Predicting Encounters in Opportunistic Networks

Radu Ioan Ciobanu
Faculty of Automatic Control and Computers
University Politehnica of Bucharest
Bucharest, Romania
radu.ciobanu@cti.pub.ro

Ciprian Dobre
Faculty of Automatic Control and Computers
University Politehnica of Bucharest
Bucharest, Romania
ciprian.dobre@cs.pub.ro

ABSTRACT

An opportunistic network is composed of human-carried mobile devices that interact in a store-carry-and-forward fashion. A mobile node stores data and carries it around; when it encounters another node, it may decide to forward the data if the encountered node is the destination or has a better chance of bringing the data closer to the destination.

In order to obtain efficient routing in such a network, we should be able to predict the future behavior of a node. This would help the algorithm decide if the data contained by the node should be further carried or forwarded, and to which node it is to be forwarded. In this paper, we present a mobile interaction trace collected at the University Politehnica of Bucharest in the spring of 2012, and analyze it in terms of the predictability of encounters and contact durations. We show that there is a regular pattern in the contact history of a node and then we prove that, by modelling the time series as a Poisson distribution, we can efficiently predict the number of contacts per time unit in the future. These assumptions are demonstrated both on the trace presented in this paper, as well as on a different trace recorded in another type of environment, showing that predictability doesn't happen only in strict and controlled situations.

Categories and Subject Descriptors

C.2.2 [**Computer-Communication Networks**]: Network Protocols—*routing protocols*; G.3 [**Mathematics of Computing**]: Probability and Statistics—*distribution functions, statistical computing, time series analysis*

General Terms

Theory, Experimentation, Verification

Keywords

Entropy, mobile devices, networking, opportunistic, Poisson distribution, predictability

HP-MOSys'12, October 25, 2012, Paphos, Cyprus.

1. INTRODUCTION

In recent years, mobile devices have become more and more ubiquitous. Thus, ad-hoc wireless networks formed over these devices have been analyzed thoroughly by researchers. One type of such a network is an opportunistic network, which consists of human-carried mobile devices that communicate with each other using a paradigm entitled store-carry-and-forward, without any infrastructure. In opportunistic networks, disconnections and highly variable delays caused by human mobility are the norm. The solution consists of dynamically building routes, as each node acts according to the store-carry-and-forward paradigm. Thus, contacts between nodes are viewed as an opportunity to move data closer to the destination. Such networks are therefore formed between nodes spread across the environment, without any knowledge of a network topology. The routes between nodes are dynamically created, and nodes can be opportunistically used as a next hop for bringing each message closer to the destination. Nodes may store a message, carry it around, and forward it when they encounter the destination or a node that is more likely to reach the destination.

The challenge in opportunistic networks is knowing when and to whom should a node pass a message in order to obtain the best possible hit rate and latency. Thus, predicting the future encounters of a device is paramount to implementing a good opportunistic routing algorithm. In this paper, we propose a way to predict the future behavior of a node in the opportunistic network by analyzing its past encounters and approximating the time series as a Poisson distribution. We use a mobile trace gathered at the University Politehnica of Bucharest in 2012 to prove our suppositions. Aside from the fact that it is a relatively small enclosed space with a lot of participants, the other great advantage of an opportunistic network implemented over an academic environment is that the contacts are inherently regular: students attend the same classes with the same professors every week. We thus show that the contact history of a node in terms of number of encounters may be approximated as a Poisson distribution. We then compute the Poisson probabilities for the final two weeks of the experiment and prove that they hold in practice.

The rest of the paper is structured as follows. Section 2 presents recent work in the area of opportunistic networking and event detection and predictability. Section 3 offers information regarding the tracing experiment at the University Politehnica of Bucharest. Next, Sect. 4 performs an analysis of the trace in terms of predictability, by computing entropies for the number of encounters and contact durations. Section 5 shows that approximating the trace as a Poisson distribution can lead to good predictability results. Finally, Sect. 6 presents conclusions and future work.

