An Analysis of Techniques for Opportunistic Networking

Alexandru Asandei, Ciprian Dobre
Department of Computer Science and Engineering
University POLITEHNICA of Bucharest
Bucharest, Romania
E-mails: alexandru.asandei@cti.pub.ro,
ciprian.dobre@cs.pub.ro

Princy Johnson
School of Engineering, Technology and Maritime Operations
Liverpool John Moores University
Liverpool , UK
E-mail: P.Johnson@ljmu.ac.uk

Abstract— **Opportunistic networks, as opposed to classic ones, rely on human-carried devices. This leads to a high mobility of all the nodes in the network and a very dynamic topology. Developing good forwarding algorithms for optimum routing in opportunistic networks presents a specific set of challenges. The most notable forwarding solutions rely on social human mobility patterns. They depend on several key factors and assumptions. The purpose of this paper is to provide an increased understanding on how these factors can influence routing performance. For this purpose we focus on the number of copies for each generated message, altruism and frequency of exchanged messages. The analysis is based on data collected from a social tracing experiment conducted at the University Politehnica of Bucharest. The setup and implementation of this experiment will also be presented in this paper.**

Keywords-opportunistic networking, social networks, wireless communication, mobile networks, routing performance, academic trace

I. INTRODUCTION

Opportunistic mobile networks consist of human-carried mobile devices that communicate with each other in a store-carry-and-forward fashion, without any infrastructure [1]. Compared to the classical networks, they present distinct challenges. In opportunistic networks, disconnections and highly variable delays caused by human mobility are the norm rather than an exception. 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 opportunities 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.

In order to design more robust routing algorithms, real-life traces can be used to extract information on the patterns of movement that people carrying mobile devices follow. Traces are collected using different types of mobile devices for various kinds of communication as well as for different scenarios. Such traces can provide valuable information for designing a good dissemination algorithm for opportunistic networks. In this paper **we present a study designed to analyze the efficiency of several popular routing algorithms.** The analysis is based on the data collected during a social tracing experiment that took place between March and May 2012 at the University Politehnica of Bucharest, Romania.

In an opportunistic network, the nodes are people that carry mobile devices. These people are organized into *communities*, based on their shared profession, workplace, hobbies etc. Generally, members of the same community interact with each other more often than with members from outside the community [17], so the community organization should be considered when designing the algorithms for opportunistic networks. Recently, increased research efforts have been focused on looking at how the social organization of humans can be used when making routing decisions in opportunistic networks. Most common approaches consist of routing messages towards more popular nodes (forward messages to people having more connections or being more mobile than the current node, therefore having an increased chance of meeting other nodes closer to the messages' destinations). The network formed between nodes, where the links are the similarity in terms of communication between them, is known as the *social graph*. We argue and demonstrate through this experimental research that such an approach has a greater tendency of leading to congestion, as popular nodes can quickly become crowded with messages. In this work we demonstrate this tendency in academic environments with medium-density crowds such as classrooms. The situation is similar in places such as airports, train stations, shopping malls, etc. In such cases, opportunistic routing algorithms trying to increase the chances of a message being successfully delivered can easily fail because of induced congestions. Here we also present alternative solutions to solve this. In addition to congestions, the human's altruism can also constitute a bottleneck for message delivering. Therefore, we consider it appropriate to also analyze the effect of altruism, as an inherent property of humans carrying the mobile devices.

The rest of the paper is organized as follows. Section 2 presents a brief overview of existing work in the area of opportunistic routing. Section 3 evaluates the BUBBLE Rap opportunistic routing algorithm along with its modified versions. In Section 4 we present the tracing experiment, and analyze the results obtained by applying several opportunistic techniques, including the BUBBLE Rap and its modified versions to the data. In Section 5 we analyze the effects of altruism on these techniques. Finally, in Section 6 we conclude the paper and present our proposed future work.

II. Related Work

An opportunistic networking algorithm faces a huge challenge: choosing the next-hop node that might forward the message one step closer to its destination, considering that the message may also be discarded along its path. Previous studies demonstrated this to be a scheduling hard problem [11]. Therefore, this research focuses on the algorithms designed to *increase the probability* of messages being successfully transferred to their destinations (*increased hit rate*) [2], while keeping the time needed to complete the transmission low (*decreased transmission delay*) [3]. More recently, authors of [4] proposed an additional dimension: preserving the success of transmission when nodes/people also behave with selfish attitude (*impact of the altruism* as a measure of fault tolerance).

Authors of [9] propose Socio-Aware Overlay, an opportunistic dissemination algorithm that creates an overlay for publish/subscribe systems (where some nodes registered to receive specific events from other nodes). The overlay is composed of nodes with high values of *centrality* (a measure of the probability of being part of communication paths). Socio-Aware Overlay is socially-aware, having its own community detection methods. The authors also propose two algorithms for distributed community detection, named Simple and k-CLIQUE.

