Pervasive and Mobile Computing 9 (2013) 472–483

Contents lists available at SciVerse ScienceDirect

Pervasive and Mobile Computing

journal homepage: www.elsevier.com/locate/pmc

Human-mobility enabled wireless networks for emergency communications during special events

Angela Sara Cacciapuoti [a], Francesco Calabrese [b], Marcello Caleffi [a,*], Giusy Di Lorenzo [b], Luigi Paura [a,c]

[a] Department of Biomedical, Electronics and Telecommunications Engineering (DIBET), University of Naples Federico II, Naples, Italy

[b] IBM Research, Dublin, Ireland

[c] Laboratorio Nazionale di Comunicazioni Multimediali (CNIT), Naples, Italy

ARTICLE INFO

Article history:
Received 23 August 2011
Received in revised form 17 January 2012
Accepted 18 January 2012
Available online 28 January 2012

Keywords:
Measurement
Human mobility
Connectivity
Small world
Urban

ABSTRACT

During social gatherings or emergency situations, infrastructure-based communication networks have difficulty operating given either increased traffic demand or possible damage. Nevertheless, current communication networks still rely on centralized networking paradigms. The adoption of a peer-to-peer communication paradigm would be better adapted to these needs, especially if it relies on the mobile phones that people normally carry, since they are automatically distributed where the communication needs are. However a question arises: can the spatio-temporal distribution of mobile phones enable a partially-connected ad hoc network that allows emergency communications to happen with an acceptable delay? To try to answer this question, we defined a methodology composed of three steps. First, the positions of seven hundred humans, spread over a metropolitan area, have been anonymously traced during a special gathering event. Then, with a multi-disciplinary approach, we have inferred the contact events from the humans' traces. Finally, we have assessed the effectiveness of an ad hoc network established by the mobile phones to disseminate emergency information to the population in a timely fashion. The results reveal that the humans' mobility can effectively enable emergency communications among a significant subset of mobile phones, although the connectivity of the network strictly depends on the number of cooperating devices and on the maximum allowed delay.

1. Introduction

The need for disaster relief communications has been recognized by the United Nations Foundations as one of the great challenges of the early 21st century: "When disasters strike, people need food, shelter, blankets, and medicine but without an effective communications network, supplies are left undelivered, and relief workers are unable to do their jobs" [1].

During social gatherings or disaster events, traditional communication networks such as cellular ones operate with difficulty, due to increased congestion levels and possible damage. Indeed, during such events, tremendous stress is placed on networks due to the sudden rise in traffic demand for both usual and emergency calls. This effect is shown in Fig. 1, which reports the density of mobile phones engaged in calls in Rome both during an ordinary weekday and right after a metro

* Correspondence to: Department of Biomedical, Electronics and Telecommunications Engineering (DIBET), via Claudio 21, I-80125 Napoli, Italy.
E-mail addresses: angelasara.cacciapuoti@unina.it (A.S. Cacciapuoti), fcalabre@ie.ibm.com (F. Calabrese), marcello.caleffi@unina.it (M. Caleffi), giusydil@ie.ibm.com (G. Di Lorenzo), paura@unina.it (L. Paura).

1574-1192/$ – see front matter © 2012 Elsevier B.V. All rights reserved.
doi:10.1016/j.pmcj.2012.01.002

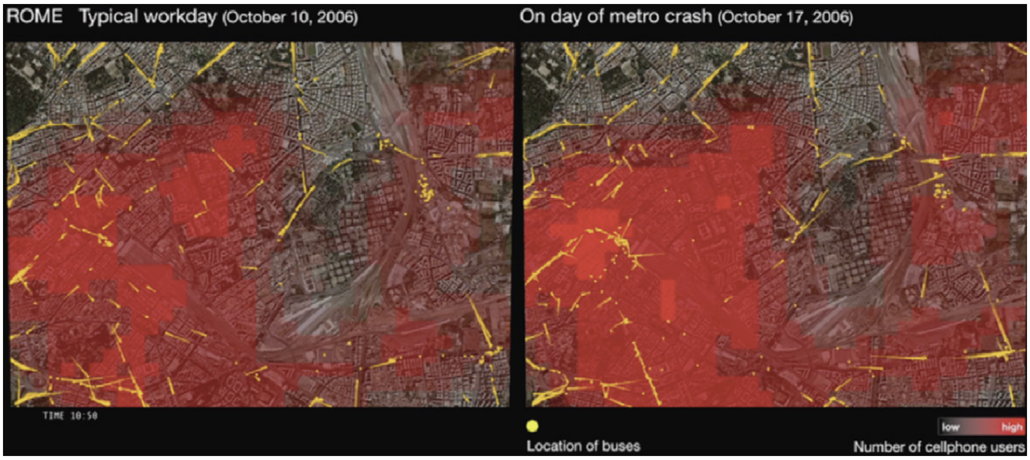

Fig. 1. Density of mobile phones engaged in calls (red) and location of buses (yellow) in Rome. The image on the left shows the cell-phone activity on a normal weekday at 10:50. The image on the right shows more activity happening around the train station at the time of a metro crash on October 17, 2006. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

crash.[1] Nevertheless, current communication networks still rely on centralized infrastructures and traditional networking paradigms.

In the last ten years, a new paradigm for wireless networks based on node mobility has gained attention by the research community. This paradigm, usually referred to as *Pocket Switched* [2], *Delay* or *Disruption Tolerant* [3,4], assumes that at each instant end-to-end paths may not exist in the network as a consequence of intermittent connectivity, low node density or the presence of isolated well-connected regions. In such a paradigm, the end-to-end connectivity is thus provided by node mobility and store-and-forward communications: intermediate nodes store the data waiting for transmission opportunities provided by mobility to deliver the data towards the final destination.

In this paper, we evaluate the feasibility of the delay tolerant paradigm to provide end-to-end connectivity in the considered scenario, namely to enable emergency communications during large scale events when traditional communication infrastructures are congested. Clearly, in such a scenario the communication devices can be the mobile phones that we normally carry, since they are automatically distributed where the communication needs are. However, the question arises: can the spatio-temporal distribution of mobile phones enable emergency communications with an acceptable delay?

To try to answer to this question, we carried out an experiment in which the positions of seven hundred mobile phone users have been anonymously traced during a special gathering event in the Boston metropolitan area: the Boston Independence Day Celebration on July 4th, 2009. Then, the mobility patterns have been estimated from the humans' traces. Finally, the effectiveness of such patterns to build a Wi–Fi delay tolerant network, able to disseminate emergency information over the population in the absence of network infrastructures and in a timely fashion, has been assessed. The Independence Day Celebration event has been chosen as concerts and fireworks were organized around the Charles river between the Harvard and Longfellow Bridges (see Fig. 2), and people congregated around the river from the early hours of the day for the celebrations starting at 8 pm. We particularly focused on the time frame 5–10 pm right before and during the concert.

