Worst-case Latency of Broadcast in Intermittently Connected Networks

Mikael Asplund

Linköping University,
SE-581 83 Linköping, Sweden,
mikael.asplund@liu.se

Simin Nadjm-Tehrani

Linköping University,
SE-581 83 Linköping, Sweden,
simin.nadjm-tehrani@liu.se

Abstract Worst-case latency is an important characteristic of information dissemination protocols. However, in sparse mobile ad hoc networks where end-to-end connectivity cannot be achieved and store-carry-forward algorithms are needed, such worst-case analyses have not been possible to perform on real mobility traces due to lack of suitable models. We propose a new metric called *delay expansion* that reflects connectivity and reachability properties of intermittently connected networks. Using the delay expansion, we show how bounds on worst-case latency can be derived for a general class of broadcast protocols and a wide range of real mobility patterns. The paper includes theoretical results that show how worst-case latency can be related with delay expansion for a given mobility scenario, as well as simulations to validate the theoretical model.

Keywords: Latency, Connectivity, Partitioned networks, Graph expansion, Dynamic networks, Real-time, Delay-tolerant networks, Analysis

Biographical notes: Mikael Asplund is an assistant professor in the Real-time Systems Group at Linköping University, Sweden. In 2011-2012 he worked one year as a Research Fellow at Trinity College Dubin. He received his MSc degree in computer science and engineering in 2005 and his PhD in computer science in 2011 both from Linköping University. His PhD thesis focused on design and analysis of partition-tolerant distributed systems, including development of middleware services for maintaining consistency and information dissemination algorithms for disaster area networks. His current research interests include dependable distributed systems, mobile and vehicular computing and real-time systems.

Simin Nadjm-Tehrani received her BSc degree (with honours) from Manchester University, UK, and her PhD degree in Computer Science at Linköping University, Sweden, in 1994. Prior to her PhD studies she worked for six years at the International companies Deloittes and PriceWaterhouse in early 80s. In 2006-2008 she was a full professor of dependable real-time systems with University of Luxembourg, and is currently back at Department of Computer and Information Science at Linköping University where she has led the Real-time Systems Laboratory since 2000. Her research interests are in dependable distributed systems with resource constraints.

1 Introduction

Wireless networks are emerging as the dominant technology for connecting devices and people together. Most wireless systems are infrastructure-based, but there are cases, such as disaster area management and vehicular communication, where ad hoc communication becomes a possible alternative. However, such networks are likely to be intermittently connected when deployed in a large geographic area, requiring the use of store-carry-forward techniques. Still, in a post-disaster scenario timely communication is vital for the relief efforts. It is needed for fast and efficient assessment of the situation and coordination of the rescue actions. Disaster area communication is therefore both mission-critical and time-critical. The same applies to car-to-car communication which is actively being developed mainly for safety-related applications.

For such applications it is of great importance to be able to analyse under what conditions a sufficient quality of service can be provided. In particular, for time-critical applications, an important aspect is to be able to guarantee message delivery within a certain amount of time. For example, a warning message might need to be propagated within some given time frame in order to be relevant. While simulation-based approaches can provide useful insights into performance under some typical scenarios, it is difficult to draw any general conclusions. As a complement to simulation, having an analytical model of the system allows deriving guarantees on system performance given that some basic assumptions are met.

Unfortunately, mobile ad hoc networks are inherently difficult to capture in analytical models. This is partly due to the fact that the actual node mobility varies considerably between application scenarios, and even worse, their characteristics are largely unknown to the research community. In addition, even with simple mobility models such as random waypoint (RWP), analytical models can become very complex. In reality, mobility is heterogeneous and complex, so we need a model that can incorporate such phenomena while still providing enough abstraction for analytical reasoning.

In this paper we address the issue of worst-case latency in intermittently connected networks for a wide class of mobility models. For this purpose, we provide an abstract description of node connectivity in a mobile network. This model taken together with generic properties of broadcast protocols will be combined to derive the worst-case latency in this setting. The novelty of our method lies in being able to analyse actual mobility traces rather than just theoretical models of mobility.

The work flow of our approach is shown in Figure 1. The basic idea is to be able to take some model of mobility and some broadcast protocol and be able to give an upper bound on the latency (given knowledge of the system load). We introduce the notion of *delay expansion* function ($D(\epsilon, s)$ in the figure) to capture the basic connectivity characteristics of a given mobility model. The delay expansion function can be derived from a trace file using an algorithm described in this paper. Using this function, and properties of a given protocol P ("Queue parameters" in the figure) we then create the *spread-time* function $T_P(x, y)$ which can be used to calculate the worst-case latency for spreading a message across partitions. Intuitively this term describes how long it takes for a protocol P to spread a message from x nodes to y nodes.

There are four contributions in this paper:

- A delay expansion function that characterises the connectivity of disconnected networks
- An algorithm to derive the delay expansion from mobility traces
- A proof of the upper bound on the worst-case latency for a class of broadcast algorithms with bounded queues
- Validation of the approach on an example protocol and mobility models including a real-life mobility trace

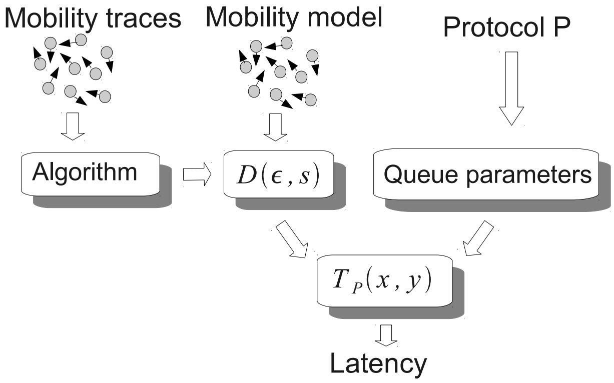

Figure 1 Overview of the approach

The rest of this paper is organised as follows. Section 2 briefly covers the graph-theoretical constructs that we use in the paper and also covers related work. Section 3 describes our system model and defines the delay expansion function that we use in the rest of the paper. Section 4 shows how to derive the actual worst-case latency of a protocol. In Section 5 we show how the delay expansion can be derived from a typical mobility model description and from real mobility traces. The theoretical model is validated using a network simulated in Section 6 with real-life mobility data. Finally, Section 7 concludes the paper.

2 Related Work

In this section we will relate our work to existing literature with regard to connectivity modelling and analytical derivations of message latency in intermittently connected networks.

2.1 Connectivity Modelling

Although, intermittently connected networks have been studied for quite some time (Davidson et al., 1985; Saito and Shapiro, 2005; Asplund et al., 2009), it was not until the introduction of delay-tolerant networks (DTN) that the problem of characterising mobile connectivity became an important problem. The large body of work on mobile ad hoc networks usually deals with connectivity in terms of link and path stability, path diversity etc. (Samar and Wicker, 2004; Yawut et al., 2008).

The connectivity of ad hoc networks has been extensively studied using percolation theory (Penrose and Pisztora, 1996; Santi and Blough, 2003), connected components (Dousse et al., 2006), temporal clustering coefficients (Tang et al., 2009) as well as other metrics (Ovalle-Martínez et al., 2005; Scellato et al., 2011).

A common assumption for analysing connectivity of delay-tolerant networks have been that the inter-meeting time is exponentially distributed, which has also been shown to be true for some common mobility models like random waypoint and random walk (Groenevelt et al., 2005). However recent work has shown that this is not true for real mobility traces (Chaintreau et al., 2007; Karagiannis et al., 2010), and that it in fact is not homogeneous, but heterogeneous and correlated (Ciullo et al., 2011; Hossmann et al., 2011; Passarella and Conti, 2011)

In this paper we will derive a global connectivity metric by considering the actual connectivity pattern of a system. Modelling dynamic connectivity as a graph can be done in two basic ways, often each node is represented by one vertex and there is one edge for every contact (Kempe and Kleinberg, 2002; Balasubramanian et al., 2007). The other alternative, which we use in this paper, is to let each node be represented by multiple vertices, one for each network configuration or time step (Merugu et al., 2004). Kong and Yeh Kong and Yeh (2008), use a concept of *long-term connectivity graph* whose edges are all the pairs of nodes with a finite expected meeting time.

Acer et al. (2010) propose a metric where a space-time connectivity graph is represented as a reachability tensor from which a specific graph measure is extracted. The authors show experimentally that this metric is correlated with the expected hitting time (time to reach a particular node for the first time) for a random walk in the network. However, no theoretical analysis is done to explain this correlation. Chen et al. (2011) present another classification of intermittently connected networks based on node-pair characteristics.

2.2 Latency in Intermittently Connected Networks

To the best of our knowledge, this is the first paper to provide a model for finding the worst-case broadcast latency in intermittently connected networks with arbitrary mobility and limited bandwidth. On a similar theme, Uddin et al. (2010) analyses the worst-case end-to-end deadlines of data flows for networks where the mobility is recurring. There is also a rich body of theoretical results on information dissemination in dynamic networks (O'Dell and Wattenhofer, 2005; Kuhn et al., 2010; Prakash et al., 2011), although most works assume connected networks. There are a number of results on the complexity of flooding for a class of connectivity models called edge-Markovian dynamic graphs where links appear and disappear according to a fixed probability (Clementi et al., 2010; Baumann et al., 2011). However, none of these works are able to model and reason about an actual mobility trace.

