

Impact of Mobility on Epidemic Broadcast in DTNs

Francesco Giudici, Elena Pagani, Gian Paolo Rossi

Information Science and Communication Department, Università degli Studi di Milano, Italy,
e-mail: {fjudici, pagani, rossi}@dico.unimi.it

Abstract The broadcast diffusion of messages in Delay Tolerant Networks (DTNs) is heavily dependent on the mobility of the nodes, since protocols must rely on contact opportunities among devices to diffuse data. This work is the first effort of studying how the dynamics of nodes affect both the effectiveness of the broadcast protocols in diffusing the data, and their efficiency in using the network resources. The paper describes three simple self-adaptive control mechanisms that keep the broadcast overhead low, while ensuring high node coverage. Those mechanisms characterize a family of protocols able to achieve some awareness about the surrounding environment, and to use this knowledge in order to improve performances. Simulation results allow to identify the winning mechanisms to diffuse messages in DTNs under different conditions.

1 Introduction

The mobile nodes of a delay and disruption tolerant network, or DTN [6], experiment intermittent connectivity, network partitions and poor radio link quality. So far, in such a critical scenario, the research mainly focused on the problem of providing unicast communications (e.g., [5, 12, 13, 1, 18]). By contrast, the one-to-all communication scheme has not received the same attention despite the fact that its service is strategic to support protocols at both application and routing levels. For instance, a broadcast service is required to diffuse scoped advertisements – e.g. about available services or events – and summaries [14], to support podcasting [15], to upload software patches or new parameter settings – e.g. in environmental observation systems – or to diffuse acknowledgements, or cure, packets [9]. In a DTN, broadcast can be designed by adopting one of the gossip-based mechanisms that have been proposed in the literature in a few slightly different alternatives by starting from the following basic scheme: when a node has in its cache a message m to diffuse, it forwards m to one or more (and possibly all) neighbors it happens to encounter while moving. The forwarding, elsewhere called infection or epidemic, can be either performed periodically [17] or whenever the contact occurs [21]. Infection can continue up to the message life time or up to a given hop/copy count. It is easy to show [8] that this simple PUSH-based algorithm provides an effective broadcast service, by achieving a node coverage arbitrarily close to 1 with a low latency, but fails in doing this ef-

ficiently. The inefficient use of the network resources seriously limits the practical use of the protocol. 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. There are several, growing levels of knowledge a node can achieve about the neighbors state and to approximate the global system state. They range from zero knowledge, as for the sketched stateless PUSH-based algorithm [21] and for the SA-BCAST described in [8], to full knowledge, which can be approximated by maintaining some log of encounters and by enabling the log exchange among encountering nodes, [16].

The primary focus in the design of a deployable, i.e. both efficient and effective, broadcast protocol is to increase the node likelihood of delivering the message only to uninfected nodes. However, whatever is the followed approach to achieve efficiency, the performances of the algorithms are greatly influenced by the mobility patterns that nodes follow, [2], and a mechanism properly working when nodes move according to the random waypoint model may be totally useless when they move according to a different mobility model. The comparative analysis of the effects of mobility on DTN protocols deserves more attention and, as far as we know, has never been applied to broadcast. This paper moves into this research track and provides some interesting contributions to understand broadcast delivery over DTNs under different mobility conditions. Firstly, the paper defines a family of broadcast protocols obtained by starting from the above mentioned PUSH-based forwarding protocol and by leveraging its behaviors through the incremental introduction of autonomic and adaptive mechanisms whose purpose is the improvement of the node awareness about the level of infection in the neighborhood. Secondly, the performances of the different mechanisms are analyzed in 3 different mobility models: the classical random way point, RWP, mobility model [2], the swarm mobility model [10], SWR, and the aggregation model, AGG, in which nodes move throughout aggregation points according to some spatial or functional mobility law. The main contribution of the paper is twofold: (i) it identifies and characterizes the winning mechanisms to diffuse messages in DTNs under different conditions; (ii) it is the first attempt to move toward the design of an autonomic and situational algorithm able of autonomously deriving the mobility context the node is moving through, and adapting its parameters accordingly in order to optimize performances.

