On the Impact of a Road-Side Infrastructure for a DTN Deployed on a Public Transportation System

Sabrina Gaito, Dario Maggiorini, Christian Quadri, and Gian Paolo Rossi

Università degli Studi di Milano
Via Comelico 39, 20135 Milano, Italy
`firstname.lastname@unimi.it`

Abstract. In these last few years we observed a tremendous raise in the field of personal communication and location- context-based services. Unfortunately, in regard to these kind of services, standard connectivity, such as 3G, is limited by scalability issues and its strictly pull-based service model. As an alternate solution, Delay-Tolerant Networks (DTN) have already been proposed as a way to obtain a scalable and efficient urban backbone by leveraging an already available public transportation system. To bound delay of DTNs when deploying in an urban environment the most cost-effective solution is the partial involvement of road-side infrastructure. This paper addresses the problem of infrastructure allocation in a real city environment by analyzing the topology and the timetable of a real urban setting to propose an algorithm suitable for reducing deployed equipment. By means of extensive simulations we show that, a very limited number of exchange points leads to a significant performance improvement and helps planning for a bounded delivery delay.

Keywords: Vehicular networks, Delay tolerant networks, Opportunistic networks, Routing protocols.

1   Introduction

In these last few years we observed a tremendous raise in the demand of personal communication and ubiquitous data access; we now think to mobile services as a part of our everyday life. Following this trend, a wide set of services addressing mobile users are now offered based on location and context and have an urban-wide scope. Among these services we may find both diffusion services, such as public utility services, local news, commercial advertisements, recommending systems, and interactive services such as participatory sensing and computing.

As a result of this hunger for location-dependent data, standard 3G connectivity is suffering from poor scalability and is limited by its pull-based service model. Moreover, many cell-phone providers will offer the same service from within their core network, making it difficult the interoperation between users of competing providers. To overcome these limitations, many service and data providers (and users) are currently looking for novel connectivity solutions that

complement the 3G network services and offer provider-independent, location-based, and cheaper than 3G massive data distribution. Opportunistic Networks (ONs), as a special case of Delay Tolerant Networks (DTNs), have attracted growing interest because of their ability to provide flexible and robust multi-hop connectivity by leveraging opportunistic contacts among mobile radio devices. Previous work already suggested to deploy an ON over the Public Transportation System (PTS) operating in a urban area (see [12,18,3,19,8,9,7]). In particular, recent researches [9,7] have assessed the potential of this solution and shown that its application to a city-wide environment ensures an high delivery ratio and an acceptable delivery delay to meet the requirements of a large class of mobile services. All these results are leading to a vision where ONs coexist with other wireless technologies in a multi-platform urban backbone.

Anyway, when deploying this kind of solution on a very large scale, the problem of ensuring scalability must be addressed. Since size and shape of bus trips are limited by human and organizational factors, delay will ramp up exponentially with the covered area due to the increasing number of hops each packet must traverse. Needless to say, keeping the delay below a certain threshold is a key requisite for an effective service provisioning. Currently, to bound delay when extending service to a metropolitan area, the most cost-effective solution is the involvement of a road-side infrastructure. Placing radio equipped devices at bus stops provides a few benefits. It mainly allows to exploit existent infrastructures managed by the public transportation authority. As a consequence, the hardening of the devices is simple and energy supply ensured. Furthermore, bus stops have been natively designed to favor line intersection and offer bus-line to bus-line connectivity with high intra-contact time between bus and road-side units.

This paper addresses the problem of infrastructure deployment/allocation in a real city environment. We analyze the topology and the timetable of a real urban setting and then we propose an algorithm which is suitable to reduce the number of road-side units. By means of extensive simulations we show that a very limited number of road-side units can lead to a significant performance improvement in term of delivery delay.

In previous work [7] we already proposed a single-copy routing protocol for DTNs in urban environment. Here we aim to raise the bar and define a methodology to demonstrate that a DTN can be effectively deployed on a PTS by introducing a smart, lightweight, and non invasive road-side infrastructure. Such a DTN will be finally able to provide mobile services based on location and context while satisfying quality of service requirements typical of urban-wide information services.