Authors of [12] propose a dissemination technique called ContentPlace that attempts to deal with data dissemination in resource-constrained opportunistic networks by making the content available only in regions where interested users are present, without overusing available resources. In order to optimize content availability, ContentPlace exploits learned information about the users' social relationships, in order to decide where to place the user data. ContentPlace's design is based on two facts: users can be grouped together logically according to the type of content they are interested in, and their movement is driven by social relationships. In order to select data from an encountered node, nodes from ContentPlace use a utility function by means of which each node can associate a utility value to any data object. When a node encounters a peer, it computes the utility values of all the data objects stored in the local and in the peer's cache. Then, it selects the set of data objects that maximizes the local utility of its cache.

The BUBBLE Rap algorithm [2] uses *local community structures* and the *centrality* of nodes to make the forwarding decision. It makes the assumptions that each node belongs to a community, and that every node has a different ranking (its social popularity) within the global and the local community. The algorithm forwards each message to the encountered nodes with ranks higher than the node currently storing them. Local ranking is used when encountering nodes from inside the local community. The global ranking is used when encountering a node outside the local community.

These and other similar opportunistic algorithms use the social graph that models the relations between the participants as it has proven to be a more reliable source to identify *local communities* and *centrality*. We refer to the two terms as they were described in the BUBBLE Rap algorithm [2]. Social information has been used to a high extent to construct routing schemes based on local communities and centrality. The authors of [2] provide a solution to increase the hit rate by using a high number of message copies. The solution provided by the Spray-Select-Focus scheme [3] solves the rapid-flooding problem by keeping only a relatively small number of fixed copies for each message.

The social graph also groups nodes/participants in local communities, which are further used by these algorithms: they rely on some community detection algorithm to discover these communities. However, they also do not use the user's feedback to validate the community detection decision. It has been shown in [5] that using the social graphs provided by the users can actually increase routing performance. However, the social relations between users do not necessarily mean only connections. The authors in [4] present the impact of several distributions of altruism in *social network models* and *human mobility traces*. It also describes useful communication patterns, which is also used in our analysis.

Since it is one of the most complex and well-studied algorithms in terms of performance opportunistic algorithms proposed so far, we will refer to the BUBBLE Rap algorithm as the bench mark in order to compare the performance of the techniques discussed in this paper. . In particular, we analyze its main flaw: in the classical BUBBLE Rap algorithm messages tend to stack in nodes having the highest rankings. Also, the original carrier deletes the message from its buffer only if the message is forwarded to its destination community. Until this happens, the node will continue to forward copies of the message, which negatively impacts the performance of the overall network when messages are exchanged frequently between users. Such results were also observed in our experimental analysis as demonstrated in the following sections.

III. Opportunistic Networking Techniques

The techniques used in our analysis are based on the routing principles of one of the most popular opportunistic approaches, namely BUBBLE Rap [2]. This routing algorithm, in fact, is the result of many popular studies on opportunistic networks [6][13][14].

Our analysis is based on two proposed algorithms, along with the classic BUBBLE Rap (for comparison). In order to make an equal comparison between these methods, the classic method (hereafter referred as the *classic* BUBBLE Rap approach) was slightly modified to use a fixed number of copies (in all experiments we consider the same number of messages exchanged in the network) of messages. Introducing a fixed number means that when a message reaches its copy limit it is always deleted from the nodes' buffer after it is forwarded to another node.

Our proposed methods are as follows. First, the *outer* algorithm assumes that sending a single copy of each message inside a community has a higher chance to reach nodes in that community, considering the route-to-more-popular-nodes BUBBLE Rap approach. As a result, it tries to push copies of the messages having destinations outside the community directly to the nodes from that community. After a first copy is forwarded inside the local community, all other copies will be reserved for encounters with outer nodes. Most forwarding decisions are still made according to the BUBBLE Rap algorithm, with the difference that this decision is not always taken (if the inner or outer number of copies has reached its limit with regard to the current message) and that messages not destined for the local community are forwarded to the first node encountered from another community regardless of the node's rank. Finally, the *lowerEPS* algorithm introduces a 10% random factor into the classic BUBBLE Rap algorithm. This is based on the observations that messages having lower ranked nodes as destinations in the classic BUBBLE Rap have to go all the way to the nodes with higher rank and then back to their destinations (just like in a hierarchical social tree). However, in the *lowerEPS* approach, when encountering a node with a lower rank, the algorithm has a 10% chance of forwarding the message to that node regardless of its rank. This probability was introduced to deal with the uncertainties [10] in computing the rank/popularity of each node (just because we have no observations we cannot make the assumption that two nodes are not connected).

Another important aspect that differentiates the algorithms is the method used to detect the local communities. The original BUBBLE Rap uses the k-clique algorithm to statistically learn the node's community based on the history of encounters. In our previous studies [5] we have shown that a social graph can deliver better results. Therefore, in our current approach only the BUBBLE Rap uses the k-CLIQUE, while the newly proposed algorithms rely on the social graph provided by the users to detect communities and the node centrality.

The local communities and centrality are two parameters used by many forwarding algorithms [2][13]. We also analyze them in our experiments. We also make two assumptions. First, we consider the use of a fixed number of copies (differently distributed in the experimental scenarios) for each message. Secondly, we assume nodes generate messages according to a community-biased communication pattern [4], with a higher probability to generate messages destined for nodes in the same local community as the originator.