The rest of this paper is organized as follows. Section 2 discusses related work, and Section 3 describes the experiment. In Section 4 we comment on the results of the experiment, while in Section 5 we discuss the limitations. Finally, conclusions are drawn in Section 6.

2. Related work

Several experiments have been conducted in the past ten years to collect human mobility data and to evaluate its impact on routing protocols.

A first small experiment was conducted in 2005 by collecting human mobility traces in a conference environment [2]. Following this, several studies have focused on comparing different datasets and studying the impact of the derived human mobility patterns on the design of routing protocols [5–7]. The study of these mobility-enabled networks, such as the one presented in [2], has also lead to the definition of new temporal distance metrics to quantify the speed of information diffusion processes, see for instance [8]. In [9], the authors studied an urban setting (Cambridge, UK) but relied on a rather small group of students.

[1] See more information at http://senseable.mit.edu/realtimerome.

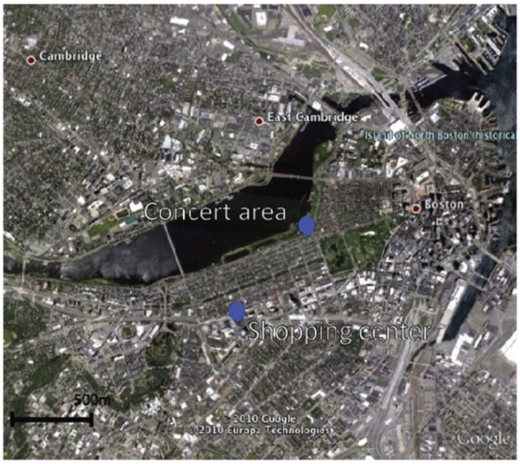

Fig. 2. Boston metropolitan area. Blue circles represent the starting locations of emergency alarms (with 100 m radius). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

The main difference between our work and the cited works lies in the traced participants selection. Our approach does not rely on volunteers who agreed to be tracked, who usually are students or conference participants who move in a small scale scenario. We instead use data that is generated by anyone using their mobile phones and connected to a particular telecom operator. This allows us to indirectly monitor the human mobility of a wider range of people and over larger temporal and spatial intervals.

Inferred traces have been considered in [10] using urban transport data, and routine behavior has been exploited for media sharing. Mobile phone data were considered in [11] to infer contact events between roughly two thousand people in the Boston metropolitan area. However, none of the studies have considered large gathering events. Moreover, none of the works presented above have studied the impact of human mobility patterns for enabling emergency communications when traditional communication infrastructures fail.

In this paper, we instead evaluate the feasibility of human mobility: (i) at large temporal and spatial scale; (ii) not relying on volunteers who agreed to be tracked; (iii) to disseminate emergency information in absence of network infrastructures and in a timely fashion. To the best of our knowledge, this is the first work in which a real-world wide-scale data set has been considered to enable emergency communications thorough a partially-connected ad hoc network.

3. Description of the experiment

To evaluate the feasibility of a partially connected ad hoc network established by mobile phones for disseminating emergency information, we have conducted a data collection experiment involving one million mobile phone users of a US telecom operator in the Boston metropolitan area, whose positions have been anonymously traced during the month of July 2009. The large user set allows us to limit biases in the obtained results, which are found in previous datasets where participation was restricted to students or conference attendees who volunteered to be tracked (see Section 2).

In this paper, we select as an event representative of a social gathering event the Boston Independence Day Celebration on July 4th, 2009,[2] when people usually congregate around the Charles river (the area depicted in Fig. 2) to attend the concert and watch fireworks organized by the city administration. The choice of Independence Day is not arbitrary, since: (i) it is a crowd gathering event, and thus it allows us to trace human mobility during a social event; (ii) although it is not an emergency event, it is a major US holiday, and so the human mobility traces are likely not to be affected by routine behaviors.

In the following, first we present the procedure adopted for measuring real world human positions. Then, we describe how the contact events have been inferred from the mobility patterns. Finally, we illustrate how the emergency alarm has been spread among the people.

3.1. Tracing human mobility during a social gathering

To trace the human positions, a set of 200 million anonymous location measurements collected by AirSage[3] has been considered. This dataset covers a region spread over 8 counties in east Massachusetts (Middlesex, Suffolk, Essex, Worcester, Norfolk, Bristol, Plymouth, Barnstable) with an approximate population of 5.5 million people, and it involves one million mobile phones (corresponding to a share of 20% of the population approximately), traced for roughly a month.

2 http://www.july4th.org.
3 http://www.airsage.com.

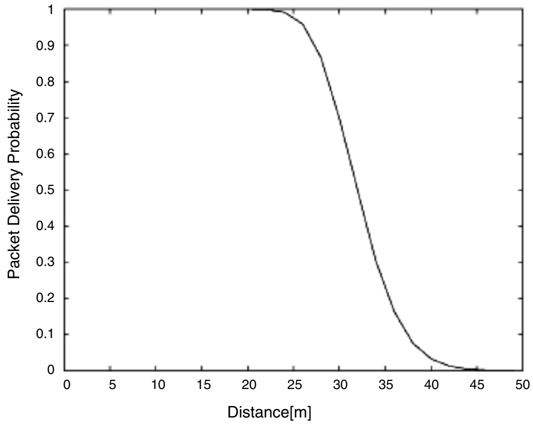

Fig. 3. Packet delivery probability as a function of the distance between two mobile phones.

The location measurements have been recorded each time a device connected to the cellular network, i.e., (i) when a call was placed or received, both at the beginning and at the end of the call; (ii) when a short message was sent or received; (iii) when the device connected to the Internet. Each location measurement represents the position, i.e., the latitude and longitude, of a certain device estimated through triangulation.

To infer human mobility patterns related to the considered event, we pre-process the dataset of location measurements by selecting the users who, during July 4th, 2009 between 01:00 pm and 23:59 pm, (i) made at least one connection to the cellular network every hour; (ii) were located in that area under analysis, namely the Boston City area, during at least 80% of the calls. The number of selected users is 703.

The selection of the users has a twofold purpose. It limits the computational complexity of the data analysis, which becomes unfeasible for greater numbers of selected users. Moreover, it improves the statistical meaning of the analysis results by selecting people for whom a sufficient number of location measurements is available. We note that the user selection, based on the number of available location measurements, does not introduce bias in our results, since it does not account for the connectivity properties of the selected users.