Other works on latency in intermittently connected networks can be broadly divided in three categories. (1) Graph exploration of a DTN graph to do optimal routing. (Xuan et al., 2003; Jain et al., 2004). Common for this and similar works (Merugu et al., 2004) is that the delay is calculated for one node at a given single time point, rather than as a general characteristic of the entire system. (2) Asymptotic best-case analyses using simple homogeneous mobility models (Grossglauser and Tse, 2002; Neely and Modiano, 2005). Although results from these studies provide upper bounds on achievable latencies in wireless networks, they do not help in analysing properties of specific networking algorithms for more general classes of mobility models. (3) Probability-based analyses for specific protocols where the inter-meeting time follows a given homogeneous and independent probability distribution (Altman et al., 2010; Balasubramanian et al., 2007; Groenevelt et al., 2005; Resta and Santi, 2012; Spyropoulos et al., 2009). Most of these concentrate on the expected delivery ratio, rather than the worst case.

3 Connectivity Models

In this section we will describe our system model and introduce the delay expansion that we later use to determine the worst-case latency.

3.1 System Model

We use a space-time graph model to describe the dynamic connectivity of the mobile network. The system is composed of a set of processes (using terminology from distributed systems). We will also use the term node synonymously with process. Given any two processes p and p' in the system with an uninterfered link between them, we denote by T_m the maximum transmission time for any message m in the network. A link (x, y) is interfered if some neighbour of y other than x is transmitting. We assume that there is a unique time for every event in the system. That is, no two messages are delivered at *exactly* the same time. This assumption allows more straightforward definitions relating to time and information spreading, but does not affect the length of the worst case spreading time due to the arbitrary interleavings of possible events. Moreover, the participating nodes have no information of this global time. Table 1 at the end of the paper summarises the symbols and notation we use throughout.

Formally we define the connectivity model C as a sequence of topologies G^i:

Definition 3.1: A *connectivity model* $C = \langle\langle G^0, T^0\rangle, \langle G^1, T^1\rangle, \ldots\rangle$ is a (possibly infinite) sequence of graphs $G^i = (V, E^i)$, $i \geq 0$, each graph representing the ith topology of the network and lasting over a duration $T^i \geq T_m$. The set V represents the the processes in the system, and E^i represents the links in the system in topology G^i.

A given mobility pattern together with a range of radio characteristics will result in a certain connectivity model. Thus, we consider node mobility in an abstract fashion by considering the connectivity pattern that results from that mobility. Figure 2 shows an example of a connectivity model for a system with six processes (a to f). There are four different topologies, G^0 to G^3. Note that while there might be topologies of shorter duration than T_m, we do not include these in the connectivity model in order to ensure that at least one message can be sent during the duration of a topology. This can also be seen as an underlying assumption on the interesting connectivity models, meaning that topologies with shorter connectivity than that needed for sending a message are considered as less relevant due to weak connectivity. We will also need a notion of a continuous time line with time points modelled as positive real numbers. This lets us define the start time of a given topology as follows:

$$start(G^i) = \sum_{0 \leq j \leq i-1} T^j$$

where T^j is the duration of the topology G^j (the start of the first topology thus evaluates to $start(G^0) = 0$).

We can now use this model to express connectivity strengths. For example, the majority partition assumption requires that for any topology G^i there is a set $M \subseteq V$ where $|M| \geq \lceil \frac{|V|}{2} \rceil$ so that for any $x, y \in M$ there is a path from x^i to y^i in G^i.

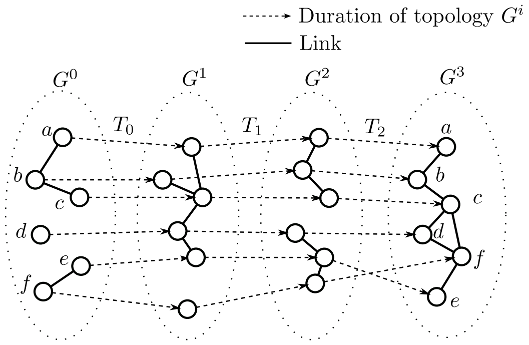

Figure 2 Example Model

Our focus in this paper is on a much weaker type of connectivity than majority which we call *space-time connected*. This model is quite weak, but there is a bound on the ability for each process to reach another either by a hop in the current topology or by waiting for another topology. This model is equivalent to network liveness (Vollset and Ezhilchelvan, 2005), which excludes the existence of permanent partitions.

Definition 3.2: A connectivity model $C = \langle \langle G^0, T^0 \rangle, \langle G^1, T^1 \rangle, \ldots \rangle$ is *space-time connected* if for any topology $G^i = (V, E^i)$ there is a finite sequence $G^i, \ldots, G^{i+n}$ of successive topologies starting with G^i, so that the graph $G = (V, E^i \cup \ldots \cup E^{i+n})$ is connected.

The model shown in Figure 2 is space-time connected since every topology, if merged with the successor topology, results in a connected graph.

In a space-time connected connectivity model it is possible to bound the delay of a dissemination protocol. Unfortunately, this definition is not very helpful in itself. First of all, given a certain physical system, how can one know the bound? Secondly, due to limited bandwidth and collisions a protocol will not be able to deliver a message within a bound computed from the duration of topologies just because there was a path. We will deal with these issues in the coming sections.

3.2 Delay expansion

We will now introduce the first of the main contributions in this paper, which is a way of expressing relevant properties of the connectivity model that is needed to derive worst-case latency bounds. Intuitively, in an intermittently connected network, latency is decided by the amount of "spreading" a message can perform in one time step. In a network that is well-connected, the spreading-rate is high, so the message will disseminate quickly. If on the other hand, the message is able to reach at most a few nodes in every step, the dissemination will be slow.

Graph expansion is a measure used in graph theory which intuitively tells how well-connected a graph is. For more details on graph expansion and related concepts we refer the reader to the excellent survey by Hoory et al. (2006). We will use the concept of expansion and add to it the notion of time delay. The time delay decides how long we must wait before the mobility model is able to generate a graph with the required expansion. Using this time and expansion factors we will later be able to derive bounds on the delivery time.

However, before introducing the delay expansion, we need to define a particular kind of graph expansion that suits our specific needs. For a given set of processes in our system we not only want to know how many of them have a neighbour to which they can send a message, but also how many such neighbours there are. This will be captured by the following definition:

Definition 3.3: Given a graph $G = (V, E)$, let $S \subseteq V$ be a non-empty subset of V and let $\overline{S} = V \setminus S$, then the *least neighbour expansion* of S is defined as:

$$e(G, S) = \min \left(\frac{|\Gamma(S) \setminus S|}{|S|}, \frac{|\Gamma(\overline{S}) \setminus \overline{S}|}{|S|} \right)$$

where $\Gamma(S)$ denotes the neighbour set of S (i.e. $\Gamma(S) = \{x | \exists y : (x, y) \in E, y \in S\}$).

Intuitively, this notation captures how many outside neighbours a given set of nodes have in a graph. We will later use this to capture how many new nodes can be reached with a message given that S is the set of nodes that have received the message so far.

Figure 3 shows a simple example of least neighbour expansion where the edges between nodes in S and outside S are drawn. Assuming that there are 10 members in the set S, the least neighbour expansion would be $e(G, S) = \min(3/10, 4/10) = 3/10$ since there are four border nodes in $\overline{S}$, 3 border nodes in S.

Figure 4 shows the result of averaging the least neighbour expansion for different set sizes in 1000 randomly generated connected graphs. Each graph contained 100 nodes and the node degree in the graphs were between 1 and 4. Although the least neighbour expansion exists

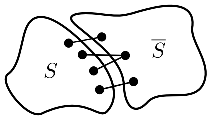

Figure 3 Least neighbour expansion

for all subsets of V, as we can see in the graph it will be very low for set sizes larger than set size 50 (i.e., larger than $|V|/2$). The explanation for this is that the set size is found in the denominator in the expression for $e(G, S)$. However, we do not need to use the least neighbour expansion for sets S with more than half the nodes since symmetry allows us to consider the expansion of the complement set $\overline{S}$ which can be expressed as $e(G, \overline{S}) = (e(G, S) \cdot |S|)/|\overline{S}|$.

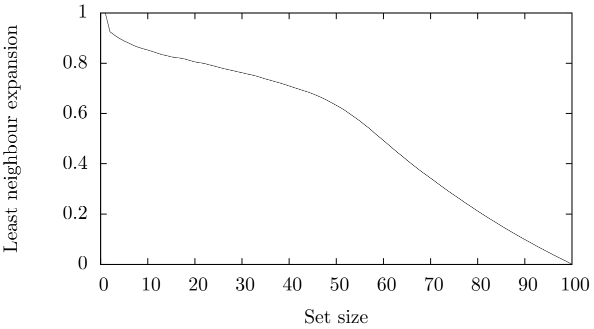

Figure 4 Least neighbour expansion vs. set size

Based on Definition 3.3, a topology that represents a partitioned network will have a least neighbour expansion of zero for at least some subset of its vertices S. However, if we consider the future encounters for nodes over an interval of time, and add an edge to a disconnected current topology for all future encounters of the nodes, we may well get a connected graph (and thereby non-zero least neighbour expansion). One can think of it as accumulating all the links that occur within this time in one graph. The following definition formally describes how such a graph is constructed from a sequence of topologies for a given time point and duration.