2 System and Mobility Models

2.1 System Model

The scenario we consider in this paper includes people walking in a limited urban area, such as a *campus area*, and equipped with wireless portable devices. No base stations are assumed and the communication between a source s and a destination

d may eventually occur through either direct contact, when, for instance, node d moves into the range of s , or indirect contact, when one or more relaying nodes help to create the multi-hop path towards the destination and the last of them finally enters the range of d . The devices have a unique identifier ID, are not required to have positioning capabilities on board and, to meet resource saving requirements, are supposed to adopt a short radio range to communicate. This latter point, together with the fact that devices can be sparsely distributed over a large area, makes high the probability of network partitions and link disruption. Throughout the paper we only assume that each mobile device, or node, periodically broadcasts a *beacon* message in its radio cell. Beacons are used to discover other devices in the neighborhood and their content is limited to the device identifier. In such a scenario, people mobility might follow either a Random Waypoint model [2] or a more structured motion, as described in the next subsection.

2.2 Mobility Models

A great deal of research is currently ongoing in order to characterize mobility models suitable for opportunistic networks. Mobility could be extracted from traces of movements of nodes in real settings; several traces are for instance provided by the CRAWDAD community [4]. The use of traces with simulators creates, however, some problems. Their timescale is hardly scaled to the simulation time and they generally model the specific behavior of a given mobility scenario thus losing the generality required during the protocol design phase. Traces are more likely useful during the validation process than during the design and performance analysis phase.

In this work, we analyze how three basic movement models affect the performance of the epidemic protocols described in sec.3. The models are: Random Waypoint (RWP), aggregation (AGG) and swarm (SWR). The classical RWP model is not realistic, but it is simple, often provided within network simulators, and the most commonly used in the literature. In the AGG model, a node moves toward an aggregation point (*ap*), chosen from a set according to a certain probability distribution P_{AGG} , and once there it pauses for a time t_{AGG} before selecting the next *ap*. This model reproduces mobility of users who may group in interest points according to some spatial or functional rule. Parameters of this model are also the number and position of *aps* over the area, the speed range $[v_{min}, v_{max}]$ and the radius of the *aps*. The arrival point of a node inside an *ap* can be determined according to a Gaussian distribution centered at the center of the *ap* and with standard deviation σ_{AGG} . In the SWR model, nodes move in a coordinated way. Each swarm has a logical center, which moves toward a destination chosen randomly. Once there, the nodes in the swarm stop for a pause time t_{SWR} before moving to a new destination. Each swarm has a number of nodes determined by a probability distribution with mean μ_{SWR} and standard deviation σ_{SWR} . Nodes in a swarm move randomly around its logical center, within maximum distance d_{SWR} . If the regions of two swarms over-

lap, nodes in the intersection may choose to migrate to the other swarm according to a probability M_{SWR} .

The RWP model reproduces sporadic encounters of two (or a few) nodes. The AGG model reproduces the encounters of many nodes, with nodes experiencing relevant neighborhood changes every time they enter an *ap*. The SWR model reproduces nodes maintaining the same neighborhood for a long time, with sporadic encounters with other (groups of) nodes. Analyzing these three basic models allows to bring into evidence the effects each of them separately has on message diffusion. However, mobility of people in an opportunistic scenario is more likely modeled by the combination of the above three patterns. At the best of our knowledge, the research community has not yet produced a synthetic model that addresses these issues or adopts statistical distributions inferred from real traces [4].

Whatever is the adopted mobility model, it should be assumed that the following *mobility assumption* applies: when a contact occurs, the reciprocal speed is such that the two nodes can set up a communication channel and a significant amount of data is exchanged before they become disconnected. This assumption is reasonable according to results achieved by observations reported in [19, 11], but can be occasionally violated in our simulations.

