



EURECOM
Department of Mobile Communications
Campus Sophia Tech
Les Templiers
450 route des Chappes
B.P. 193
06410 Biot
FRANCE

Research Report RR-12-272

**Delay Analysis of Epidemic Schemes in Sparse and Dense
Heterogeneous Contact Environments**

July 20th, 2012
Last update October 4th, 2012

Pavlos Sermpezis and Thrasyvoulos Spyropoulos

¹EURECOM's research is partially supported by its industrial members: BMW Group, Cisco, Monaco Telecom, Orange, SAP, SFR, STEricsson, Swisscom, Symantec.

Tel : (+33) 4 93 00 81 00

Fax : (+33) 4 93 00 82 00

Email : {pavlos.sermpezis,thrasyvousos.spyropoulos}@eurecom.fr

Delay Analysis of Epidemic Schemes in Sparse and Dense Heterogeneous Contact Environments

Pavlos Sermpezis and Thrasyvoulos Spyropoulos

Abstract

Epidemic algorithms have found their way into many areas of computer science, such as databases and distributed systems, and recently for communication in Opportunistic and Delay Tolerant Networks (DTNs). To ensure analytical tractability, existing analyses of epidemic spreading predominantly consider homogeneous contact rates between nodes. However, this assumption is generally not true, especially in opportunistic networks between mobile nodes. In this paper we consider two classes of contact/mobility models with heterogeneous contact rates. We derive simple closed form approximations and bounds for the basic epidemic spreading step, and prove conditions for their accuracy. To demonstrate the utility of this result in practice, we use it to derive the delay of epidemic routing and variants (gossiping and 2-hop routing) as well as to construct a network size estimator based on local measurements. We use simulation results based on synthetic heterogeneous models to validate our analysis, as well as real traces to demonstrate that our expressions can still be useful even in scenarios with significantly more complex structure.

Index Terms

Mobile opportunistic networks, epidemic algorithms, heterogeneous contact dynamics, performance modeling

Contents

1	Introduction	1
2	Approximation and Bounds for Epidemic Step	2
2.1	Heterogeneous Contact Networks with Full Mixing	4
2.2	Heterogeneous Contact Networks with Sparse Mixing	8
2.3	Bounding the Expected Step Time	9
2.3.1	Lower Bound	9
2.3.2	Upper Bound	10
3	Applications	11
3.1	Epidemic Routing	11
3.2	Probabilistic Routing - Gossiping	12
3.3	2-hop Routing	12
3.4	Network Size Estimator	12
4	Model Validation	14
4.1	Model Validation	14
4.2	Real Mobility Traces	15
4.2.1	Message Spreading	18
4.2.2	Network Size Estimation	18
5	Related Work	21
6	Conclusions	21
7	Appendix	22
7.1	Approximation for Expected Unicast Delay under Epidemic Routing	22
7.2	Delay Approximation under 2-hop Routing	23
7.3	Asymptotical Behavior of the Network Size Estimator	23

List of Figures

1	Epidemic spreading over a homogeneous network with N nodes . . .	3
2	Epidemic spreading in a heterogeneous network with 4 nodes . . .	3
3	Delay of epidemic spreading (broadcast) in Dense (Full Mixing) networks. Contact rates are distributed according to (a): $F(\lambda) \sim$ Log-normal, $\mu_\lambda = 0.33 \cdot 10^{-3} \text{sec}^{-1}$, $\frac{\sigma_\lambda}{\mu_\lambda} = 0.5$, (b): $F(\lambda) \sim$ Exponential, $\mu_\lambda = 1 \cdot 10^{-3} \text{sec}^{-1}$, $\frac{\sigma_\lambda}{\mu_\lambda} = 1$ and (c): $F(\lambda) \sim$ Pareto, $\mu_\lambda = 2.8 \cdot 10^{-3} \text{sec}^{-1}$, $\frac{\sigma_\lambda}{\mu_\lambda} = 2.23$	16
4	Delay of epidemic spreading (broadcast) in Sparse (Poisson Mixing) networks, with different meeting probabilities. Contact rates are distributed according to (a): $F(\lambda) \sim$ Pareto, $\mu_\lambda = 2.8 \cdot 10^{-3} \text{sec}^{-1}$, $\frac{\sigma_\lambda}{\mu_\lambda} = 2.23$, (b): $F(\lambda) \sim$ Log-normal, $\mu_\lambda = 2.8 \cdot 10^{-3} \text{sec}^{-1}$, $\frac{\sigma_\lambda}{\mu_\lambda} = 2.23$ and (c): $F(\lambda) \sim$ Exponential, $\mu_\lambda = 1 \cdot 10^{-3} \text{sec}^{-1}$, $\frac{\sigma_\lambda}{\mu_\lambda} = 1$	17
5	Delay of two unicast routing protocols in a full mixing network. (a) Epidemic routing and (b) Probabilistic routing with $p_r = 0.2$. Contact rates are distributed according to $F(\lambda) \sim$ Log-normal, $\mu_\lambda = 0.33 \cdot 10^{-3} \text{sec}^{-1}$, $\frac{\sigma_\lambda}{\mu_\lambda} = 1.86$	18
6	Box-plots of the unicast (i.e. message delivery) delay under epidemic and 2-hop routing. On each box, the central horizontal line is the median, the edges of the box are the 25th and 75th percentiles, the whiskers extend to the most extreme data points not considered outliers, and outliers are plotted individually as crosses. The thick lines represent the theoretical values predicted by our model.	19
7	Network size estimator. (a): Estimator compared to real values for each sub-network of the Cabspotting trace, (b): Estimator compared to real values for traces generated by three mobility models.	20
8	Network size estimator. ECDF of estimations among all nodes for (a): a Sub-network of 100 nodes of the Cabspotting trace, (b): a trace of 100 nodes, generated by the SLAW mobility model	20

1 Introduction

Epidemic spreading is probably one of the most popular bio-inspired principles that have made their way into computer engineering. Epidemic algorithms and variants (e.g. gossip) have been used for communication in distributed systems, synchronization of distributed databases, content searching in peer-to-peer systems, etc. Recently, epidemic-based schemes have also been proposed for routing and data dissemination in *Opportunistic Networks*. *Opportunistic* or *Delay Tolerant Networks* (DTNs) are envisioned to support communication in case of failure or lack of infrastructure (disaster, censorship, rural areas), but also to enhance existing wireless networks (e.g., offload cellular traffic, novel applications).

In Opportunistic networks, when two mobile devices (smartphones, laptops, etc) are in proximity (*in contact*) they can exchange data or information using local wireless communication (e.g. Bluetooth or WiFi). Message dissemination can be end-to-end or content-centric, yet neither the existence nor the knowledge of an end-to-end path is assumed. Nevertheless, over a sequence of node encounters, messages can get copied to many nodes as well as forwarded over multiple hops. In the *epidemic routing* case [1], any node that has a message X (is “infected”) will forward it to any node encountered that does not have it yet (is “susceptible”). While this guarantees that every node in the network will eventually receive the message, it comes with a high resource overhead. Numerous variants have been proposed to improve the resource usage of epidemic routing while maintaining good performance (see [2, 3] for a detailed survey).

Since the mobility process of nodes involved (e.g. humans or vehicles carrying the devices) is, in most cases, not deterministic, the performance of epidemic algorithms (and variants) heavily depends on the underlying contact patterns between nodes. To this end, epidemic algorithms have been extensively studied through both simulations and analytical models. While simulations with state-of-the-art synthetic models or real mobility traces can provide more reliable predictions for the *specific* scenario tested, analytical models can give quick, qualitative results and intuition, answer “what-if” questions, and help optimize protocols (e.g. choosing the number of copies in [4], or gossip probability [5]).

For the sake of tractability, state-of-the-art analytical models for epidemic spreading mainly rely on simple mobility assumptions (e.g. Random Walk, Random Waypoint), where node mobility is stochastic and independent, identically distributed (IID) (see e.g. [5–7]). Nevertheless, numerous studies of real mobility traces [8–11] reveal a different picture. One key finding is that contact rates between different pairs of nodes can vary widely. Furthermore, many pairs of nodes may never meet. This puts in question the accuracy and utility of these models’ predictions. Yet, departures from these assumptions [9, 12–14] seem to quickly increase complexity and/or limit the applicability of results.

This raises the question: can we derive *useful and accurate* closed form expressions for the performance of epidemic schemes, under more generic mobility assumptions? To this end, in this paper we consider a large class of contact/mobility

models with heterogeneous contact rates. At first, we assume every pair of nodes $\{i, j\}$ meets according to a random process with a different *mean* contact rate λ_{ij} , drawn from an arbitrary probability distribution $F(\lambda)$ with known mean μ_λ and variance σ_λ^2 . We then further relax our assumptions and allow pairs of nodes to never meet each other. Our contributions can be summarized as follows:

- We derive simple closed form approximations for the expected delay of the basic epidemic algorithm that only involve 1st and 2nd moments of the contact rate distribution $F(\lambda)$. We prove that these approximations become exact asymptotically on the network size. We also derive upper and lower bounds for finite network sizes. To our best knowledge, this is the first closed-form result for the two contact classes considered (Section 2).
- To demonstrate how these results could be used in practice, we derive closed form expressions for the delay of various epidemic based protocols and use our theory to propose an efficient *local* network size estimator for heterogeneous contact networks (Section 3).
- We validate these results against various synthetic simulation scenarios belonging to the above contact classes, and show that their accuracy is significant even for moderate network sizes (Section 4.1).
- Finally, we test our results and estimator on real mobility traces, and show that our simple expressions can still be of use even in scenarios with significantly more complex structure than the one assumed (Section 4.2).

As a final note, while our focus here will be in the domain of opportunistic and delay-tolerant networking, we believe our framework has more general applicability to epidemic algorithms in many different contexts, as long as the contact process between nodes fits our contact classes. For example, one could imagine the probabilistic spread of malware over an email or chat network [15], where the spread is possible when two nodes decide to communicate. For this reason, we present our analysis in the general context of epidemic spreading. We focus on opportunistic networks when we discuss applications and compare our results to real scenarios.

2 Approximation and Bounds for Epidemic Step

We consider a network composed of N nodes. Each pair of nodes $\{i, j\}$ “meets” at random intervals with an average rate λ_{ij} . Some nodes might never come in contact, in which case $\lambda_{ij} = 0$. A contact network can then be represented by a matrix $\Lambda = \{\lambda_{ij}\}$ with zero and non-zero entries.