The raw location measurements of the selected users, which consist of a sequence of locations with arbitrary sampling rate, are then processed to remove fictitious trips due to localization errors and to obtain location traces with a sampling rate equal to 10 min. Finally, the set of processed locations is used as waypoints to infer the mobility patterns by assuming uniform mobility between two consecutive waypoints. The interested readers are referred to [11,12] for further information.

3.2. Inferring contact events from human mobility

To evaluate the feasibility of a partially connected ad hoc network established by the mobile phones for disseminating emergency information, we have supposed that no cellular infrastructure is working. As a consequence, the devices can communicate only via Wi–Fi ad hoc communications. The assumption of Wi–Fi enabled devices is justified by the latest market research analyses, which estimate that 144 million Wi–Fi enabled mobile phones were shipped worldwide in 2009 [13] and predict that the Wi–Fi enabled phone penetration will quadruple by 2015, reaching 66% of all mobile phone shipments [14].

To infer the contact events from the human mobility patterns derived in Section 3.1, we import the user traces into a network simulator, Network Simulator 2 (ns-2) [15]. The adoption of ns-2 allows us to simulate the effects of all the involved layers (physical, data link and networking) on the alarm dissemination process, which were not considered in previous work using similar traces [11].

To take into account the Wi–Fi technology properties, we have assumed a transmission range shorter than 50 m, as shown in Fig. 3, where the probability of correct packet reception between two neighboring devices is given as a function of the distance. The adopted transmission range takes into account the effective indoor transmission ranges of the different Wi–Fi standards, and further details about the adopted physical and data link layers, as well as the channel model, are described in [16]. According to this, in a certain time slot two users experience a transmission opportunity with a probability of packet delivery whose value depends on their mutual distance.

We note that the location measurements could be affected by a localization error greater than the adopted transmission range. However, since the localization error of the measurements can be assumed independent and identically distributed with zero mean, no biases are introduced in the statistics of inferred contact events, as shown in [11].

3.3. Disseminating emergency information

In our experiment, we consider a scenario in which an emergency occurs in a place with high population density and an alarm message must be broadcast to all the people located within a certain distance from the emergency location.

Table 1
Device classification summary.

Active	A device allowed to re-broadcast an alarm
Passive	A device never allowed to broadcast an alarm
Enabling	An active device that starts the alarm broadcasting
Infected	A device that has received an alarm
Target	A device located inside the target area

For alarm dissemination, we have assumed that a fraction of the devices, namely the *active* ones, cooperate in spreading the alarm by adopting a simple store-and-broadcast protocol: once an active node receives an alarm message (i.e., once it becomes *infected*), it keeps broadcasting the message every τ s.[4] The remaining devices, namely the *passive* ones, can simply receive the alarm messages from the active infected nodes but they never re-broadcast such messages.

The classification of nodes into active and passive, summarized in Table 1, has a twofold purpose.

First, it is unreasonable to assume that all the mobile phones actively participate in spreading the alarm. In fact, it can happen that only some special users (like police officers or firefighters) are authorized to forward alarms, and in this case the active nodes represent the *trustable* users. But it can also happen that some users could exhibit selfish behavior by disabling the broadcasting functionality, i.e., by acting as passive nodes. Or even some devices could be passive as a consequence of power or bandwidth management. As we do not have any information on the mobile phone users that could let us differentiate them, in our experiment we randomly assign a certain percentage of users as active, and the remaining ones as passive. We then repeat the experiments 50 times, each with a different selection of active users.

Moreover, by classifying nodes as active and passive, we can assess the alarm spreading as a function of the number of active nodes, and therefore indirectly as a function of the system requirements, such as bandwidth and minimum number of participating nodes.

In the following, we identify two circular areas centered on the emergency location and with different radii: the *enabling area* and the *target area*.

The enabling area is the region where the emergency takes place and where people are close enough to be aware that an emergency is occurring without the need to receive any message. As a consequence, alarm messages are originated by the active devices that, at a certain time, are located in this area, namely the *enabling devices*. Clearly, all the active devices contribute to diffusing the alarm by moving and by re-broadcasting a received alarm message, but only the enabling ones are allowed to start the spreading.

Finally, the target area is the region where the alarm messages must be spread so that people can leave the area or take a certain action to contain the emergency effects. Clearly, due to human mobility and message broadcasting, the alarm messages can also spread outside the target area and, based on the case, this effect can be suitable or not. However, since the target area is defined as the area where the emergency actually affects people, in this work we are mainly interested in characterizing the alarm diffusion in such an area.

4. Results of the experiment

In this section, we analyze and discuss the results obtained by means of numerical simulations with Network Simulator 2 based on the mobility patterns inferred by the experiment described in Section 3.

4.1. Scenario

In our experiment, we consider an emergency originated in one of two locations in Boston (see Fig. 2): (i) concert area, Hatch Shell; (ii) shopping center, Prudential Center. For both the locations, we consider an enabling area radius equal to 100 m or 200 m and two target areas with radii of 1 km and 2 km, respectively. The enabling area radii and the target area radii are modeled according to the lethal ranges and the minimal evacuation distances provided by the US Bureau of Alcohol, Tobacco and Firearms [17], respectively.

We adopt the *infection rate*, i.e., the rate between the number of infected target devices and the total number of target devices, as the performance metric. Clearly, due to human mobility some devices could enter or exit the target area during the time. However, in our scenario a device remains a 'target' also after it leaves the target area, giving a conservative measure of the performance that can be obtained. This assumption is reasonable, since each human that entered the target area at least once is a potential target of the emergency. As an example, the emergency could be a terrorist bio-attack and the emergency message could contain instructions on how to limit the injuries.

Since the speed of the alarm message diffusion deeply affects the amount of damage and/or number of injuries caused by the emergency, we consider 30 min as the period of time available for alarm spreading and analyze the infection rate as a function of time.

[4] In our experiment $\tau = 1$ s since such an interval time is commonly used in ad hoc routing protocols for state information or route discovery signaling.

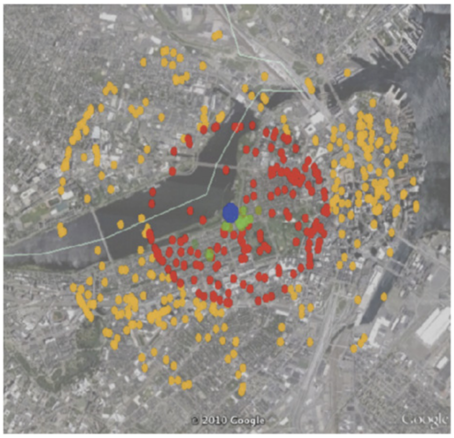

(a) 100 active nodes, after 10 min.