3 A Family of Broadcast Protocols for DTNs

Given a general DTN scenario as described in the previous section, purpose of this paper is the comparative analysis of a family of topology-independent broadcast protocols under different mobility conditions. Protocols are obtained by the incremental introduction of adaptive and situational mechanisms that progressively augment the protocol capability to adapt to changing conditions. Purpose of the paper is the evaluation of the role played by the different mechanisms and parameter settings to ensure broadcast *effectiveness*, i.e. the capability of the protocol to eventually achieve node coverage arbitrarily close to 1, and *efficiency*, i.e. the capability of the protocol to keep the generated broadcasts-per-message as close as possible to $O(n \ln n)$ [3]. Each broadcast message m is supposed to have a “scope” that is defined by the source and specified through a *lifetime*; when this time expires, a node deletes the copy of m and stops its diffusion. In the following, we assume long lived messages to better understand the broadcast behavior independently of other constraints.

The basic broadcast protocol to start with is represented by the above mentioned PUSH-based algorithm: a node p , holding a message m , starts a forwarding of m with probability $Prob_p = 1$ whenever a node enters its radio range. We indicate this protocol with the name P-BCAST. It is possible to improve P-BCAST behaviors by adding some autonomic capability that extracts an approximation of global state from locally observed data.

The first improvement of P-BCAST is to let p start the forwarding of m when the percentage of neighborhood change with respect to the previous diffusion exceeds

a threshold Nth . To this purpose, each node maintains a local view of the neighborhood changes by exploiting the underlying beaoning. To avoid broadcast storms, the nodes in range adopt a random mechanism that desynchronizes transmissions and suppresses duplicates. This way, p , and the nodes in range of p , control the duplicate generation and maximize the probability of infecting new nodes by adopting a membership-driven infection. This *reactive* mechanism is supposed not to be influential when moving according to RWP in sparse conditions, which is very close to model a single encounter scenario; by contrast, it should positively affect performances when moving according to AGG.

A second, *adaptive* mechanism can be adopted to let a node p to be able to tune the value of $Prob_p$ according to the delivery status of m in the area where p is moving. In fact, when most nodes in the neighborhood have delivered m , p should reduce the probability $Prob_p$ of a new diffusion accordingly, and vice versa. In a zero knowledge paradigm, the node p is unable to achieve this type of knowledge; however, p can derive the symptoms of the delivery status from sensing the events generated by its encounters. The number of received duplicates of m is the first of these symptoms; it reveals that other encounters have delivered m and is helpful to decrease, according to a given function \mathcal{F} , the local value of $Prob_p$. In [8], it has been shown that this simple mechanism works properly to limit the number of duplicates under high coverage conditions and, with the addition of some further control mechanism, it is also able to promptly increase $Prob_p$ when p moves from a covered area to another where m has not been widely delivered. Both the *reactive* and the *adaptive* mechanisms are part of a new protocol that we will call self-adaptive broadcast, or SA-BCAST.

In [8], the authors have observed that SA-BCAST is unable to terminate before the message lifetime expires. This is intuitively motivated by observing that the algorithm has to ensure a drip feed of message transmissions to manage node joins and temporary partitions. A non-terminating protocol has severe effects on the efficiency that can be only mitigated by adding more knowledge to the nodes about their neighborhood. The introduction of an encounters' history can provide the extra information required and moves the protocol from a zero-knowledge paradigm to a full knowledge one.

To verify the benefits of a history-based algorithm, we implemented the following simple mechanism: each node maintains a *history* of the nodes to which it forwarded m and from which it received a duplicate. In order to speed up the learning process, we allow the *history* exchange among encounters by piggybacking it on m . If a node p encounters a node it has in the local history, then p suppresses the message forwarding. The history mechanism can be added to both the above mentioned protocols to obtain HP-BCAST and HSA-BCAST.

