

Impact of History on Epidemic Broadcast in DTNs

Giorgio Gamberini, Francesco Giudici, Elena Pagani, Gian Paolo Rossi
Information Science and Communication Department
Università degli Studi di Milano, Italy
e-mail: g.gamberini@gmail.com, {fgiudici,pagani,rossi}@dico.unimi.it

Abstract—In this work, we present an analysis of the gain achieved in terms of efficiency with PUSH-based epidemic diffusion of messages, when maintaining and exchanging information about the infection status. Different policies for the exchange of partial status information are compared, with the aim of optimizing bandwidth usage. We analyze the effects of adding a copy-count mechanism to the maintenance of status with respect to the capability of efficiently stopping the diffusion once full coverage is achieved; to this purpose, we adopted an analytical model proposed in the literature to determine the copy-count threshold. We perform initial measures to highlight how the mobility model affects the algorithm performance.

I. INTRODUCTION

The broadcast of messages over *delay tolerant networks* (DTNs) has several applications that include the diffusion of scoped advertisements – e.g. about available services or events – and summaries [12], podcasting [13], the upload of software patches or new parameter settings – e.g. in environmental observation systems – or the diffusion of acknowledgements, or *cure*, packets [11]. The presence of intermittent node connectivity, node sparsity and topology dynamics imposes to move apart from flooding to adopt encounter-based epidemic approaches that maintain the required topology-independent and stateless nature but better fit with the sparse and intermittent nature of a DTN. An effective, best effort, broadcast can be designed by starting from the following basic PUSH-based algorithm [16], [8], which we indicate as P-BCAST:

P-BCAST – *when a node has in its cache a message m to diffuse, it forwards m to one or more (and possibly all) encountered neighbors. The forwarding, elsewhere called infection or epidemic, can be either performed periodically [14] or whenever the contact occurs [16].*

In this work, the encounter-driven mechanism is considered.

This algorithm guarantees node coverage approximately close to 1 with low latency but generates a redundant load of duplicates. This is mainly motivated by the fact that nodes perform epidemic forwarding with a very limited knowledge about the state of the encountered nodes and, as a consequence, they often happen to forward the message to already infected nodes. The primary focus in the design of a broadcast protocol for DTNs is to reduce the duplicates while keeping high the node coverage. To achieve this goal, the broadcast design must respond to the following questions: (i) how to provide a node with information about the state of its neighbors to enable the

forwarding control? (ii) How to ensure the termination of the forwarding algorithm when the message has been delivered to the entire population of nodes? A viable approach to simply answer to the second point is to stop the forwarding when a message lifetime or a given hop/copy count has been reached. The choice of this bound influences the capability of achieving high coverage and can be obtained from the analysis of the mobility traces of the target environment [7] or, under restrictive assumptions, derived analytically [8]. By contrast, the forwarding control, point (i) above, can properly be achieved by allowing a node to maintain a local history of encounters and to exchange it with the other encounters. A history mechanism has been recently adopted either to compute utility functions to control unicast forwarding [2], [3], [10], or as the log of infection events to control broadcast forwarding. Several aspects have yet to be explored. There is the need to define the amount of history a node should maintain and/or exchange, to evaluate the real performance advantages it guarantees, to verify its validity under mobility models other than the random waypoint (RWP) [4], to design mobility-independent *stop conditions* and to explore the scalability issues the history involves. The main contribution of this paper is to consider part of the above arguments and to provide initial hints for the design of a history-based epidemic broadcast in DTNs.


II. PROBLEM DESCRIPTION AND ASSUMPTIONS

Let us consider a DTN of n independent and mobile nodes that are supposed to be sparsely distributed over a large area and to use a short radio range to communicate. Throughout this paper, we assume that each node knows the nodes' cardinality, n , and we leave to another paper the more general problem of unknown cardinality. Moreover, a node is aware of the subset of nodes, the *encounters*, in its radio range. Encounters are identified through *beacons*, which are periodically broadcast at the underlying layer [5]. In this scenario, we can easily obtain a basic history-based epidemic broadcast, denoted as HP-BCAST, by adding a history of contacts to P-BCAST:

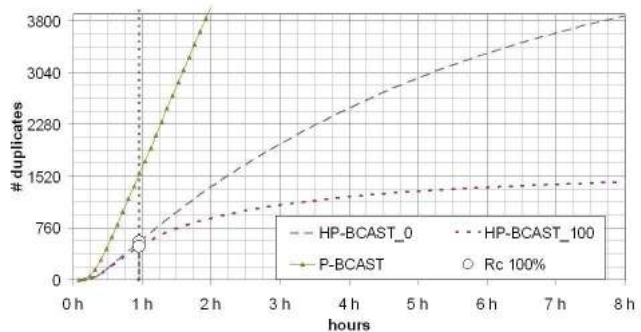
HP-BCAST – *whenever a node p forwards or receives the message m , it adds to a local data structure, $history_p$, the list of its current encounters or m 's sender respectively. Initially, $history_p$ only contains p itself. p broadcasts a message m if its current neighborhood has nodes other than those held in $history_p$. It skips forwarding otherwise. p terminates the epidemic algorithm when $history_p$ contains n entries.*

In the basic HP-BCAST (indicated as HP-BCAST₀), each node

This work has been partially funded by the Italian Ministry of University and Research in the framework of the "Context-Aware Routing Over Opportunistic Networks (CARTOON)" PRIN Project.



(a)



(b)

Fig. 1. (a) Coverage, and (b) cumulative number of duplicates in the RWP model

maintains a local history. The simple evolution of this algorithm is obtained by enabling a node to exchange its local history with the encounters. This policy accelerates the node's awareness about the epidemic diffusion and should provide a more effective forwarding control. We denote this new algorithm as HP-BCAST₁₀₀. It can be obtained from HP-BCAST₀ by adding the following statement: *when p forwards m then it piggybacks (the 100% of) its history _{p} on m . When p receives m from q , it merges history _{p} with history _{q} .*

It is worth observing that the described HP-BCAST₁₀₀ algorithm is not as aggressive as it could be in its attempt of propagating the infection knowledge. In fact, nodes could achieve a much quicker awareness about the evolution of the system infection by piggybacking the history on top of the beacon messages. In sec.III, we will consider all these issues and compare the above algorithms to identify their contribution to achieve the control on forwarding. We mainly focus on the protocol's capability of maximizing node *coverage* (i.e. the number of nodes that deliver m), while minimizing both latency, bandwidth and energy consumption. In sec.IV, we discuss the issue of obtaining a *stop condition* and compare different approaches. In sec.V, we consider different policies to exchange the history and provide preliminary scalability considerations.

A. Simulation Environment

The protocol evaluation is performed by using GloMoSim [15]. The simulation setting considers a system of 100 nodes sparsely distributed over a 1000×1000 mt. area. Nodes move at a speed in $[1, 2]$ m/s, thus reproducing a pedestrian environment. They are equipped with a low power 802.11 radio device with 10 m. communication range and DCF at the MAC layer. Beacons are performed every 1 sec. We consider long lived broadcasts, with simulations lasting up to 6 hours. Results are averaged over 50 simulations performed with variable random seed.

The protocol behavior is highly influenced by the underlying mobility models. The RWP model is largely used in the literature because it is simple and is provided by most network simulators. However, several experiences of collecting mobility traces from workplaces or campus area have shown that the

distribution underlying the human mobility are heavy-tailed distributions with the Pareto law as a common shape [6]. A more recent research [1] has observed that the mobility traces could be more accurately described through a distribution law of inter-contact times following the shape of an exponential variable followed by a Pareto variable. This is coherent with the combined coexistence of a base Brownian motion, which leads to unplanned and short contact opportunities, with intentional mobility patterns that, at a given time, trigger oriented trajectories of one or more nodes towards common points in the space, or aggregation points (AP). With the aim of analyzing the protocol behavior in realistic conditions, we performed simulations under both RWM and AP mobility conditions (sec.IV).

III. EVALUATION OF PUSH- AND HISTORY-BASED ALGORITHMS

We measured the performances of the considered algorithms in the described scenario. In fig.1, we report the node coverage (the white circle (Rc 100%) evidences the point where 100% coverage is reached in the worst case), and the cost paid in terms of redundant, or *duplicate*, messages. The coverage slope over time also shows how quickly an algorithm approximates the full coverage and, as a consequence, the latency. P-BCAST is the best performing algorithm in terms of both coverage and latency but it generates unacceptable redundant traffic. By contrast, it is clear that the use of history provides a twofold advantage: on one hand, it does not affect coverage and latency and, on the other hand, it provides an effective mechanism to control the forwarding. This is obtained by simply providing a node with some approximation of the global system state and, of course, the higher is the global state awareness (HP-BCAST₁₀₀), the lower the amount of redundant messages. It is worth noticing that P-BCAST has no information about the state of its encounters and this leads to a very poor forwarding control (fig.1(b)).