2   Related Work

Despite the fact that DTNs have been proposed only recently, in 2003 [6], a considerable effort has been devoted by the scientific community to devising reliable routing strategies on these kinds of architectures [11,15,13]. In particular, [13] is

a multi-copy forwarding scheme where probabilities are calculated on a per-bus base and adjusted with an aging policy. As special cases, in [10,4,17,21] the contact opportunities between nodes are exploited to build opportunistic networks. Among ONs, networks built on top of public transportation systems have been attracting attention in recent years.

The first contributions [16,5] focused on rural environments in developing regions where a number of villages are spread over a large territory linked by buses. The common goal of all the mentioned projects is to provide network access for elastic non real-time applications so that the local population may enjoy basic Internet services (e.g., e-mail and non-real time web browsing).

Campus bus networks (e.g., [3,19,2]) are designed to serve students and faculties who commute between colleges or from/to nearby towns. These kinds of services are usually characterized by a relatively small number of nodes when compared to a fully fledged urban environment. The main contribution in this direction is represented by [3], where five colleges are linked with nearby towns and to one another over an area of 150 square miles. The authors of the aforementioned paper propose a multi-copy routing algorithm, namely MaxProp, based mainly on message priorities. These priorities depend on the path likelihoods to destination nodes on the basis of historical data and other complementary mechanisms. By means of simulation MaxProp is shown to outperform oracle-based protocols hinged on knowledge of deterministic meetings between peers. This research was been extended in [19], where inter-contact time distribution was analyzed both at bus and line levels. On the same bus network, a system of throwboxes [20] was deployed to enhance the capacity of the DTN. The last case ([2]), models the routing system as an allocation problem and try to optimize a specific routing metric such as worst-case delivery delay or the fraction of packets that are delivered within a deadline.

Scaling up, we find bus networks in urban environment (e.g., [12,1,18]) which are usually characterized by a considerable number of densely deployed lines, many contact opportunities, and frequent contacts. In [12], authors propose a commercial application called *Ad Hoc City*. Ad Hoc City is based on a multi-tier wireless ad-hoc network architecture and provides elastic Internet access by means of Access Points which are responsible for a geographical area. Authors verified the validity of the proposed approach against real movement traces by *King County Metro bus system* in Seattle, WA. Using the same real data as for [12], authors of [1] propose a cluster-based routing algorithm for intra-city message delivery. In [1] an efficient large-scale clustering methodology is devised: nodes are clustered on the basis of encounter frequency and multi-copy forwarding takes place between members of the same cluster hosting the destination node. The contribution from [18] uses data from the public transportation system of Shanghai to test the performance of a single-copy forwarding mechanism based on intra-contact times.

A recent work about performance analysis for deployment at urban scale is [9]. In this work, authors study the inter-contact times of the transportations systems of Zurich and Amsterdam, discovering that they follow an exponential

distribution. Starting from this result, they predict the performance of an epidemic routing using a Markov chain model. These results are comparable, as average value, to the ones obtained from our simulations, demonstrate the feasibility of using a PTS as a DTN backbone.

3   The Urban Environment

The urban environment we used to evaluate our proposal is the city of Milan (Italy) and its PTS. Milan is a medium size town (typical for many European cities) and its PTS is a complex system extending above and below ground. The ground PTS of Milan, spans 69 lines over 49 square miles for a total paths length of 683 miles (13.85 miles for every square mile).

Fig. 1. PTS layout in the city of Milan

As it can be observed from Fig. 1, the overall city structure is clearly not manhattan-like; crosses between bus lines may occur at any time and there is no constant space between intersection points. Moreover, buses can run along three kinds of routes: they may span across the city, run around the center, or cover only a peripheral section making a relatively small loop. While manhattan-like topologies have been extensively studied in the past, uneven topologies like the one we just described are less addressed. This is why Milan represents an interesting case study: for its uneven topological structure and the frequent contacts between nodes.