The described mechanisms lead to a family of broadcast protocols whose behavior has been mainly designed to properly work under both RWP and AGG mobility models. By contrast, the SWR mobility generates some new issues that deserve specific attention. In fact, with this setting, we will show in the next section that the 100% coverage is hard to reach, while a high number of duplicates might be generated. In order to achieve high effectiveness while maintaining an acceptable effi-

ciency, a mechanism of *local duplicate suppression* has been added to HP-BCAST (thus obtaining LP-BCAST) and HSA-BCAST (LSA-BCAST). When a beacon from a new neighbor is received, a node p sets a timer slightly larger than the beacon period. When the timer expires, the message is sent with the history, *including all the nodes that are new neighbors*. If another infected node with the timer set receives this message, it schedules its own transmission only if it has new neighbors not included in the history. This way, diffusion is not prevented for nodes that connect two swarms or two different groups in a swarm, but multiple infections of the same nodes are avoided. The local suppression mechanism leads to obtain the complete family of protocols shown in fig.1.

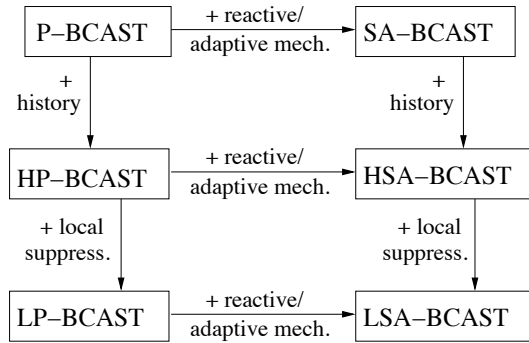


Fig. 1 Family of adaptive protocols

4 Performance Evaluation

4.1 Simulated Conditions and Performance Indexes

We implemented the described protocols in the framework of the GloMoSim [20] simulation environment. The simulation setting considers a system of 50 nodes sparsely distributed over a 1000×1000 m. 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. Beaconing is performed every 1 sec.; after 3 missing beacons, the corresponding neighbor is removed from the neighbor list.

The simulations run different values of N_{th} and two different functions \mathcal{F} : a linearly decreasing function (or **Lin10**) and an inverse exponential function (or **InvExp**). When an infected node p receives a duplicate from a node that is infected from less than 3 min., then p sets $Prob_p$ to $MAXP=1$. Otherwise, $Prob_p$ is either decremented of 0.1 or halved in line with functions **Lin10** or **InvExp**, respectively. $Prob_p$ has a lower bound defined by $MINP=0.01$. We consider long lived broadcasts, with simulations lasting up to 6 hours. All simulation results are averaged

over 50 simulations performed with variable random seed. The mobility models are provided by BonnMotion [7]; in order to allow movements to reach a steady state, the first 1000 sec. of the traces are not considered for the measures. The parameters of the different mobility models are presented in the next subsections.

Performance indexes are the *coverage*, i.e. the percentage of nodes infected, the *duplicate messages* (a message is a duplicate if it is received by an already infected node), and the *target ratio* $T = (msgrecv - dups) / msgrecv$, with *msgrecv* the total number of messages received and *dups* the total number of duplicates among them. T is a measure of efficiency in using the network resources. Of course, T is optimized by $dups = 0$ and is affected by the number of the encounter nodes and by the progress of the infection in the neighborhood. In fact, packets are broadcast to the nodes in range, let us say k ; so that, for any message sent, we count $msgrec = k$ and the *dups* value depends on the level of infection among the k nodes.

4.2 Random Waypoint vs. Aggregation Model

In the measurements presented in this section, the pause time for the RWP model is 0. With AGG, we performed experiments with a number of *aps* variable from 3 to 10, uniformly distributed in the area. A node stops in an *ap* for 10 minutes. The next *ap* is chosen according to a uniform probability distribution. The distance of a node from the *ap* center is determined by σ_{AGG} in $[0, 15]$, and follows the probability distribution reported in Table 1.