Depending on the scenario considered, this contact might be two smartphones coming within (Bluetooth) transmission range, two users initiating an online chat,




Figure 1: Epidemic spreading over a homogeneous network with N nodes




Figure 2: Epidemic spreading in a heterogeneous network with 4 nodes

etc. During a contact, a message (e.g. a file, a virus, or a trending video) currently on one of the nodes *could* be forwarded to (“infect”) the other node as well. In the basic epidemic scheme (*epidemic routing* in the context of DTNs [1]), a message starts from a source node, and a message transfer occurs at *every* contact opportunity involving a node with the message and one without it. We are usually interested in the time until either all nodes (broadcast) or a specific destination node (unicast) get infected. This time heavily depends on Λ .

To simplify analysis, it is commonly assumed (in both DTNs and biology) that *all* inter-contact times are *independent and exponentially distributed* with the same rate $\lambda_{ij} = \lambda$ (IID) [5, 6]. This assumption allows one to model epidemic spreading with a pure-birth Markov chain, as depicted in Fig. 1.

We say that the process is at *step* k (or state k), when k nodes, including the source, have the packet. The random variable $T_{k,k+1}$, is the time it takes to move from *step* k to *step* $k+1$, and can be easily shown in the IID case to be exponentially distributed with rate $k(N - k)\lambda$ and mean value $E[T_{k,k+1}] = \frac{1}{k(N-k)\lambda}$. The expected delay until *all* nodes are infected is then $\sum_{k=1}^{N-1} E[T_{k,k+1}]$.

While the expected delay of the basic epidemic step $E[T_{k,k+1}]$ is easy to derive in the IID contact case, this assumption is rather strong. Our goal in this paper is to raise this assumption and derive closed form expressions that are both accurate and simple. Specifically, we first assume that pairwise rates are not identical (Section 2.1). Then we also allow an arbitrary number of pairs to never meet (Section 2.2).

2.1 Heterogeneous Contact Networks with Full Mixing

We first consider the following class of contact models.

Definition 1 (Heterogeneous Full Mixing). *Inter-contact times between a given pair of nodes i and j are exponentially distributed with a different non-zero rate λ_{ij} , and independent of each other. λ_{ij} is drawn from an arbitrary distribution $F(\lambda)$, $\lambda \in (0, \infty)$, with known mean μ_λ and variance σ_λ^2 .*

In other words, we allow different pairs to contact each other more frequently than others (but, we assume *for the moment* all contact rates to be non-zero). While this contact class is far from exhaustive, $F(\lambda)$ can describe a significantly broader range of scenarios. For example, large σ_λ^2 values imply that some pairs will rarely contact each other while others much more often. An $F(\lambda)$ symmetric around μ_λ (e.g. uniform distribution) implies a balanced number of high and low rates, while a right-skewed $F(\lambda)$ (e.g. Pareto) describes a network with most pairs having large intercontact times, but few meeting very frequently. Small μ_λ values could correspond to slow moving nodes, e.g. pedestrians, (or large geographical areas). Finally, multi-modal $F(\lambda)$ functions might approximate scenarios with some hierarchical structure.

However, we retain the assumption of exponential and independent pair inter-contact times. While inter-meeting rates cannot be expected to be always and exactly exponential (although some studies do claim this to be the case, at least for a subset of node pairs [9, 10]), detailed studies of the same real traces [16, 17] show that (aggregate) intermeeting times exhibit an exponential tail. This exponential tail is also supported by known results about the hitting times of random walks [18]. Finally, an interesting recent study makes the case that the usually studied aggregate inter-meeting time distribution might be non-exponential, even if the individual pair meetings are exponential but with different rates [19], which is consistent with the above mobility class.

Besides theoretical reasons, there are also important practical reasons for this assumption. The few studies we are aware of that assume more general inter-contact times [13, 20] only deal with asymptotic behavior (i.e. whether delay is finite or infinite). We therefore choose to retain this assumption and provide more expressive results. We then use real traces as well as one scenario where inter-contact times are explicitly designed to be non-exponential (but “heavy-tailed”) [21], to see when these results are or are not useful.

Unlike the case of IID contacts, where we only need to track the *number* of nodes k that are infected, in this heterogeneous case we need to know *which nodes exactly* are included in the k infected nodes. As is shown in Fig. 2, complexity increases quickly, even for a simple 4-node network. The state space explodes with $\binom{N}{k}$ different starting states for *step* k , each with a different probability. While keeping track of these probabilities could be done recursively, the task becomes intractable and such a chain can only be solved numerically. To avoid this compli-

cation, the main idea behind our results is to prove that, in the limit of large N , all such starting states become *statistically equivalent*, and then collapse them.

Before we proceed, we prove a simple proposition, that we will need later, which follows from Chebyshev's inequality.

Proposition 2.1. *Let $X = \sum_{i=1}^N \lambda_i$, where the λ_i are IID random variables with finite mean μ_λ and finite non-zero variance σ_λ^2 . Let further c_μ be a real positive number. Then, if μ_X is the mean value of X ,*

$$P\left(\frac{|X - \mu_X|}{\mu_X} \geq c_\mu\right) \leq \frac{1}{N} \frac{\sigma_\lambda^2}{c_\mu^2 \mu_\lambda^2}$$

Proof. As X is a sum of N IID random variables, it holds that $\mu_X = N\mu_\lambda$ and $\sigma_X^2 = N\sigma_\lambda^2$ [22], where σ_X^2 denotes the variance of X . We can apply Chebyshev's inequality for X to show that $P\left(\frac{|X - \mu_X|}{\sigma_X} \geq c\right) \leq \frac{1}{c^2}$, for any real positive number c . By setting $c = \frac{c_\mu \mu_X}{\sigma_X}$ and replacing μ_X and σ_X with the expressions above, we end up with the result. \square

We now state our first result that gives a simple closed form approximation for the delay of an epidemic step. The accuracy of this approximation improves as (i) the number of nodes in the network increases, (ii) the mean contact intensity μ_λ increases, or (iii) the variance σ_λ^2 decreases. Furthermore, our simulation-based evaluation (Section 4.1) reveals that the accuracy of this result, in practice, is satisfactory even for moderate network sizes (e.g. 100 nodes).

Theorem 2.2. *For all contact models belonging in the Heterogeneous Full Mixing class, the expected time from step k to step $k+1$ in epidemic spreading can be approximated by:*

$$E[T_{k,k+1}] = \frac{1}{k(N-k)\mu_\lambda} + \frac{\sigma_\lambda^2}{[k(N-k)]^2 \mu_\lambda^3}. \quad (1)$$

Proof. There are three sources of randomness at play when calculating $E[T_{k,k+1}]$:

(i) A network is initially created according to $F(\lambda)$. In other words, $N(N-1)/2$ contact rates λ_{ij} are drawn independently from $F(\lambda)$. The resulting graph or (symmetric) contact rate matrix $\Lambda = \{\lambda_{ij}\}$ is a contact network *instance*.

(ii) When at step k , there are k nodes with the message. Conditioned on Λ , $T_{k,k+1}$ is a random variable whose distribution will also depend on the actual set of k nodes that have the message, and their contact rates with the remaining nodes. Let \mathbf{C}_m^k denote this set, where m is an integer indicating one of the $\binom{N}{k}$ possible sets of infected relays at step k .

(iii) Finally, conditional on both the network instance Λ and \mathbf{C}_m^k , $T_{k,k+1}$ will also depend on the randomness of the inter-contact times involved.

Let node $i \in \mathbf{C}_m^k$ and node $j \notin \mathbf{C}_m^k$ and denote as t_{ij} the next time node i will meet node j , after the k^{th} node received the packet. Then, $T_{k,k+1} = \min_{i \in \mathbf{C}_m^k, j \notin \mathbf{C}_m^k} \{t_{ij}\}$.

As t_{ij} are independent, exponentially distributed random variables $T_{k,k+1}$ is also exponentially distributed with rate $\sum_{i \in \mathbf{C}_m^k} \sum_{j \notin \mathbf{C}_m^k} \lambda_{ij}$. Hence,

$$E [T_{k,k+1} | \mathbf{C}_m^k, \Lambda] = \frac{1}{\sum_{i \in \mathbf{C}_m^k} \sum_{j \notin \mathbf{C}_m^k} \lambda_{ij}} \quad (2)$$

Using the properties of conditional expectation, we get the expected time from *step* k to *step* $k+1$:

$$E [T_{k,k+1} | \Lambda] = \sum_{m=1}^{\binom{N}{k}} \frac{1}{\sum_{i \in \mathbf{C}_m^k} \sum_{j \notin \mathbf{C}_m^k} \lambda_{ij}} \cdot P\{\mathbf{C}_m^k | \Lambda\} \quad (3)$$

The problem in Eq.(3) is that keeping track of the probabilities $P\{\mathbf{C}_m^k | \Lambda\}$ is exceedingly complex, and even if we did (e.g. recursively) it would not lead to a useful expression.

Instead, we take the expectation over all possible Λ

$$E [T_{k,k+1}] = \int_{\Lambda} \sum_{m=1}^{\binom{N}{k}} \frac{1}{\sum_{i \in \mathbf{C}_m^k} \sum_{j \notin \mathbf{C}_m^k} \lambda_{ij}} \cdot P\{\mathbf{C}_m^k | \Lambda\} P\{\Lambda\} d\Lambda. \quad (4)$$

In Eq.(4), we first pick the network Λ and its elements, and then consider all possible partitions of nodes into two sets of k and $N - k$. We repeat over all networks and average. Since the network creation process (i.e. populating the elements of matrix Λ) is *ergodic*, we could exchange this order to get¹:

$$E [T_{k,k+1}] = \sum_{m=1}^{\binom{N}{k}} E \left[\frac{1}{X} \right] \cdot P\{\mathbf{C}_m^k\}, \quad (5)$$

where X is a random variable equal to $\sum_{i \in \mathbf{C}_m^k} \sum_{j \notin \mathbf{C}_m^k} \lambda_{ij}$, and the expectation of $\frac{1}{X}$ is taken over Λ . Furthermore \mathbf{C}_m^k now simply denotes one of the possible partitions of nodes into k and $N - k$ nodes. Since the different elements in Λ are IID, the $k(N - k)$ contact rates λ_{ij} included in X are distributed according to $F(\lambda)$ for any partition \mathbf{C}_m^k . By taking the expectation over all Λ , there is no dependence anymore on the exact partition and

$$E [T_{k,k+1}] = E \left[\frac{1}{X} \right] \sum_{m=1}^{\binom{N}{k}} P\{\mathbf{C}_m^k\} = E \left[\frac{1}{X} \right]. \quad (6)$$

Calculating $E[\frac{1}{X}]$ is still no easy task. Since it involves a $k(N - k)$ -th order convolution (X is a sum of $k(N - k)$ terms). Instead, we will estimate it, by expressing

¹This can be rigorously shown, in the limit of a large network. However, the proof is technical and is omitted from this report.

the function $f(X) = \frac{1}{X}$ as a Taylor series expansion, centered at μ_X , the mean value of X .

$$T_f(X) = \sum_{n=0}^{\infty} \frac{f^{(n)}(\mu_X)}{n!} (X - \mu_X)^n = \sum_{n=0}^{\infty} \frac{(-1)^n (X - \mu_X)^n}{(\mu_X)^{n+1}} \quad (7)$$

We can approximate f by taking the first m terms of the Taylor series. That will result in:

$$f(X) \approx \sum_{n=0}^m \frac{(-1)^n}{(\mu_X)^{n+1}} (X - \mu_X)^n \quad (8)$$

An approximation for the mean value of $f(X)$ results after taking the expectation of both sides in the last equation.

$$E[f(X)] \approx \sum_{n=0}^m \frac{(-1)^n}{(\mu_X)^{n+1}} M_n \quad (9)$$

where $M_n = E[(X - \mu_X)^n]$ is the n^{th} central moment.