2. RELATED WORK

A thorough review of opportunistic networking is presented in [6]. The analysis, developed in the context of the EU Haggle project, highlights the properties of main networking functions, including message forwarding, security, data dissemination and mobility models. The authors also propose various solutions for communication in opportunistic networks and introduce HCMM, a mobility model that merges the spatial and social dimensions. Moreover, several well-known opportunistic forwarding algorithms are also presented, such as BUBBLE Rap [11], PROPICMAN [19] and HIBOp [2]. Other dissemination techniques for opportunistic routing include Socio-Aware Overlay [22], Wireless Ad Hoc Podcasting [16] or ContentPlace [3]. Of these algorithms, only ContentPlace performs a prediction of the future encounters of a node, based on the history of previous contacts. A taxonomy for data dissemination algorithms is proposed in [4].

Jain et al. [15] propose a framework for evaluating routing algorithms for delay-tolerant networks and analyze the performance of several such algorithms in terms of the amount of knowledge about the network that they require. The authors reach the conclusion that the algorithms that use the least knowledge tend to perform poorly, but that efficient algorithms can be constructed using less than the complete global knowledge. A set of abstract knowledge oracles are proposed for the framework, that are able to answer questions about the environment (i.e. the network). There are four types of oracles: Contacts Summary Oracle (which provides the average waiting time until the next contact), Contacts Oracle (which answers questions regarding contacts between two nodes at any point in time), Queuing Oracle (which gives information about buffer queuing) and Traffic Demand Oracle (which can answer any question regarding the present or future traffic demand). Several algorithms are then proposed which gradually use all the oracles, and the authors show that the smarter algorithms (i.e. with more knowledge) perform better than the others both in terms of delay and delivery ratio, and the performance differences become more clear when the network is more intermittent. However, the conclusion is that in networks with a lot of communication opportunities (such as a network formed over a small enclosed space with lots of participants, like the one presented in this paper) the need for routing algorithms with a great amount of knowledge is minimal. On the other hand, in resource-limited environments, where contact opportunities, bandwidth and storage are low, smarter algorithms would provide a significant benefit, especially those which use the Queuing and Traffic Demand Oracles. However, these two are the hardest to implement in real life and in many cases it may not even be worth it.

Islam and Waldvogel [14] state that opportunistic routing protocols in the literature are dependent on the history of devices for extracting routing information, but the size of this routing information is limited, which may introduce inaccuracies and thus a weaker message delivery. The authors analyze the predictive qualities of history-based routing algorithms using extensive simulation on real CRAWDAD traces [7], and reach the conclusion that the repetitive nature of a path is proportional to the mobility extent of the devices and thus contact history obtained from dense opportunistic networks can be reliable.

In [20], Song et al. analyze the predictability of human behavior and mobility on user traces obtained from mobile carriers. They use several variants of the entropy function as the most fundamental quantity that captures the degree of probability that characterizes a time series. Ihler et al. [12] model a normal periodic behavior of a time series by a time-varying Poisson process model and then modulate it using a hidden Markov process, in order to account for bursty events. They show that using a Poisson model is significantly more accurate at detecting future behavior and known events than a traditional threshold-based technique. This model can also be used to investigate periodicity in the data, such as systematic weekday and time of day effects. Since contact information in an opportunistic network is also a time series, we believe that it can also be approximated as a Poisson distribution, as shown in Sect. 5.

3. TRACING EXPERIMENT

For gathering traces of human mobility, an experiment was performed at the University Politehnica of Bucharest in the spring of 2012 [17]. For this experiment, an application entitled HYCCUPS Tracer was implemented with the purpose of collecting contextual data from Android smartphones. It runs in the background and collects availability and mobile interaction information such as usage statistics, user activity, battery statistics or sensor data, but what really interests us is the fact that it collects information about encounters between devices. This is done in two ways: Bluetooth and AllJoyn [13]. Bluetooth interaction scans for paired devices in the immediate vicinity and stores contact information such as the ID of the encountered device and the time of contact, as well as the duration of the contact. The information stored by AllJoyn tracing is similar, but is based on the construction and deletion of wireless sessions using the AllJoyn framework based on WiFi. The difference between Bluetooth and WiFi is that WiFi consumes more battery life, but it is more stable [8]. Tracing was executed periodically with a predefined timeout for Bluetooth, and asynchronously for AllJoyn interactions.