Table 1 Distribution of distance of nodes from *ap* center

σ	68.3%	99.0%	99.7%
0	0 m.	0 m.	0 m.
5	5 m.	12.9 m.	15 m.
10	10 m.	25.8 m.	30 m.
15	15 m.	38.7 m.	45 m.

When switching to the AGG model, SA-BCAST fruitfully uses the contact opportunities in *aps* to speed up the infection (fig.2). The collateral effect is that SA-BCAST, although able of smoothing down the generated traffic as soon as a high coverage is reached (fig.2), is too aggressive in diffusing when nodes are in an *ap*, thus generating a high number of duplicates (fig.3). Some improvements can be achieved with a more stringent \mathcal{F} , as shown in fig.3 for $\sigma_{AGG} = 10$. With $N_{th} > 100$ some contact opportunities may be missed, thus further reducing traffic; coverage is anyway achieved, although with a higher latency, thanks to the existence of multiple relays. The number of *aps* also has impact: when it tends to ∞ , the RWP and AGG models coincide. By contrast, with 3 *aps* we observed more duplicates, because

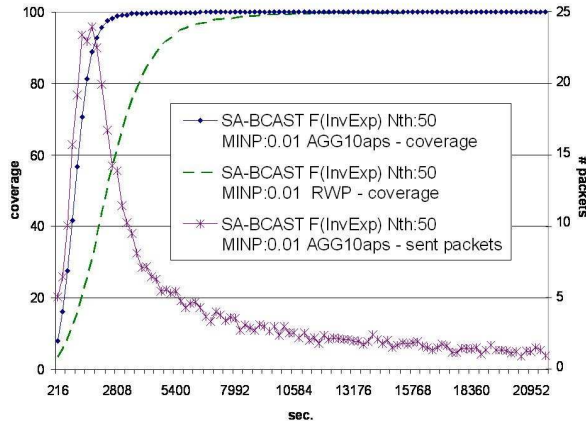


Fig. 2 Coverage and number of generated packets vs. time for SA-BCAST in either the RWP or the AGG model, with $\sigma_{AGG} = 10$




Fig. 3 Cumulative number of duplicates after 6 hours, for different mobility models, \mathcal{F} , Nth , with $\sigma_{AGG} = 10$

(i) nodes are more dense in *aps*, and (ii) there is a higher probability of encounters during movements between two *aps*, which are used for dissemination.

An interesting aspect is shown in fig.4, where the network coverage is reported for a single simulation; in order to emphasize the behavior, with $\mathcal{F} = \text{InvExp}$ the $Prob_p$ is decreases by dividing the current value by 20. With 10 *aps*, where encounters during movements are more sporadic, coverage increases in steps, which correspond to relevant membership changes in the *aps*. Steps become less high with the progress of the simulation, because the probability of entering an *ap* with already infected nodes increases. This behavior is much more evident for $Prob_p$ decreasing

more quickly. It confirms that SA-BCAST, although nodes are unaware of being either in an *ap* or on the road, allows to effectively exploit node density to increase the coverage. However, this is achieved at the expenses of efficiency. The size of the aggregation area also has impact (fig.5): with a larger σ_{AGG} , nodes in an *ap* are not all in mutual communication range. When a node enters in an *ap*, the messages exchanged do not affect all nodes and, at the same time, reduce $Prob_p$ thus preventing excessive diffusion. Yet, efficiency is still far from that achieved in the RWP model. The lesson learnt is thus that:

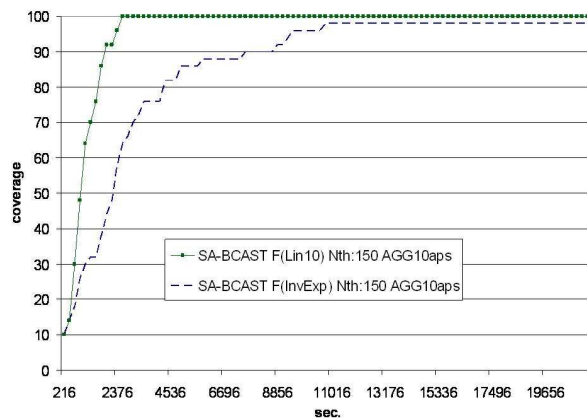


Fig. 4 Progress of coverage vs. time in a single simulation, with $\sigma_{AGG} = 10$