This method, of approximating a function with a finite Taylor sum and taking the expectation of it for evaluating the mean value of the function, is widely known as the *delta method* [23, 24].

As the final step, we need to derive the expected value of X , μ_X (as well as higher moments, depending on the number of Taylor series terms we use). To derive Eq.(1) we consider only the first three terms in Eq.(9). We thus only need the 1st and 2nd moments of X (sum of independent variables), given by [22]

$$\mu_X = k(N - k)\mu_\lambda, \quad (10)$$

$$\sigma_X^2 = k(N - k)\sigma_\lambda^2. \quad (11)$$

Replacing them in Eq.(9) gives us Eq.(1).

Convergence of approximation: The more terms of Eq.(9) we take into account, the larger the accuracy we can achieve. This accuracy depends on the distribution of X and on the linearity of f . The more probability mass is concentrated around the mean value μ_X and the more linear f is over the range of data, the less terms are needed [23, 24]. Using Proposition 2.1, the probability that X is not within a certain interval around the mean, $\pm c_\mu \mu_X$, is inversely proportional to the number of nodes in the network and to the mean intermeeting rate μ_λ . It is also proportional to the variance of $F(\lambda)$, which means that large differences between different pair contact rates will require larger network sizes for the approximation to reach a given level of accuracy. However, what our result implies is that, if our network is large enough, the statistics over a *single* network instance will converge to their ensemble statistics of $F(\lambda)$. In practice, there is usually no need to consider more than a few terms. \square

In the following analysis we use the above approximation that depends on the first two moments of the distribution $F(\lambda)$. Our choice is a tradeoff between usability and expressibility of the result, and its accuracy. The above theorem provides the guidelines for exploring various tradeoffs.

2.2 Heterogeneous Contact Networks with Sparse Mixing

We have so far required *all* pairs of nodes to meet with non-zero rate, i.e. $\lambda_{ij} > 0$. This is an undesirable restriction. Intuition as well as studies of real traces suggest that many pairs of nodes in fact never meet. We now relax this assumption and extend our results to sparse contact networks.

Definition 2 (Sparse Contact Graph with Heterogeneous Mixing). *For each pair of nodes i and j the following holds: (i) with probability $1 - p$ they never contact each other, (ii) with probability p they regularly contact with rate λ_{ij} . For the latter pairs, λ_{ij} is drawn from an arbitrary distribution $F(\lambda)$, $\lambda \in (0, \infty)$, with known mean μ_λ and variance σ_λ^2 , and the inter-contact times are exponentially distributed.*

In other words, we now first create a *Poisson (or Erdős-Renyi) graph* [25] between nodes. We then assign rates λ_{ij} , as before, but only to the existing links. With the parameter p , we can now also capture arbitrarily sparse scenarios, where each node meets only an a percentage of all nodes².

Theorem 2.3. *In scenarios belonging in the contact class defined by Definition 2, the approximation for the expected time from step k to step $k+1$ in epidemic spreading is:*

$$E[T_{k,k+1}] = \frac{1}{k(N-k)p\mu_\lambda} + \frac{(1-p) + \frac{\sigma_\lambda^2}{\mu_\lambda^2}}{[k(N-k)p]^2\mu_\lambda} \quad (12)$$

Proof. The proof follows similar steps as that of Theorem 2.2. We therefore only sketch here the key differences. We can again define, at *step* k , the set \mathbf{C}_m^k of “infected” nodes. However, we now also need to define for each node $i \in \mathbf{C}_m^k$ the set $\mathbf{D}_m^k(i) = \{j : j \notin \mathbf{C}_m^k \text{ and } \lambda_{ij} > 0\}$, which is the set of the nodes j that have not received yet the packet and can meet node i . In the full mixing case, the size of the set $\mathbf{D}_m^k(i)$ is $\|\mathbf{D}_m^k(i)\| = (N - k)$. However, in the sparse case, $0 \leq \|\mathbf{D}_m^k(i)\| \leq N - k$.

The total transition rate from state k to state $k + 1$, conditional on \mathbf{C}_m^k is then $\sum_{i \in \mathbf{C}_m^k} \sum_{j \in \mathbf{D}_m^k(i)} \lambda_{ij}$. Using similar arguments as in Theorem 2.2, we can show that $E[T_{k,k+1}] = E\left[\frac{1}{X}\right]$, where X is a sum of $\sum_{i \in \mathbf{C}_m^k} \|\mathbf{D}_m^k(i)\|$ of IID random variables λ_{ij} , and the expectation is taken over all network instances (i.e. over all Poisson graphs with parameter p and link weights according to $F(\lambda)$).

²We do assume that the probability p is large enough for connectivity to be achieved. In practice, the theory of Poisson graphs tells us that connectivity can be achieved with arbitrary low p as long as N is large enough (with percolation occurring at an average degree as low as 1 in the limit) [25].

But $\sum_{i \in \mathbf{C}_m^k} \|\mathbf{D}_m^k(i)\|$ is itself a random variable. By the definition of a Poisson graph, for each node $i \in \mathbf{C}_m^k$, there are $N - k$ other nodes $j \notin \mathbf{C}_m^k$ that are its neighbors each with probability p , and independent of all other links. Hence, each $\|\mathbf{D}_m^k(i)\|$ is binomially distributed as

$$P\{\|\mathbf{D}_m^k(i)\| = x\} = \binom{N-k}{x} p^x (1-p)^{(N-k)-x}, \quad (13)$$

Its expectation and variance are given by $E[\|\mathbf{D}_m^k(i)\|] = (N-k)p$ and $\sigma_{\|\mathbf{D}_m^k(i)\|}^2 = (N-k)p(1-p)$, respectively. Since $\|\mathbf{D}_m^k(i)\|$ are also independent for different i we can use the rule for expectations of random sums of IID variables [22] to get the mean and variance of random variable X ,

$$\mu_X = E\left[\sum_{i \in \mathbf{C}_m^k} \|\mathbf{D}_m^k(i)\|\right] \cdot E[\lambda_{ij}] = k(N-k)p \cdot \mu_\lambda, \quad (14)$$

$$\begin{aligned} \sigma_X^2 &= E\left[\sum_{i \in \mathbf{C}_m^k} \|\mathbf{D}_m^k(i)\|\right] \sigma_\lambda^2 + \mu_\lambda^2 \sigma_{\left(\sum_{i \in \mathbf{C}_m^k} \|\mathbf{D}_m^k(i)\|\right)}^2 \\ &= k(N-k)p [\sigma_\lambda^2 + \mu_\lambda^2(1-p)] \end{aligned} \quad (15)$$

With these first and second moments of X in hand, we can go ahead and derive an approximation for $E[T_{k,k+1}] = E\left[\frac{1}{X}\right]$. The analysis is the same as in Full Mixing class. We only need to replace μ_X and σ_X^2 from Eq.(14) and Eq.(15). \square

2.3 Bounding the Expected Step Time

While our approximation is provably accurate in the limit of large networks, it is of interest to know what kind of errors to expect in theory for finite sized networks. We will do this by deriving lower and upper bounds for the mean step delay. Eq.(6) suggests that the derivation of $E[T_{k,k+1}]$ boils down to the derivation of the expectation of $E\left[\frac{1}{X}\right] = E[f(X)]$ (see proof of Theorem 2.2). Yet $f(X)$ is a convex function and X a sum of $(k(N-k))$ random variables. We can thus base our analysis on known bounds for the expectation of convex functions of random variables. We will express our bounds for the more generic case of Poisson mixing. The full mixing result can be derived by setting $p = 1$.

2.3.1 Lower Bound

A lower bound for $E[f(X)]$ can be taken by the classic Jensen's inequality, and is equal to $f(\mu_X)$. Therefore the lower bound for the step time is:

$$E[T_{k,k+1}] \geq \frac{1}{k(N-k)p\mu_\lambda} \quad (16)$$

Note that this is also equal to the first order approximation using the Delta method in Section 2.1. The tightness of this bound, as in the case of the approximation, increases under the same conditions as Theorem 2.2 (larger network size, etc.).

2.3.2 Upper Bound

Bounding $E[T_{k,k+1}]$ from above is more crucial than bounding it from below, as it provides us with a worst case estimate (in the stochastic sense).

An upper bound for a convex function $f : [a, b] \rightarrow \mathbb{R}$ of a random variable X is given by the Edmundson-Madansky inequality [26]. As proven in [27], more tight bounds can be estimated by using higher-order moments of the random variable X . These upper bounds are given by the next equation, where $EM_n \leq EM_{n-1}$ and as $n \rightarrow \infty$ the upper bound converges to $E[f(X)]$.

$$EM_n = \sum_{i=0}^n \binom{n}{i} \frac{E[(X-a)^i(b-X)^{n-i}]}{(b-a)^n} f\left(a + \frac{i}{n}(b-a)\right) \quad (17)$$

Eq.(17) can give very bounds by considering a large number of terms. Nevertheless, this leads to a complex expression and we might not always know all the moments of X , $E[X^n]$. We will thus focus only on the 2nd order bound.