Definition 3.4: Given a connectivity model C, a time point t, a duration d, and a transmission time T_m, the *delay neighbourhood* $N(C, t, d)$ is a graph obtained from the union of topologies $G^i, \ldots, G^{i+n}$ within C such that $start(G^{i+1}) \geq t + T_m$ and $start(G^{i+n}) + T_m \leq t + d$.

Intuitively, this notion captures the potential one-hop reach of a network at a time point and within an interval d in the future. That is, if a message resides at a node x at time point t, then during an interval of length d it has the possibility to spread to a node y such that the edge $(x, y) \in E$ where E is the edge set of the delay neighbourhood $N(C, t, d)$. Note that the order in which the edges in E appear does not matter since we do not consider multiple hops in one delay neighbourhood.

The reason for the slightly convoluted time requirements is that we can only include the topologies that overlap enough with the specified time interval d. Figure 5 shows the idea, where the four topologies G^i to G^{i+3} will be included in $N(C, t, d)$ since there is enough time for a message to be sent in the first and last of the four topologies.

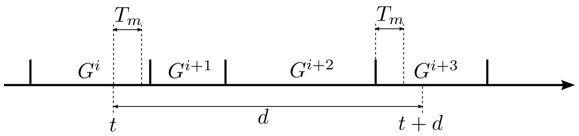

Figure 5 delay neighbourhood

We will now proceed with one of the more central definitions of the paper, where we combine the least neighbour expansion with the delay neighbourhood to define the *delay expansion* $D(\epsilon, s)$. This notion is a measure of the delay required to reach an expansion of ϵ for all subsets of size s. The intuition is that a connectivity model with a small delay expansion will take a small amount of time until the number of stable links that exist during this period will be enough to let a message spread to a given number of nodes.

Definition 3.5: Let C be a space-time connected model with node set V, ϵ be a positive real number, and s a positive integer. The *delay expansion* $D(\epsilon, s) = \min(\{d : \forall S \subseteq V, |S| = s, \forall t \geq 0, e(N(C, t, d), S) \geq \epsilon\})$ is the minimum delay d so that all delay neighbourhoods have a least neighbour expansion of ϵ for sets of size s.

Note that the definition allows for different ϵ for different set sizes. Moreover, the function $D(\epsilon, s)$ is defined for all space-time connected connectivity models and all $s \in [1, n)$ (n being the number of nodes) for the extreme case where $\epsilon = 1/n$, since the only requirement on the delay neighbourhood graph is that it is connected. We will now present a symmetry theorem that relates the delay expansion of a set size s with the delay expansion of $\overline{s} = n - s$, where n is the number of nodes. This will turn out to be useful since it means that we only need to know the delay expansion for all set sizes $s \leq n/2$.

Theorem 1: *Given any connectivity model C with n nodes for which $D(\epsilon, s)$ is defined, let $\overline{s} = n - s$ and $\overline{\epsilon} = \epsilon \cdot s/\overline{s}$, then $D(\overline{\epsilon}, \overline{s}) = D(\epsilon, s)$.*

The intuition behind this theorem is that the time taken for a small set S to encounter new neighbours outside S is the same time taken for the complement set $\overline{S}$ to encounter nodes in S. This is a consequence of how the delay expansion is defined. All proofs are provided in the appendix.

4 Deriving Worst-case Latency

So far, we have been concerned with the properties of the mobility and connectivity models of the system, and introduced the delay expansion as an abstract notion of connectivity. Now we turn to analysing protocols in this setting. We will first introduce the function $T_P(s, s')$ for capturing the actual time required for a broadcast protocol to reach a given number of nodes. We will show a couple of properties of this function and demonstrate how it behaves for an ideal spreading protocol. Then we will consider what happens in a system with limited bandwidth and send queues. Finally, we derive an analytical expression of the worst-case latency under certain assumptions.

4.1 Introducing the Spread-Time

We consider broadcast-type protocols whose purpose it is to spread a message to all nodes in the network. We begin by introducing some notation that will be used for describing protocol behaviour. We will use a style which is similar to that of timed I/O-automata. A process will be denoted as being either *informed* or *uninformed* of a given message. If a process performs the action *receive*(m) at time t, then it will be informed of m at all time points $t' \geq t$ (i.e., a node never ceases to be informed of a message). If a process x performs an action *send*(m) at time t, then any process y with an uninterfered link from x during the interval $[t, t + T_m]$ will perform the action *receive*(m) no later than time $t + T_m$. Finally, a *run* of a protocol P is a sequence of actions and times that follow the specification of the protocol (e.g. excluding spurious send and receive actions) in a given connectivity model.

We now introduce the spread-time function $T_P(s, s')$ that tells us the worst-case time taken for a network to go from s nodes being informed to s' nodes being informed (where $s' > s$) when running protocol P. Clearly, if we are able to find an expression for $T_P(s, s')$ then, in a system with n nodes, $T_P(1, n)$ will give us the worst-case broadcast latency of P. Before proceeding with the definition we introduce a help function $t(r, s, m)$ which denotes the first time-point of run r when at least s nodes are informed of message m.

Definition 4.1: Consider a protocol P, a connectivity model C, and two non-zero set sizes s and $s' > s$. Let $R \neq \emptyset$ be the set of all pairs $\langle r, m\rangle$, where r is a run of P and m is a message, such that $t(r, s, m)$ exists. The *worst-case spread time* from s to s' with protocol P is then defined as $T_P(s, s') = \max_{\langle r,m\rangle \in R} t(r, s', m) - t(r, s, m)$.

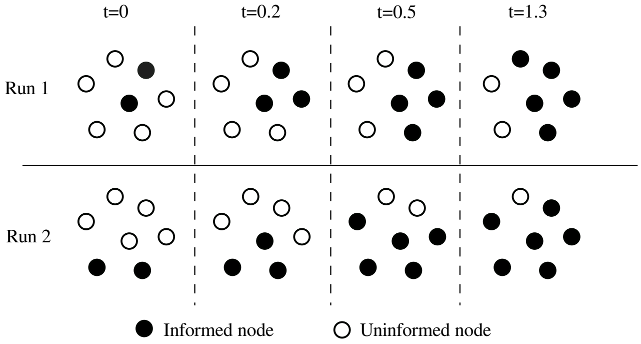

Figure 6 Spread time example

Figure 6 shows an example of a protocol P with two runs for a given connectivity model (illustrated with some specific time points). The spread time from 2 to 5 nodes is then $T_P(2, 5) = max\{1.3, 0.5\} = 1.3$.

Whether or not $T_P(s, s')$ is well-defined and finite depends on the protocol as well as the connectivity model. More specifically, whenever R is non-empty, the protocol-mobility pair has the *potential* to spread a message to s' nodes. We will in the coming subsections show some conditions under which the function is defined and where we can give upper bounds for it.

4.2 Properties of $T_P(s, s')$

Now we have come to the point where we can derive an expression of $T_P(s, s')$ based on knowledge of the delay expansion function. Let us first consider an ideal protocol I with a theorem that states that given a bound on $D(\epsilon, s)$ then the ideal protocol will spread to $\lceil \epsilon \cdot s \rceil$ more nodes in less than $D(\epsilon, s)$ time units. The ideal protocol has the following properties: (1) a *send*(m) action is always followed by a *receive*(m) action on neighbouring nodes (within T_m), even if the link is interfered (2) a node which is informed of message m at time t will perform a *send*(m) action at all times $t' > t$. This idealised description of a protocol can be useful in cases where the system load is considerably less than the available bandwidth so that queuing times can be neglected.

Theorem 2: *Given a space-time connected model C for which $D(\epsilon, s)$ is defined, and an ideal protocol I, for any integer $1 \leq k \leq \lceil \epsilon \cdot s \rceil$, $T_I(s, s + k) \leq D(\epsilon, s)$.*

Recall that we are interested in finding an expression for $T_P(1, n)$ since this is the latency for a protocol to spread a message from one node to all nodes in the system. However, the above theorem only gives the spread-time for spreading in smaller increments (i.e. smaller than $\lceil \epsilon \cdot s \rceil$). We will now show how to derive $T_P(1, n)$ given that we have expressed $T_P(s, s')$. The intuition here is that we should be able to simply add the durations. That is, the time taken to reach n nodes is the time taken to first reach $s < n$ nodes plus the time taken to reach n nodes starting from s informed nodes.

Theorem 3: *For any protocol P and any connectivity model C such that $T_P(s_1, s_3)$ is defined, we have $T_P(s_1, s_3) \leq T_P(s_1, s_2) + T_P(s_2, s_3)$, for all s_2 such that $1 \leq s_1 < s_2 < s_3 \leq |V|$.*

By combining Theorems 2 and 3, it is possible to derive an expression for the worst-case latency $T_I(1, n)$.

However, we will instead turn the focus to less idealised protocols that account for limited bandwidth and then derive the worst-case latency for those protocols in Section 4.4.

4.3 Dealing with Limited Bandwidth and Queues

In order to account for the effects of limited bandwidth and interference, we need to model protocols which are less powerful and more realistic than the ideal protocol. We call this class of protocols Q-b-fair. The intuition of a Q-b-fair protocol is that (1) for each message there is a bound Q on the number of other messages that will be sent before m is sent to an uninformed node, and (2) the minimal number of messages that can be transmitted over a link is b (i.e., this is an abstract notion of bandwidth). We will not provide an actual algorithm for such a protocol but rather give their specification. We will also argue that this specification can be met by a real implementation (at least under fault-free conditions).