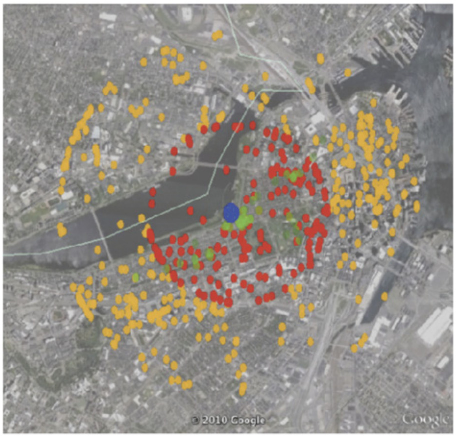

(b) 100 active nodes, after 20 min.

(c) 100 active nodes, after 30 min.

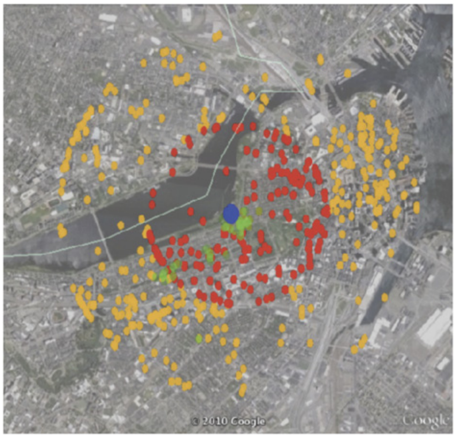

(d) 700 active nodes, after 10 min.

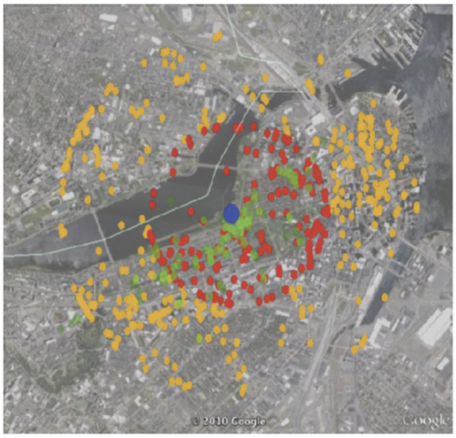

(e) 700 active nodes, after 20 min.

(f) 700 active nodes, after 30 min.

Fig. 4. Concert area, 5:00 pm: visualization of the alarm spreading in the target area for one simulation with the enabling radius equal to 100 m. The enabling area is colored in blue. Red circles represent nodes within 1 km of the alarm center, orange circles represents nodes within 2 km of the alarm center and green ones represent infected nodes. The radius of each circle is 50 m. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Moreover, as the population's spatial distribution and mobility change over time (as shown, for instance, in [12]), we consider two different starting times: (i) 5:00 pm; (ii) 9:00 pm. The first time corresponds to when people started gathering around the river, and the second time is in the middle of the concert and right before the fireworks.

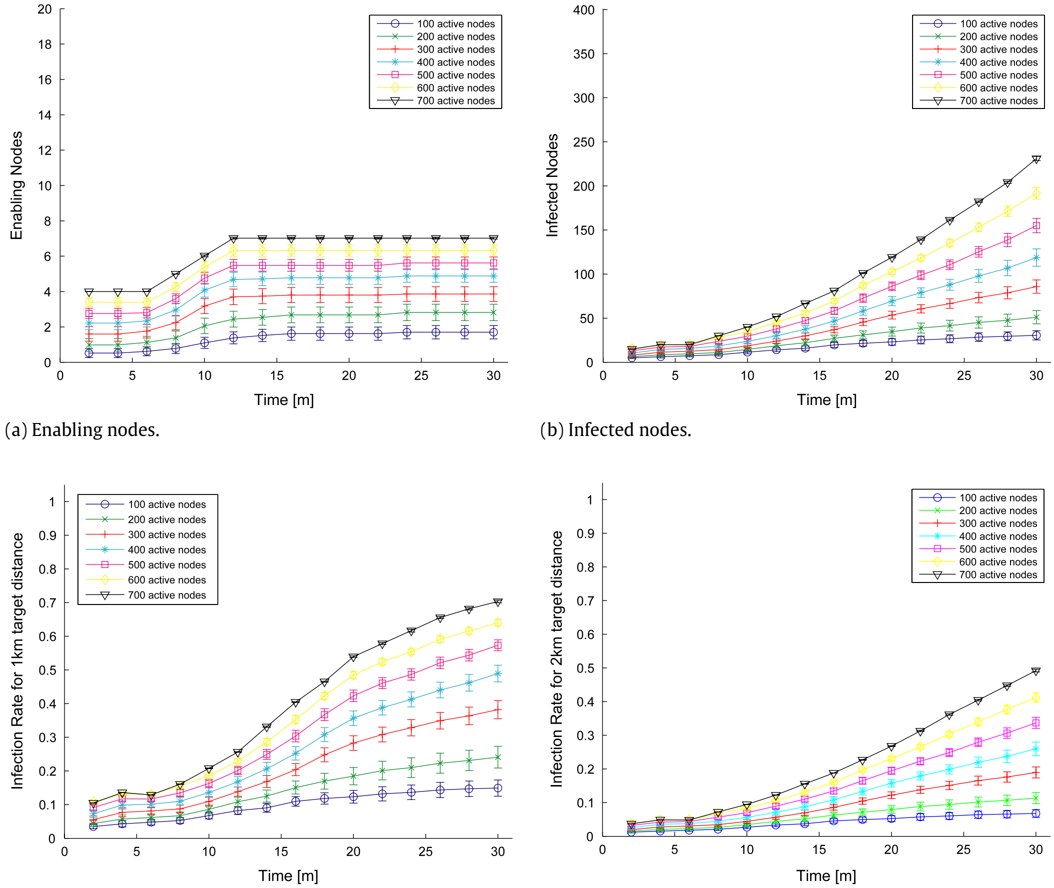

(a) Enabling nodes.

(b) Infected nodes.

(c) Infection rate for 1 km target area.

(d) Infection rate for 2 km target area.

Fig. 5. Concert area, 5:00 pm: average and 99% confidence interval of 50 simulations with the enabling radius equal to 100 m.

Thus, there are four considered scenarios for the experiment: (i) concert area, 05:00 pm, enabling radius equal to 100 m; (ii) concert area, 05:00 pm, enabling radius equal to 200 m; (iii) shopping center, 05:00 pm, enabling radius equal to 100 m; (iv) concert area, 09:00 pm, enabling radius equal to 100 m. For each scenario we ran fifty simulations, thus involving over 10^7 contact events, and in each simulation a different set of active nodes, randomly selected, is considered, i.e., we do not impose any specific mobility patterns on the active nodes. We note that, for the sake of completeness, in one of the experiments (see Fig. 5) we report both the average values and the 99% confidence intervals. It is easy to see that the number of simulated contact events is large enough to assure a high reliability for the estimation process.