4   Infrastructure Deployment

In this section we are going to introduce the road-side network infrastructure and describe our algorithm to select infrastructure nodes position.

4.1   Road-Side Exchange Points

Our road-side infrastructure will be made of apparatus we baptized Road-Side eXchange Points (RSXPs). RSXPs are wireless-equipped appliances located at bus stops and working as fixed disconnected relays. Operating in a urban area,

RSXPs are supposed to be permanently attached to power supply and can adopt complex hardware to manage substantial data flows and store a large amount of data in the internal memory.

An RSXP will build an internal database of lines passing nearby. When a bus comes into transmission range with an RSXP the following will happen:

1. The RSXP will declare to the bus all known lines.
2. The bus may decide, based on its forwarding policy, to send to the RSXP packets whose next hop is known.
3. The RSXP will send to the bus all packets for which it is a valid next hop.

As we can see, the routing decision sits only on the bus, and the RSXP will act as an external cache; changing routing policy will not require any upgrade or management on the road-side infrastructure. Moreover, RSXPs are not known to the buses: they are exploited when encountered in the same fashion as a contact opportunity with another bus. This way, a faulty RSXP will not jeopardize data forwarding on the DTN.

4.2   Reduction Algorithm Overview

The reduction algorithm we propose is composed of three phases.

1. Bus stops covering the same radio area are merged.
2. Relevant bus stops are identified from lines encounter probability.
3. Redundancy is eliminated by solving a Set-Covering optimization problem.

In the following Sections we describe in detail the three phases of the algorithm.

4.3   Sites Selection

In order to select deployment sites for RSXPs we start from bus stop positions. At bus stops we have scheduled pulls up where a longer intra-contact time is experienced and a significant data transfer can take place. Moreover, intersecting lines share the same bus stop, making the place a good location for a packet waiting for a ride. Of course, placing an RSXP at every bus stop will create a huge overlap and, as a consequence, a huge waste of resources.

In order to reduce the number of RSXP we implement the following greedy algorithm.

step 1. An RSXP is created at each bus stop and is tagged with all bus lines stopping there.

step 2. For every RSXP its neighborhood is analyzed and tagged lines are merged. Two RSXPs are neighbors if and only if they are in mutual transmission range and in line of sight.

step 3. RSXPs are sorted in descending order, based on the number of the distinct lines the RSXP has in sight. The list is then scanned starting from the first element; the algorithm removes all other RSXPs which are neighborhood of current RSXP.

step 4. All RSXPs tagged with only one line are discarded because they will not offer any real contact opportunity.

This algorithm guarantees that resulting RSXPs have all declared lines in-sight, because if two stops of different lines are in-sight, also buses belonging to those lines have both stops in-sight.

When this approach is applied to the city of Milan, the total number of RSXPs is reduced from 3741 to 561 with a reduction of 85% for the number of sites.

4.4   Required RSXPs

The number of RSXP can be further reduced based on the lines encounter probability. We want to use the road-side infrastructure to forward data only between those pairs of lines which have a low encounter probability. The encounter probability between pairs of lines is calculated as in [7], where a Binomial model based on history of buses encounters is used. This way, we will boost probability only for those hops who are less likely to happen.

We define an RSXP to be required if and only if there is at least a pair of lines which can be reached by the RSXP whose encounter probability is below a certain threshold, τ. Only pairs of lines having a non-zero encounter probability are considered, to avoid the case when two lines never encounter. More formally, given L the set of all lines in the system, and L_x the set of lines reachable from RSXP x ($L_x \subseteq L$) we can define the set U as the set of all pairs of line whose encounter probability is lower than τ.

$$U = \{\{u, v\} \mid p(u, v) \leq \tau, u \in L, v \in L, \tau \in (0, 1]\}$$

RSXP x is required if a pair $\{u, v\} \in U$ such that u and $v \in L_x$ exists.

By applying this criterion we can tune our system to the best value for τ with respect to the observed performances.

For the city of Milan, the number of required RSXPs depending on τ is reported in Tab. 1, second column.