Every time two nodes meet for a sufficient amount of time, there is an opportunity to exchange messages, we say that there is a number of *send opportunities.* However, since there are potentially many messages in the queue that need to be sent, not all messages can take advantage of a given send opportunity. One can think of Q as the worst-case length of a virtual global queue. Moreover, every node will have at least b send opportunities (i.e. the possibility to send b messages) at every meeting. By meeting we here mean that two nodes are in contact while a topology lasts. We also assume that no message is ever lost in transit. We now state these properties formally:

Definition 4.2: A protocol P is said to be Q-b-fair if the following holds:

1. If there is a topology G^i where x is informed of a message m at time $start(G^i)$, y is uninformed of m at time $start(G^i)$ for $(x, y) \in E^i$, then x will have at least b *send opportunities* for m during the interval $[start(G^i), start(G^{i+1}) - T_m]$

2. If a message m has had a sequence of $Q \geq 1$ consecutive send opportunities, then a $send(m)$ action is performed at the time of at least one of those opportunities.

3. If process x performs a $send(m)$ action at time $t \in [start(G^i), start(G^{i+1}) - T_m]$ and $(x, y) \in E^i$, then process y will perform $receive(m)$ no later than time $t + T_m$.

In order to successfully implement a Q-b-fair protocol, the protocol running at a node will need to send b messages at every meeting. One would therefore need to have a MAC protocol that guarantees each node at least some time to communicate with one of its neighbours during every topology. Most wireless medium access protocols such as 802.11 do not provide such real-time guarantees. However, this is possible to achieve, at least probabilistically (Yang and Kravets, 2006). Naturally, a protocol would also need to have a neighbourhood discovery protocol to keep track of its current neighbours.

In practice, we can approximate most broadcast protocols to be Q-b-fair. The parameter b is simply the number of messages that can be exchanged in a meeting between two nodes. One way of providing a probabilistic bound on Q is to let the queuing order at a node be randomised (Haas et al., 2006). Thus, at every meeting there is a certain probability of the message being sent; and given enough such meetings the probability of a successful send will be high. More detailed analysis on how to appropriately determine Q can be performed but is outside the scope of this paper.

Theorem 4 will now give a bound on the time it takes for a Q-b-fair protocol to inform k additional nodes starting from s informed nodes. We will assume here that the delay expansion time is less than a bound D for all set sizes between s and $s + k$. We will call the act of informing one new uninformed node a *successful send.*

Theorem 4: *Let D and ϵ be positive real numbers and s, k integers where $1 \leq k \leq \lceil \epsilon s \rceil$. Given a Q-b-fair protocol P, and a space-time connected model C for which $D(\epsilon, s') \leq D$ for all $s' \in [s, s + k]$, the spread time from s informed nodes to k new nodes $T_P(s, s + k)$, is bounded by $\left\lceil \frac{Q}{b} \right\rceil \cdot D$.*

Now we have all the tools we need to actually derive an expression for the worst-case latency for a Q-b-fair protocol and some knowledge of the delay expansion of a given connectivity model.

4.4 Worst-case Latency of Broadcast

We are now at the stage where we can calculate the bound on worst-case latency from one node to all n nodes using the sum $T_P(1, n) \leq T_P(1, x_1) + \ldots + T_P(x_n - 1, n)$. In order to find the appropriate values for x_i, we can use the following iterative formula until $x_i = n$:

$$T_P(1, x_i) \leq T_P(1, x_{i-1}) + T_P(x_{i-1}, x_i) \qquad (1)$$

Where the next x_i is calculated as $x_i = \min(n/2, \lceil x_{i-1} \cdot \epsilon \rceil)$ if $x_{i-1} < n/2$ and as the biggest set size x_i so that $x_{i-1} + \lceil \epsilon(n - x_i + 1) \rceil \geq x_i$ otherwise. The size of each new term can then be calculated using the appropriate theorem (i.e., Theorems 2 or 4).

However, for convenience, it is sometimes worthwhile to have a single closed-form expression, which we now proceed to present as a theorem for the case when the least neighbour expansion $\epsilon = 1/u$, $u \in \mathbb{Z}^+$. Considering Figure 4 it does not seem unreasonable to consider a single bound ϵ for all set sizes less than half the number of nodes. For the cases where these assumptions are not appropriate we suggest using the iterative approach.

Theorem 5: *Let D be a positive real number and n, Q, b, u be positive integers. Given a network of size n, a Q-b-fair protocol P (with $Q \geq b$), and a space-time connected model C where $D(1/u, s) \leq D$ for all set sizes $s \leq n/2$, the worst-case latency of broadcast for P is:*

$$T_P(1, n) \leq 2 \left\lceil \frac{Q}{b} \right\rceil \left(\left\lceil \log_2 \frac{n}{2u} \right\rceil + 1 \right) \cdot D \cdot u$$

As an example, consider a system with the following parameters: maximum global queue length $Q = 100$, minimum number of send opportunities on a stable link $b = 50$, number of nodes $n = 50$, minimum least neighbour expansion $\epsilon = 1/u = 1/3$, maximum delay to get the required expansion $D = 5[s]$ (where [s] denotes seconds). Then the worst-case broadcast latency of any message sent with a Q-b-fair protocol will be:

$$T_P(1, 50) \leq 300[s]$$

This expression provides a quick method of getting a bound on latency, and lets us see how the different parameters affect the final result. However, it is less exact than the iterative formula (1), which for the same example results in $T_P(1, 50) \leq 180[s]$.

5 Deriving the Connectivity Model

So far we have introduced the delay expansion $D(\epsilon, s)$ as a theoretical concept to capture mobility and connectivity properties of the nodes. We now consider how to actually find this function. In this section we show how this can be done using real mobility traces. Such traces could for example be obtained by tracking node movements using GPS devices in a real scenario or exercise, or by using a mobility trace generator. A trace is typically composed of a large number of time-stamped records that specifies the location and possibly the velocity for a given node at that time point. Using such a description, one can then extract a sequence of topologies (i.e., a connectivity model) from which the delay expansion can be calculated.

Unfortunately, finding the maximal expansion of a graph is a hard problem. This leaves us with two options, either we try to find a suboptimal bound that is safe (i.e., not optimistic) but that might provide very pessimistic results, or we try to find a way to approximate the expansion and risk the possibility that the result is optimistic. The former of these is appropriate if we want a theoretically correct bound on the latency and the latter if we are interested in an approximate figure that is closer to reality.

We will start by exploring the first option by finding a safe bound on the expansion (and thereby, on the latency) in the rest of this section. We have also tried the second option as explained in Section 6. The presentation of the method to obtain a pessimistic bound will rely on the fact that there is a close relationship between least neighbour expansion and the *all-multicommodity flow problem*. Therefore, in Section 5.0.1, we will first introduce the concept of flows and describe the all-multicommodity flow problem, followed by its application in Section 5.0.2.

5.0.1 Multicommodity Flows

Given a graph $G = (V, E)$, a source s, and a sink t a *flow* is a mapping $f : V \times V \to \mathbb{R}$ such that (1) for all $x, y \in V$, $f(x, y) = -f(y, x)$, (2) for all $x \in V \setminus \{s, t\}$, $\sum_{y \in V} f(x, y) = 0$, and (3) if $(x, y) \notin E$ then $f(x, y) = 0$. One can think of flows as commodities that are to be shipped from one node to another through the network. Using this analogy, the all-multicommodity flow problem states that every node wants to ship some distinct amount of commodity to every other node in the network, but where each link can handle only so much flow at the same time. A little more formally we can express this as (adapted from (Hoory et al., 2006)):

Definition 5.1: Given an n-vertex input graph $G = (V, E)$, an *all-multicommodity flow assignment* of size δ is a set of $n \cdot (n - 1)/2$ distinct flows, one for each pair of nodes (source and sink) such that:

- For every flow i, $\sum_{y \in V} f_i(s_i, y) = \delta$, where s_i is the source of the distinct flow, and δ is the flow size.
- Every edge has a maximum capacity of 1, so for all edges $(x, y) \in E$, $\sum_i |f_i(x, y)| \leq 1$

The max all-multicommodity flow problem is to find the largest possible flow size δ. The solution to this problem provides the total capacity of the network assuming that each node wants to ship something to every other node. The multi-commodity flow problem has indeed been used to find the capacity of wireless (and other) networks (e.g. (Garetto et al., 2007)). However, our reason for using this concept is not the capacity itself, but rather the connection with graph expansion (Hoory et al., 2006), as will be shown in the next section.

5.0.2 Finding the Delay Expansion from a Trace File

We will now use an all-multicommodity flow assignment to bound the least neighbour expansion of a graph. We present a theorem which provides a bound on the least neighbour expansion given some flow assignment.

Theorem 6: *Given a graph $G = (V, E)$, $n = |V|$ and an all-multicommodity flow assignment of size δ (according to Definition 5.1), for every node x, let F_x be the number of flows such that for some y: $|f_i(x, y)| > 0$ and let $F_{max} = \max(F_x)$. Then for every set $S \subset V$, $|S| = s \leq n/2$, the least neighbour expansion is:*

$$e(G, S) \geq \frac{n - s}{F_{max}}$$

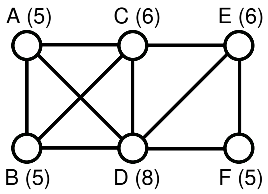

Flows:		
A-**B**	**B**-**D**	**C**-**E**
A-**C**	**B**-**C**	**C**-E-**F**
A-**D**	**B**-D-**E**	**D**-**E**
A-C-**E**	**B**-D-**F**	**D**-**F**
A-D-**F**	**C**-**D**	**E**-**F**

Figure 7 Flow assignment example, each node is marked with the number of flows going through it. The table shows the nodes that each flow passes through, with the sink and source marked as bold.