If $F(\lambda)$ is defined in a closed interval $[a, b]$ then the computation of the upper bound EM_2 is trivial as Eq.(17) requires only the values of the mean and variance of X , which are already derived. However, if $F(\lambda)$ is defined in $(0, \infty)$, we cannot apply Eq.(17) as it is. For these cases, we can still derive an upper bound that holds for a fraction of all the possible random networks of size N that result from $F(\lambda)$.

Theorem 2.4. *Let $F(\lambda)$ ($0 < \mu_\lambda, \sigma_\lambda < \infty$) be an arbitrary distribution of contact rates defined in $(0, \infty)$. Then, $\forall \epsilon > 0, \exists N_0$: an upper bound (EM_2) for the mean step delay exists for a fraction $p_0 \geq 1 - \epsilon$ of all possible networks instances in $F(\lambda)$ of size $N > N_0$. The bound is given by*

$$E[T_{k,k+1}] \leq \frac{1}{k(N-k)p\mu_\lambda} \left[1 + \frac{\frac{c^2+1}{2} \left((1-p) + \left(\frac{\sigma_\lambda}{\mu_\lambda} \right)^2 \right)}{k(N-k)p - c^2 \left((1-p) + \left(\frac{\sigma_\lambda}{\mu_\lambda} \right)^2 \right)} \right] \quad (18)$$

and $N_0 = \lceil 1 + \frac{1}{\epsilon} \left(\frac{\sigma_\lambda}{\mu_\lambda} \right)^2 \rceil$, where $c = \frac{1}{\sqrt{\epsilon}}$.

Proof. Chebyshev's inequality states that the values of X are within the interval $[\mu_X - c\sigma_X, \mu_X + c\sigma_X]$ with a probability $p_0 \geq 1 - \epsilon$, where $\epsilon = \frac{1}{c^2}$. As the value of X is determined by the network (size and first and second moments of meeting rates across pairs), it follows that this will hold for at least a fraction p_0 of the networks. To apply the bound of Eq.(17), we can now consider X only inside the interval $[\mu_X - c\sigma_X, \mu_X + c\sigma_X]$. We can thus replace $\alpha = \mu_X - c\sigma_X$ and $\beta = \mu_X + c\sigma_X$, and use the values of μ_X and σ_X in Eq.(17). Considering only EM_2 leads us to the desired Eq.(18), after some calculations.

Furthermore, from Proposition 2.1, we set $c \cdot \sigma_X = c_\mu \cdot \mu_X$, where c_μ is a constant in $(0, 1)$ because X is always positive. Then it holds that $c \cdot \frac{\sigma_X}{\mu_X} < 1$ or equivalently $\left(c \cdot \frac{\sigma_X}{\mu_X} \right)^2 < 1$. Replacing $\epsilon = \frac{1}{c^2}$, μ_X and σ_X^2 we take $\frac{1}{\epsilon} \frac{(1-p) + \left(\frac{\sigma_\lambda}{\mu_\lambda} \right)^2}{k(N-k)p} <$

1. The previous expression takes its max value for $k = 1$, thus to be satisfied for every step k it must always hold that $N > 1 + \frac{1}{\epsilon} \frac{(1-p) + \left(\frac{\sigma_\lambda}{\mu_\lambda}\right)^2}{p} = N_0$. \square

From Theorem 2.4 it follows that an upper bound that holds for a fraction $1 - \epsilon$ of all the possible network instances cannot be always found for arbitrarily small N . The smaller the ϵ , the larger N is demanded. For very large networks the upper bound holds for every network instance. However, for little skewed contact mean rates distributions the bound exists and ϵ can be selected to be small even for moderate network sizes.

3 Applications

Having found the necessary approximations and bounds for individual epidemic steps in heterogeneous scenarios, we turn our attention to applications of these results, focusing on opportunistic networks. We first use the basic building blocks of our analysis to predict the “end-to-end” delay for epidemic broadcast and unicast routing [1], probabilistic flooding (gossiping) [5], and 2-hop routing [6]. We also use our result to derive a local estimator of network size N , and to analyze its convergence properties.

3.1 Epidemic Routing

The *broadcast* delay of epidemic spreading (i.e. the total time to deliver the message to every node), can be easily derived by adding up the expected delays of all steps (due to the linearity of expectation). Hence,

$$E[T_{epid}^{br}] = \sum_{k=1}^{N-1} E[T_{k,k+1}]. \quad (19)$$

In *unicast* routing, we are interested instead on the time until a given destination node receives the message. Assuming the selection of the source and destination nodes is random, the probability that the destination node is the k^{th} node to receive the message is the same for every k . Consequently,

$$E[T_{epid}^{uni}] = \frac{1}{N-1} \sum_{n=2}^N \sum_{k=1}^{n-1} E[T_{k,k+1}] = \frac{1}{N-1} \sum_{k=1}^{N-1} (N-k) E[T_{k,k+1}]. \quad (20)$$

For $E[T_{k,k+1}]$ we can use either the approximate value of Eq.(1) or Eq.(12), depending on the type of scenario at hand (Full or Poisson mixing). We can also derive lower and upper bounds using the respective results from Section 2. As an example, the 2nd order approximation for epidemic unicast routing in a Full mixing scenario is

$$E[T_{epid}^{uni}] \approx \frac{\ln(N)}{N} \frac{1}{\mu_\lambda} + \frac{1.65N + 2\ln(N)}{N^3} \frac{\sigma_\lambda^2}{\mu_\lambda^3}. \quad (21)$$

and the details for its derivation are given in Appendix 7.1.

3.2 Probabilistic Routing - Gossiping

In probabilistic routing, when two nodes meet, the node that carries the message forwards it to the other node with probability p_r (except if the other node is the destination node). This is equivalent to *thinning* the (Poisson) contact process between two pairs [22]. The meeting events of a node pair will be a Poisson process with decreased rate $\lambda'_{ij} = p_r \lambda_{ij}$. Thus, the mean value and variance of the rates become $\mu'_\lambda = p_r \mu_\lambda$ and $\sigma'^2_\lambda = p_r^2 \sigma_\lambda^2$. We can replace these values in any of the equations in Section 2 to get the mean step delay $E[T_{k,k+1}]$. We can then use the epidemic equation above (Eq.(20)) to get approximations or bounds for the unicast delay. Replacing μ'_λ and σ'^2_λ in Eq.(21) gives us a 2nd order approximation for probabilistic routing: $E[T_{rnd}^{uni}] = \frac{E[T_{epid}^{uni}]}{p_r}$.

3.3 2-hop Routing

In the 2-hop routing scheme, the source sends the message to every node it meets, like in epidemic routing. However, other nodes receiving the message can only give it directly to the destination, when and if they encounter it. We present here the delay approximation in Poisson mixing scenarios (see Appendix 7.2).

$$E[T_D^{2hop}] = \left(\frac{1}{p\mu_\lambda} + \frac{(1-p) + \left(\frac{\sigma_\lambda}{\mu_\lambda}\right)^2}{p^2\mu_\lambda(N-1)} \right) \sum_{k=1}^{N-1} \frac{(N-2)!}{(N-k-1)! (N-1)^k} \quad (22)$$

3.4 Network Size Estimator

As a final example of an application of our analysis, we turn our attention to the problem of distributed estimation of the network size N . In opportunistic networks, a node might want to know the network size for application reasons (e.g. to join the largest network around) or to tune specific protocol parameters (e.g. TTL). As one example, in Spray and Wait routing the number of copies needed in order to achieve a given performance relative to the optimal one depends on the network size N [4]. If a new node joins the network, or N changes over time due to nodes joining and leaving, it will need to derive and maintain an estimate of N .

One way would be to count the *unique* number of nodes met. However, this process may lead to very long delays as *all* nodes must be encountered. While node collaboration could speed up this process, it comes with a communication overhead and the need to keep N unique IDs and compare them at every meeting. For this reason, approximate methods have been proposed [4, 28], based on either purely local [4] or distributed algorithms [28].

The advantage of purely local approaches [4] is much faster convergence with less overhead (albeit at a cost of a small statistical error). However, this approach

applies to stable, homogeneous networks only and cannot be directly applied to heterogeneous networks. The approach we propose here requires each node to maintain a *locally measured* estimate related to contact times. If node i just met some node, then it will measure the time until it meets *any* other node. Let us denote such consecutive samples as $t_{1,n}^i$ and the estimator $\hat{T}_1^i = \frac{1}{m} \sum_{n=1}^m t_{1,n}^i$. Then, it holds that

$$\lim_{m \rightarrow \infty} E \left[\frac{1}{m} \sum_{n=1}^m t_{1,n}^i \right] \rightarrow \frac{1}{\sum_j \lambda_{ij}}. \quad (23)$$

In other words, the local estimator \hat{T}_1^i converges to $E[T_{k,k+1}]$, *as measured from the perspective of node i* : each $t_{1,n}^i$ is exponentially (and independently) distributed with rate $\sum_j \lambda_{ij}$. In practice, this estimator would be implemented with an exponential weighted moving average (as in [4]).

\hat{T}_1^i depends on the size of the network N . We could thus try to build an estimator out of it. Let us assume full mixing and define μ_λ^i as $\mu_\lambda^i = \frac{1}{N-1} \sum_j \lambda_{ij}$. Then it holds that

$$E \left[\hat{T}_1^i \right] = \frac{1}{(N-1)\mu_\lambda^i}. \quad (24)$$

If we had a perfect estimator for \hat{T}_1^i and knew μ_λ^i we could use the following estimator $\hat{N} = \frac{1}{\hat{T}_1^i \mu_\lambda^i} + 1$. However, measuring μ_λ^i requires seeing *all* nodes *multiple* times in order to get reliable estimates for the λ_{ij} rates involved, as well as maintaining IDs. In addition, as nodes join and leave this value change over time, possibly faster than the estimation process.

Instead, of measuring and using the *actual* value of μ_λ^i each node could use the expected value μ_λ , the mean value of $F(\lambda)$. This is a *structural property* of the network, independent of network size and instance, and could be known a priori or slowly estimated with less overhead per time. Based on these observations, the following theorem derives an estimator for the network size \hat{N} based on \hat{T}_1^i , μ_λ (and p for sparse mixing). The proof follows a similar methodology as those in Section 2 and can be found in Appendix 7.3.