4.5   Further Reduction

One last level of refinement can be reached: it is possible that the same pair of lines is covered by more than one RSXP. This means that potentially we might deploy more RSXPs than the ones needed to cover the set U.

This last reduction has been obtained by solving a set covering problem on U. Given R as the set of required RSXP, the incidence matrix A can be defined as:

$$A = [a_{i,j}] \qquad i \in U, j \in R$$

$$a_{i,j} = \begin{cases} 1 & \text{if RSXP } j \text{ can see both buses of pair } i \in U \\ 0 & \text{otherwise} \end{cases}$$

Then we must find

$$\min \sum_{j \in R} X_j$$

with

$$\sum_{j \in R} a_{i,j} X_j \geq 1 \qquad \forall i \in U, X_j \in \{0, 1\}$$

In the formulas, X_j represents the binary variable associated to RSXP j.

The third column of Tab. 1 reports the number of RSXPs for the city of Milan when applying the complete algorithm.

Table 1. Number of RSXPs for several values of τ

τ	Required RSXPs	Deployed RSXPs
0.1	30	17
0.2	78	47
0.3	128	79
0.4	168	101
0.5	198	110

5   Simulation Setup

In this section we describe the simulation environment which has been set up to test how DTN performances benefit from the presence of a road-site infrastructure. To simulate data transfer over the PTS we have used URBeS (Urban Routing Backbone Simulator) simulator [7], which allowed us to consider a real urban topology and bus schedules.

In our experiments we considered IEEE 802.11b technology. Available bandwidth is 10 Mpbs with a radio range of 100 meters. Communication takes urban canyons into account; we consider only line-of-sight contacts.

Data traffic generation is performed continuously during working hours: from 8 A.M. to 8 P.M. During simulations, unicast traffic is generated from <*one bus of a bus line*> to <*any bus of the destination bus line*>, following a rationale discussed in [19]. Each operating bus generates 64 KB packets at a constant rate; each packet will be delivered to a distinct, randomly chosen, bus line. Data packet generation will vary from 20 to 400 packets per hour from every bus. The total number of packets coming into play is thus between $138,364$ (20 packets per hour) and $2,766,965$ (400 packets per hour).

When a packet is generated it is placed in the node local buffer until forwarding becomes possible; an isolated bus will keep accumulating packets while no contacts are experienced. When an encounter happens all the packets are checked for forwarding. Forwarding is, of course, subject to bandwidth and buffer limitations. Bandwidth is accounted using a token bucket mechanism while buffer space availability is simply checked before transmission. In case of contention a first-in-first-out policy is applied. The size of the application buffer on board at

every bus has been set to 4 GB. Packets are then forwarded in accordance with the adopted routing policy.

When a bus reaches the end of the line it may or may not queue up and wait for another scheduled departure. If the bus keeps a place in line it will hold all its data and will keep generating packets while waiting. If, on the other hand, the bus leaves service all content will be pushed to the first bus waiting in line. If there is no bus in line (because there are no more scheduled departures) all stored packets are dropped and will be considered lost.

When a packet is queued for forwarding, the adopted routing policy is used to identify the next hop. Packet forwarding can take place under three different conditions: (i) the bus which is carrying the packet encounters a bus belonging to the next-hop line, (ii) the carrying bus encounter an RSXP which has in its database the next-hop line for the packet, and (iii) if the packet is queued on an RSXP and a bus belonging to the next-hop line shows up. In our previous work [14] we have shown the scalability problems of the multi-copy routing algorithms Rapid [2] and MaxProp [3]. We have also shown the scalability of Op-HOP (Opportunistic Hopping on Probability), a single-copy routing algorithm which have been proposed in [7] and is designed to maximize the delivery probability. According to these results, we adopt Op-HOP as routing policy in order to better asses the actual gain provided from the infrastructure.

