

Routing in Quasi-Deterministic Intermittently Connected Networks

Paolo Giaccone¹, David Hay¹, Giovanni Neglia², and Leonardo Rocha²

¹ Dipartimento di Elettronica, Politecnico di Torino, Italy,
`name.surname@polito.it`

² Maestro project-team, INRIA Sophia-Antipolis Méditerranée, France
`name.surname@sophia.inria.fr`

Key words: Intermittently Connected Networks, Delay Tolerant Networks, Routing

Extended Abstract

Some of the recent applications using wireless communications (wildlife monitoring, inter-vehicles communication, battlefield communication,...) are characterized by challenging network scenarios. Most of the time there is not a complete path from a source to a destination (because the network is sparse), or such a path is highly unstable and may change or break while being discovered (because of nodes mobility and time-variations of the wireless channel). Networks under these conditions are usually referred to as Intermittently Connected Networks (ICNs) or Delay Tolerant Networks (DTNs). In such scenarios information delivery is then based on the store-carry-forward paradigm: a mobile node first stores the routing message from the source, carries it from a physical location to another and then forwards it to an intermediate node or to the destination. Typical examples of ICNs are those where nodes are intrinsically mobile (independently from data transfer purpose): vehicular networks [2] (in which data is carried over cars and buses), “pocket area networks” [2] (in which data is carried by people carrying small devices like PDAs), mixed ground/satellite networks and networks of sensors attached to animals [9]. Also some scenarios in which some nodes are mobile and some nodes are fixed (e.g, mobile devices with fixed gateways) present the same challenges.

The wide range of applications, promising performance results and concise modeling have led to an extensive research on ICNs during the last few years (e.g, [5, 6, 7, 3]). At the core of this research line are *routing and scheduling algorithms*: at any given time, each node should find when and where to forward the data stored in its buffer so that it reaches the destination in a timely manner. Moreover routing for ICNs is not only limited to *forwarding schemes*, where a single copy of each packet is present in the network [1], but it also include *replication schemes*, which send many copies of the same data packet across the network. A prime example of replication schemes are *epidemic* routing algo-

rithms (a.k.a *flooding algorithms*) in which each node sends each packet to *all* its neighbors. Replication improves performance in terms of delivery probability and delivery delay when contacts cannot be predicted or when transmissions are unreliable, but at the same time it implies higher costs in terms of required bandwidth, transmission energy and buffer requirements (see [10]).

Most of the research on routing in ICNs has focused on two extreme cases: 1) when contacts among nodes are deterministic and known in advance (e.g. in the case of space communications among satellites, probes and earth or space stations [8]) or 2) when they cannot be predicted (e.g. for human and animal mobility [2, 9]) and are supposed to obey to some generic random mobility model, like random way-point, random direction or brownian models. Many interesting scenarios do not fall in any of these two cases: even complex mobility patterns often exhibit some form of periodicity or in other cases the underlying node mobility is known in advance, but it can be modified by random effects. A clear example is that of a vehicular network carrying data over public transportation (e.g., buses): the predictions of the contact times are derived from the schedule and routes of the buses; on the other hand, delays in bus operations clearly change the contact times or even prevent contacts to occur, implying that the predictions are not necessarily accurate.

Our preliminary investigation suggests that there is currently no framework to study comprehensively all the range of possible scenarios between deterministic contacts and unpredictable random contacts. For this reason, we have decided to investigate a specific class of networks characterized by *small deviations* from the deterministic contact model. We refer to such networks as “quasi-deterministic” ICNs.

Our current research consider bus networks as a case study for the general problem of routing in quasi-deterministic networks. Besides being an interesting application scenario itself, this specific network scenario will allow us to understand the key issues our models and our algorithms need to consider.

We envision that the infrastructure of bus-enabled data network is formed by (some) buses and bus stations equipped with wireless devices, e.g. based on WiFi technologies like in Dieselnet [2]. When two of them come within transmission range of each other, they can transfer data. Some access points at bus stops can also be connected to the Internet. Passengers on a bus (/waiting at a stop) can use this infrastructure through their mobile devices, associating to the bus (/stop) access point.