IV. Evaluation Using Real-World Traces

In order to obtain trace information regarding the mobility of the members inside an academic environment a real-world tracing experiment has been performed at the University Politehnica of Bucharest (UPB).

For the UPB trace (hereafter called UPB2012), the participants were asked to run an Android application called HYCCUPS Tracer (http://cipsm.hpc.pub.ro/hyccups), which was designed to collect contextual data from smartphones. The application runs in the background and collects features pertaining to availability and mobile interaction, namely the sensor data (accelerometer and proximity sensor readings), Alljoyn interactions (construction and deletion of wireless sessions by means of the Alljoyn framework [15]), Bluetooth interactions (paired devices in the immediate vicinity), Wi-Fi Hotspots (periodical Wi-Fi scan results including name, MAC address and signal strength of wireless APs in the close surroundings). Tracing is executed both periodically (using a predefined timeout), as well as asynchronously on certain events (such as Alljoyn interactions or user events). In addition, all the participants were asked to provide access to their social graph via their Facebook account.

We applied the proposed opportunistic algorithms on top of this trace, emulating their behavior. For the generation of messages we consider the use of a community-biased communication pattern. This means that periodically half of the total number of devices generate new messages, resulting in n/2 new messages where n is the total number of devices. It is community-biased because the destination is chosen from the local community of the node generating the message with a high probability (according to a normal distribution). We also assume that each node has a message buffer that can store up to 2048 messages. The size of the buffer was chosen by considering an average message size of 10KB which would lead to a total of 20MB at maximum load. Most Android (newer generation) devices have a memory limit for each application of 16 to 24 MB [16]. Implicitly the message buffer of a node cannot exceed this limit.

In the end we evaluated several performance metrics. The *hit rate* is an important metric for an opportunistic network, as it represents the ratio between data objects that have successfully arrived at requesting nodes, and the total number of requests. It suggests the efficiency of a dissemination algorithm, showing the fraction of requests that can be served by a dissemination algorithm. The *delivery cost* represents the ratio between the total number of exchanged messages during the course of the experiment and the number of generated messages. It should be as low as possible as it represents the congestion of the network. The *latency* values represent the time (in milliseconds) elapsed between generating a message and delivering it to the destination. In an opportunistic network, which is a type of delay tolerant network, delivery latency is not as important, nonetheless it should be improved when possible.

The x-axis in Figures 3-5 below represents the time (in days) when new messages are generated. This ranges from messages generated only once in the lifespan of the experiment (at the beginning) - which is the value 60 - to 0.02 days - which means new messages are generated every half hour.

The UPB2012 experiment was conducted over a period of 9 weeks, from early March to early May 2012. Most participants were bachelor students in their final year of study. There were a total of 52 participants. Amongst the participants 2 were members of the faculty staff, and 2 were Master students.

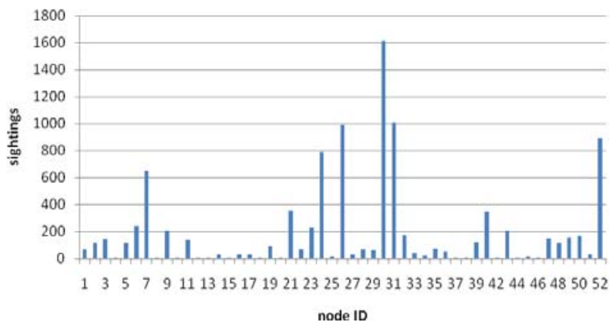

Fig. 1. Sighting distribution for the UPB 2012 trace.

The HYCCUPS application was designed to scan for peer-to-peer interaction using both Bluetooth and WiFi (to compensate for the deficiencies of each technology). The period between two successive periodical discoveries was user-configurable - ranging from 10 to 60 minutes, with a default value of 15 minutes used by the majority of participants.

When encountering another participant a wide variety of information is logged and stored in the device's memory. There are two main information groups: the identity of the two devices having an interaction and data regarding the state of each device which is collected regardless of any interactions. For the Bluetooth connections, MAC address and device name are used to uniquely identify devices. For WiFi connections we used the IMEI code of the phone (if available) or the Alljoyn ID. WiFi interactions also logged the disconnection timestamp, not only the initial contact timestamp (as was the case for Bluetooth).

A contact is considered to start at the first time a device is seen and end when it is not rediscovered in the time interval of the periodic discovery. Throughout the duration of the experiment, there were a total of 12003 contacts with an average contact time of 27 minutes and an average inter-contact time of 35.85 hours. The contact time represents the duration of a contact between two devices. Inter-contact time represents the intervals between successive encounters between the same devices. Contact distribution, which is the number of times a certain device is encountered, can be observed in Figure 1.

Social communities are an important aspect of the experiment. The average degree in the social graph is 11, the minimum is 1 and the maximum number of social connections is 34. The social graph can be observed in Figure 2. Users within the same community tend to interact more with each other, and previously [5] we demonstrated that the use of this information can lead to better performance. Several human mobility models previously proposed, such as [7] and [8] also highlight the interaction patterns observed in the experimental data better.