To illustrate the application of this theorem consider the graph G in Figure 7. There are $n = 6$ nodes, so there are $n(n-1)/2 = 15$ flows in the graph, one for each pair of nodes. In the table, a possible flow assignment is shown and the resulting number of flows going through each node (denoted F_x in Theorem 6) is shown in parentheses beside each node. The maximum number of flows going through any node is $F_{max} = \max\{F_x\} = F_D = 8$. We can now apply the theorem to see that the least neighbourhood expansion for all subsets S of size 3 must be:

$$e(G,S) \geq \frac{n-s}{F_{max}} = \frac{3}{8}$$

This means that any subset of size 3 must have at least $\lceil 3 \cdot 3/8 \rceil = \lceil 9/8 \rceil = 2$ neighbours outside the set itself. Similarly, a subset of size 2 will have a least neighbour expansion $e(G,S) = 1/2$, so the nodes in set will have at least one neighbour outside the set.

We will now use Theorem 6 to construct an algorithm as seen in Listing 1 that will in polynomial time heuristically find the delay D and expansion ϵ such that $D(\epsilon, s) \leq D$. The input to the algorithm is a space-time connected model C as defined in Section 3.1, Definition 3.2, and, a minimum link time T_m, and a set size s. In order to get the connectivity model C from a trace file, one needs to extract all contacts (e.g. using a tool such as cbm (Khelil et al., 2005)) and remove all links that are too short.

The algorithm works in two steps. The first step (lines 1 to 7) will find the worst-case delay D required to achieve delay neighbourhood graphs $N(C, start(G^i), D)$ so that any subset of size s is guaranteed to have external neighbours (lines 4-6). Finally, the resulting D is found by taking the maximum (i.e., worst) of all D_i (line 7). We add an extra T_m at the end since we will consider only the delay neighbourhood at the start time of each topology.

In the second step (lines 8-13), a bound on the least neighbour expansion ϵ is calculated so that $D(\epsilon, s) \leq D$ for all values of s. Again, we iterate over all topologies of C and create the delay neighbourhood $N(C, start(G^i), D_i)$ (line 9). We then create a flow assignment using the CreateFlowAssignment function (line 10) which can be any polynomial-time heuristic algorithm that provides an all-multicommodity flow assignment (we have used a greedy algorithm when testing the algorithm). The number of flows going through each node is calculated (line 11), and by applying Theorem 6 the resulting least neighbour expansion is found (line 12). Finally, we take the minimum of all the ϵ_i to get the worst-case bound (line 13).

Listing 1 GetDelayExpansion

```
Input:   Space-time connected model $C$,
         minimum link duration $T_m$
         set size $s$
Output:  Delay $D$, least neighbour expansion bound $\epsilon$

GetDelayExpansion
       //Get delay:
 1     for every topology $G^i \in C$:
 2         $D_i \leftarrow 0$
 3         $N \leftarrow$ the $N(C, start(G^i), D_i)$
 4         while a subset of size $s$ might be disconnected:
 5                 $D_i \leftarrow D + T_m$
 6                 $N \leftarrow N(C, start(G^i), D_i)$
 7     $D \leftarrow max\{D_i\} + T_m$

       //Get expansion:
 8     for every topology $G^i \in C$:
 9         $N \leftarrow N(C, start(G^i), D_i)$
10         CreateFlowAssignment($N$)
11         $F_x \leftarrow$ the number of flows passing node $x$
12         $\epsilon_i \leftarrow (n - s)/(\max\{F_x\})$
13     $\epsilon \leftarrow \min\{\epsilon_i\}$

14     return $D$, $\epsilon$
```

Due to Theorem 6, we know that the algorithm gives a proper least-neighbour expansion for every delay neighbourhood. And since we go through every topology, we will get the worst-case delay required to achieve this expansion.

To summarise this section, we have provided an algorithm that takes as input a connectivity model (e.g., a mobility trace file), a link duration and a set size s and it will then return D and ϵ so that the delay expansion $D(\epsilon, s)$ is less than or equal to the bound D. By running the algorithm for all set sizes s up to half the number of nodes, and using the methods explained in Section 4.4, these numbers can then be used to derive the worst-case latency of broadcast for this particular trace.

6 Validation

In order to assess the validity of our assumptions and to put the latency results in perspective we have performed some simulation-based studies. The main result of the paper is an upper bound on the worst-case latency for broadcast in intermittently connected networks (as presented in Section 4). In this section we will demonstrate that the upper bound is meaningful, and the intermediate tools to compute the upper bound can be used to derive a useful approximation of the upper bound in practice.

We will start by demonstrating that the bound computed according to the method described in Section 4 is tight. That is, we can give an example of a node movement trace where the resulting broadcast latency is equal to the upper bound given by our calculations. We created a mobility trace where 8 nodes were positioned in a circle with a radius of 100 meters. Every 200 seconds, a number of pairwise node meetings were created by letting the nodes move to the centre of the circle. The pattern with which the nodes met formed a connected graph with least neighbour expansion $1, 1, 2/3, 1/2$ for set sizes $1, 2, 3, 4$ respectively. Figure 8 shows the resulting latency for a protocol with parameters $Q = 2, b = 1$ (i.e., maximum queue length of 2). The simulated curve was produced using the network simulator One (version 1.4) (Keränen et al., 2009) with a ideal broadcast protocol, and the theoretical upper bound using the iterative approach described in Section 4.4, equation (1). We can see that the two curves follow each other exactly, which demonstrates that there is no way of getting a less pessimistic bound in the general case.

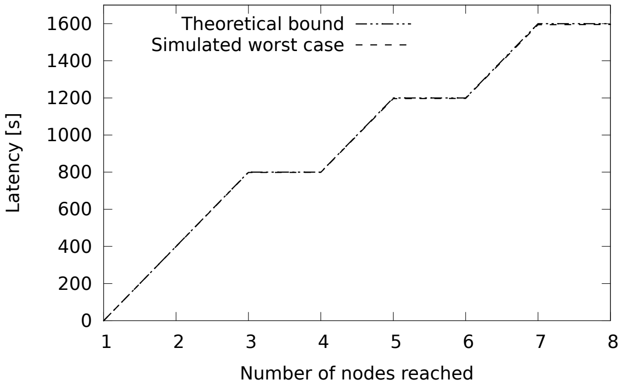

Figure 8 Tightness of upper bound on worst case broadcast latency

Next, we move on to a larger and more realistic model, with 50 nodes moving in an area of $500m \times 500m$. The mobility model was a location-based model intended to mimic a disaster scenario where people cluster around certain places but with some movement between the locations as well. There were 20 locations, and 20 of the nodes were fixed at one of the locations. The 30 mobile nodes moved at walking speed (1-2 m/s). Each node moved in a circular area with a radius of 10 meters using the random waypoint model around one of the 20 locations location. Every 40 seconds a mobile node will change to a new location to which it will travel at vehicular speed (36-90 km/h). The transmission range of the devices was 20 meters, and the simulation lasted for 2000 seconds. The traffic was generated by randomly choosing a sender among all the nodes every 10-100 seconds. Again the protocol is modelled in an abstract manner by determining the Q and b parameters. In the simulations the parameters were $Q = 2$ and $b = 1$ (thus, we let a message queue for 10 seconds before being sent during 10 seconds).

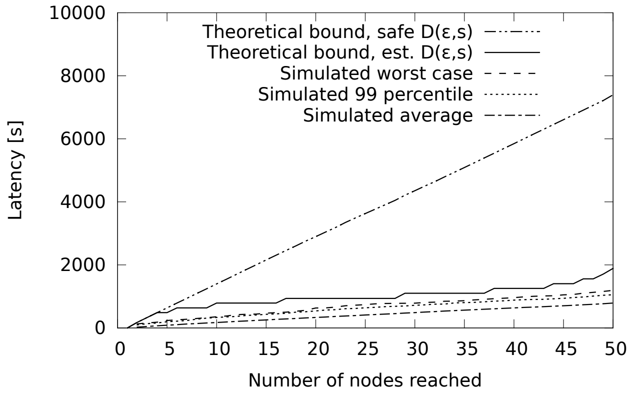

Figure 9 Latency for disaster area scenario

Figure 9 shows the worst-case latency as a function of the number of nodes reached. The lowest three curves show the actual simulation-based latency based on 50 runs with the worst case, the 99 percentile and the average.

The middle curve shows a theoretical bound of the worst-case latency. This is based on the results of this paper together with an estimation of the delay expansion. The estimation of the delay-expansion was done by using the *Get delay* part of the algorithm Listing 1, and the following procedure for calculating ϵ. We first randomly created 750 connected subsets, and then calculated the least neighbour expansion for each subset using Definition 3.3. Finally, we let ϵ be the worst (smallest) of these values. The result is fairly close to the simulated one, and provides a practical upper bound on the latency for a given mobility.