Here we focus on the simple case when we want to transfer some data from a bus or a bus stop to a remote bus or a remote stop. We envision two possible applications even for this simple unidirectional scheme. First, the data could be some non-time-critical information collected by sensors on the bus/stop for operation/management purpose that needs to be transferred to the bus system central operation point via a bus stop connected to the Internet. Second, the data could be destined to a passenger. We can think about possible hybrid systems, where, for example, the user requests its emails or a file through the standard

cellular data connection and then get the reply through the DTN, that could offer a cheap data transfer service.

There are different options in designing such a system. First, the system could rely only on forwarding -i.e. a single copy of the data is propagated along a path- or could take advantage of multiple copies spread in the network to increase delivery probability and reduce delivery time. A second choice is between exploiting only transmission opportunities between buses and bus stops or exploiting also direct transmission opportunities between different buses. In the latter case we can expect the system capacity to increase, but at the same time meetings between buses are more unreliable, not only in terms of the time they are going to occur, but also in terms of their existence itself, being that delays can prevent buses to miss a meeting opportunity.

In order to gain a better feeling of reasonable modeling assumptions and of possible design choices we have started considering the actual public bus network in Turin, Italy, that has about 50 frequency based bus lines (up to 12 buses per hour) and 3000 bus stops. Our current contributions follow.

1. Analyzing real bus traces, we have characterized the statistical properties of bus delays at stops.
2. In [4], some of us have determined optimal routing schemes under deterministic contacts. We have then evaluated to which extent these schemes are robust to noise in the meeting process, i.e. how performance decrease when routing is based on predicted contact times ignoring the presence of noise.
3. Given the contact predictions and *a priori* statistical information on the noise process, we have developed a *multi-hop routing and scheduling algorithm* and evaluated its performance as a matter of throughput, delay and delivery probability.

References

1. Aruna Balasubramanian, Brian Levine, and Arun Venkataramani. DTN routing as a resource allocation problem. *SIGCOMM Comput. Commun. Rev.*, 37(4):373–384, 2007.
2. John Burgess, Brian Gallagher, David Jensen, and Brian Neil Levine. MaxProp: Routing for vehicle-based disruption-tolerant networks. In *IEEE INFOCOM*, 2006.
3. Delay tolerant networking research group.
4. David Hay and Paolo Giaccone. Optimal routing and scheduling for deterministic delay tolerant networks. In *Wireless On-Demand Network Systems and Services, 2009. WONS 2009. Sixth International Conference on*, pages 27–34, Feb. 2009.
5. Sushant Jain, Kevin Fall, and Rabin Patra. Routing in a delay tolerant network. In *ACM SIGCOMM*, pages 145–158, 2004.
6. Cong Liu and Jie Wu. Scalable routing in delay tolerant networks. In *ACM MobiHoc*, pages 51–60, 2007.
7. Cong Liu and Jie Wu. Routing in a cyclic mobispace. In *ACM MobiHoc*, pages 351–360, 2008.

8. Lloyd Wood, Will Ivancic, Wes Eddy, Dave Stewart, James Northam, Chris Jackson, and Alex da Silva Curiel. Use of the delay-tolerant networking bundle protocol from space. In *Proc. of the 59th International Astronautical Congress*, September 2008.
9. Pei Zhang, Christopher M. Sadler, Stephen A. Lyon, and Margaret Martonosi. Hardware design experiences in zebranet. In *SenSys '04: Proceedings of the 2nd international conference on Embedded networked sensor systems*, pages 227–238, New York, NY, USA, 2004. ACM.
10. Xiaolan Zhang, Giovanni Neglia, Jim Kurose, and Don Towsley. Performance modeling of epidemic routing. *Comput. Netw.*, 51(10):2867–2891, 2007.