The duration of the tracing experiment was 64 days, between March and May 2012 and had 66 participants. The participants were chosen so that they covered as many years as possible from both Bachelor and Master courses. Thus, there were: one first year Bachelor student, one third year Bachelor student, 53 fourth year Bachelor students, three Master students, two faculty members and six external participants (from an office environment). We were interested only in the participants at the faculty, so we eliminated the external nodes. We also eliminated some nodes that had too little contact information, because they were irrelevant to our experiment. Such nodes belonged to students that did not keep their Android application on at all times when they were at the faculty as they were instructed, or who haven't gone to many courses in the experiment period. In the end, we remained with 53 nodes in the experiment that had useful information.

An academic environment is a natural situation where an opportunistic network can bring benefits [18, 9, 10], therefore we consider the participants in this experiment as part of such a network. The advantages of opportunistic routing and data dissemination are accentuated in an academic environment mainly because it is an enclosed space with lots of potential device carriers located at very close distances to each other. Provided we have a good dissemination algorithm, this means that data will circulate quickly and reach all desired destinations with a good probability.

4. TRACE ANALYSIS IN TERMS OF PREDICTABILITY

In order to verify if the behavior of nodes in an opportunistic network built in an academic environment is predictable, we have analyzed in detail the mobile trace presented in Sect. 3. This section presents an overview of this analysis, in regard to node encounters and contact duration. Moreover, we present the entropy values obtained from the trace using two different probability functions and show that certain information regarding future contacts can be predicted.

4.1 Encounters and Contact Duration

Because the nodes in the opportunistic network presented in this paper are students and teachers at the University Politehnica of Bucharest, we believe that their behavior is predictable. This should happen because the participants have a fixed daily schedule and interact with each other at fixed times in a day. For example, a teacher and the students from a group interact when the students attend the teacher's class, which happens regularly each week. Likewise, two students in the same group would interact at almost all times when they are at the faculty.

We tried to prove that this supposition holds by analyzing the traces presented in Sect. 3. The first metric used was the total number of encounters between a node and the other nodes. An encounter is considered to begin at the first moment when a node is in range of the current node and to end when the two nodes are not in range anymore for a certain period of time. The number of encounters specifies the popularity of a node (the more encounters it has, the more popular it is). The second metric we used was the contact duration of every encounter of a node in a given time interval. Similar to the number of encounters, it suggests the popularity of a node, but also its mobility. If a node has many encounters in a time interval, but all the encounters are short in terms of duration, it means that the node is very mobile and it doesn't stay in the same place for long periods of time.

Figure 1 shows the total number of daily encounters of a randomly chosen node with each of the other nodes for the entire duration of the experiment. It can be seen that on Tuesdays, Wednesdays and Thursdays, the node has regular encounters with fairly the same nodes. The number of contacts per day sometimes differs, which probably happens because there were short periods of time when the nodes were not in contact (for example, a student went out of the faculty grounds), but the nodes encountered are basically the same every week. This shows (on a purely intuitive level for now) that there may be a certain amount of predictability in the behavior of our nodes. It should be noted that not all charts for other nodes in the experiment look as regular as the one shown in Fig. 1, but for most of them the behavior appears predictable.

Figure 2 presents contact durations per day for the same

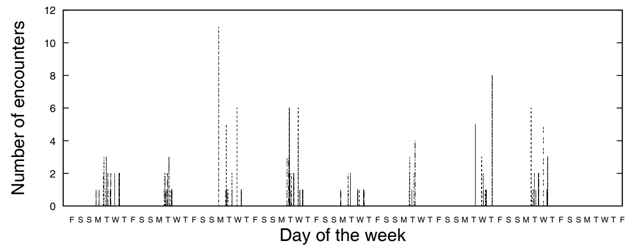

Figure 1: Total encounters per day for a random node.