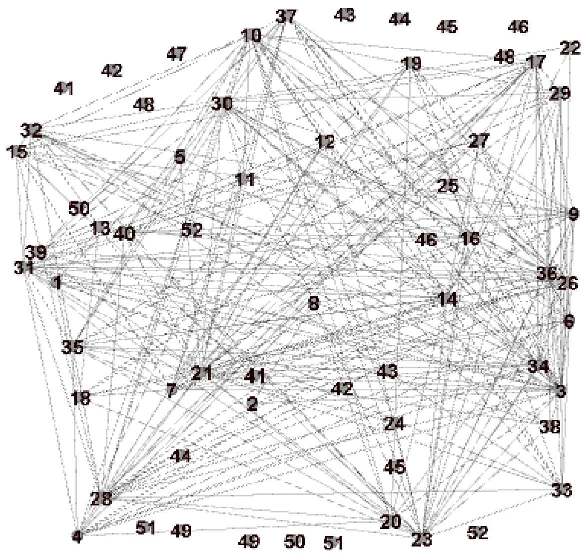

Fig. 2. The social graph, UPB 2012 trace.

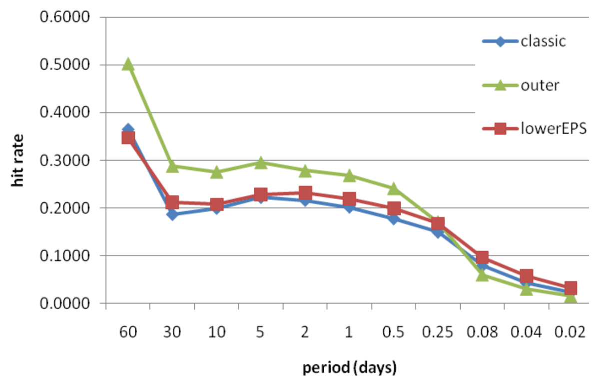

Fig. 3. The results for hit rate, UPB 2012 trace, 2 copies.

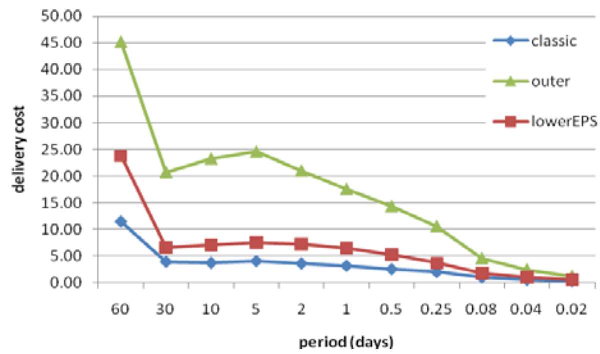

Fig. 4. The results for delivery costs, UPB 2012 trace, 2 copies.

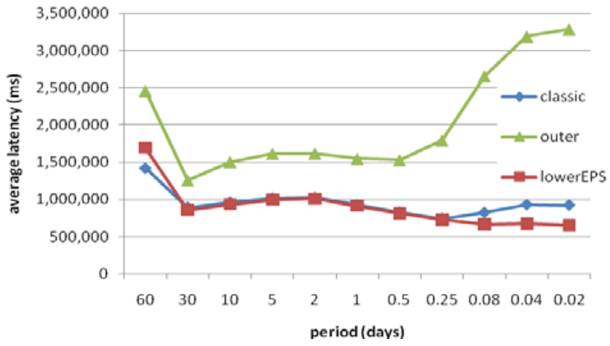

Fig. 5. The results for latency, UPB 2012 trace, 2 copies.

Using this experimental setup we emulated the proposed algorithms. First, we limited the number of copies per message to only 2 copies. In this case, the obtained results are presented in Figures 3-5.We notice that *outer* has the best average hit rate. After the 0.25 mark (new messages generated every 6 hours), when *outer* starts losing its advantage all algorithms begin to have a very similar hit rate and delivery cost. In this scenario, the classic BUBBLE Rap algorithm has the poorest performance compared to the other solutions.

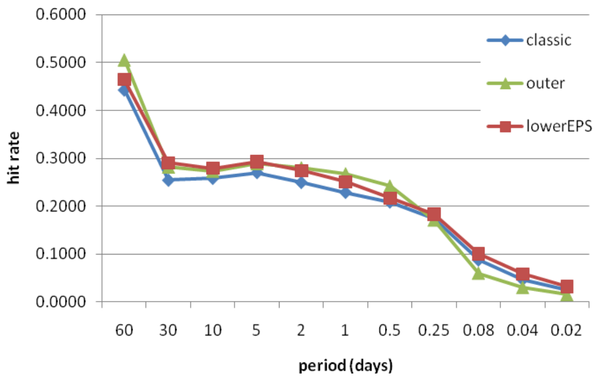

Fig. 6. The results for hit rate, UPB 2012 trace, 8 copies.