6   Simulation Results

In this section we are going to discuss the outcome of our simulation experiments. We will present Op-HOP performance improvements with respect to various values of τ, including 0 and 1. When τ is 0, pure ad hoc mode will be adopted: no RSXP will be deemed as required and only bus-to-bus forwarding will take place. This value will provide a baseline for performance comparison. On the opposite, if τ is 1, all RSXPs will be required and all forwarding operations will be performed through the infrastructure: bus-to-bus forwarding is forbidden. Experiments with this value of τ will provide a performance upper bound.

6.1   Data Delivery

Unsurprisingly, exploiting the network infrastructure is reducing the average delivery delay, as reported in Fig. 2(a). From this figure, we can observe a significative improvement, up to one hour (i.e., 30%), even with low values of τ. Interestingly enough, increasing the offered load seems to make the situation better thanks to packets sitting on RSXPs and leaving in bursts for the next hop. This behavior can also be partially explained observing the average delivery ratio (Fig. 2(b)). By combining Fig. 2(a) and Fig. 2(b) we can see that packets following long paths gets discarded more frequently. Discarding packets running long distances reduces the average delivery time for the ones who are able to reach the destination.

Moreover, lower values of τ seem to make delivery ratio worst while for $\tau \geq 0.3$ the situation improves. This behavior indicates that, depending on the PTS,
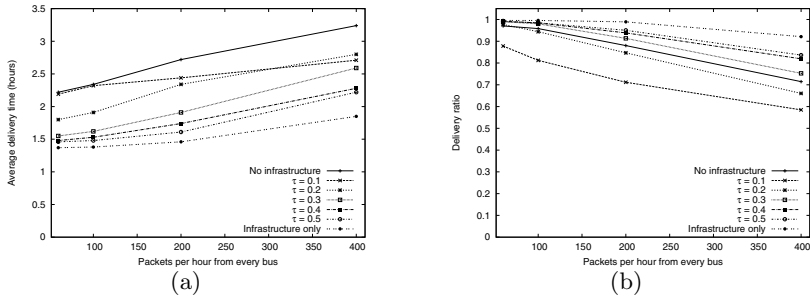
Fig. 2. Performances at increasing offered load with multiple value of τ: average delay (a) and average delivery ratio (b)

there is a lower bound for the number of RSXPs to deploy. Below this bound, the infrastructure is not able to manage all the offered load. The explanation is that overloaded RSXPs will not be able to spool out all queued packets during a single bus contact, increasing significantly the buffering time and leading to considerable packet drops at the end of the day.

Furthermore, we detected an uneven usage for the RSXP buffers. As an example, in the worst case, with $\tau = 0.1$ (17 RSXPs) at peak traffic (6:30 P.M.) we have 14% of RSXPs whose buffer is empty and 23% of them who are not able to spool all queues in a single bus encounter. On the other hand, the most uniform situation can be observed with $\tau = 0.5$. As a matter of fact, for the PTS of Milan, a good tradeoff value for τ seems to be 0.4. In all the examined cases the system seems to perform better than the infrastructure-less version and setting τ to 0.5 does not give a significative improvement over 0.4.

6.2   Number of Hops

The average number of hops for delivered packets is reported in Fig. 3(a). As we can see in the picture, for low values of τ the use of an infrastructure makes the paths longer. This is confirming the performance degradation outlined in Sec. 6.1. Greater values of τ improve the situation but do not outperform the infrastructure-less version in a sensible way when the traffic is very high.

Obviously, if we consider only the lines traversed by a packet, the paths are actually becoming much shorter, as depicted by Fig. 3(b). This is due to the fact that short paths traversing the infrastructure benefit from contacts opportunity happening for sure (soon or later). These short paths will no longer get penalized in favor of longer, and more probable, ones. Nevertheless, here we can observe that the reduction of the number of involved lines is not proportional to τ or the number of RSXPs. Once again, the answer can be found in an uneven distribution of the network load: full buffers at RSXPs will prevent some buses to use the shortest, planned, route and will make them wait for a bus-to-bus forwarding opportunity.
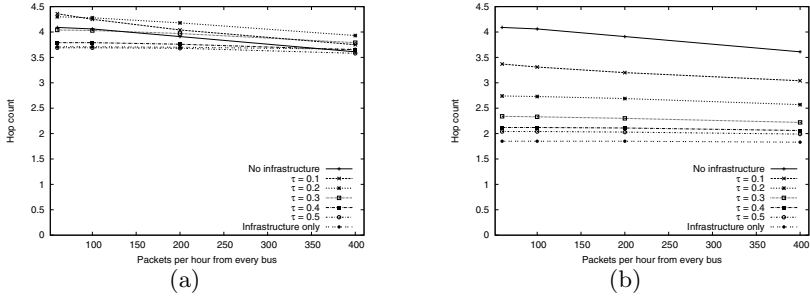
Fig. 3. Statistics on delivered packets: average number of hops (a) and average number of involved lines (b)