4.2. Performance analysis

In the first scenario we consider an alarm originating at the concert area at 05:00 pm with an enabling radius equal to 200 m, and we analyze the results of the experiment as a function of the number of active nodes.

Fig. 4 shows a qualitative comparison (result of a single simulation) of the alarm diffusion performance with 100 and 700 active nodes. In the figure, the enabling area is denoted with a blue circle, the 1 km target nodes with red circles and the 2 km ones with orange circles, while the green circles represent the infected nodes. Obviously, the alarm spreading over time is more evident when 700 nodes are active. Moreover, when 700 nodes are active, the alarm spreads in the larger (2 km) target area, while the alarm is limited to the smaller (1 km) target area if only 100 nodes are active. Two factors concur in limiting the alarm diffusion: the low device density (with respect to the transmission range) and the limited mobility.

Therefore, in Fig. 5 we conduct a quantitative comparison of four different metrics as a function of time: (i) the number of enabling nodes (Fig. 5-a); (ii) the number of infected nodes (Fig. 5-b); (iii) the infection rate for the 1 km target area (Fig. 5-c); (iv) the infection rate for the 2 km target area (Fig. 5-d).

Clearly, the number of enabling nodes (Fig. 5-a) depends on the number of active nodes. However, even if almost all the nodes cooperate in the alarm diffusion, the average enabling node number is lower than eight. Moreover, the maximum number of enabling nodes is reached at the beginning of the alarm spread and it remains steady, meaning that the alarm spread is mainly due to broadcasting.

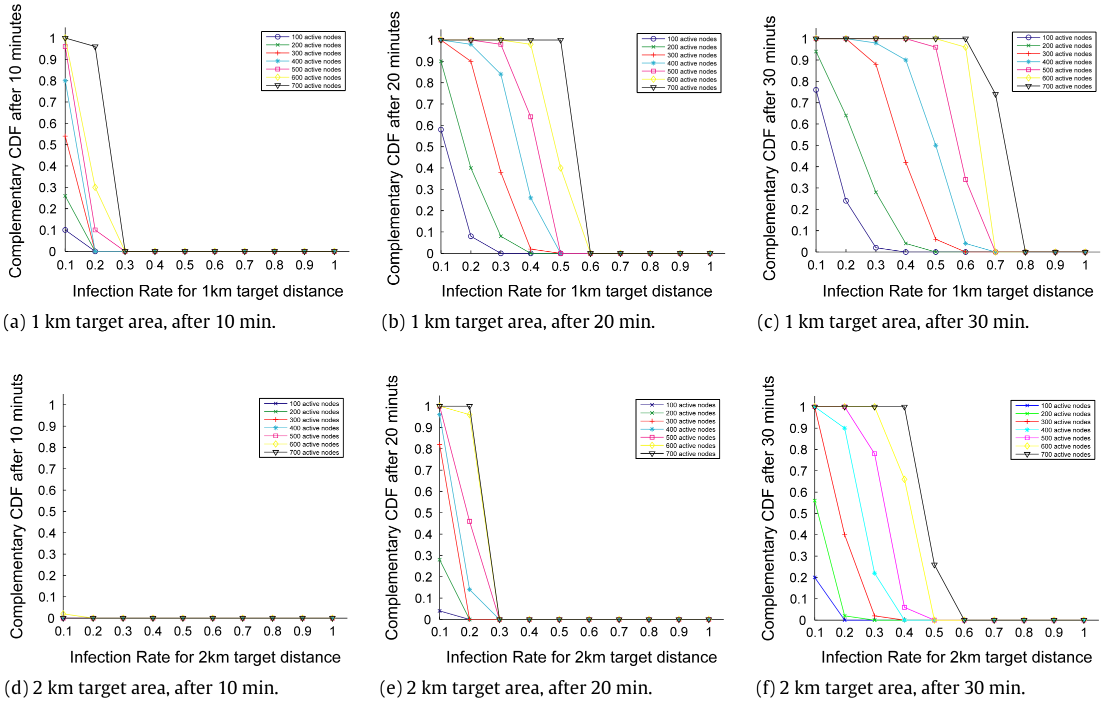

(a) 1 km target area, after 10 min. (b) 1 km target area, after 20 min. (c) 1 km target area, after 30 min.

(d) 2 km target area, after 10 min. (e) 2 km target area, after 20 min. (f) 2 km target area, after 30 min.

Fig. 6. Concert area, 5:00 pm: complementary CDF for the infection rate with the enabling radius equal to 100 m.

By comparing this value with the total number of infected nodes (Fig. 5-b), we have that the ratio between the infected and the enabling nodes is always greater than 10, thus proving that the alarm diffusion is mainly due to ad hoc communications. However, the results also show that the infected node number is deeply affected by the number of active nodes and that only a fraction (between 5% and 40%) of the devices are able to receive the alarm in a timely fashion.

To exclude from our analysis nodes that in the considered time interval are far away from the emergency location (the furthest are roughly 100 km distant), in the following two figures we limit our attention to the infection rate for the target areas. With regard to the smaller area (Fig. 5-c), we have that, at the beginning of the alarm spread, the infection rates increase smoothly, while after ten minutes all the curves start exhibiting a steep slope. Although less evident, this behavior is still present in the results for the larger area (Fig. 5-d). With reference to the smaller target area, the fraction of infected nodes is remarkable when a significant amount of nodes cooperate in alarm spreading, reaching after 30 min on average 50% of infected nodes when half of the nodes are active and a value greater than 70% in the most favorable case. In most cases, rates are roughly halved for the larger target area.

Finally, in Fig. 6 we analyze the distribution of the infection rate among the different simulations by means of the complementary cumulative distribution function (CDF). The complementary CDF for a random variable X denotes the probability $P(X \geq x)$. The results show approximate normal distributions. Moreover, they show that when only a fraction of the nodes are involved in alarm forwarding, the infection rates deeply depend on how these nodes are selected. For instance, if we consider the scenarios in which roughly half of the nodes cooperate for alarm diffusion (300 and 400 active nodes respectively), we have that after 30 min the infection rate varies between 20% and 70% for the smaller area and between 10% and 40% for the larger one.

In the second scenario, we consider again an alarm originating at the concert area at 05:00 pm, but with an enabling radius equal to 200 m. Clearly, the larger the enabling area is, the higher the infection rates are, as shown by Figs. 7 and 8. However, by comparing Fig. 5 with Figs. 6 and 7 with Fig. 8, respectively, it is easy to see that the behavior of the considered metrics does not change significantly when the enabling radius is doubled.