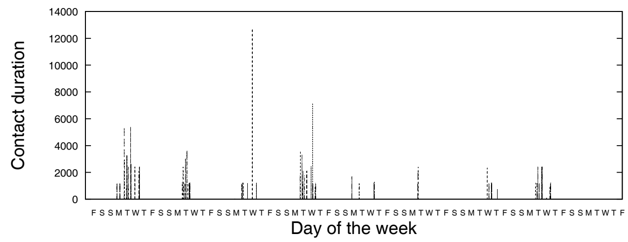

Figure 2: Total contact durations per day for a random node.

node as before. Just as in Fig. 1, it can be seen that on Tuesdays, Wednesdays and Thursdays the contact durations are similar. The conclusion here is that contact durations are also seemingly predictable, albeit the charts are less regular than the ones for encounters.

4.2 Entropies

In order to verify if a node's behavior in the opportunistic network is predictable, we used Shannon's entropy, which is basically a measure of predictability (the lower the entropy, the higher the chances are of a prediction being successful). When the entropy is 0, it means that a node's behavior is 100% predictable. The formula for entropy is: $H(X) = -\sum_{i=1}^{n} p(x_i) \ln p(x_i)$, where X is a discrete random variable with possible values in the interval $x_1, ..., x_n$ and $p(X)$ is a probability mass function for X. The obvious thought here would be to use the probability of encountering node N at the next time interval as $p(X)$. However, the sum of probabilities in this case would not be 1 as it should be, because a node might be in contact with more than one other node at a given time. Therefore, we have split the entropy computation into two parts: predicting that the next encounter will be with node N, and predicting if a contact will take place at the next time interval. In theory, combining these two values will result in a prediction of the time of an encounter with a given node.

The first step was to compute the entropies for contacts with a given node N. The probability function in this case was computed as the ratio between the total number of times node N was encountered and the total number of contacts with any nodes during the experiment. The second entropy function was computed as a ratio between the number of time units a node was in contact with another node for, and the entire duration of the experiment. We used different values for the time unit (one second, one minute and one hour) and the results were similar. Figure 3 shows the cumulative distribution functions for the two entropies. It can be

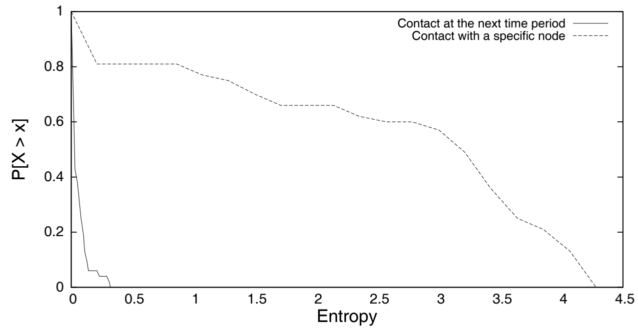

Figure 3: Entropy values for predicting the time of contact with any node and contacts with given nodes.

observed that having a contact at the next period of time is mostly predictable, because the entropy is always lower than 0.35. However, predicting the node that will be seen at the next encounter is not so easily done based solely on the history of encounters, and this is shown by the high entropy values (as high as 4.25, meaning that a node may encounter on average any one of $2^{4.25} \simeq 19$ nodes).

In the rest of this paper, we will focus on attempting to predict the time of the next encounter with any node and the total number of encounters at that moment, and leave the prediction of contacts with a certain node for future work. As we have previously seen, further information is required in order to predict exactly what nodes will be encountered, such as knowledge about the social relationships between the participants in the opportunistic network. As shown in [5], there is a higher chance of a node interacting with nodes that are part of its social circle than with other nodes, and this information can be used in designing an efficient algorithm.