Fig. 4. Average humber of hops between buses and RSXPs for delivered packets

Another interesting consideration comes from Fig. 4, where the average number of hops not taking place between buses is reported. While it is obvious that increasing τ we increase the number of RSXP and, in turn, the degree of involvement of the infrastructure; it is less obvious that there is an upper bound for the number of hops provided by the infrastructure regardless of the number of RSXPs. In our case, the average infrastructure involvement is not greater than 1.75 hops for any value of τ; unless, of course, we forbid data forwarding between buses.

By comparing Fig. 3(a) and Fig. 4 we can observe that, for $\tau \geq 0.3$, 50% of the hops are from bus to RSXP (or vice versa). Looking to Fig. 3(b), for the same values of τ, the average number of lines traversed by a packet is 2. Putting these information together we can state that, after a certain density of RSXPs (i.e., a certain threshold for τ) the average path converges toward a sequence like "*source line* $\rightarrow$ *RSXP* $\rightarrow$ *destination line*", and deploying RSXPs above such density is a potential waste of resources. This is the explanation for the upper bound in Fig. 4.

From the figures above, in the case of Milan, we can confirm that setting τ to 0.4 gives a good tradeoff between performances and infrastructure deployment.

7   Conclusion

In this paper we addressed the allocation problem for a network infrastructure to augment the performance of an Opportunistic Network deployed on the Public Transportation System of a real city. The proposed solution consists in the creation of a road-side network infrastructure by means of disconnected relays we call Road-Side eXchange Points. Using an algorithm, we identify a proper subset of bus stops candidate to take part in fixed infrastructure.

By means of extensive simulations on the topology and the timetable of the city of Milan, Italy, we shown the impact of RSXPs density on the DTN performances. Results indicate that delivery delay will benefit from the presence of the infrastructure but, in order to improve the delivery ratio, a minimum density is required. The number of traversed hops indicates that the infrastructure involvement in the forwarding process is upper bounded and deployment of RSXPs above a certain density is a potential waste of resources. Using these results, we demonstrate that a lightweight, and non invasive, road-side infrastructure is a feasible way to deploy a DTN over a PTS to provide mobile services based on location and context on a urban scale. By adopting a single-copy protocol, we can also ensure scalability and satisfying quality of service requirements typical of urban-wide information services.

References

1. Ahmed, S., Kanhere, S.: Cluster-based forwarding in delay tolerant public transport networks. In: 32nd IEEE Conference on Local Computer Networks, LCN 2007, pp. 625–634 (October 2007)
2. Balasubramanian, A., Levine, B., Venkataramani, A.: Dtn routing as a resource allocation problem. In: Proceedings of the 2007 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications, SIGCOMM 2007, pp. 373–384. ACM (2007)
3. Burgess, J., Gallagher, B., Jensen, D., Levine, B.N.: Maxprop: Routing for vehicle-based disruption-tolerant networks. In: Proceedings of the 25th IEEE International Conference on Computer Communications, INFOCOM 2006, pp. 1–11 (April 2006)
4. Cai, H., Eun, D.Y.: Aging rules: what does the past tell about the future in mobile ad-hoc networks? In: MobiHoc 2009: Proceedings of the Tenth ACM International Symposium on Mobile Ad Hoc Networking and Computing, pp. 115–124. ACM (2009)
5. De Oliveira, C.T., Braga, R.B., Taveira, D.M., Fern, N.C., Duarte, O.C.M.B.: A predicted-contact routing scheme for brazilian rural networks. In: Electrical Engineering Program, COPPE/UFRJ (2008)
6. Fall, K.: A delay-tolerant network architecture for challenged internets. In: SIGCOMM 2003: Proceedings of the 2003 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications, pp. 27–34. ACM (2003)
7. Gaito, S., Maggiorini, D., Rossi, G.P.: Bus Switched Networks: an Ad Hoc Mobile Platform Enabling Urban-Wide Communications. Ad Hoc Networks Journal (to appear, 2012)

