following, the traces for a fixed day are analyzed. According to relatively short delays requested by people surveyed, we implemented a simple greedy algorithm that computes over the trace the *optimal in latency* diffusion tree for every source. The tree provides statis-