5. PREDICTING FUTURE NODE ENCOUNTERS

As we have seen in the previous section, the entropy for predicting if a contact will take place at the next time interval is lower than 1, which means that the behavior of a node in terms of encounters with any other nodes is highly predictable. Since there are only two cases in this prediction (i.e. having a contact or not having a contact), we can model a node's behavior as a Bernoulli distribution, which is a particular case of a binomial distribution. However, simply knowing if there will be a contact at a given time may not be enough for a good opportunistic routing algorithm. Therefore, it would be good if we were able to know exactly how many contacts will there be in the specified time interval, and a binary distribution such as Bernoulli does not offer such information. Consequently, we believe that a Poisson distribution might be suitable for this situation, because it expresses the possibility of a number of events (in our case encounters with other nodes) to occur in a fixed time interval. We will show in the rest of this section that the Poisson distribution applies to the trace presented in Sect. 3.

The probability mass function of a Poisson distribution is $P(N,\lambda) = \frac{e^{-\lambda}\lambda^N}{N!}$, where in our case $P(N,\lambda)$ represents the probability of a node having N contacts at a given time interval. In order to prove that a Poisson distribution applies to our trace, we used Pearson's chi-squared test [21], which tests a null hypothesis stating that the frequency distribution of mutually exclusive events observed in a sample is consistent with a particular theoretical distribution (in our case Poisson). We applied the chi-squared test for every node in the network individually.

The time interval chosen for applying the Poisson distribution and the chi-squared test was one hour. We tried to choose this interval in order to obtain a fine-grained analysis of the data. Choosing a smaller interval (such as a minute) and estimating the next contact incorrectly may lead to missing it completely (in the sense of using it appropriately). When we have an interval such as an hour, we can predict that in the next hour there will be a certain number of contacts with a higher rate of succes, and the opportunistic routing algorithm can be ready for those contacts in the respective hour.

The first step of the chi-squared test was to count the frequency distribution of contacts per hour for the entire duration of our trace. We then stated the null hypothesis, namely that the number of encounters a node has per hour follows a Poisson distribution. The λ parameter can either be included in the hypothesis or it can be estimated from the sample data (as it was in our case). We computed it using the max likelihood method by averaging the number of encounters per hour over the entire experiment. Knowing λ, we were then able to find out the probability for having N encounters at the next time interval according to the Poisson distribution. Using this probability, we finally performed the chi-squared test according to the formula $\chi^2_{k-p-1} = \sum_k \frac{(f_o - f_e)^2}{f_e}$, where f_o is the observed frequency, f_e is the expected frequency (computed using the Poisson distribution), k is the number of classes (which depends on the way the number of encounters are distributed for each node) and p is the number of parameters estimated from the data (in this case 1, the λ value).

We used the 0.05 level of significance for proving the hypothesis using a chi-squared table, and the results can be seen in Fig. 4 (1). We also included the nodes that have not had any encounters in the "Accepted" category, since a distribution with only zeros is a valid Poisson distribution. As can be seen from Fig. 4, only 20.75% of the hypotheses were

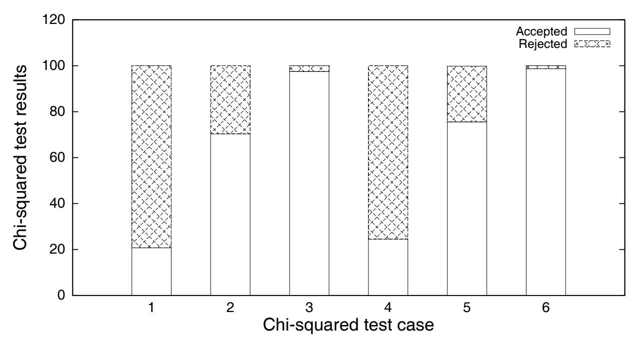

Figure 4: Results of chi-squared tests for various scenarios. Datasets 1, 2 and 3 are computed using the total number of encounters and varying the max likelihood (1 - for the entire experiment, 2 - per weekday, 3 - per hour of a day of the week). Datasets 4, 5 and 6 are computed using unique encounters.