Finally, the upper curve in the figure represents a safe upper bound on the latency as dictated by equation (1) in Section 4. The difference to the previous curve is that instead of trying to find a good estimate of $D(\epsilon, s)$ from the trace file, we use the algorithm Listing 1 to get a safe upper bound on the delay expansion. We then use the same calculations as for the middle curve to derive the bound on worst-case latency. As we can see, as opposed to the result in Figure 8, the bound is very pessimistic in this case. The for this is that the worst-case delay expansion can be very high during some time periods in the mobility trace file. While this bound is guaranteed to be safe, the estimated worst-case is perhaps the most useful outcome of the work in practical scenarios.

Finally, the last mobility trace we have analysed is a large scale real-life trace based on the movement of taxis in the San Francisco area. The trace was collected by Piorkowski et al. (2009) based on data made available by the cabspotting project during May 2008. We used a subset of 272 cars that were in contact with some other node in this group at least every 15 minutes. We assume that each taxi has a wireless device with a range of 300 meters, and simulated 10,000 seconds with each node sending a packet every 10-500 seconds. Figure 10 shows the results for the simulation (worst-case, 99 percentile and average) at the bottom and the theoretical bound at the top. We estimate $D(\epsilon, s)$ in the same way as for the disaster scenario, but calculating a safe bound turns out to be very time consuming for such a large scale network. This further confirms our assessment that the estimation-based approach is the more practically useful application of our results.

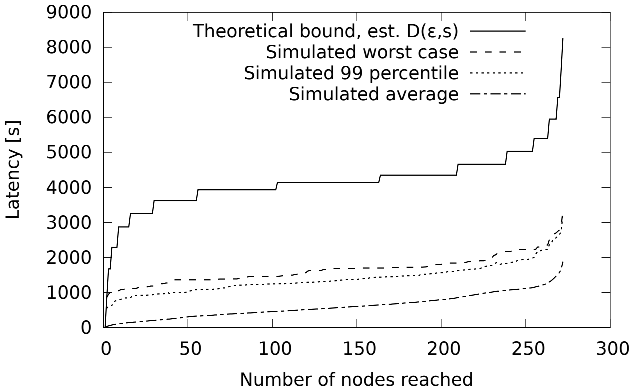

Figure 10 Latency for San Francisco cab scenario

7 Conclusions

In this paper we have presented an analytical method for deriving the worst-case latency of broadcast in intermittently connected networks. A key strength of the approach is that it is applicable to any kind of mobility movement as opposed to other proposed schemes that usually assume node movement to be homogeneous. Another benefit is that we decouple the analysis of the connectivity model (which results in a more abstract representation in the form of a delay expansion function) from the properties of the protocol. Moreover, we have shown that bounds on the delay expansion can be derived from an arbitrary trace file of the mobility movements.

We believe that our method can be useful to protocol and application designers in determining the envelope in which the application is *guaranteed* to provide the desired quality of service. Using this approach the application designer can state that if the connectivity of the network is well behaved (i.e., with a delay expansion less than some ϵ), and the system load is so that the queue length does not exceed Q, then the protocol is guaranteed to deliver a message within X seconds. The key novelty is that we provide a metric to characterise dynamic and intermittently connected networks that can actually be extracted from mobility traces.

Naturally, there are some limitations. First of all, the delay expansion function does not rule out adversarial mobility movements. This means that we must consider some very improbable cases to really get the worst-case latency. Although this is exactly what we intended (otherwise one would perform an average-case analysis), one might be able to find tighter bounds for a given mobility model.

Moreover, we have assumed a medium access control layer that can provide some guarantees on message delivery over a stable link. As future work, it would be interesting to have a more fine-grained model of interference and allocation of the medium, as well as other ways to model message queues. This could also lead to more detailed models of different protocols, thereby allowing this method to help in determining the suitability of a given protocol to some given mobility scenario.

We believe that there are many ways that one can build upon the theoretical framework we have presented in this paper. The delay expansion concept is well suited for quickly characterising node mobility. We have already found the concept interesting for application to probabilistic analyses, providing an alternative to the common uniform inter-meeting time assumption (Asplund and Nadjm-Tehrani, 2012).

Acknowledgements

This work was supported by the Swedish Research Council (VR) grant 2008-4667. During the final stages of preparing the manuscript the first author was supported, in part, by Science Foundation Ireland grant 10/CE/I1855.

Table 1 Notation Summary

Symbol	**Description**
T_m	Maximum transmission time of a message
V	Set of nodes in the system
S	Subset of nodes ($S \subseteq V$)
$\overline{S}$	Complementary set ($\overline{S} = V \setminus S$)
C	Connectivity model (Definition 3.1)
G^i	Topology graph (Definition 3.1)
E^i	Set of links in topology G^i (Definition 3.1)
T^i	Duration of topology i (Definition 3.1)
$e(G, S)$	Least neighbour expansion (Definition 3.3)
ϵ	Bound on least neighbour expansion
$N(C, t, d)$	Delay neighbourhood (Definition 3.4)
$D(\epsilon, s)$	Delay expansion (Definition 3.5)
$T_P(s, s')$	Worst case spread time (Definition 4.1)
Q	Bound on queue length (Section 4.3)
b	Bandwidth (Section 4.3)

References

U. G. Acer, P. Drineas, and A. A. Abouzeid. Random walks in time-graphs. In *Proc. The Second International Workshop on Mobile Opportunistic Networking (MobiOpp)*. ACM, 2010. doi: 10.1145/1755743.1755761.

E. Altman, T. Basar, and F. D. Pellegrini. Optimal monotone forwarding policies in delay tolerant mobile ad-hoc networks. *Perform. Eval.*, 67(4):299–317, 2010. ISSN 0166-5316. doi: 10.1016/j.peva.2009.09.001.

M. Asplund and S. Nadjm-Tehrani. Analysing delay-tolerant networks with correlated mobility. In X.-Y. Li, S. Papavassiliou, and S. Ruehrup, editors, *11th International Conference on Ad-Hoc Networks and Wireless (AdHocNow)*, volume 7363 of *LNCS*. Springer, 2012.

M. Asplund, S. Nadjm-Tehrani, and K. Zagar. Middleware extensions that trade consistency for availability. *Concurr. Comput.-Pract. Exp.*, 21(9):1181–1203, June 2009. doi: 10.1002/cpe.1383.

A. Balasubramanian, B. Levine, and A. Venkataramani. DTN routing as a resource allocation problem. *SIGCOMM Comput. Commun. Rev.*, 37(4):373–384, 2007. ISSN 0146-4833. doi: 10.1145/1282427.1282422.

H. Baumann, P. Crescenzi, and P. Fraigniaud. Parsimonious flooding in dynamic graphs. *Distributed Computing*, pages 1–14, 2011. ISSN 0178-2770. doi: 10.1007/s00446-011-0133-9.

A. Chaintreau, P. Hui, J. Crowcroft, C. Diot, R. Gass, and J. Scott. Impact of human mobility on opportunistic forwarding algorithms. *IEEE Trans. Mobile Comput.*, 6(6):606–620, June 2007. doi: 10.1109/TMC.2007.1060.

Y. Chen, V. Borrel, M. H. Ammar, and E. W. Zegura. A framework for characterizing the wireless and mobile network continuum. *Computer Communication Review*, 41(1):5–13, 2011.

D. Ciullo, V. Martina, M. Garetto, and E. Leonardi. Impact of correlated mobility on delay-throughput performance in mobile ad hoc networks. *Networking, IEEE/ACM Transactions on*, 19(6):1745 –1758, dec. 2011. ISSN 1063-6692. doi: 10.1109/TNET.2011.2140128.

A. E. F. Clementi, C. Macci, A. Monti, F. Pasquale, and R. Silvestri. Flooding time of edge-Markovian evolving graphs. *SIAM J. Discrete Math.*, 24(4):1694–1712, 2010. ISSN 0895-4801. doi: 10.1137/090756053.

S. B. Davidson, H. Garcia-Molina, and D. Skeen. Consistency in a partitioned network: a survey. *ACM Comput. Surv.*, 17(3):341–370, 1985. ISSN 0360-0300. doi: 10.1145/5505.5508.

O. Dousse, M. Franceschetti, and P. Thiran. A case for partial connectivity in large wireless multi-hop networks. In *Proc. UCSD - ITA workshop, San Diego CA*, 2006.

M. Garetto, P. Giaccone, and E. Leonardi. Capacity scaling in delay tolerant networks with heterogeneous mobile nodes. In *Proc. 8th ACM international symposium on Mobile ad hoc networking and computing (MobiHoc)*, pages 41–50. ACM, 2007. ISBN 978-1-59593-684-4. doi: 10.1145/1288107.1288114.

R. Groenevelt, P. Nain, and G. Koole. The message delay in mobile ad hoc networks. *Perform. Eval.*, 62(1-4):210–228, 2005. ISSN 0166-5316. doi: 10.1016/j.peva.2005.07.018.

M. Grossglauser and D. Tse. Mobility increases the capacity of ad hoc wireless networks. *IEEE/ACM Trans. Netw.*, 10(4):477–486, Aug. 2002. doi: 10.1109/TNET.2002.801403.

Z. Haas, J. Halpern, and L. Li. Gossip-based ad hoc routing. *IEEE/ACM Trans. Netw.*, 14(3):479–491, June 2006. doi: 10.1109/TNET.2006.876186.

S. Hoory, N. Linial, and A. Wigderson. Expander graphs and their applications. *Bull. Amer. Math. Soc.*, 43: 4319–561, 2006.

T. Hossmann, T. Spyropoulos, and F. Legendre. Putting contacts into context: Mobility modeling beyond intercontact times. In *Twelfth ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc 11)*, Paris, France, May 2011. ACM.