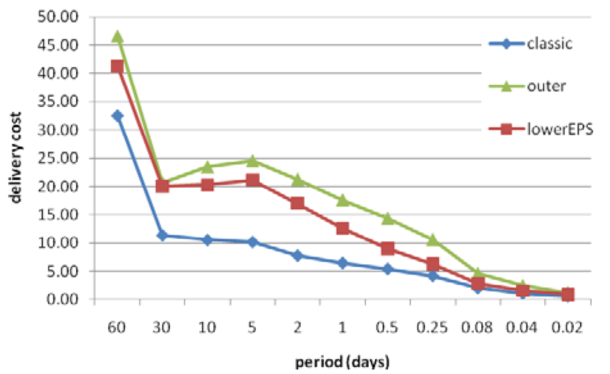

Fig. 7. The results for delivery costs, UPB 2012 trace, 8 copies.

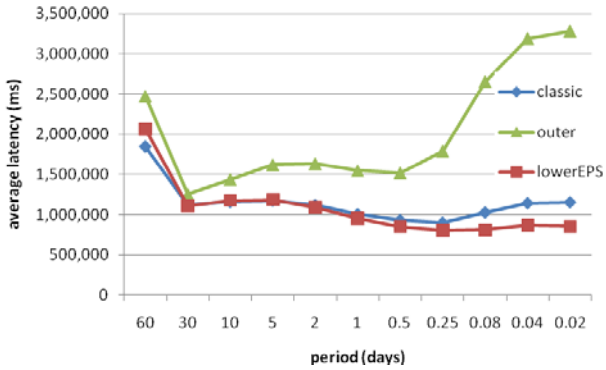

Fig. 8. The results for latency, UPB 2012 trace, 8 copies.

While the hit rate tends to become indistinguishable between the algorithms as the message generation period decreases, the average latency shows important changes in behavior. It is clear that outer becomes very inefficient generating a high latency but it is also noticeable that the exploration component provided by *lowerEPS* contributes to a lower average latency as the 0.25 mark is reached. This happens while having a slightly higher hit rate as the *classic* and *outer* methods, meaning that the messages that are indeed delivered are delivered through a faster route.

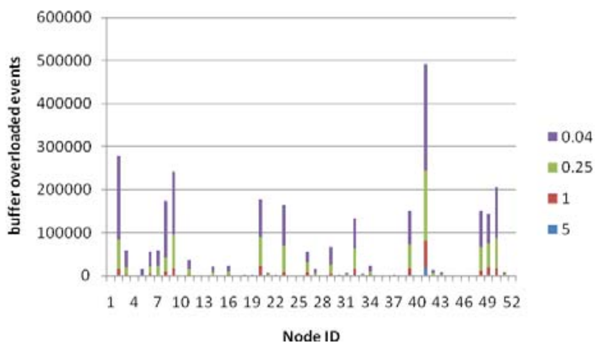

Fig. 9. The distribution of buffer saturation for *classic*, UPB 2012 trace, 8 copies.

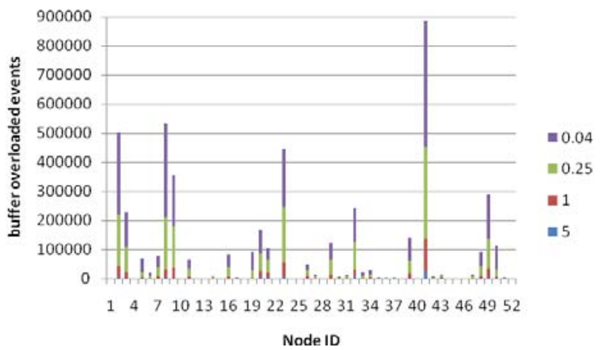

Fig. 10. The distribution of buffer saturation for *outer*, UPB 2012 trace, 8 copies.

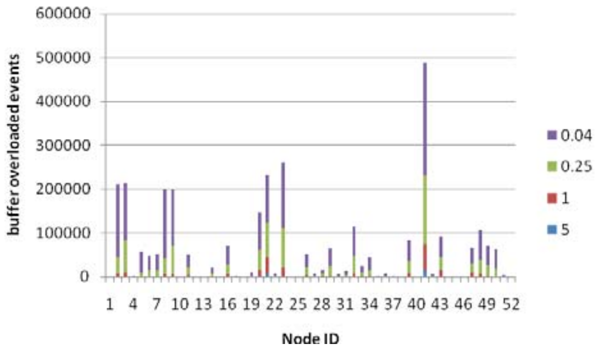

Fig. 11. The distribution of buffer saturation for *lowerEPS*, UPB 2012 trace, 8 copies.

When the number of copies is increased to 4 the behavior observed is somewhere between 2 and 8 copies. Thus we present in Figures 6-8 the results for 8 copies to have a more contrasted view on the trends of the observed metrics.

As can be observed in Figure 6 the average hit rate of *outer* remains almost the same, gaining little benefit from the increased number of copies. The other two algorithms thus manage to reach a hit rate almost as good as *outer*.