accepted in this case. However, we have observed in Sect. 4 that a node's encounter history has a somewhat repetitive pattern for days of the week, so we then attempted to compute λ as the averaged number of contacts in the same day of the week. Therefore, we ended up with a larger number of chi-squared hypotheses to prove (53 nodes $\times$ 7 days in a week) for each node, but also with a much finer-grained approximation of the data. The results for this situation can be seen in Fig. 4 (2), with only 29.65% of the hypotheses being rejected. Still we went one step further, knowing that students at a faculty generally follow a fixed schedule in given days of the week and thus we computed the max likelihood value as an average per hour per day of the week (consequently having 53 $\times$ 7 $\times$ 24 hypotheses per node). Thus, the results obtained were very good, with only 2.49% of all the hypotheses rejected.

The previous results were computed for the total number of encounters in an hour. However, if the Android tracer application misbehaved at some point in the experiment and instead of logging a long contact between two nodes, logged a large number of very short contacts, the results of applying a Poisson probability may be wrong. Because of this situation, we also applied the chi-squared tests described above using only unique contacts. Therefore, the number of contacts in a hour will be equal to the number of different nodes encountered in that hour. The results are shown in Fig. 4 (4,5,6). For the first test case (with λ computed over the entire experiment), 75.47% of the hypotheses were rejected. In the test that uses the average per day of the week, 24.26% of all chi-squared hypotheses were rejected and finally just 1.31% of distributions were not Poisson according to the chi-squared test for computing the max likelihood value per hour of a weekday.

Although these results look good, it might be that this situation only happens on this particular trace. Consequently, we ran all the tests presented in this section on another trace [5] which was also performed at the University Politehnica of Bucharest, but for a shorter duration (35 days) and with fewer participants (22 students and teachers). However, the results for the λ-per-hour test with unique contacts are even better than for the current trace, since only 0.11% of the hypotheses were rejected.

However, the traces performed in the campus of our faculty may be a particular case, meaning that we cannot generalize our assumptions yet. Therefore, we tested our theory on a different type of trace, entitled St. Andrews [1], which was collected using a mobile sensor network with Tmote Invent devices carried by 27 members of the University of St. Andrews for a period of 79 days, in which the participants were asked to carry their devices and to keep them on at all times, whether in or out of the town of St. Andrews. The Invent devices were able to detect and store information about encounters between each other within a radius of 10 metres, and were programmed to send discovery beacons at every 6.67 seconds. This trace corresponds to a different situation that the traces performed at the University Politehnica of Bucharest, since this is a larger and more open space, with less regularity, less contact opportunities and smaller contact durations. The results we obtained were good, with only 12.39% of the chi-squared hypotheses being rejected.

To further prove our assumptions, we have eliminated the last two weeks from the trace and computed the Poisson distribution probabilities for each hour per day of the week

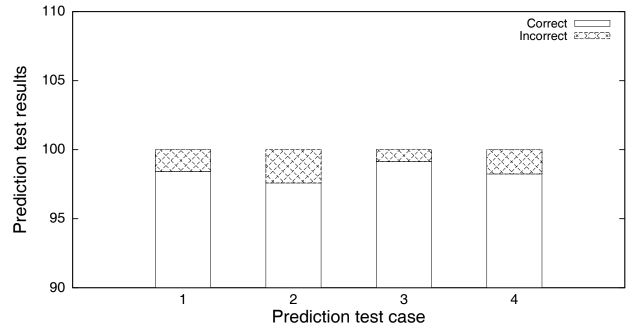

Figure 5: Prediction success of the Poisson distribution.

on the remaining series. We compared the value that had the highest Poisson probability (i.e. the most likely value according to the distribution) with the real values. If the Poisson predictions were to be correct, then the two values should be equal. The results of this test for both total and unique contacts are shown in Fig. 5. It can be seen that for total encounters 97.59% of the Poisson-predicted values are correct for the next to last week, and 98.42% for the last. When taking into account individual encounters, the predictions are even better (98.24% and 99.14% respectively).