Theorem 3.1. *For a network with N nodes, in which the nodes meet according to the Heterogeneous Contact Class of Definition 2 - with parameters μ_λ , σ_λ and p - an estimator for the total number of nodes in the network is*

$$\hat{N} = \frac{1}{p} \cdot \frac{1}{\hat{T}_1^i \cdot \mu_\lambda} + 1, \quad (25)$$

where $\hat{T}_1^i = \frac{1}{m} \sum_{k=1}^m t_{1,k}^i$. For this estimator, it holds that

$$E[\hat{N}] \approx 1 + (N - 1) \cdot \frac{1}{1 + \frac{(1-p) + \frac{\sigma_\lambda^2}{\mu_\lambda}}{p(N-1)}} \cdot \left[1 + \frac{1}{m} \frac{1 + 3 \frac{(1-p) + \frac{\sigma_\lambda^2}{\mu_\lambda}}{p(N-1)}}{\left(1 + \frac{(1-p) + \frac{\sigma_\lambda^2}{\mu_\lambda}}{p(N-1)} \right)^2} \right], \quad (26)$$

that is, it is asymptotically unbiased as the number of samples m and network size increase.

Remark: We stress here that it is beyond the scope of this paper to derive the best estimator possible, or show how it should be implemented. Our goal here was only to show how our model could be used to derive a reasonable estimator for heterogeneous networks and analyze its performance.

4 Model Validation

In order to validate the accuracy of our results we first compare them (in Section 4.1) against simulations of various *synthetic* scenarios belonging to the Heterogeneous Contact Classes of Definitions 1 and 2.

4.1 Model Validation

Since our theoretical results are only asymptotically accurate in N , the network size, we use Monte Carlo simulations to examine the accuracy of our various analytical expressions (i) in finite size networks and (ii) as a function of other parameters of interest (e.g. statistics of the network generating function $F(\lambda)$).

In each simulation, we create a network of N nodes and a contact pattern by generating a $N \times N$ matrix $\mathbf{\Lambda} = \{\lambda_{ij}\}$. Each entry λ_{ij} characterizes the contact process of the pair of nodes i and j : it is zero with probability $1 - p$ (in Poisson mixing), otherwise it takes values drawn from a chosen distribution $F(\lambda)$ with mean μ_λ and variance σ_λ^2 . Then for the pairs that meet ($\lambda_{ij} > 0$) we generate a sequence of contact events with exponentially distributed intercontact times with rate $\lambda_{ij} > 0$.

For every simulation point (network size / mobility pattern pair), we run 10000 simulations and calculate the average message delay. We have considered various contact rate distributions $F(\lambda)$ and parameters. We compare the simulation value to the approximations, lower and upper bounds derived in Section 2. Since in most cases $0 < \lambda < \infty$, we apply Theorem 2.4 for the upper bound (which holds for 95% of the possible networks, in the sense of this theorem).

Fig. 3 shows the delay for epidemic broadcast in dense networks (Full Mixing scenarios). We present three different mobility cases with similar inter-meeting time

mean rate but different variance. There are a couple of interesting things to observe in these figures: (i) the accuracy of the approximation and the tightness of the bounds is significant even for small networks of less than 100 nodes, (ii) accuracy improves with network size and lower variance, as our theory predicts, and (iii) in Fig. 3(c) the upper bounds for 95% of the cases exist only after a network size N_0 , in the sense of Theorem 2.4.

In Fig. 4, we present the simulated delay and the analytical results for epidemic broadcast in sparse networks (Poisson Mixing scenarios). The previous observations (Fig. 3) hold also in these cases. Additionally, we can observe that for sparse networks analytical values converge slower (i.e. for larger network sizes) compared to the dense networks³. Especially, the more sparse a network is (i.e. smaller p), the slower the analytical values converge.

In Fig. 5, we also present some results for unicast routing schemes, namely for the delay of epidemic (unicast) routing and probabilistic routing, as described in Section 3. The plots again suggest significant accuracy for our analytical results already for reasonable network sizes, and convergence properties consistent with the theory.

4.2 Real Mobility Traces

The results of the previous section show that our analytical predictions achieve significant accuracy even in finite networks whose mobility patterns fall under the Full or Poisson Mixing heterogeneous contact classes. While these contact classes are rather broad, whether they capture “real” scenarios, and to what extent, depends on the application setting, contact scenario, etc.

In the context of opportunistic networks, some mobility traces collected in real experiments and/or networks do exist. Arguably, the size of most of them is small and they represent each only a single instance of the random mobility process at play, often with a number of measurement complications and errors. Nevertheless, it is of interest to see how our performance predictors behave in some of these scenarios, and whether they can capture the quantities of interest (even if qualitatively), despite the considerably higher complexity (e.g. community structure) of such scenarios, and departures from the assumptions for which our predictors are designed.

To this end, we use the following sets of real mobility traces: (i) *Cabspotting* [29], which contains GPS coordinates from 536 taxi cabs collected over 30 days in San Francisco, and (ii) *Infocom* [30], which contains traces of Bluetooth sightings of 78 mobile nodes from the 4 days iMotes experiment during Infocom 2006. We also generated mobility traces with three recent mobility models that have been shown to capture well different aspects of real mobility traces, namely, TVCM [31], HCMM [32] and SLAW [21]. In order to compare with analysis, we parse each trace and estimate the mean contact rate for all pairs $\{i, j\}$. We then produce

³The absence of delay values for small networks is due to lack of connectivity of the contact graph.

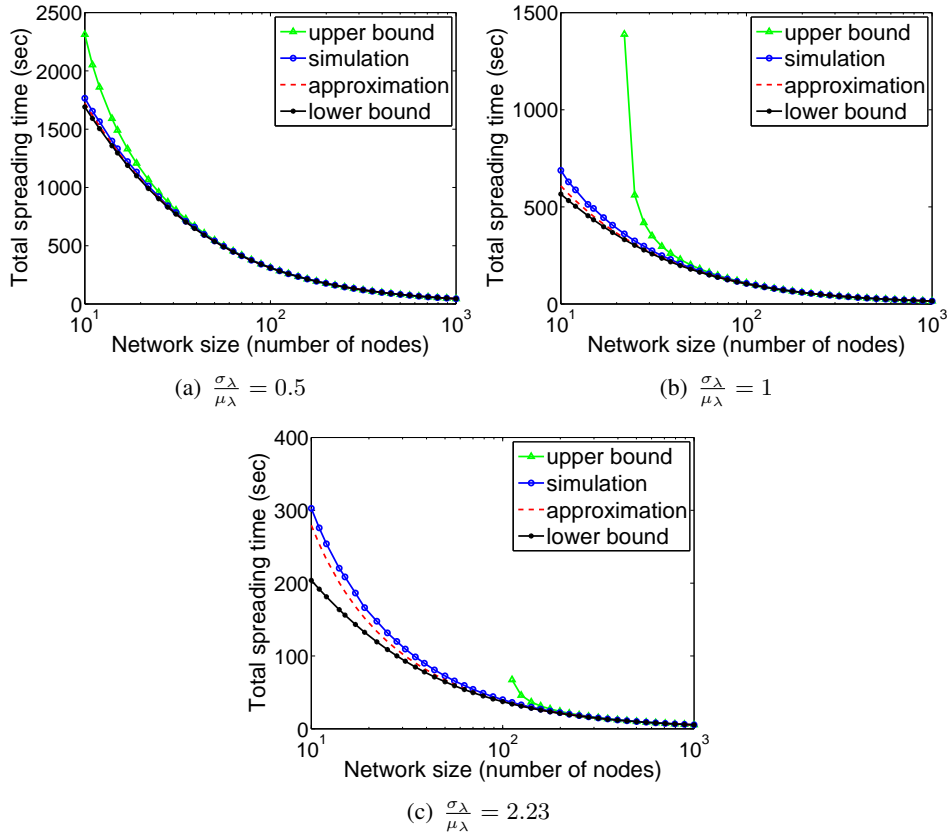


Figure 3: Delay of epidemic spreading (broadcast) in Dense (Full Mixing) networks. Contact rates are distributed according to (a): $F(\lambda) \sim \text{Log-normal}$, $\mu_\lambda = 0.33 \cdot 10^{-3} \text{sec}^{-1}$, $\frac{\sigma_\lambda}{\mu_\lambda} = 0.5$, (b): $F(\lambda) \sim \text{Exponential}$, $\mu_\lambda = 1 \cdot 10^{-3} \text{sec}^{-1}$, $\frac{\sigma_\lambda}{\mu_\lambda} = 1$ and (c): $F(\lambda) \sim \text{Pareto}$, $\mu_\lambda = 2.8 \cdot 10^{-3} \text{sec}^{-1}$, $\frac{\sigma_\lambda}{\mu_\lambda} = 2.23$.

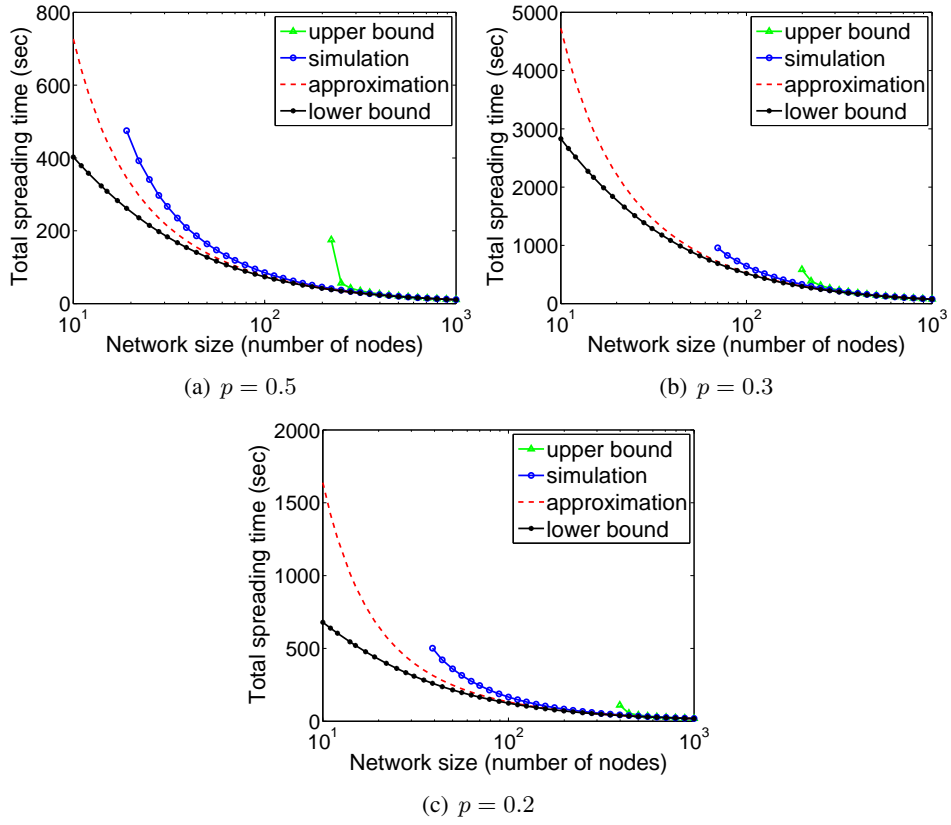


Figure 4: Delay of epidemic spreading (broadcast) in Sparse (Poisson Mixing) networks, with different meeting probabilities. Contact rates are distributed according to (a): $F(\lambda) \sim \text{Pareto}$, $\mu_\lambda = 2.8 \cdot 10^{-3} \text{sec}^{-1}$, $\frac{\sigma_\lambda}{\mu_\lambda} = 2.23$, (b): $F(\lambda) \sim \text{Log-normal}$, $\mu_\lambda = 2.8 \cdot 10^{-3} \text{sec}^{-1}$, $\frac{\sigma_\lambda}{\mu_\lambda} = 2.23$ and (c): $F(\lambda) \sim \text{Exponential}$, $\mu_\lambda = 1 \cdot 10^{-3} \text{sec}^{-1}$, $\frac{\sigma_\lambda}{\mu_\lambda} = 1$.