In the third scenario, we consider an alarm originating at the shopping center at 05:00 pm, and we compare the results for the infection rates with those obtained for the first scenario. Looking at Fig. 9, we observe that, not only are the infection rate values after 30 min similar to those obtained for the concert area scenario, but also the behavior of the curves is quite similar. This behavior is reasonable since the two originating locations are close to each other and thus both the mobility patterns and the human density are quite similar.

Therefore, in the fourth scenario we consider an alarm originating at the concert area at 09:00 pm. Here we have that all the infection rate values are higher than those obtained in the previous scenarios (Fig. 10). Moreover, we have that the infection rates increase rapidly at the beginning of the alarm spread. The main reason for this is that the characteristics of the mobility patterns of this scenario slightly differ from those of the previous two, mainly because people congregate around the concert area at 09:00 pm to attend the concert.

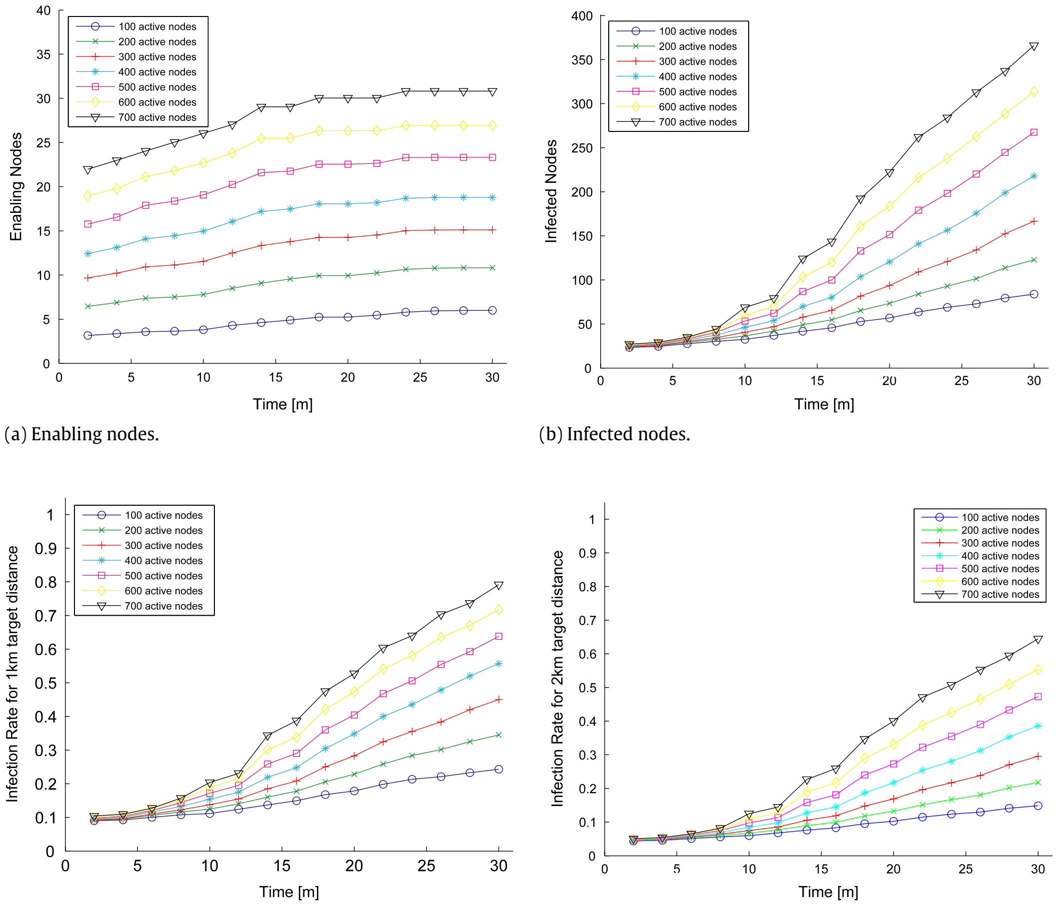

(a) Enabling nodes.

(b) Infected nodes.

(c) Infection rate for 1 km target area.

(d) Infection rate for 2 km target area.

Fig. 7. Concert area, 5:00 pm: average of 50 simulations with the enabling radius equal to 200 m.

To confirm this reasoning, in the last figure (Fig. 11) we analyze the distribution of the number of nodes infected by each active node (at the end of the alarm spread) by means of the complementary CDF with logarithmic scale. Here, the complementary CDF for a value *x* denotes the probability that an active node is responsible for having infected more than *x* nodes.

We found that in general a few active nodes are responsible for most of the information spreading. In fact, for both the scenarios the distributions approximate a power-law (straight lines in logarithmic scale), confirming the results obtained in smaller experiments. Moreover, by comparing the two scenarios we have that during the 09:00 pm scenario (Fig. 11-b) the distributions of infected (and thus met) nodes are higher than those of the 05:00 pm scenario (Fig. 11-a). Finally, when the number of active nodes is smaller, the distributions are rather flat, because each active node has the possibility of infecting many passive nodes. This is in contrast with what happens for a larger number of active nodes, where the distribution decreases more sharply.

5. Experimental limitations

Only one large scale event happened in the Boston metropolitan area during the considered time interval. We plan to collect data for other events in other cities in order to compare performance based on the type of event.

The number of nodes considered is a subset of the real number of mobile phone users at the event. This limitation is derived from the requirement to select users for which location information is sufficiently frequent to infer traces. This may create biases in the selected users. We expect that increasing the number of users would improve routing performance.

Finally, it is important to note that the large scale event analyzed did not involve any emergency situation. While we agree that human behavior might change in the case of emergencies [18], we believe that our analysis is still of interest as it relies on real human spatial distribution over large areas, and simulates communication opportunities during special events not involving large changes in mobility patterns.