In terms of delivery cost, a similar trend can be observed with a particularity. The average delivery costs of *lowerEPS*

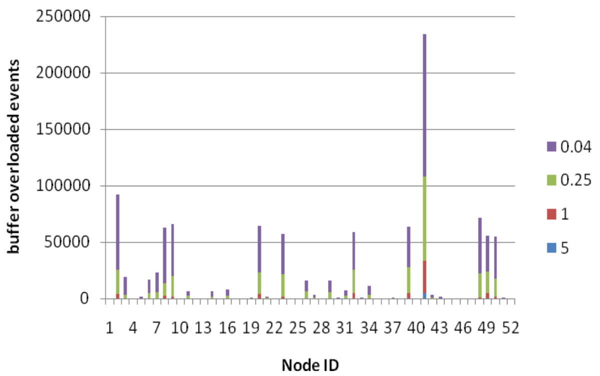

Fig. 12. The distribution of buffer saturation for *classic*, UPB 2012 trace, 2 copies.

have increased more than *classic* although the increase in hit rate is comparable.

While the 0.25 mark remains an important pivot point for most metrics, in the case of *lowerEPS* the exploration component exhibits its impact on latency sooner. As soon as the daily mark is reached the trend in latency of *lowerEPS* is to distance itself from that of *classic*. This is a strong argument for the benefits of a small random exploratory component.

Another important aspect to be considered when analyzing these solutions is how the network becomes saturated with messages. In Figures 9-11 we show the distribution of buffer saturation for a number of 8 copies per message. The legend shows the values for the considered message generation periods (in days).

Buffer saturation is distributed almost the same for all solutions. This shows that the basic forwarding strategy of BUBBLE Rap causes messages to stack in the nodes with higher centrality, thus noticeably overloading them as the daily mark is hit for the generation of messages.

Having a trace with a significant number of participants implies that our results are relevant in conditions of saturation. Even though it occurs slightly later and to a small extent the saturation of the network is significant even for only 2 message copies, as depicted in Figure 12.

V. The Role of Altruism in Opportunistic Routing

Selfishness appears when a node refuses to forward any messages except those generated by itself. Altruism is the opposite of this. Such behavior of nodes can occur in real-life for a variety of reasons. For example, based on the preferences of each user, they might have a higher availability to forward messages generated from inside their local community. Altruism can also be used to model node failures due to network issues.

Two altruism models were implemented, in terms, to outline the tolerance of the proposed algorithms. From the models presented and studied in [4], normal distribution and community-biased distribution will be used in our analysis as they were considered to be the most realistic.

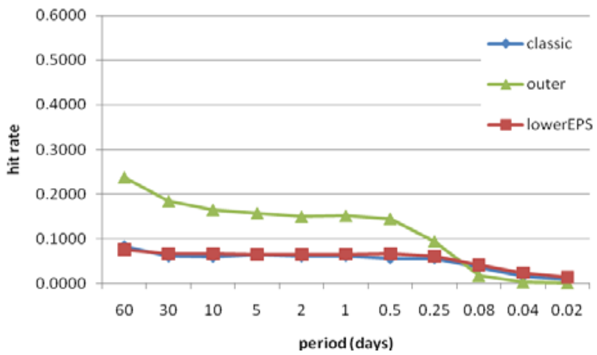

Fig. 13. The results for hit rate, Normal distribution altruism, UPB 2012 trace, 2 copies.

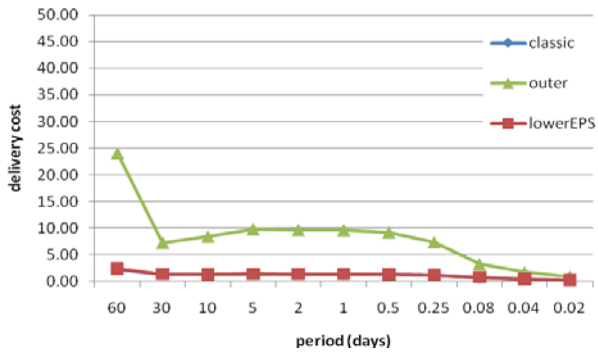

Fig. 14. The results for delivery costs, Normal distribution altruism, UPB 2012 trace, 2 copies.

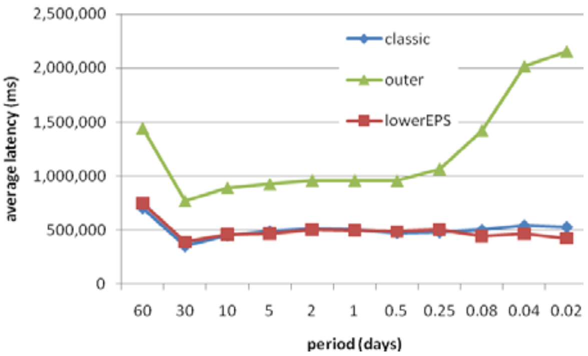

Fig. 15. The results for latency, Normal distribution altruism, UPB 2012 trace, 2 copies.