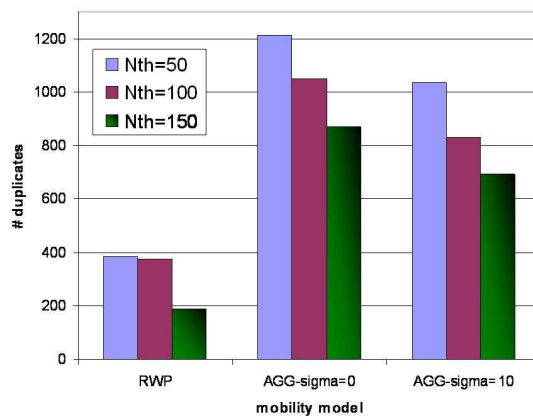


Fig. 5 Cumulative number of duplicates after 6 hours for SA-BCAST with $\mathcal{F} = \text{Lin10}$, variable *Nth* and mobility model, and 10 *aps*

- with sporadic encounters, monitoring both the neighborhood and the infection state allows to guarantee both effectiveness and efficiency; but,
- in aggregation points, slowing down diffusion when duplicates are received is a late reaction. Moreover, when the nodes are all infected, encounters in densely populated regions generate a large amount of useless traffic, due to the lack of knowledge about the system state. The need of a stop condition is much more evident.

We then measured the performance achieved with HP-BCAST and HSA-BCAST. In our measures, we assume that each node has enough memory to record up to 50 encounters. The history actually implements a stop condition because, when




Fig. 6 Cumulative number of generated duplicates vs. time with and without history, in AGG model with $\sigma_{AGG} = 10$

each node has the identifiers of all the other nodes in its own history, no message is anymore generated (fig.6), while the number of messages with SA-BCAST diverges. Its effect on the coverage is that the history alone (HP-BCAST) reduces the latency; this effect is also achieved in HSA-BCAST because nodes are able to react to the encounter with a node never seen before. By contrast, the efficiency achieved with HSA-BCAST with *aps* is even higher than that achieved by SA-BCAST in the RWP model (fig.7), that is, the history is effective in suppressing duplicate transmissions. We are currently analyzing the trade-off between efficiency gain and memory overhead for history storage and exchange, in case a partial history is maintained or exchanged.




Fig. 7 Target ratio vs. coverage with and without history, in RWP and AGG model with $\sigma_{AGG} = 10$

4.3 Aggregation vs. Swarm Model

The SWR model used in the experiments has pause time of 10 minutes, $\sigma_{SWR} = 1.73$, μ_{SWR} of 4 or 15, $d_{SWR} = 15$ m., and $M_{SWR} = 0.2$. With this model, coverage can be incremented when two swarms partially overlap, one of which has already been infected. On the other hand, once a swarm has been infected, the nodes belonging to it should refrain from transmitting again till the swarm membership does not change. Hence, better performance is achieved with Nth low, which promptly detects swarm overlapping. In fig.8, the performance of SA-BCAST is reported for $\mathcal{F} = \text{Lin10}$; with $\mathcal{F} = \text{InvExp}$ the coverage achieved is worse. Yet, in the latter case a lower number of duplicates is generated (fig.9). Hence, the InvExp function has been adopted for experiments with HSA-BCAST and LSA-BCAST.