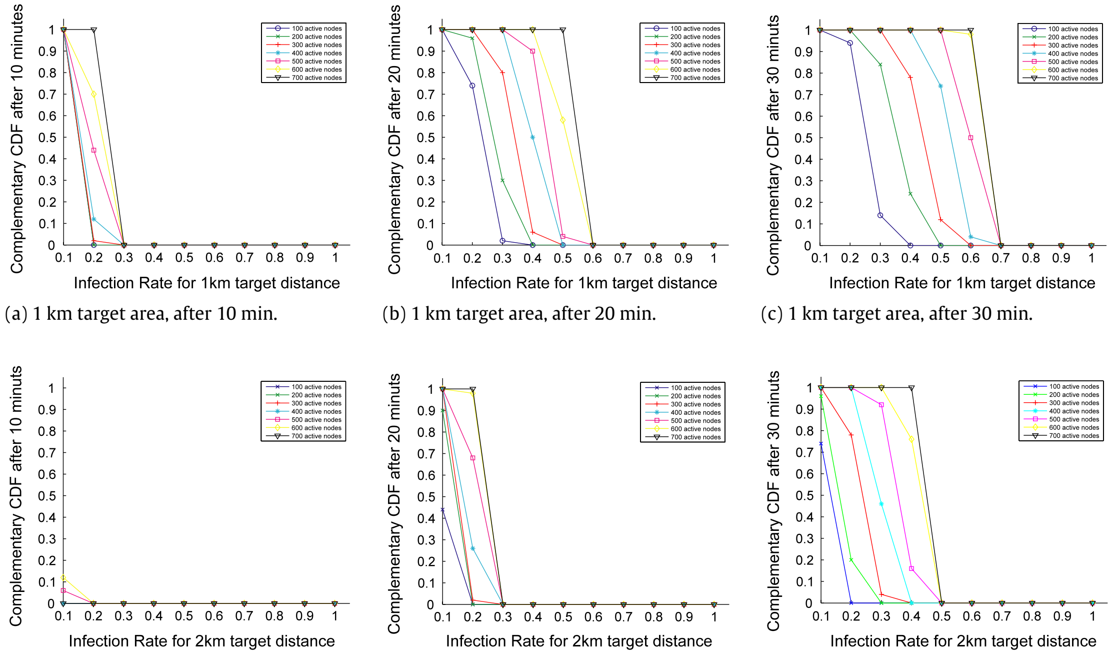

(a) 1 km target area, after 10 min. (b) 1 km target area, after 20 min. (c) 1 km target area, after 30 min.

(d) 2 km target area, after 10 min. (e) 2 km target area, after 20 min. (f) 2 km target area, after 30 min.

Fig. 8. Concert area, 5:00 pm: complementary CDF for the infection rate with the enabling radius equal to 200 m.

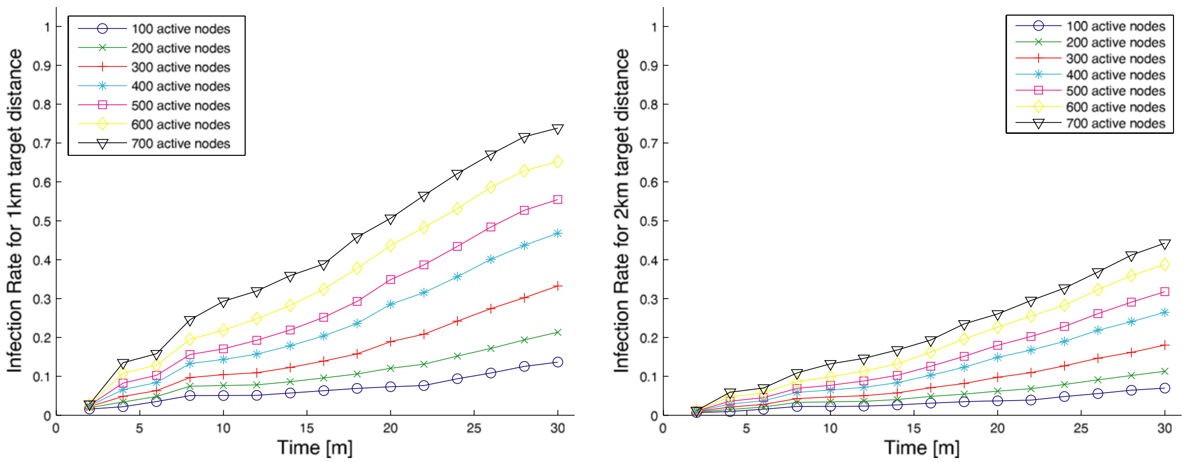

(a) Infection rate for 1 km target area. (b) Infection rate for 2 km target area.

Fig. 9. Shopping center, 5:00 pm: average of 50 simulations with the enabling radius equal to 100 m.

6. Conclusions

In this paper we have presented the results of an experiment in which the positions of seven hundred humans have been anonymously traced during a social gathering event to evaluate the possibility to disseminate alarms in a human mobility-enabled wireless network. During our experiments we found that the considered paradigm can effectively enable emergency communications among a significant subset of nodes, although the connectivity of the network strictly depends on the number of cooperating devices and on the maximum allowed delay. We hope that the encouraging results of this work will stimulate new data collection campaigns to further study large scale event scenarios for ad hoc communications. These results could be useful as a baseline for other experiments involving different mobility traces (e.g. from volunteers) during similar scenarios (large scale events), and for actual deployment of social communication based networks. As a future work, we are planning to analyze the mobility and inter-contact characteristics of the infecting nodes and design alarm dissemination strategies able to exploit these characteristics to minimize the overhead and resource requirements.

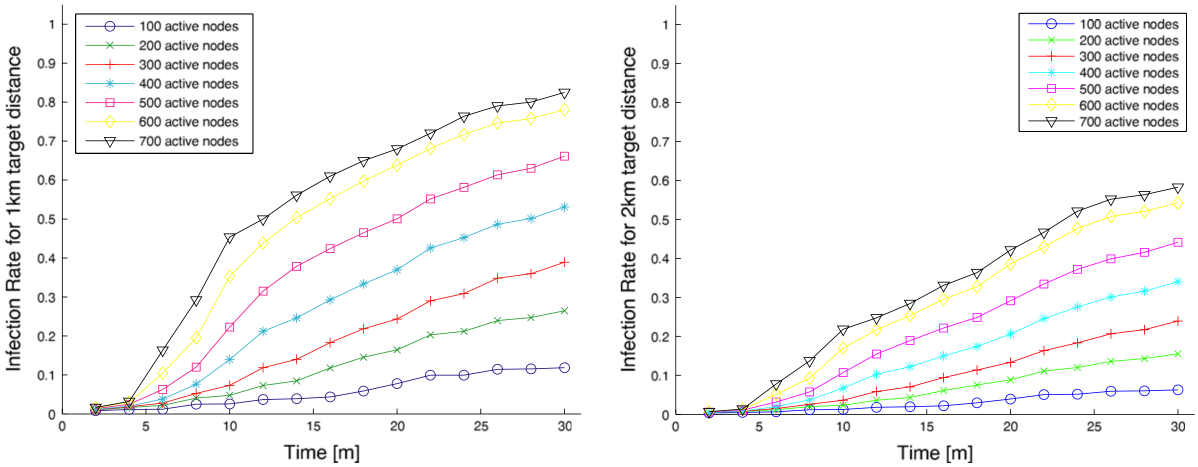

(a) Infection rate for 1 km target area.