The positive effects of the history's use grows along the algorithm run. In fact, in HP-BCAST₁₀₀ the number of duplicates each node generates in the second hour is around 4; in the third hour it drops to 2. However, the amount of duplicates is not a fully expressive measure of the history ability of identifying a

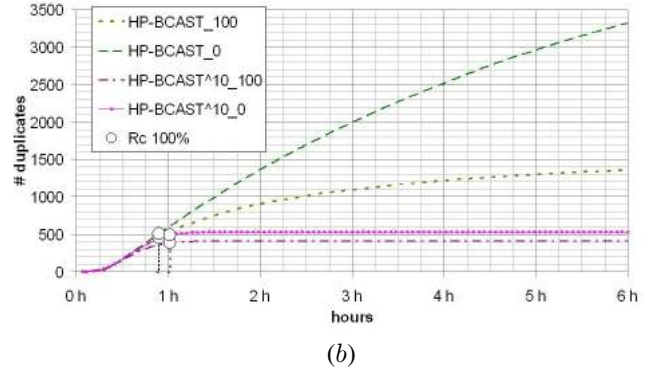
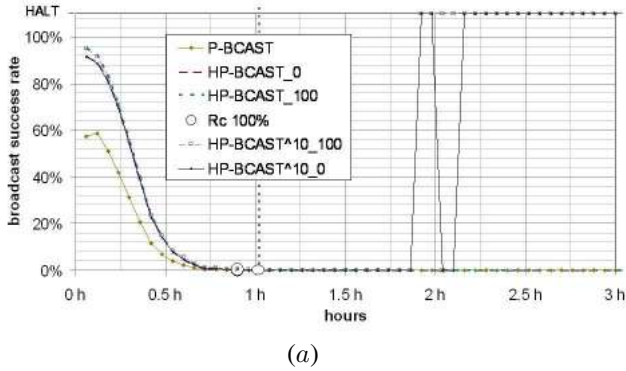


Fig. 2. (a) Broadcast success rate, and (b) cumulative number of duplicates in the RWP model

non infected node in the neighborhood. In fact, it is affected by the broadcast nature of the radio channel where, to correctly deliver m to the uninfected node q , m may be duplicated in a node p which happens to be in range. A good forwarding control should identify the presence of uninfected encounters and refrain from forwarding otherwise. The *broadcast success rate* bsr more properly captures this ability; it is defined as the ratio *hitting broadcasts/total broadcasts*, where the “hitting broadcast” is a broadcast which delivered m to at least one uninfected node. Of course, $bsr = 1$ indicates that all the broadcasts hit the mark. The index drops to 0 when the full coverage is reached (after nearly 1 hour in our experiments). The behavior of bsr (fig.2(a)) confirms the remark above. Both the history-based algorithms show identical values for both the bsr and the amount of duplicates (fig.1(b)). This indicates that the knowledge provided by the history about the system is useless when the coverage is less than 100% because nodes are more likely to encounter new nodes than nodes registered in their history. History comes into play when the full coverage is almost reached, and $bsr = 0$. From this point on, the history helps nodes to control forwarding and, of course, the larger the knowledge, the more effective the control. However, the node’s knowledge does not grow as quickly as the nodes infection and this influences the efficiency of the forwarding control. In HP-BCAST₀ the history slowly grows at the rate of contacts and eventually tends to a full history condition if the contact opportunity is equally distributed between each pair of nodes (i.e. in a RWP). HP-BCAST₁₀₀ has a quicker learning process because nodes are allowed to exchange their local system view when a message is broadcast. The fastest mechanism to propagate the system view would be enabling the nodes to share their local history through beacons, but this would be paid in terms of a larger, and probably unaffordable, resource consumption.

The above arguments lead to the following considerations. Whenever it is not possible to adopt a policy that ensures the algorithm termination after the full coverage has been reached (see, next Section), the history becomes an efficient and effective solution to control forwarding in the region with $bsr = 0$. By contrast, when a termination condition is

achievable then the history has a marginal role and a simple local history mechanism can be adopted.