The local suppression mechanism does not provide benefits in the aggregation model. However, in the SWR model (figs.8 and 9) it is able to improve both coverage and – above all – efficiency. This derives from the small delay before diffusing: if two swarms A and B are overlapping, such that nodes in A own m while nodes in B do not, an infected node in swarm A is likely to observe a sequence of new neighbors appearing at a short interval one after another. One “late” transmission allows to infect more new neighbors at one time. At the same time, the history mechanism allows infection propagation in swarm B: the newly infected nodes in the intersection have empty histories. They see all their neighbors in swarm B as not being in their histories, thus starting message diffusion. This repeats recursively till the whole swarm B is infected.

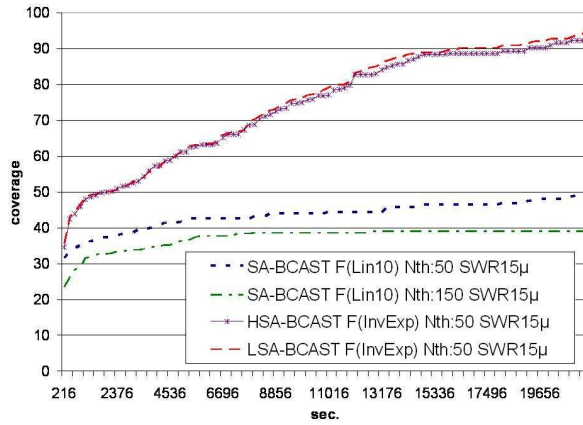


Fig. 8 Coverage vs. time for different protocols

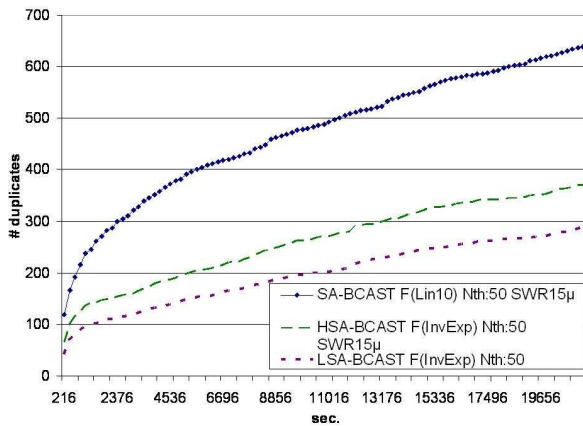


Fig. 9 Cumulative number of duplicates vs. time for different protocols

5 Conclusions and Future Work

A family of protocols – adopting different mechanisms to adapt the infection aggressiveness – has been measured in different mobility models. The AGG model brings into evidence the importance of maintaining some knowledge about the system in order to improve the usage of resources and approximate a stop condition, in highly mobile but not too sparse environments. The SWR model emphasizes the need of both suppressing duplicates when no changes in the neighborhood occur, and being prompt in diffusing when partitions happen to be merged.

Our future work is developing in three directions. We are analyzing the trade-off between performance obtained and the amount of knowledge maintained and exchanged. We are designing a dynamically adaptive protocol that monitors neighborhood in order to estimate the mobility model it currently is in, in order to dynamically change its own diffusion policy. We are developing a *hybrid* mobility model, in which nodes can assume and alternate several possible human behaviors, in order to (i) reproduce a realistic environments for an accurate evaluation of the diffusion protocols, and (ii) test the adaptive algorithm under deployment.

Acknowledgements 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.