In the case of normal distribution, the nodes that cease forwarding messages are chosen according to a normal distribution between 0 and 1 with a threshold of 0.5. Whereas community-biased distribution adds a constraint that allows altruism to manifest only when forwarding messages that are destined for another community. In other words, a node will never refuse to forward messages destined for its local community.

First we conduct the same experiments using normal distribution altruism and present the results in Figures 13-15.

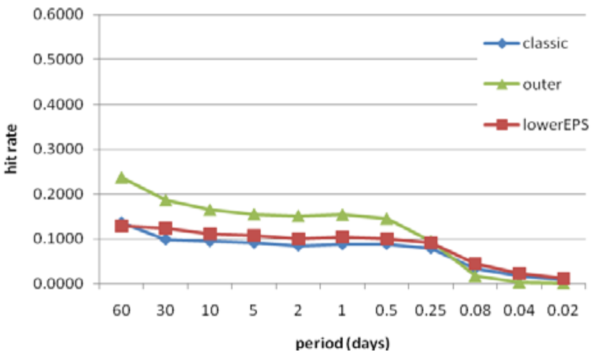

Fig. 16. The results for hit rate, Normal distribution altruism, UPB 2012 trace, 8 copies.

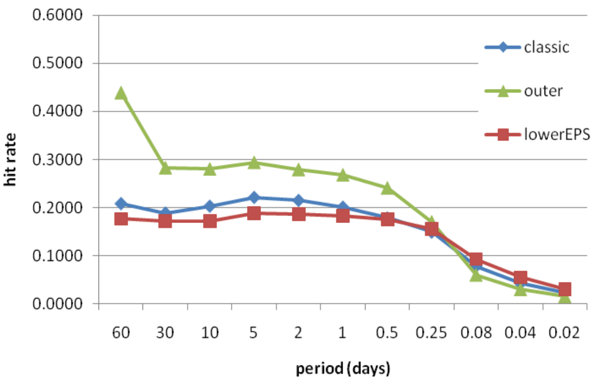

Fig. 17. The results for hit rate, Community-biased distribution altruism, UPB 2012 trace, 2 copies.

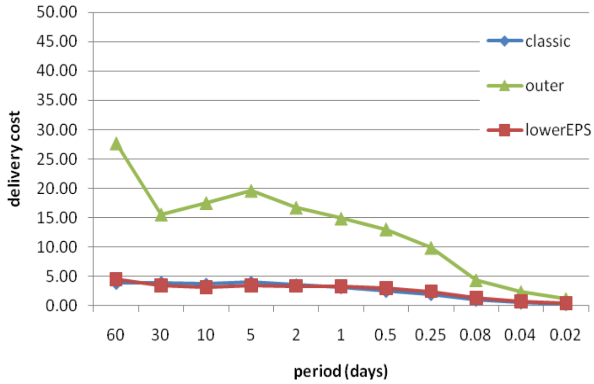

Fig. 18. The results for delivery costs, Community-biased distribution altruism, UPB 2012 trace, 2 copies.

The most noticeable modification is the downshift of the average hit rates to approximately 50% of their previous values. This decrease is also exhibited in delivery cost and latency and is explained by the direct proportionality with the hit rate. However, the general behavior remains the same for all metrics monitored.

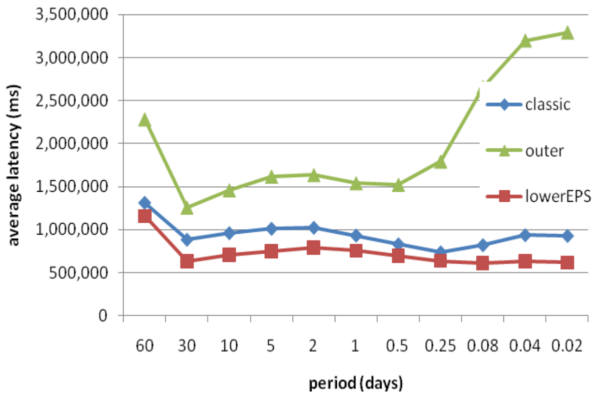

Fig. 19. The results for latency, Community-biased distribution altruism, UPB 2012 trace, 2 copies.

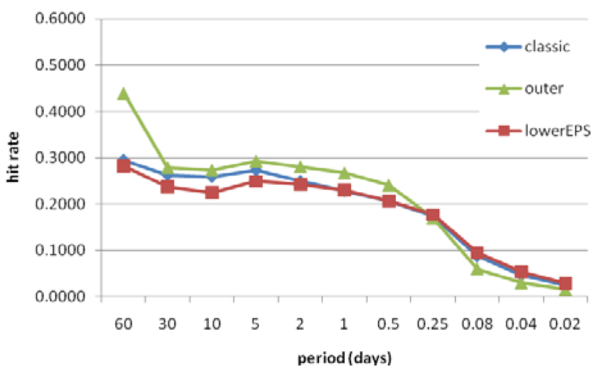

Fig. 20. The results for hit rate, Community-biased distribution altruism, UPB 2012 trace, 8 copies.