8. Gaito, S., Maggiorini, D., Pagani, E., Rossi, G.P.: Distance vector routing for public transportation vehicular networks: performance evaluation on a real topology. In: Proc. of the 2nd IFIP Wireless Days Conf., pp. 338–342 (2009)

9. Hartog, L.D., Spyropoulos, T., Legendre, F.: Using public transportation as a dtn backbone: Mobility properties and performance analysis. In: Proceedings of AOC 2010. IEEE, Montreal (2010)

10. Hui, P., Chaintreau, A., Scott, J., Gass, R., Crowcroft, J., Diot, C.: Pocket switched networks and human mobility in conference environments. In: WDTN 2005: Proceedings of the 2005 ACM SIGCOMM Workshop on Delay-tolerant Networking, pp. 244–251. ACM (2005)

11. Jain, S., Fall, K., Patra, R.: Routing in a delay tolerant network. In: SIGCOMM 2004: Proceedings of the 2004 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications, pp. 145–158. ACM (2004)

12. Jetcheva, J., Hu, Y.C., PalChaudhuri, S., Saha, A., Johnson, D.: Design and evaluation of a metropolitan area multitier wireless ad hoc network architecture. In: Mobile Computing Systems and Applications, pp. 32–43 (October 2003)

13. Lindgren, A., Doria, A., Schelén, O.: Probabilistic routing in intermittently connected networks. SIGMOBILE Mob. Comput. Commun. Rev. 7(3), 19–20 (2003)

14. Maggiorini, D., Gaito, S., Rossi, G.P., Quadri, C.: On the scalability of Delay-Tolerant routing protocols in urban environment. In: IFIP Wireless Days Conference (WD 2011), Niagara Falls, Ontario, Canada (October 2011)

15. Nelson, S.C., Bakht, M., Kravets, R.: Encounter-based routing in dtns. In: INFOCOM 2009, pp. 846–854. IEEE (April 2009)

16. Pentland, A., Fletcher, R., Hasson, A.: Daknet: rethinking connectivity in developing nations. Computer 37(1), 78–83 (2004)

17. Sandulescu, G., Nadjm-Tehrani, S.: Opportunistic dtn routing with window-aware adaptive replication. In: AINTEC 2008: Proc. of the 4th Asian Conference on Internet Engineering, pp. 103–112. ACM (2008)

18. Sede, M., Xu, L., Da, L., Min-You, W., Minglu, L., Wei, S.: Routing in large-scale buses ad hoc networks. In: Wireless Communications and Networking Conference, pp. 2711–2716 (2008)

19. Zhang, X., Kurose, J., Levine, B.N., Towsley, D., Zhang, H.: Study of a bus-based disruption-tolerant network: mobility modeling and impact on routing. In: MobiCom 2007: Proceedings of the 13th Annual ACM International Conference on Mobile Computing and Networking, pp. 195–206. ACM (2007)

20. Zhao, W., Chen, Y., Ammar, M., Corner, M., Levine, B., Zegura, E.: Capacity enhancement using throwboxes in dtns. In: Proc. IEEE Intl. Conf. on Mobile Ad hoc and Sensor Systems (MASS), pp. 31–40 (2006)

21. Zignani, M., Gaito, S.: Extracting human mobility patterns from gps-based traces. In: Proceedings of the 3rd IFIP Wireless Days Conference 2010 (2010)