IV. STOP CONDITION

The previous arguments motivate the need of identifying a condition that enables the nodes to stop the algorithm run. This condition can be either derived from the observation of the neighborhood or by exploiting the underlying mobility model. The most intuitive stop condition can be associated to full history condition. Unfortunately, even under the favourable condition of well-known node cardinality and of equally distributed contact opportunities, the full history condition is eventually achieved too late with respect to the time the full coverage condition actually occurs. Our simulations show that the full coverage condition is achieved after 1 hour of simulated time. At this time, each node on average knows the 15% of the infected nodes with HP-BCAST₀ and the 65% with HP-BCAST₁₀₀ (fig.4(b)). Moreover, with HP-BCAST₀ only 2 packets/hour overall are sent after 24 hours, and all nodes stop within 47 hours. With HP-BCAST₁₀₀, only 1 packet/hour is sent after 24 hours, and all nodes stop within 41 hours. The elapsed time between the real and the approximated event has impact on the efficiency of any broadcast protocol in DTN because it represents a large time interval in which nodes keep on forwarding few but useless messages. This means that the full history condition is not a representative condition of the global system state and a different stop condition must be searched that ensures an earlier algorithm termination.

The approach proposed in [8] moves in this direction. Under the condition of a well known cardinality and RWP, the authors derive that each node can stop the forwarding when a copy count of $\tau = 2 \lceil \ln n + \gamma \rceil$ broadcasts is reached, with γ the Euler-Mascheroni’s number. This approach has been enhanced in [9], by adding a history either locally maintained by nodes or completely shared; let us indicate with HP-BCAST _{α} ^{τ} this approach, with $\alpha = 0$ in the former case and $\alpha = 100$ in the latter. The resulting algorithm becomes:

HP-BCAST _{α} ^{τ} – when a node p encounters a node q , it increments a local copy count; if q is already in its history, p skips forwarding.

In our simulation conditions, $\tau = 10$, while the algorithms




Fig. 3. (a) Coverage and (b) cumulative number of duplicates in the aggregation model

studied so far correspond to having $\tau = \infty$. Both HP-BCAST_0^{10} and $\text{HP-BCAST}_{100}^{10}$ achieve full coverage, although with a slightly higher latency than in the case $\tau = \infty$ due to nodes that exhaust their available copy count before all nodes have been infected, thus decreasing the number of relays and slowing down the diffusion. Yet, thanks to the bound on the number of diffusions, the HP-BCAST^τ approach is optimal in RWP and is thus much more effective than HP-BCAST^∞ in limiting useless traffic (fig.2(b)): full coverage is reached after 3672 sec., and message diffusion stops after 6696 sec. with $\text{HP-BCAST}_{100}^{10}$ and after 7560 sec. with HP-BCAST_0^{10} . The *bsr* index before full coverage is the same whether τ is infinite or not, confirming the effectiveness of the history in suppressing useless transmissions (fig.2(a)). However, with $\tau = \infty$ nodes continue performing transmissions till their histories are full; with bounded τ no useless transmission is anymore generated after all nodes halt. In the latter case, though, with HP-BCAST_0^{10} , the last infected nodes could issue sporadic transmissions before halting.

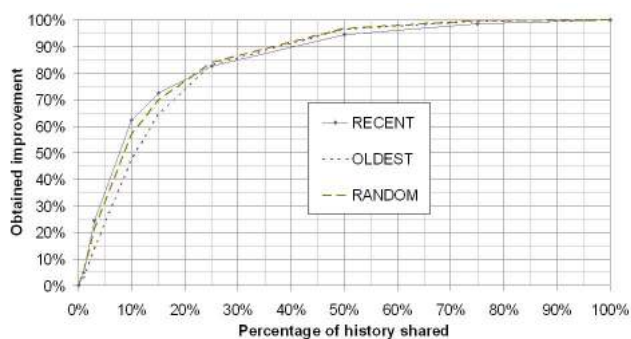
The weakness of the τ bound is its dependence on the uniform distribution of contacts as in RWP. We simulated $\text{HP-BCAST}_\alpha^\tau$ in the aggregation model described in sec.II-A, with 10 aggregation areas of 30 m. radius each, and nodes that remain in an area for 10 minutes. In this mobility scenario, HP-BCAST^τ suffers multiple contacts between the same pair of nodes (fig.3(a)); this does not lead to duplicate generation but forces the nodes to waste their broadcast “budget” every time they re-encounter a node seen in the past. As a consequence, the algorithm is too conservative and full coverage is not achieved: in the conditions shown in figure, all nodes terminate the algorithm when, on average, the 2% of them is still uninfected; the minimum coverage observed is 94%. By exasperating the non-uniformity of contacts till considering a swarm mobility, the coverage drops to 17% in the worst case. By contrast, under the same conditions the history-based algorithms guarantee full coverage independently of the mobility model although this result is paid with a higher number of duplicates (fig.3(b)).