References

1. Burgess, J., Gallagher, B., Jensen, D., Levine, B.N.: MaxProp: Routing for Vehicle-Based Disruption-Tolerant Networks. In: Proc. IEEE INFOCOM (2006).
2. Camp, T., Boleng, J., Davies, V.: A Survey of Mobility Models for Ad Hoc Network Research. In: Wireless Communication and Mobile Computing (WCMC): Special issue on Mobile Ad Hoc Networking: Research, Trends and Applications, 2(5), pp. 483-502 (2002).
3. Cooper, D.E., Ezhilchelvan, P., Mitrani, I.: High Coverage Broadcasting for Mobile Ad-hoc Networks. In: Proc. NETWORKING 2004, Lecture Notes in Computer Science, vol. 3042, pp. 100-111 (2004).
4. CRAWDAD – A Community Resource for Archiving Wireless Data At Dartmouth. <http://crawdad.cs.dartmouth.edu/>
5. Davis, J.A., Fagg, A.H., Levine, B.N.: Wearable computers as packet transport mechanisms in highly-partitioned ad-hoc networks. In: Proc. 5th IEEE Intl. Symp. on Wearable Computers (2001).
6. Delay Tolerant Networking Research Group. <http://www.dtnrg.org/wiki>
7. de Waal, C., Gerharz, M.: BonnMotion – A mobility scenario generation and analysis tool. <http://web.informatik.uni-bonn.de/IV/Mitarbeiter/dewaal/BonnMotion/>
8. Giudici, F., Pagani, E., Rossi, G.P.: Stateless and self-adaptive broadcast in delay and disruption tolerant networks. Technical Report, Università degli Studi di Milano (2008).
9. Harras, K.A., Almeroth, K.C., Belding-Royer, E.M.: Delay Tolerant Mobile Networks (DTMNs): Controlled Flooding in Sparse Mobile Networks. In: Proc. IFIP Networking (2005).
10. Hong, X., Gerla, M., Pei, G., Chiang, C.: A group mobility model for ad hoc wireless networks. In: Proc. 2nd ACM international Workshop on Modeling, Analysis and Simulation of Wireless and Mobile Systems, pp. 53-60 (1999).
11. Hui, P., Chaintreau, A., Scott, J., Gass, R., Crowcroft, J., Diot, C.: Pocket switched networks and human mobility in conference environments. In: Proc. ACM SIGCOMM Workshop on Delay Tolerant Networking (2005).
12. Jones, E.P.C., Li, L., Ward, P.A.S.: Practical routing in delay-tolerant networks. In: Proc. ACM SIGCOMM Workshop on Delay-tolerant networking (WDTN), pp. 237-243 (2005).
13. Juang, P., Oki, H., Wang, Y., Martonosi, M., Peh, L., Rubenstein, D.: Energy-efficient computing for wildlife tracking: Design tradeoffs and early experiences with ZebraNet. In: Proc. ASPLOS (2002).

14. Lee, U., Magistretti, E., Zhou, B., Gerla, M., Bellavista, P., Corradi, A.: MobEyes: Smart Mobs for Urban Monitoring with a Vehicular Sensor Network. In: IEEE Wireless Communications, 13(5) (2006).
15. Lenders, V., Karlsson, G., May, M.: Wireless Ad Hoc Podcasting. In: Proc. 4th IEEE Conf. SECON, pp. 273-283 (2007).
16. Lindgren, A., Doria, A., Schelen, O.: Probabilistic routing in intermittently connected networks. In: Proc. 1st Intl. Workshop on Service Assurance with Partial and Intermittent Resources (SAPIR), Lecture Notes in Computer Science, Vol. 3126, pp. 239-254 (2004).
17. Montessor, A., Jelasity, M., Babaoglu, O.: Gossip-based Aggregation in Large Dynamic Networks. In: ACM Transactions on Computer Systems, 23(3), pp.219-252 (2005).
18. Spyropoulos, T., Psounis, K., Raghavendra, C.S.: Single-copy routing in intermittently connected mobile networks. In: Proc. 1st IEEE SECON, pp. 235-244 (2004).
19. Su, J., Chin, A., Popinova, A., Goely, A., de Lara, E.: User Mobility for Opportunistic Ad-Hoc Networking. In: Proc. 6th IEEE Workshop on Mobile Computing Systems and Applications (2004).
20. UCLA Parallel Computing Laboratory: GloMoSim – Global Mobile Information Systems Simulation Library. University of California at Los Angeles. <http://pcl.cs.ucla.edu/projects/glomosim/>
21. Vahdat, A., Becker, D.: Epidemic Routing for Partially Connected Ad Hoc Networks. Technical Report CS-200006, Duke University (2000).