We have shown in this section that, by knowing the history of encounters between the nodes in an opportunistic network in every hour of every day of the week, we can predict the future behavior of a device in terms of number of contacts per time unit. Having this knowledge helps the opportunistic routing algorithm decide what to do with the data it carries (data that has either been generated locally, or that is currently only being carried around by the node). If the node knows that it will have few contacts this hour, but the number will increase for the next hour, it may choose to keep the data bundles until then, instead of forwarding and then deleting them. Likewise, if the next time intervals are known to have few encounters, the node might choose to forward data while it still has the chance, lest the messages are delayed for long periods of time. An entire set of heuristics can be applied in such a situation, and this is something we wish to study in detail as future work.

6. CONCLUSIONS AND FUTURE WORK

We have presented here a mobile trace performed over an academic environment at the University Politehnica of Bucharest for two months in the spring of 2012. We have analyzed it from the point of view of node encounters and contact durations and we have also shown that the history of contacts in terms of the number of encounters can be modelled as a Poisson distribution with very good predictability results. We have proven these assumptions to hold for two other traces [5, 1] that have different conditions then the one presented in this paper. Having the knowledge about future encounters is basically an implementation of the Contacts Summary Oracle, according to Jain et al. [15]. Because the network presented here offers a lot of contact opportunities, a good routing algorithm can easily be created with only this information. Thus, we believe that expanding routing algorithms with knowledge about future node encounters can lead to better results in terms of hit rate and latency.

However, adding even more knowledege in the shape of Jain's Contacts Oracle, i.e. knowing which nodes will be encountered in a given time interval, will improve the routing algorithm further. Therefore, we plan to attempt a prediction of the contact durations and particular nodes that a device will be in contact with in the future, using similar methods. It is also important to note that, since nodes in an opportunistic network are devices belonging to human carriers, the social aspect should be taken into account. Since it has been shown that nodes tend to interact more often with other members of their social communities [5], if the predictions are made not only by simply analyzing the history of events, but also by combining this history with knowledge about social relationships and communities, a very efficient opportunistic routing algorithm can be implemented.

Having the Contacts Summary Oracle and the Contacts Oracle, we plan to develop our own opportunistic routing algorithm and test its efficiency on the traces presented here. We wish to focus on the advantages of the social aspects of opportunistic networking, since we believe this is the future in pocket-switched networks.

7. ACKNOWLEDGMENTS

This work was partially supported by project "ERRIC - Empowering Romanian Research on Intelligent Information Technologies/FP7-REGPOT-2010-1", ID: 264207, and by national project "TRANSYS - Models and Techniques for Traffic Optimizing in Urban Environments", Contract No. 4/28.07.2010, Project CNCSIS-PN-II-RU-PD ID: 238. The work has been co-funded by the Sectoral Operational Programme Human Resources Development 2007-2013 of the Romanian Ministry of Labour, Family and Social Protection through the Financial Agreement POSDRU/89/1.5/S/62557.

8. REFERENCES

[1] G. Bigwood, D. Rehunathan, M. Bateman, T. Henderson, and S. Bhatti. Exploiting self-reported social networks for routing in ubiquitous computing environments. In *Proceedings of the 2008 IEEE International Conference on Wireless & Mobile Computing, Networking & Communication*, pages 484–489, Washington, DC, USA, 2008. IEEE Computer Society.

[2] C. Boldrini, M. Conti, J. Jacopini, and A. Passarella. HiBOp: a History Based Routing Protocol for Opportunistic Networks. In *World of Wireless, Mobile and Multimedia Networks, 2007. WoWMoM 2007. IEEE Int. Symp. on a*, pages 1–12, 2007.

[3] C. Boldrini, M. Conti, and A. Passarella. Exploiting users' social relations to forward data in opportunistic networks: The HiBOp solution. *Pervasive Mob. Comput.*, 4:633–657, 2008.