The above arguments lead to conclude that, under the favourable condition of well-known cardinality, a copy count

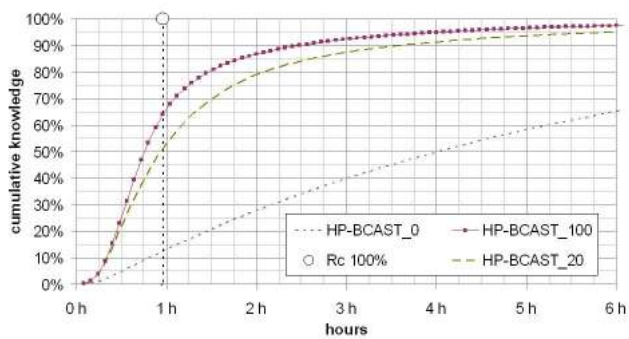
bound can be profitably adopted to stop the forwarding. This choice ensures to achieve good performances with RWP mobility model and acceptable performances with AP. In this case, the combined use of history slightly influences the performance and may be optimized with the use of history sharing mechanisms, as described in the next Section. When nodes are unaware of the system node cardinality, the above mechanisms fail and the real challenge becomes the design of an autonomic mechanism that helps nodes to control forwarding by learning the system state through the neighborhood observation.

V. POLICIES FOR KNOWLEDGE SHARING

Since DTNs must exploit contact opportunities – which could also be sporadic – in order to forward messages, the bandwidth optimization is a fundamental goal. To this purpose, a trade-off should be defined between the amount of history exchanged and the bandwidth saved in suppressing useless diffusions. We studied three policies to extract entries from the local history to be piggybacked onto the message: the node maintains its history ordered so as to record from the most to the least recent encounters. Entries can be extracted (i) in *random* order, (ii) *most recent first*, or (iii) *least recent first*. The first policy is used as benchmark. The second policy aims at notifying an encounter about the infection status of nodes in the surroundings. The third policy aims at notifying an encounter about the infection status of nodes in a region it could be directed to. In fig.4(a), the three policies are compared for different percentage of history shared. All policies show the same behavior: by indicating with 100% the improvement – in terms of generated duplicates – achieved with HP-BCAST_{100} , all policies obtain an improvement around 80 – 85% with the exchange of just 20 – 25% of the history. Moreover, for a small portion of history shared, the best behavior is achieved with the *most recent first* policy. In fact, by comparing the amount of knowledge about the system, and by indicating with 100% the full history condition, the knowledge growth – by sharing 20% of the history (HP-BCAST_{20}) – well approximates the behavior achievable with the full sharing (fig.4(b)); this effect is obviously reproduced by the number of generated duplicates (fig.5).



(a)



(b)

Fig. 4. (a) Comparison among sharing policies. (b) Cumulative knowledge about the coverage status

VI. CONCLUDING REMARKS

In this work, we compare by simulations different PUSH-based epidemic diffusion algorithms, which may use a history of encounters and a copy-count bound, analytically determined, in order to suppress useless transmissions. The results bring into evidence that the model proposed in [9] must be extended in order to accurately estimate an upper bound on the diffusions a node can perform, adequate to settings where node mobility is more similar to that observed in real environments. The history is effective in suppressing duplicates, but does not allow to promptly detect when full coverage has been reached.

In this work, we assume that nodes know the number n of nodes in the system and have enough memory to record a history of size n . In case of unknown or varying membership, nodes have the problem of appropriately dimensioning the history. Moreover, the analytical model becomes useless because it is unable to determine τ . Currently, we are studying the trade-off between the maintenance of a partial history – applicable e.g. when nodes do not know the system cardinality – and the performance improvement achieved in comparison with P-BCAST. We are studying different policies to determine what encounters to remember in the history, and we are also evaluating the possibility of finding some heuristics to implement a stop condition given a partial history. These policies will be compared with the benchmark obtained when nodes collect knowledge about the system by exchanging the history with every beacon.