S. Jain, K. Fall, and R. Patra. Routing in a delay tolerant network. *SIGCOMM Comput. Commun. Rev.*, 34(4):145–158, 2004. ISSN 0146-4833. doi: 10.1145/1030194.1015484.

T. Karagiannis, J.-Y. L. Boudec, and M. Vojnović. Power law and exponential decay of intercontact times between mobile devices. *IEEE Transactions on Mobile Computing*, 9:1377–1390, 2010. ISSN 1536-1233. doi: 10.1109/TMC.2010.99.

D. Kempe and J. Kleinberg. Protocols and impossibility results for gossip-based communication mechanisms. In *Proc. 43rd Annual IEEE Symposium on Foundations of Computer Science (SFCS)*, pages 471–480, Nov. 2002. doi: 10.1109/SFCS.2002.1181971.

A. Keränen, J. Ott, and T. Kärkkäinen. The ONE simulator for dtn protocol evaluation. In *Proceedings of the 2nd International Conference on Simulation Tools and Techniques*, Simutools '09, pages 55:1–55:10, ICST, Brussels, Belgium, Belgium, 2009. ICST. ISBN 978-963-9799-45-5. doi: 10.4108/ICST.SIMUTOOLS2009.5674.

A. Khelil, P. Marron, and K. Rothermel. Contact-based mobility metrics for delay-tolerant ad hoc networking. In *Proc. 13th IEEE International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems*, pages 435–444, 27–29 Sept. 2005. doi: 10.1109/MASCOTS.2005.20.

Z. Kong and E. M. Yeh. On the latency for information dissemination in mobile wireless networks. In *Proc. 9th ACM international symposium on Mobile ad hoc networking and computing (MobiHoc)*, pages 139–148. ACM, 2008. ISBN 978-1-60558-073-9. doi: 10.1145/1374618.1374638.

F. Kuhn, N. Lynch, and R. Oshman. Distributed computation in dynamic networks. In *Proceedings of the 42nd ACM symposium on Theory of computing*, STOC '10, pages 513–522, New York, NY, USA, 2010. ACM. ISBN 978-1-4503-0050-6. doi: 10.1145/1806689.1806760.

S. Merugu, M. Ammar, and E. Zegura. Routing in space and time in networks with predictable mobility. Technical Report GIT-CC-04-07, College of Computing, Georgia Institute of Technology, 2004.

M. Neely and E. Modiano. Capacity and delay tradeoffs for ad hoc mobile networks. *IEEE Trans. Inf. Theory*, 51(6):1917–1937, June 2005. doi: 10.1109/TIT.2005.847717.

R. O'Dell and R. Wattenhofer. Information dissemination in highly dynamic graphs. In *Proceedings of the 2005 joint workshop on Foundations of mobile computing*, DIALM-POMC '05, pages 104–110, New York, NY, USA, 2005. ACM. ISBN 1-59593-092-2. doi: 10.1145/1080810.1080828.

F. J. Ovalle-Martínez, I. Stojmenović, F. García-Nocetti, and J. Solano-González. Finding minimum transmission radii for preserving connectivity and constructing minimal spanning trees in ad hoc and sensor networks. *J. Parallel Distrib. Comput.*, 65(2):132 – 141, 2005. ISSN 0743-7315. doi: 10.1016/j.jpdc.2004.09.010.

A. Passarella and M. Conti. Characterising aggregate inter-contact times in heterogeneous opportunistic networks. In J. Domingo-Pascual, P. Manzoni, S. Palazzo, A. Pont, and C. Scoglio, editors, *NETWORKING 2011*, volume 6641 of *Lecture Notes in Computer Science*, pages 301–313. Springer Berlin / Heidelberg, 2011. ISBN 978-3-642-20797-6. doi: 10.1007/978-3-642-20798-3_23.

M. D. Penrose and A. Pisztora. Large deviations for discrete and continuous percolation. *Adv. Appl. Probab.*, 28(1):29–52, 1996.

M. Piorkowski, N. Sarafijanovic-Djukic, and M. Grossglauser. A parsimonious model of mobile partitioned networks with clustering. In *First International Conference on Communication Systems and Networks (COMSNETS)*, pages 1 –10. IEEE, jan. 2009. doi: 10.1109/COMSNETS.2009.4808865.

R. Prakash, Y. Sasson, M. Mohsin, D. Cavin, and A. Schiper. The impact of mobility on the time complexity for deterministic broadcasting in radio networks. *Int. J. of Ad Hoc and Ubiquitous Computing*, 8(3), 2011. doi: 10.1504/11.42351.

G. Resta and P. Santi. A framework for routing performance analysis in delay tolerant networks with application to noncooperative networks. *Parallel and Distributed Systems, IEEE Transactions on*, 23(1):2 –10, jan. 2012. ISSN 1045-9219. doi: 10.1109/TPDS.2011.99.

Y. Saito and M. Shapiro. Optimistic replication. *ACM Comput. Surv.*, 37(1):42–81, 2005. ISSN 0360-0300. doi: 10.1145/1057977.1057980.

P. Samar and S. B. Wicker. On the behavior of communication links of a node in a multi-hop mobile environment. In *Proc. 5th ACM international symposium on Mobile ad hoc networking and computing (MobiHoc)*, pages 145–156. ACM, 2004. ISBN 1-58113-849-0. doi: 10.1145/989459.989478.

P. Santi and D. Blough. The critical transmitting range for connectivity in sparse wireless ad hoc networks. *IEEE Trans. Mobile Comput.*, 2(1):25–39, 2003. ISSN 1536-1233. doi: 10.1109/TMC.2003.1195149.

S. Scellato, I. Leontiadis, C. Mascolo, P. Basu, and M. Zafer. Evaluating temporal robustness of mobile networks. *IEEE Transactions on Mobile Computing*, 2011. To appear.

T. Spyropoulos, T. Turletti, and K. Obraczka. Routing in delay-tolerant networks comprising heterogeneous node populations. *IEEE Trans. Mobile Comput.*, 8(8):1132–1147, Aug. 2009. ISSN 1536-1233. doi: 10.1109/TMC.2008.172.

J. Tang, M. Musolesi, C. Mascolo, and V. Latora. Temporal distance metrics for social network analysis. In *Proc. 2nd ACM workshop on Online social networks (WOSN)*, pages 31–36. ACM, 2009. ISBN 978-1-60558-445-4. doi: 10.1145/1592665.1592674.

M. Y. S. Uddin, F. Saremi, and T. F. Abdelzaher. End-to-end delay bound for prioritized data flows in disruption-tolerant networks. In *Real-Time Systems Symposium (RTSS), 2010 IEEE 31st*, pages 305 –316, 30 2010-dec. 3 2010. doi: 10.1109/RTSS.2010.39.

E. Vollset and P. Ezhilchelvan. Design and performance study of crash-tolerant protocols for broadcasting and reaching consensus in MANETs. In *Proc. 24th IEEE Symposium on Reliable Distributed Systems (SRDS)*, pages 166–175, 2005. doi: 10.1109/RELDIS.2005.15.

B. Xuan, A. Ferreira, and A. Jarry. Computing shortest, fastest, and foremost journeys in dynamic networks. *Int. J. Found. Comput. Sci.*, 14(2):267–285, 2003.

Y. Yang and R. Kravets. Achieving delay guarantees in ad hoc networks through dynamic contention window adaptation. In *Proc. 25th IEEE International Conference on Computer Communications INFOCOM 2006*, pages 1–12, Apr. 2006. doi: 10.1109/INFOCOM.2006.213.

C. Yawut, B. Paillassa, and R. Dhaou. Mobility metrics evaluation for self-adaptive protocols. *J. Networks*, 3 (1):53–64, Jan. 2008. doi: 10.4304/jnw.3.1.53-64.

Appendix

Proof of Theorem 1: **(1)** $D(\overline{\epsilon}, \overline{s}) \leq D(\epsilon, s)$: Consider the connectivity model C with $D(\epsilon, s) = d$ and any time point t, and let $N = N(C, t, d)$ be the delay neighbourhood at that time. By Definition 3.5 we then know that $e(N, S) \geq \epsilon$ for all subsets S of size s. Now recall that we can express the least neighbour expansion for the complement set $\overline{S}$ as $e(N, \overline{S}) = e(N, S) \cdot s/\overline{s}$, meaning that $e(N, \overline{S}) \geq \epsilon \cdot s/\overline{s} = \overline{\epsilon}$. Again by Definition 3.5, $D(\overline{\epsilon}, \overline{s})$ is the *minimal* delay for which this inequality is true, meaning that $D(\overline{\epsilon}, \overline{s})$ is defined and that $D(\overline{\epsilon}, \overline{s}) \leq d = D(\epsilon, s)$. **(2)**$D(\overline{\epsilon}, \overline{s}) \geq D(\epsilon, s)$: By (1) we know that $D(\overline{\epsilon}, \overline{s})$ is defined. The rest of this step is analogous to (1). □