Increasing the number of copies shows no major changes in behavior as opposed to the scenario with no altruism. As can be seen in Figure 16 the only effect is of slowly shifting the hit rate up again, compensating for some of the negative effects of altruism. This compensation, although present, is insufficient and cannot compensate to a high enough degree.

Secondly, we analyze the impact of community-biased altruism on the UPB2012 trace (see Figures 17-19). The decrease in hit rate is not as dramatic as normal distribution altruism, but it is still significant. In terms of hit rate the *lowerEPS* algorithm is more affected. This could indicate that transfers with nodes having lower ranks (that are usually explored by *lowerEPS* and not by *classic*) tend to be with other communities.

An increased number of copies has the same effect as in the case of normal distribution altruism by shifting the hit rate back up, without changing previous behavior.

VI. CONCLUSIONS

In this paper we analyzed several techniques for opportunistic routing and the impact they have on key communication metrics. The observations made from this research work provide useful information on how social-based routing algorithms for opportunistic networks can be fine-tuned to better cope with limitations such as number of copies or altruism.

The main conclusions are as follows. When network saturation occurs, forcing messages that are not destined for the local community to be pushed faster out of the community can have a positive effect on the hit rate, although it does not greatly benefit from an increased number of copies. We have demonstrated that introducing even a small exploration factor can decrease the average latency for frequently exchanged messages while maintaining the same hit rate. When dealing with altruism a reasonable increase in the number of copies cannot sufficiently compensate for the negative effects exhibited on the hit rate.

This analysis was done in conditions of network saturation. In future it would be interesting to study the point at which an opportunistic network becomes saturated, depending on the number of nodes and explore the limits of social-based routing.

ACKNOWLEDGEMENT

This work was partially supported by project "ERRIC - Empowering Romanian Research on Intelligent Information Technologies/FP7-REGPOT-2010-1", ID: 264207. 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. This paper has benefited from the collaborative research efforts of the EU Green-Net group.

REFERENCES

[1] Kevin Fall, 'A delay-tolerant network architecture for challenged internets', in Proc. of the 2003 conference on Applications, technologies, architectures, and protocols for computer communications, SIGCOMM '03, pp. 27–34, New York, NY, USA, (2003). ACM.

[2] Hui, P., Crowcroft, J., Yoneki, E.: BUBBLE Rap: social-based forwarding in delay tolerant networks. In: Proceedings of the 9th ACM international symposium on Mobile ad hoc networking and computing. MobiHoc '08, New York, NY, USA, ACM (2008) 241–250.

[3] E. Jenefa JebaJothi, V. Kavitha, T. Kavitha, 'Contention Based Routing in Mobile Ad Hoc Networks with Multiple Copies', in Journal of Computing, Volume 2, Issue 5, May 2010.

[4] Kuang Xu, Pan Hui, Victor O. K.Li, Jon Crowcroft, Vito Latora, Pietro Lio, 'Impact of altruism on opportunistic communications', in Proceeding ICUFN'09 In Proc. of the first international conference on Ubiquitous and future networks, IEEE Press Piscataway, NJ, USA, 2009.

[5] Radu Ciobanu, Ciprian Dobre, 'Social Aspects to Support Opportunistic Networks in an Academic Environment', In Proc. of ADHOC-NOW 2012, Springer-Verlag Berlin Heidelberg 2012, LNCS 7363, pp. 69–82, 2012.

[6] Elizabeth M. Daly and Mads Haahr. Social network analysis for routing in disconnected delay-tolerant manets. In Proceedings of ACM MobiHoc, 2007.

[7] Musolesi, M., Mascolo, C.: Designing mobility models based on social network theory. SIGMOBILE Mob. Comput. Commun. Rev. 11 (July 2007) 59–70.

[8] Boldrini, C., Passarella, A.: HCMM: Modelling spatial and temporal properties of human mobility driven by users' social relationships. Comput. Commun. 33 (June 2010) 1056–1074.

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

[10] Heckerman, D., Mamdani, A., Wellman, M.P.: Real-world applications of Bayesian networks. Commun. ACM 38, 3 (March 1995), 24-26.

[11] Greifenberg, J., Kutscher, D.: Efficient publish/subscribe-based multicast for opportunistic networking with self-organized resource utilization. In Proc. of the 22nd International Conference on Advanced Information Networking and Applications - Workshops. AINAW '08, Washington, DC, USA, IEEE Computer Society (2008) 1708–1714.

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

[13] F. Guidec and Y. Maheo. Opportunistic content-based dissemination in disconnected mobile ad hoc networks. In Proc. of UBICOMM '07, 49-54.

[14] C. Boldrini, M. Conti, and A. Passarella, "Modelling data dissemination in opportunistic networks," Proceedings of ACM Workshop on Challenged Networks (CHANTS), 2008.

[15] Alljoyn official website, last accessed May 12, 2012, from http://www.alljoyn.org/

[16] Official site, last accessed May 10, 2012, from http://code.google.com/p/cmpt371t1/wiki/SoftwareEngineeringProcesses

[17] S. Okasha. Altruism, group selection and correlated interaction. British Journal for the Philosophy of Science, 56(4):703–725, December 2005.