REFERENCES

- [1] Apolloni B., Bassis S., Rossi G.P., Pagani E., *A Mobility Model fitting Experimental Traces in Opportunistic Scenarios*. Technical Report, Università degli Studi di Milano, 2008.
- [2] Boldrini C., Conti M., Iacopini L., Passarella A., *HiBOP: a History Based Routing Protocol for Opportunistic Networks*. Proc. IEEE WoWMoM 2007.
- [3] Burns B., Brock O., Levine B. N., *MV Routing and Capacity Building in Disruption Tolerant Networks*. Proc. IEEE Infocom 2005.
- [4] Camp T., Boleng J., Davies V., *A Survey of Mobility Models for Ad Hoc Network Research*. Wireless Communication and Mobile Computing (WCWC): Special issue on Mobile Ad Hoc Networking: Research, Trends and Applications, 2(5), pp. 483–502, 2002.
- [5] Cerf V., Burleigh S., Hooke A., Torgerson L., Durst R., Scott K., Fall K., Weiss H., *Delay-Tolerant Networking Architecture*. RFC 4838, April 2007. Work in progress.

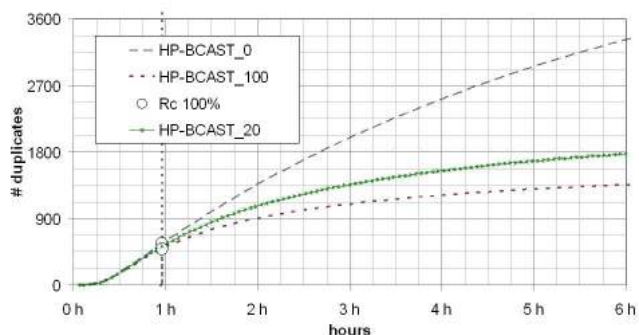


Fig. 5. Cumulative number of duplicates

- [6] Chaintreau A., Hui P., Crowcroft J., Diot C., Gass R., Scott J., *Impact of Human Mobility on the Design of Opportunistic Forwarding Algorithms*. Proc. INFOCOM 2006.
- [7] Chaintreau A., Mtibaa A., Massoulié L., Diot C., *The diameter of opportunistic mobile networks*. Proc. 2007 ACM CoNEXT Conference, pp. 1–12.
- [8] Cooper D.E., Ezhilchelvan P., Mitrani I., *High Coverage Broadcasting for Mobile Ad-hoc Networks*. Proc. 3rd Intl. IFIP-TC6 Networking Conf., Lecture Notes in Computer Science (Springer), Vol. 3042, pp. 100–111, 2004.
- [9] Cooper D.E., Ezhilchelvan P., Mitrani, I., Vollset E., *Optimization of Encounter Gossip Propagation in Mobile Ad-Hoc Networks*. Proc. 13th IEEE Intl. Symp. on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS 2005), pp. 529–532, 2005.
- [10] Dubois-Ferriere H., Grossglauser M., Vetterli M., *Age Matters: Efficient Route Discovery in Mobile Ad Hoc Networks Using Encounter Ages*. Proc. 4th ACM Intl. Symp. on Mobile ad hoc networking and computing (MobiHoc 2003), pp. 257–266.
- [11] Harras K.A., Almeroth K.C., Belding-Royer E.M., *Delay Tolerant Mobile Networks (DTMNs): Controlled Flooding in Sparse Mobile Networks*. Proc. IFIP Networking 2005.
- [12] Lee U., Magistretti E., Zhou B., Gerla M., Bellavista P., Corradi A., *MobEyes: Smart Mobs for Urban Monitoring with a Vehicular Sensor Network*. IEEE Wireless Communications, 13(5) (2006).
- [13] Lenders V., Karlsson G., May M., *Wireless Ad Hoc Podcasting*. Proc. 4th IEEE Conf. SECON, pp. 273–283 (2007).
- [14] Montresor A., Jelasity M., Babaoglu O., *Gossip-based Aggregation in Large Dynamic Networks*. ACM Transactions on Computer Systems, 23(3), Aug. 2005, pp.219–252.
- [15] UCLA Parallel Computing Laboratory, *GloMoSim – Global Mobile Information Systems Simulation Library*. University of California at Los Angeles. <http://pcl.cs.ucla.edu/projects/glomosim/>
- [16] Vahdat A., Becker D., *Epidemic Routing for Partially Connected Ad Hoc Networks*. Technical Report CS-200006, Duke University, April 2000.