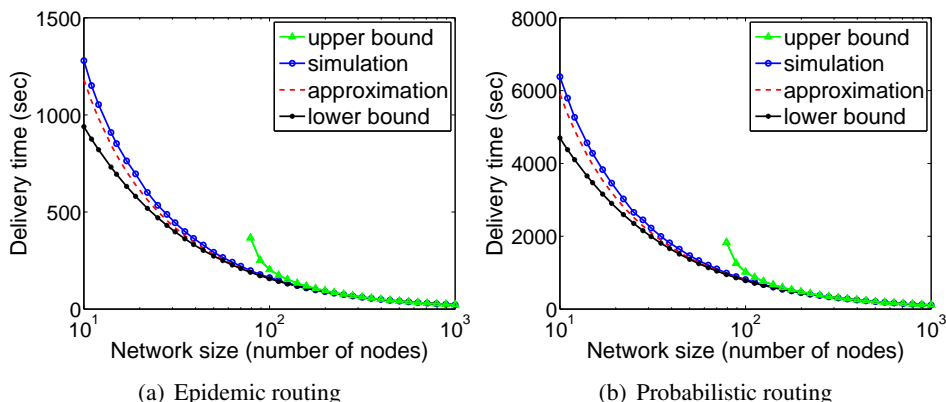


Figure 5: Delay of two unicast routing protocols in a full mixing network. (a) Epidemic routing and (b) Probabilistic routing with $p_r = 0.2$. Contact rates are distributed according to $F(\lambda) \sim \text{Log-normal}$, $\mu_\lambda = 0.33 \cdot 10^{-3} \text{sec}^{-1}$, $\frac{\sigma_\lambda}{\mu_\lambda} = 1.86$.

estimates for the 1st and 2nd moments of these rates, $\hat{\mu}_\lambda$ and $\hat{\sigma}_\lambda^2$, as well as the percentage of connected pairs \hat{p} and use them in our analytical expressions.

4.2.1 Message Spreading

As our first example, we consider unicast routing. Fig. 6 shows the message delay under epidemic and 2-hop routing. Source and destination are chosen randomly in different runs and messages are generated in random points of the trace.

The first thing to observe is that delay values span a wide range of values for different source-destination pairs. This implies a large amount of heterogeneity in the “reachability” of different nodes. Our analytical predictions are shown as thick dark horizontal lines. As it can be seen, our result is in most cases close to the median and in almost all cases between the 25th and 75th percentile of the delay observed in both the real traces and mobility models⁴.

4.2.2 Network Size Estimation

In this second example, we will check the performance of our network estimator. We assign to each node in the network an estimator that counts its inter-meeting time intervals and computes the network size according to Theorem 3.1. In Fig. 7(a), we consider the cabspotting trace. We form sub-networks of different sizes by picking a subset of nodes randomly from the original trace. Somewhat surprisingly, it can be observed that our estimator closely follows the actual network size. Note that we have used the values of μ_λ and p of the original network (536 nodes) for all the considered sub-networks. This supports our argument that

⁴We tend to underestimate the delay in SLAW, a mobility model that was designed to capture power-law characteristics of contact meetings [21]

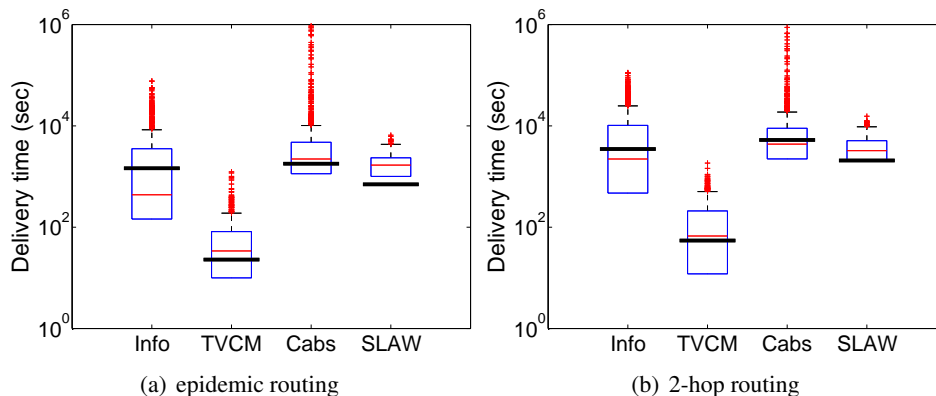


Figure 6: Box-plots of the unicast (i.e. message delivery) delay under epidemic and 2-hop routing. On each box, the central horizontal line is the median, the edges of the box are the 25th and 75th percentiles, the whiskers extend to the most extreme data points not considered outliers, and outliers are plotted individually as crosses. The thick lines represent the theoretical values predicted by our model.

such general properties of the mobility process provide sufficient knowledge to the estimator. Nodes could thus estimate them once (paying the overhead involved) and reuse them for much longer periods of time.

We have also simulated our estimator in the sythetic mobility scenarios and have found an acceptable match as well. Results are shown in Fig. 7(b). The reason for the lower accuracy of the estimations for the scenarios generated by HCMM and TVCM, compared to the scenario generated by SLAW, is that these mobility models generate networks with significant community structure [31, 32].

While nodes seem to find the correct size, *on average*, intuition suggests that, in finite size networks, some nodes will be “faster” on average than others (i.e. have a higher $\sum_j \lambda_{ij}$), due to the randomness of each pairwise rate. It can be seen from our discussion of the estimator (Section 3.4) and Eq.(24) that faster nodes will overestimate the network size while slower nodes will underestimate it.

To this end, we plot in Fig. 8 the experimental CDF of the network size estimators for two scenarios. In the SLAW mobility scenario (where the real network value is 100 and the mean value of our estimate is 94), Fig. 8(b), the individual estimations of different nodes are not spread very far from their mean value, which means that the majority of the nodes estimate quite well the size of the network. The eCDF for the cabspotting scenario, Fig. 8(a) shows a bit more spread, but the majority of the nodes estimate the size of their network with relatively low error. Nevertheless, the results make a case for an enhancing mechanism in which nodes could occasionally average their estimates over some contacts, in order to obtain even more accurate estimates.

Summarizing, it is somewhat remarkable that our estimator and delay predictor are close to the actual results (qualitatively or even quantitatively in some cases) in

a range of real or realistic scenarios; studies of these scenarios reveal considerable differences to the much simpler contact classes for which our results are derived. We should also be careful not to jump to generalizations about the accuracy of these results in all real scenarios, as we are aware of situations that could force our predictors to err significantly. Nevertheless, we believe these results are quite promising in the direction of finding simple, usable analytical expressions even for complex, heterogeneous contact scenarios.

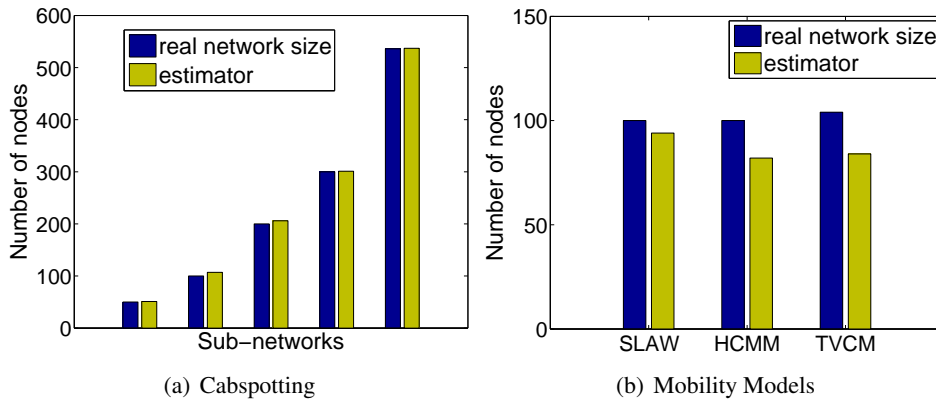


Figure 7: Network size estimator. (a): Estimator compared to real values for each sub-network of the Cabspotting trace, (b): Estimator compared to real values for traces generated by three mobility models.

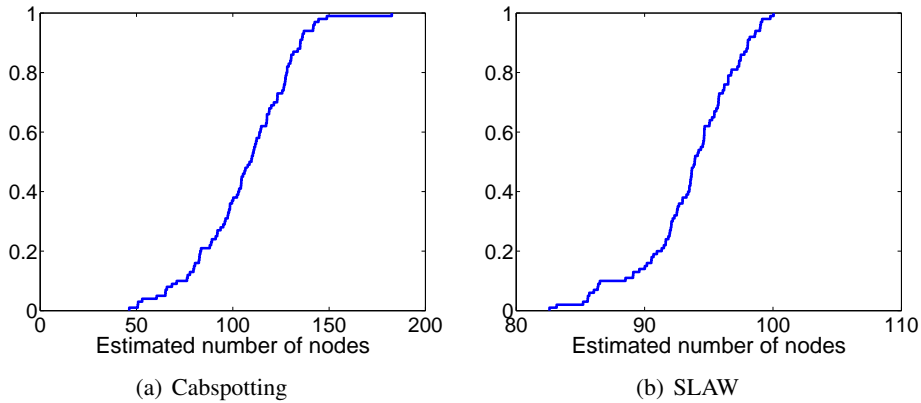


Figure 8: Network size estimator. ECDF of estimations among all nodes for (a): a Sub-network of 100 nodes of the Cabspotting trace, (b): a trace of 100 nodes, generated by the SLAW mobility model

5 Related Work

Models for epidemic spreading of diseases [25] and/or computer malware [15], were early derived, based on the well known SIR model, and studied widely. In DTNs, efforts to analyze the performance of epidemic routing and other protocols also abound. Stochastic analyses, like the one in [6], define a Markov chain as in Fig. 1, in order to give closed form results for epidemic and 2-hop routing. Fluid models [5, 7, 33], take an approach similar to the SIR model in biology, and define the number of messages in the network as a continuous function (of time). Then, ordinary differential equations (ODEs) are used to derive expressions for the total delay, delivery probability etc. While these models provide closed form results and thus can be used in tuning protocol parameters (e.g. gossiping probability [5], number of copies [4], TTL [34]), they all assume a homogeneous network with a common meeting rate for every pair of nodes.