[4] R. Ciobanu and C. Dobre. Data dissemination in opportunistic networks. In *18th Int. Conf. on Control Systems and Computer Science*, CSCS-18, pages 529–536, 2011.

[5] R. I. Ciobanu, C. Dobre, and V. Cristea. Social aspects to support opportunistic networks in an academic environment. Feb. 2012.

[6] M. Conti, S. Giordano, M. May, and A. Passarella. From opportunistic networks to opportunistic computing. *Comm. Mag.*, 48:126–139, 2010.

[7] CRAWDAD. http://crawdad.cs.dartmouth.edu/.

[8] E. Ferro and F. Potorti. Bluetooth and Wi-Fi wireless protocols: a survey and a comparison. *IEEE Wireless Comm.*, 12(1):12–26, 2005.

[9] T. Henderson, D. Kotz, and I. Abyzov. The changing usage of a mature campus-wide wireless network. In *Proc. of the 10th annual int. conf. on Mobile computing and networking*, MobiCom '04, pages 187–201, New York, USA, 2004. ACM.

[10] P. Hui, A. Chaintreau, J. Scott, R. Gass, J. Crowcroft, and C. Diot. Pocket switched networks and human mobility in conference environments. In *Proc. of the 2005 ACM SIGCOMM workshop on Delay-tolerant networking*, WDTN '05, pages 244–251, New York, USA, 2005. ACM.

[11] P. Hui, J. Crowcroft, and E. Yoneki. BUBBLE Rap: social-based forwarding in delay tolerant networks. In *Proc. of the 9th ACM int. symp. on Mobile ad hoc networking and computing*, MobiHoc '08, pages 241–250, New York, USA, 2008. ACM.

[12] A. Ihler, J. Hutchins, and P. Smyth. Learning to detect events with markov-modulated poisson processes. *ACM Trans. Knowl. Discov. Data*, 1(3), Dec. 2007.

[13] Q. I. C. Inc. Introduction to AllJoyn. HT80-BA013-1 Rev. B, 2011.

[14] M. A. Islam and M. Waldvogel. Prediction quality of contact history in opportunistic networks. In *2011 IFIP Wireless Days (WD)*, pages 1–3. IEEE, Oct. 2011.

[15] S. Jain, K. Fall, and R. Patra. Routing in a delay tolerant network. *SIGCOMM Comput. Commun. Rev.*, 34(4):145–158, Aug. 2004.

[16] V. Lenders, M. May, G. Karlsson, and C. Wacha. Wireless ad hoc podcasting. *SIGMOBILE Mob. Comput. Commun. Rev.*, 12:65–67, 2008.

[17] R. C. Marin and C. Dobre. Exploring predictability in mobile interaction. May 2012.

[18] M. McNett and G. M. Voelker. Access and mobility of wireless PDA users. *SIGMOBILE Mob. Comput. Commun. Rev.*, 7:55–57, 2003.

[19] H. A. Nguyen, S. Giordano, and A. Puiatti. Probabilistic routing protocol for intermittently connected mobile ad hoc network (propicman). *2007 IEEE Int. Symp. on a World of Wireless Mobile and Multimedia Networks*, pages 1–6, 2007.

[20] C. Song, Y. Qu, N. Blumm, and A.-L. Barabási. Limits of Predictability in Human Mobility. *Science*, 327:1018–1021, 2010.

[21] A. Stuart, K. Ord, and S. Arnold. *Kendall's Advanced Theory of Statistics, Classical Inference and the Linear Model*, volume Volume 2A (2007 reprint)). Wiley, sixth edition, 1999.

[22] E. Yoneki, P. Hui, S. Chan, and J. Crowcroft. A socio-aware overlay for publish/subscribe communication in delay tolerant networks. In *Proc. of the 10th ACM Symp. on Modeling, analysis, and simulation of wireless and mobile systems*, MSWiM '07, pages 225–234, New York, NY, USA, 2007. ACM.