(b) Infection rate for 2 km target area.

Fig. 10. Concert area, 9:00 pm: average of 50 simulations with the enabling radius equal to 100 m.

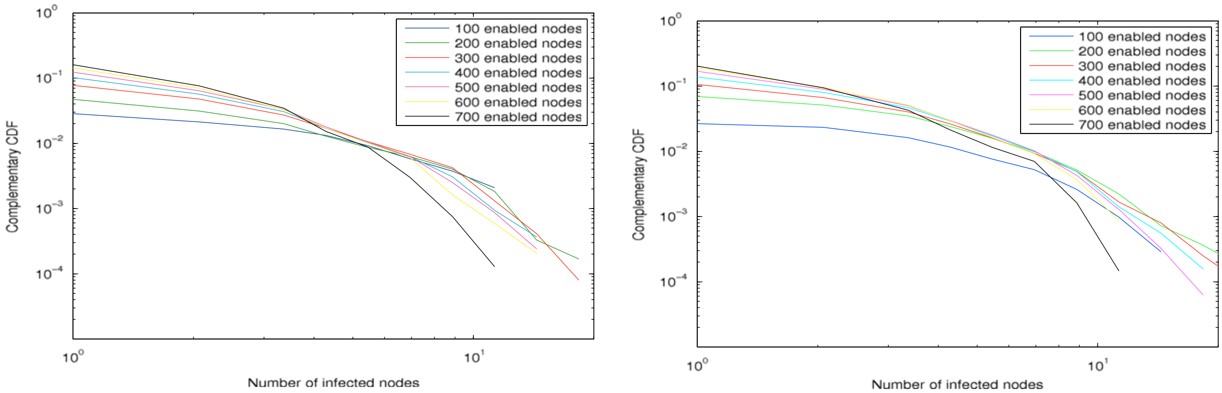

(a) 5:00 pm.

(b) 9:00 pm.

Fig. 11. Concert area: distribution of the number of nodes infected by the active ones after 30 min. Average of 50 simulations with the enabling radius equal to 100 m.

Acknowledgments

The first, third and fifth authors' work is partially supported by the project "Global & Reliable End to End e-Commerce & On Line Service Platform (GRECO)" founded by the Italian National Program *Industria 2015* and by the projects "HarBour traffIc opTimizAtion sysTem (HABITAT)" and "DRIVEr monitoring: technologies, methodologies, and IN-vehicle INnovative systems for a safe and ecocompatible driving (DRIVE IN2)" founded by the Italian National Program *Piano Operativo Nazionale Ricerca e Competitivitá 2007–2013*. The second and fourth authors' work was partially conducted while the authors were at the Department of Urban Studies and Planning, MIT.

References

[1] Tech. Rep., United Nations Foundations, 2009. http://www.unfoundation.org/what-we-do/legacy-of-impact/technology/disaster-relief-comms.html.
[2] P. Hui, A. Chaintreau, J. Scott, R. Gass, J. Crowcroft, C. Diot, Pocket switched networks and human mobility in conference environments, in: Proc. of the ACM Sigcomm Workshop on Delay-Tolerant Networking, WDTN, 2005, pp. 244–251.
[3] E.M. Daly, M. Haahr, The challenges of disconnected delay-tolerant manets, Elsevier Ad Hoc Networks 8 (2) (2010) 241–250.
[4] A.S. Cacciapuoti, M. Caleffi, L. Paura, Mobile p2p: peer-to-peer systems over delay tolerant networks, in: A. Vasilakos, Y. Zhang, T. Spyropoulos (Eds.), Delay Tolerant Networks: Protocols and Applications, CRC Press, 2011, pp. 1–35.
[5] A. Chaintreau, P. Hui, J. Crowcroft, C. Diot, R. Gass, J. Scott, Impact of human mobility on opportunistic forwarding algorithms, IEEE Transactions on Mobile Computing 6 (6) (2007) 606–620.
[6] P. Hui, J. Crowcroft, E. Yoneki, Bubble rap: social-based forwarding in delay tolerant networks, in: Proc. of the 9th ACM International Symposium on Mobile Ad Hoc Networking and Computing, MobiHoc, 2008, pp. 241–250.
[7] E.M. Daly, M. Haahr, Social network analysis for routing in disconnected delay-tolerant manets, in: Proc. of the 8th ACM International Symposium on Mobile Ad Hoc Networking and Computing, MobiHoc, 2007, pp. 32–40.
[8] J. Tang, M. Musolesi, C. Mascolo, V. Latora, Temporal distance metrics for social network analysis, in: Proc. of the 2nd ACM Workshop on Online Social Networks, WOSN, 2009, pp. 31–36.

[9] J. Leguay, A. Lindgren, J. Scott, T. Friedman, J. Crowcroft, Opportunistic content distribution in an urban setting, in: Proc. of the ACM SIGCOMM Workshop on Challenged networks, CHANTS, 2006, pp. 205–212.
[10] L. McNamara, C. Mascolo, L. Capra, Media sharing based on colocation prediction in urban transport, in: Proc. of the Annual Int. Conference on Mobile computing and networking, MobiCom, 2008, pp. 58–69.
[11] A. Cacciapuoti, F. Calabrese, M. Caleffi, G. DiLorenzo, L. Paura, Human-mobility enabled networks in urban environments: is there any (mobile wireless) small world out there? Elsevier Ad Hoc Networks (2011).
[12] F. Calabrese, G. DiLorenzo, C. Ratti, Human mobility prediction based on individual and collective geographical preferences, in: IEEE International Conference on Intelligent Transportation Systems, 2010, pp. 312–317.
[13] Wi-fi capable handsets: residential and enterprise markets for wi-fi and dual-mode handsets, Tech. Rep., ABI Research, 2009.
[14] Wi-fi enabled mobile phone handsets in the US, 2010–2015, Market Research Analysis, Coda Research Consultancy, 2010.
[15] K. Fall, K. Varadhan, The ns manual (formerly ns notes and documentation), Tech. Rep., The VINT Project, 2010.
[16] M. Caleffi, L. Paura, *M*-dart: multi-path dynamic address routing, Wiley: Wireless Communications and Mobile Computing 11 (3) (2011) 392–409.
[17] International Association of Fire Chiefs, National Fire Protection Association, Fundamentals of Fire Fighter Skills, Jones & Bartlett Publishers, 2008.
[18] J.P. Bagrow, D. Wang, A.L. Barabasi, Collective response of human populations to large-scale emergencies, PLoS ONE 6 (3) (2011).