Recent studies on real network traces [8–11] suggest that the homogeneity assumption is not true. To overcome this limitation, a number of works has introduced heterogeneity in contact network models, by allowing different meeting rates for each node pair [9, 11, 12, 14, 35]. Yet, most of these works use the heterogeneous model to design new, better protocols (e.g. multicast [9] or unicast [11]) that take heterogeneity into account, but do not analyze performance. One exception is [14], but only for the cases of direct transmission and 2-hop routing. To our best knowledge, the work closer to this paper is that of [35], where a very generic contact graph is considered. However, due to the large generality of the contact model, only upper bounds for the delay can be provided.

In our work, while we allow arbitrary link rates between nodes, as in [9, 14, 35], we restrict the underlying contact graph model, in order to derive closed-form approximations and bounds. We validated our results with synthetic simulations for the targeted contact classes, as previous work did [4–6], but also demonstrate their applicability in real networks.

As a final note, in theoretical biology and epidemiology, there are many studies trying to capture heterogeneity and model it in the context of complex networks [36, 37], using different levels of mixing [38] or stratified populations [39]. Yet, the majority of these works focus on deriving thresholds above which the epidemic will spread and their results usually consider infinite time.

6 Conclusions

In this paper, we have considered two classes of heterogeneous contact models, and have derived simple closed form approximations and bounds for epidemic spreading. From the validation of the model against synthetic models and realistic traces we can conclude that: (a) simple delay expressions, that can be used for performance prediction and protocol optimization, exist not only for the homogeneous contact case; (b) performance predictions that are accurate qualitatively, and

(somewhat more surprisingly) sometimes quantitatively also, can be made even for a number of real scenarios, despite the highly more complex structure of the latter. We believe that our methodology might be applicable in more generic contact graph structures as well (e.g. the configuration model [25]). However, we also think that a limit exists, probably related to the size of the min-cut of the contact graph, beyond which are techniques are not applicable and other methods are needed (e.g. [35]). We intend to explore these issues further as part of future work.

7 Appendix

7.1 Approximation for Expected Unicast Delay under Epidemic Routing

Substituting Eq.(1) in Eq.(20) we get

$$\begin{aligned}
E[T_{epid}^{uni}] &= \frac{1}{N-1} \sum_{k=1}^{N-1} (N-k) E[T_{k,k+1}] \\
&= \frac{\sum_{k=1}^{N-1} (N-k) \cdot \left(\frac{1}{k(N-k)\mu_\lambda} + \frac{\sigma_\lambda^2}{[k(N-k)]^2 \mu_\lambda^3} \right)}{N-1} \\
&= \frac{1}{N-1} \sum_{k=1}^{N-1} \left(\frac{1}{k\mu_\lambda} + \frac{\sigma_\lambda^2}{k^2(N-k)\mu_\lambda^3} \right) \\
&= \frac{1}{(N-1)\mu_\lambda} \left[\sum_{k=1}^{N-1} \frac{1}{k} + \frac{\sigma_\lambda^2}{\mu_\lambda^2} \sum_{k=1}^{N-1} \frac{1}{k^2(N-k)} \right]
\end{aligned} \tag{27}$$

Since $\sum_{k=1}^{N-1} \frac{1}{k}$ is an *harmonic sum* it can be approximated as [40]

$$\sum_{k=1}^{N-1} \frac{1}{k} \approx \ln(N-1). \tag{28}$$

Also, using *partial fraction decomposition*, the second sum in Eq.(27) becomes

$$\sum_{k=1}^{N-1} \frac{1}{k^2(N-k)} = \frac{1}{N^2} \left(\sum_{k=1}^{N-1} \frac{1}{k} + \sum_{k=1}^{N-1} \frac{1}{N-k} + N \sum_{k=1}^{N-1} \frac{1}{k^2} \right). \tag{29}$$

Using the approximation for the *harmonic sum* (Eq.(28)) and the approximation [40]

$$\sum_{k=1}^{N-1} \frac{1}{k^2} \approx 1.65, \tag{30}$$

in Eq.(29) we get

$$\sum_{k=1}^{N-1} \frac{1}{k^2(N-k)} = \frac{1.65N + 2 * \ln(N-1)}{N^2}. \tag{31}$$

Replacing Eq.(28) and Eq.(31) in Eq.(27) and approximating $N - 1 \approx N$ we result in Eq.(21).

7.2 Delay Approximation under 2-hop Routing

Under 2-hop routing, in *step* k there are k nodes that carry the message (the source and $k - 1$ relays). As relays can forward the message only to the destination node and the source to everyone it meets, there are $N - 1$ possible meeting events in which a message exchange can take place, i.e. (a) $N - k - 1$ possible meetings between the source and a non-infected node, other than the destination, and (b) k possible meetings between the infected nodes (including the source) and the destination. Due to randomness, the probability that the destination node will be involved in the exact next meeting event with message exchange is $\frac{k}{(N-k-1)+k} = \frac{k}{N-1}$.

Then, if $T_{k,D}^{2hop}$ is the time from *step* k till the message delivery and $T_{k,k+1}^{2hop}$ the time till the next meeting event with message exchange, it holds that

$$E[T_{k,D}^{2hop}] = E[T_{k,k+1}^{2hop}] + \left(1 - \frac{k}{N-1}\right) E[T_{k+1,D}^{2hop}],$$

and recursively we result in:

$$E[T_D^{2hop}] = \sum_{k=1}^{N-1} \frac{(N-1)!}{(N-1)^k} E[T_{k,k+1}^{2hop}]. \quad (32)$$

In Eq.(32) the expectation $E[T_{k,k+1}^{2hop}]$ corresponds to the delay of *step* k for 2-hop routing, i.e. the expected time to go from k infected nodes to $k + 1$. Compared to the epidemic routing, the difference is that now, in *step* k , the meeting events *with message exchange* are only $N - 1$, instead to $k(N - k)$. We thus need to replace the value $k(N - k)$ in the denominator of all results in Section 2 with $N - 1$ to get the respective value for $E[T_{k,k+1}^{2hop}]$. For example, the equivalent of Eq.(1) under 2-hop routing is

$$E[T_{k,k+1}] = \frac{1}{(N-1)\mu_\lambda} + \frac{\sigma_\lambda^2}{[(N-1)]^2\mu_\lambda^3}. \quad (33)$$

The previous arguments can be applied also in the Poisson Mixing cases. So if we replace in Eq.(12) the terms $k(N - k)$ with $(N - 1)$ we will get the expected step delay under 2-hop routing $E[T_D^{2hop}]$ for Poisson Mixing scenarios and if we further use the result to Eq.(32) we will end up to Eq.(22).

7.3 Asymptotical Behavior of the Network Size Estimator

Proof. The proposed estimator for the network size, i.e. the total number of nodes in the network is

$$\hat{N} = \frac{1}{p} \cdot \frac{1}{\hat{T}_1^i \cdot \mu_\lambda} + 1 \quad (34)$$

where $\hat{T}_1^i = \frac{1}{m} \sum_{k=1}^m t_{1,k}^i$.

The expectation of \hat{N} is given by

$$E[\hat{N}] = \frac{1}{p \cdot \mu_\lambda} \cdot E \left[\frac{1}{\hat{T}_1^i} \right] + 1 \quad (35)$$

and for computing it we need first to compute $E \left[\frac{1}{\hat{T}_1^i} \right]$. Thus

$$E \left[\frac{1}{\hat{T}_1^i} \right] = E \left[\frac{1}{\frac{1}{m} \sum_{k=1}^m t_{1,k}^i} \right] = m \cdot E \left[\frac{1}{\sum_{k=1}^m t_{1,k}^i} \right] \quad (36)$$

Approximating the expectation in Eq.(36) with a Taylor sum (applying the *Delta Method*, as in Section 2) we get

$$E \left[\frac{1}{\hat{T}_1^i} \right] \approx m \cdot \left(\frac{1}{E[\sum_{k=1}^m t_{1,k}^i]} + \frac{\text{var} \left(\sum_{k=1}^m t_{1,k}^i \right)}{\left(E[\sum_{k=1}^m t_{1,k}^i] \right)^3} \right) \quad (37)$$

In Eq.(36), the different samples $t_{1,k}^i$ are IID random variables, so, it holds that [22]

$$E \left[\sum_{k=1}^m t_{1,k}^i \right] = m \cdot E[t_1^i] \quad (38)$$

$$\text{var} \left(\sum_{k=1}^m t_{1,k}^i \right) = m \cdot \text{var}(t_1^i) \quad (39)$$

t_1^i is exponentially distributed with rate $\sum_j \lambda_{ij}$. If we denote $X = \sum_j \lambda_{ij}$, then the probability density function of X , i.e. f_X , is the convolution of the pdf of the different λ_{ij} involved in the sum. Applying the property of conditional expectation, we have

$$\begin{aligned} E[t_1^i] &= \int_0^\infty E[t_1^i | X] f_X(X) dX \\ &= \int_0^\infty \frac{1}{X} f_X(X) dX = E \left[\frac{1}{X} \right] \end{aligned} \quad (40)$$

where the equality between the first and the second line holds because t_1^i is exponentially distributed with rate X . With the same arguments we get also the variance of t_1^i

$$\begin{aligned} \text{var} \left(t_1^i \right) &= \int_0^\infty \text{var} \left([t_1^i | X] \right) f_X(X) dX \\ &= \int_0^\infty \frac{1}{X^2} f_X(X) dX = E \left[\frac{1}{X^2} \right] \end{aligned} \quad (41)$$