Proof of Theorem 2: Consider any space-time connected model C, any run r of I, and any message m, where there is a time t when exactly s processes are informed of m. Let S be this set of informed processes, and let $d = D(\epsilon, s)$. From Definition 3.5 we know that $N(C, t, d)$ is a graph with least neighbour expansion of *at least* ϵ. Therefore, by Definition 3.3, S will have at least $\epsilon \cdot s$ neighbours outside S in $N(C, t, d)$, and thereby at least $\lceil \epsilon \cdot s \rceil$ neighbours since $\epsilon \cdot s$ is not necessarily an integer. Moreover, since $k \leq \lceil \epsilon \cdot s \rceil$, there will be at least k such neighbours. Considering $N(C, t, d)$, we see from Definition 3.4 that these neighbours will appear some time during the interval $[t, t + d - T_m]$. By property (2) of the ideal protocol, each process in S will send the message m at all times in this interval, and by property (1) all their neighbours will have received the message by the time $t + d$, which means that at least $s + k$ nodes are informed of m at that time. Since we considered any run of I, any message m, by Definition 4.1 the spread time $T_P(s, s + k)$ is no larger than $t + d - t = d = D(\epsilon, s)$. □

Proof of Theorem 3: By Definition 4.1, there will be at least one run of P in which there is a point t_1 at which s_1 nodes are informed and within $T_P(s_1, s_3)$ of which s_3 nodes are informed. Consider any such run r, and let $t_3 = T_P(s_1, s_3)$ (i.e., the first time point when s_3 nodes must have been informed). Since we in the system model assume that there is a unique time point for every event in the system, there is also a time point t_2, $t_1 < t_2 < t_3$, where exactly s_2 processes have been informed. By Definition 4.1 we have that $t_2 - t_1 \leq T_P(s_1, s_2)$ and $t_3 - t_2 \leq T_P(s_2, s_3)$. Adding these inequalities we get: $t_2 - t_1 + t_3 - t_2 = t_3 - t_1 = T_P(s_1, s_3) \leq T_P(s_1, s_2) + T_P(s_2, s_3)$. □

Proof of Theorem 4: The proof will be done in three steps. (1) First, we will consider a *subinterval* of the entire interval required to spread from s to $s + k$ nodes. At the beginning of the subinterval we do not know how many nodes are informed, except that it is between s and $s + k$. So, we will show that for any subinterval of length D, where there are between s and $s + k$ informed nodes at the start of the subinterval, the number of send opportunities for all informed nodes will be at least $b \cdot k$. (2) We will then derive the number of send opportunities in any interval of length $\left\lceil \frac{Q}{b} \right\rceil \cdot D$. (3) Finally, by relating the number of send opportunities with the number of successful sends (i.e., resulting in one more informed node), we can show that in that interval there will be at least k successful sends.

(1) Consider any space-time connected model C with $D(\epsilon, s') \leq D$ for all s' as described in the theorem, any Q-b-fair protocol P, any message m, and any interval starting at some time t with length D (i.e., this is a subinterval of the total time interval required to spread the message to k more nodes) such that for any run of P, the number of m-informed nodes at time t is between s and $s + k$. Since the theorem assumes that there are at least s informed nodes, at least one such interval must exist. By Definition 3.5, and the fact that the expansion delay cannot get worse for set sizes larger than s, $N(C, t, d)$ is a graph whose least neighbour expansion is at least ϵ for all sets of sizes s to $s + k$. By Definition 3.4, this means that for at least $\epsilon \cdot s$ of the informed nodes, there will be some topology G^i where they are adjacent to uninformed nodes. Since the number of informed border nodes must be integer, there are at least $\lceil \epsilon \cdot s \rceil \geq k$ such nodes. By property 1 of the Q-b-fair protocol, this means that any such node will have b send opportunities. So in total, there will be at least $b \cdot k$ send opportunities.

(2) Now consider any run of P where at some time t there are s nodes informed. We will now show from t to $t' = t + \left\lceil \frac{Q}{b} \right\rceil \cdot D$ there will have been at least Qk send opportunities. We restrict our attention to the runs where there are *no more* than $s + k$ informed nodes at time t' (as in the other case, at least k nodes have been informed and there is nothing left to prove). By the first step of the proof, we know that any sub interval of length D in $[t, t']$ will result in at least $b \cdot k$ send opportunities. So the total number of send opportunities o during the interval $[t, t']$ will be:

$$o \geq \left\lfloor \frac{(t' - t)}{D} \right\rfloor \cdot b \cdot \left\lfloor \frac{\left\lceil \frac{Q}{b} \right\rceil \cdot D}{D} \right\rfloor \cdot b \cdot = \left\lceil \frac{Q}{b} \right\rceil \cdot b \cdot k$$

$$\geq \frac{Q}{b} \cdot b \cdot k = Qk$$

(3) Finally, for the same runs selected under (2) above, we will now connect the number of send opportunities with the number of new informed nodes. By property 2 of Q-b-fair protocol, having Qk consecutive send opportunities for message m during some interval of time means that there will be at least k *send*(m) actions. By property 3, all those send actions will also result in a new informed node.

Thus, we have shown that, starting from s nodes, during the time $\left\lceil \frac{Q}{b} \right\rceil \cdot D$ at least k more nodes have been informed. Since we considered any run of P, and any message m, we know that this is true for all runs, and therefore by Definition 4.1, the worst-case spread time is $T_P(s, s+k) \leq \left\lceil \frac{Q}{b} \right\rceil \cdot D$. □

Proof of Theorem 5: The proof will be done in several steps. First we will express how long it takes to double the number of informed nodes (given that less then $n/2$ nodes are informed). Then we use this to express the time taken to spread from 1 to $n/2$ nodes. Similarly for the case when more than $n/2$ nodes are informed we derive the time taken to halve the number of *uninformed* nodes, and use this to find the time to spread from $n/2$ to n nodes.

Let $k = \lceil \frac{s}{u} \rceil$, then by Theorem 4, for any set size s, such that $1 \leq s$, and $s + k \leq n/2$:

$$T_P(s, s+k) \leq \left\lceil \frac{Q}{b} \right\rceil \cdot D \tag{2}$$

This means that if we start with s informed nodes, the time to spread to k more nodes is at most $\left\lceil \frac{Q}{b} \right\rceil \cdot D$. By applying Theorem 3, we can use inequality (2) u times starting from $l \cdot u$ nodes, where $l \geq 1$ getting:

$$T_P(lu, 2lu) \leq \left\lceil \frac{Q}{b} \right\rceil \cdot D \cdot u$$

Similarly, we see that $T_P(1, u) \leq \left\lceil \frac{Q}{b} \right\rceil \cdot D \cdot u$. We can now express the time taken to spread from one node to $n/2$ nodes:

$$\begin{aligned} T_P(1, n/2) &\leq T_P(1, u) + \ldots + T_P(2^{L-1}u, n/2) \\ &\leq \left\lceil \frac{Q}{b} \right\rceil (L+1) \cdot D \cdot u \end{aligned} \tag{3}$$

where $L = \lceil \log_2 \frac{n}{2u} \rceil$. For set sizes s, where $n/2 \leq s \leq n$ we first let $\epsilon = \overline{s}/us$ so that by Theorem 1 $D(\epsilon, s) = D(\epsilon \cdot s/\overline{s}, \overline{s}) = D(1/u, \overline{s})$. Since $\overline{s} \leq n/2$, we also know that $D(1/u, \overline{s}) \leq D$ so $D(\epsilon, s) \leq D$. As before, we can now use theorems 3, and 4 u times starting from $n - 2lu$ nodes:

$$T_P(n - 2lu, n - lu) \leq \left\lceil \frac{Q}{b} \right\rceil \cdot D \cdot u \tag{4}$$

The spread time from $n - u$ to n nodes $T_P(n - u, n)$ is bounded by $\left\lceil \frac{Q}{b} \right\rceil \cdot D \cdot u$ (again by theorems 3, and 4). Using this together with (4) we can now express the time taken to spread from $n/2$ to n nodes (with L defined as above):

$$\begin{aligned} T_P(n/2, n) &\leq T_P(n/2, n - 2^{L-1}u) + \ldots + T_P(n - u, n) \\ &\leq \left\lceil \frac{Q}{b} \right\rceil (L+1) \cdot D \cdot u \end{aligned} \tag{5}$$

Finally, by combining (3) and (5), we get the result of the theorem. □

Proof of Theorem 6: Consider any cut U (i.e., set of edges) in the graph that partitions the graph into two sets S and $\overline{S} = V \setminus S$, where we can let S be the smaller of the two sets. Clearly, the total flow going through the cut is

$$\sum_{\forall i, \forall (x,y) \in U} |f_i(x, y)| = |S||\overline{S}|\delta \tag{6}$$

Now consider the set $\Gamma(S) \setminus S$ which are all the nodes in $\overline{S}$ with neighbours in S. The sum of flows going through $\Gamma(S) \setminus S$ must be less than or equal to $|\Gamma(S) \setminus S| \cdot F_{max} \cdot \delta$. Since the flow through the cut must be less than or equal to the total flow through $\Gamma(S) \setminus S$ we have:

$$|S||\overline{S}|\delta \leq |\Gamma(S) \setminus S| \cdot F_{max} \cdot \delta$$

Rewriting (recall that $|\overline{S}| = n - s$):

$$\frac{|\Gamma(S) \setminus S|}{|S|} \geq \frac{n - s}{F_{max}} \tag{7}$$

The same reasoning for the flow through $\Gamma(\overline{S}) \setminus \overline{S}$ gives:

$$\frac{|\Gamma(\overline{S}) \setminus \overline{S}|}{|S|} \geq \frac{n - s}{F_{max}} \tag{8}$$

Putting together Definition 3.3 with inequalities (7) and (8) gives the desired expression. □