To compute $E\left[\frac{1}{X}\right]$ and $E\left[\frac{1}{X^2}\right]$, Eq.(40) and Eq.(41), we will apply once more the *Delta method* and consider, as before, the first three terms of the Taylor sum. Thus,

$$E[t_1^i] \approx \frac{1}{E[X]} \cdot \left(1 + \frac{\text{var}(X)}{(E[X])^2}\right) \quad (42)$$

$$\text{var}(t_1^i) \approx \frac{1}{(E[X])^2} \cdot \left(1 + \frac{3\text{var}(X)}{(E[X])^2}\right) \quad (43)$$

Substituting the values of equations (38), (39), (40) and (41) in Eq.(37), after a few simple steps we end up to

$$E\left[\frac{1}{\hat{T}_1^i}\right] \approx \frac{E[X]}{1 + \frac{\text{var}(X)}{(E[X])^2}} \cdot \left(1 + \frac{1}{m} \cdot \frac{1 + 3\frac{\text{var}(X)}{(E[X])^2}}{\left(1 + \frac{\text{var}(X)}{(E[X])^2}\right)^2}\right) \quad (44)$$

Node i can meet $N - 1$ other nodes each one with probability p and their meeting rate is drawn from $F(\lambda)$ and is also independent of the other rates. It means that $X = \sum_j \lambda_{ij}$ is a sum a random number of IID random variables λ_{ij} , so, its expectation and variance are given by (Section 2.2)

$$E[X] = p(N - 1)\mu_\lambda \quad (45)$$

$$\text{var}(X) = p(N - 1) \left(\sigma_\lambda^2 + (1 - p)\mu_\lambda^2\right) \quad (46)$$

Now if we use the values of Eq.(45) and Eq.(46) in Eq.(44) and then apply the result to Eq.(35) we end up to Eq.(26).

For large networks, i.e. large N , the ratio $\frac{(1-p) + \frac{\sigma_\lambda^2}{\mu_\lambda^2}}{p(N-1)}$ becomes much smaller than 1 and Eq.(26) is approximately

$$E[\hat{N}] = (N - 1) \cdot \left(1 + \frac{1}{m}\right) + 1 \quad (47)$$

Now it is clear that for a large number of samples, m , the expectation of the estimator equals the real value of the network size N .

□

References

- [1] A. Vahdat and D. Becker, “Epidemic routing for partially connected ad hoc networks,” Tech. Rep. CS-200006, Duke University, 2000.
- [2] T. Spyropoulos, R. N. Rais, T. Turletti, K. Obraczka, and A. Vasilakos, “Routing for disruption tolerant networks: taxonomy and design,” *Wirel. Netw.*, vol. 16, pp. 2349–2370, Nov. 2010.

- [3] L. Pelusi, A. Passarella, and M. Conti, "Opportunistic networking: data forwarding in disconnected mobile ad hoc networks," *Comm. Mag., IEEE*, vol. 44, pp. 134–141, Nov. 2006.
- [4] T. Spyropoulos, K. Psounis, and C. S. Raghavendra, "Efficient routing in intermittently connected mobile networks: the multiple-copy case," *IEEE/ACM Trans. Netw.*, vol. 16, pp. 77–90, Feb. 2008.
- [5] X. Zhang, G. Neglia, J. Kurose, and D. Towsley, "Performance modeling of epidemic routing," *Computer Networks*, vol. 51, no. 10, pp. 2867–2891, 2007.
- [6] R. Groenevelt, P. Nain, and G. Koole, "The message delay in mobile ad hoc networks," *Performance Evaluation*, vol. 62, pp. 210–228, 2005.
- [7] Z. J. Haas and T. Small, "A new networking model for biological applications of ad hoc sensor networks," *IEEE/ACM Transactions on Networking (TON)*, vol. 14, pp. 27–40, 2006.
- [8] W. Hsu and A. Helmy, "On nodal encounter patterns in wireless lan traces," *IEEE Transactions on Mobile Computing*, vol. 9, pp. 1563–1577, 2010.
- [9] W. Gao, Q. Li, B. Zhao, and G. Cao, "Multicasting in delay tolerant networks: a social network perspective," in *Proc. of ACM MobiHoc*, 2009.
- [10] V. Conan, J. Leguay, and T. Friedman, "Characterizing pairwise inter-contact patterns in delay tolerant networks," in *Proc. ACM Autonomics*, 2007.
- [11] P. Hui, J. Crowcroft, and E. Yoneki, "Bubble rap: Social-based forwarding in delay-tolerant networks," *IEEE Transactions on Mobile Computing*, vol. 10, pp. 1576–1589, 2011.
- [12] T. Spyropoulos, T. Turletti, and K. Obraczka, "Routing in delay-tolerant networks comprising heterogeneous node populations," *IEEE Transactions on Mobile Computing*, vol. 8, pp. 1132–1147, 2009.
- [13] C. Boldrini, M. Conti, and A. Passarella, "Less is more: Long paths do not help the convergence of social-oblivious forwarding in opportunistic networks," in *Proc. ACM MobiOpp*, 2012.
- [14] C.-H. Lee and D. Y. Eun, "On the forwarding performance under heterogeneous contact dynamics in mobile opportunistic networks," *IEEE Transactions on Mobile Computing*, vol. 99, no. PrePrints, 2012.
- [15] Y. Wang, D. Chakrabarti, C. Wang, and C. Faloutsos, "Epidemic spreading in real networks: An eigenvalue viewpoint.," in *SRDS*, pp. 25–34, IEEE Computer Society, 2003.

- [16] T. Karagiannis, J.-Y. Le Boudec, and M. Vojnović, “Power law and exponential decay of inter contact times between mobile devices,” in *Proc. ACM MobiCom*, 2007.
- [17] H. Cai and D. Y. Eun, “Crossing over the bounded domain: from exponential to power-law intermeeting time in mobile ad hoc networks,” *IEEE/ACM Trans. Netw.*, vol. 17, pp. 1578–1591, Oct. 2009.
- [18] D. Aldous and J. Fill, “Reversible markov chains and random walks on graphs. (monograph in preparation.)” <http://stat-www.berkeley.edu/users/aldous/RWG/book.html>.
- [19] A. Passarella and M. Conti, “Characterising aggregate inter-contact times in heterogeneous opportunistic networks,” in *Proc. of IFIP - Volume Part II, NETWORKING’11*, pp. 301–313, 2011.
- [20] A. Chaintreau, P. Hui, J. Crowcroft, C. Diot, R. Gass, and J. Scott, “Impact of human mobility on the design of opportunistic forwarding algorithms,” in *Proc. IEEE INFOCOM*, 2006.
- [21] K. Lee, S. Hong, S. J. Kim, I. Rhee, and S. Chong, “Slaw: A new mobility model for human walks,” in *IEEE INFOCOM*, 2009.
- [22] S. M. Ross, *Stochastic Processes*. John Wiley & Sons, 2 ed., 1996.
- [23] E. Cooch and G. White, *Program MARK A gentle Introduction*. 11 ed., 2012.
- [24] G. W. Oehlert, “A note on the delta method,” *The American Statistician*, vol. 46, no. 1, pp. 27–29, 1992.
- [25] M. Newman, *Networks: An Introduction*. New York, NY, USA: Oxford University Press, Inc., 2010.
- [26] A. Madansky, “Bounds on the expectation of a convex function of a multivariate random variable,” *The Annals of Mathematical Statistics*, vol. 30, no. 3, pp. 743–746, 1959.
- [27] S. P. Dokov and D. P. Morton, “Higher-order upper bounds on the expectation of a convex function,” *Stochastic Programming EPrint Series*, 2002.
- [28] A. Guerrieri, I. Carreras, F. D. Pellegrini, D. Miorandi, and A. Montessor, “Distributed estimation of global parameters in delay-tolerant networks,” *Computer Communications*, vol. 33, no. 13, pp. 1472–1482, 2010.
- [29] M. Piorkowski, N. Sarafijanovic-Djukic, and M. Grossglauser, “CRAW-DAD data set epfl/mobility (v. 2009-02-24).” Downloaded from <http://crawdad.cs.dartmouth.edu/epfl/mobility>, Feb. 2009.

- [30] J. Scott, R. Gass, J. Crowcroft, P. Hui, C. Diot, and A. Chaintreau, "CRAWDAD data set cambridge/haggle (v. 2009-05-29)." Downloaded from <http://crawdad.cs.dartmouth.edu/cambridge/haggle>, May 2009.
- [31] W.-J. Hsu, T. Spyropoulos, K. Psounis, and A. Helmy, "Modeling spatial and temporal dependencies of user mobility in wireless mobile networks," *IEEE/ACM Transactions on Networking (TON)*, vol. 17, no. 5, pp. 1564–1577, 2009.
- [32] C. Boldrini and A. Passarella, "Hcmm: Modelling spatial and temporal properties of human mobility driven by users' social relationships," *Computer Communications*, vol. 33, no. 9, pp. 1056–1074, 2010.
- [33] A. Khelil, C. Becker, J. Tian, and K. Rothermel, "An epidemic model for information diffusion in manets," in *Proc. of ACM MSWiM*, 2002.
- [34] R. Lu, X. Lin, T. H. Luan, X. Liang, X. Li, L. Chen, and X. Shen, "Prefilter: An efficient privacy-preserving relay filtering scheme for delay tolerant networks.," in *Proc. IEEE INFOCOM*, 2012.
- [35] A. Picu, T. Spyropoulos, and T. Hossmann, "An analysis of the information spreading delay in heterogeneous mobility dtns," in *Proc. IEEE WoWMoM*, 2012.
- [36] M. J. Keeling, "The effects of local spatial structure on epidemiological invasions.," *Proc. R. Soc. B*, vol. 266, no. 1421, pp. 859–867, 1999.
- [37] M. E. J. Newman, "Spread of epidemic disease on networks," *Phys. Rev. E*, vol. 66, p. 016128, Jul 2002.
- [38] F. Ball, D. Mollison, and G. Scalia-Tomba, "Epidemics with two levels of mixing," *Annals of Applied Probability*, vol. 7, pp. 46–89, 1997.
- [39] S. Rushton and A. Mautner, "The deterministic model of a simple epidemic for more than one community," *Biometrika*, vol. 42, pp. 126–132, 1955.
- [40] I. S. Gradshteyn and I. M. Ryzhik, *Table of Integrals, Series, and Products, Seventh Edition*. Academic Press, 7th ed., 2007.