An Analysis of the Information Spreading Delay in Heterogeneous Mobility DTNs

Andreea Picu
Communication Systems Group
ETH Zürich, Switzerland
Email: lastname@tik.ee.ethz.ch

Thrasyvoulos Spyropoulos
Mobile Communications
EURECOM, France
Email: firstname.lastname@eurecom.fr

Theus Hossmann
Communication Systems Group
ETH Zürich, Switzerland
Email: lastname@tik.ee.ethz.ch

***Abstract*—Epidemic spreading is one of the most popular bio-inspired principles, which has made its way into computer networking. This principle naturally applies to Opportunistic or Delay Tolerant Networks (DTNs), where nodes probabilistically meet their neighbors thanks to mobility. Epidemic-based algorithms are often the only choice for DTN problems such as broadcast and unicast routing, distributed estimation etc. Existing analyses of epidemic spreading in various contexts only consider specific graph geometries (complete, random, regular etc) and/or homogeneous exponential node meeting rates. In addition, in wired networks, synchronous communication is usually assumed.**

In this paper, we relax these assumptions and provide a detailed analysis of epidemic spreading in DTNs with heterogeneous node meeting rates. We observe the special properties of a Markov model, describing the epidemic process and use them to derive bounds for the delay (expectation and distribution). We apply our analysis to epidemic-based DTN algorithms for routing and distributed estimation and validate the bounds against simulation results, using various real and synthetic mobility scenarios. Finally, we empirically show that the delay distribution is relatively concentrated, and that, depending on graph properties (communities, scale-freeness), the delay scales very well with network size.

I. Introduction

Opportunistic or *Delay Tolerant Networks* (DTNs) are envisioned to support communication in case of failure or lack of infrastructure (disaster, censorship, rural areas) and to enhance existing networks (e.g., offload cellular traffic), enabling novel applications. Nodes harness unused bandwidth by exchanging data when they are in proximity (*in contact*). The data is then *stored* by the nodes, *carried* through their mobility and eventually *forwarded* to destinations, thus achieving multihop communication, despite the lack of end-to-end paths.

Due to the inherent uncertainty and randomness of this type of network, the epidemic spreading principle is central to many DTN algorithms [1], [2], [3]. Here, epidemic spreading operates as follows: given a piece of information/message m, every node carrying a copy of m must further replicate the message to every node it encounters (provided the encountered node does not already have m). Thus, information/messages will spread almost like an epidemic through the network, with every node eventually receiving a copy of m. Modifications of this basic algorithm include spreading the message with a probability less than 1 at each contact ("gossip").

In DTNs, algorithms based on the epidemic spreading principle have been extensively studied, at first through lengthy and complex simulations and later through analytical models, based on Markov chains [4] and fluid approximations [5], [6]. For the sake of tractability, state-of-the-art epidemic spreading models mainly rely on simple mobility assumptions (e.g. Random Walk, Random Waypoint), where node mobility is stochastic and independent identically distributed (IID). Yet, studies of real scenarios [7], [8] reveal more complex structure, comprising considerable heterogeneity in node mobility, questioning the usefulness of these models' predictions.

A more realistic analysis of the basic principle of epidemic spreading is thus important for all the algorithms that it underlies. These range from routing to distributed estimation and including content dissemination. For example, DTN *broadcast routing* or flooding is achieved via the epidemic propagation of the broadcast message. Moreover, the first algorithm proposed for DTN *unicast routing* is the well-known epidemic routing [1] and while more sophisticated, resource-friendly routing schemes have been proposed, epidemic routing still serves as the key benchmark against which any scheme must compare in terms of delay and delivery probability. The epidemic spreading principle is also employed in the *distributed estimation of global network parameters*, both in DTNs [2] and in traditional networks [9]. Network size [10], average buffer occupancy [11] or overall network storage are just a few examples of global network parameters needed by DTN algorithms to tune themselves and provide the best performance they can.

In this paper, we analyze the delay of epidemic spreading under a much larger class of node mobility patterns than in previous studies. Previous studies model every node pair's meeting times with *identical* Poisson processes of rate λ [5], [6], [4], or use at most a few different rates corresponding to mobility classes [3]. Here, we use a set of *fully heterogeneous* Poisson processes with rates λ_{ij}, unique to each node pair (i, j). We choose to maintain, at this point, the assumption of independence (among contact processes) and the Poisson approximation, for two reasons: (i) for analytical tractability, as even a small departure from the simple IID exponential model rapidly increases complexity and (ii) we believe that the introduction of

This work is funded by ETH Zürich under Research Grant ETH−20 09−1 and by the EU Commission under the SCAMPI (FP7−258414) FIRE Project.

978-1-4673-1239-4/12/$31.00 ©2012 IEEE

generic, heterogeneous rates already facilitates analysis over a significantly broader class of mobility models, which better approximate realistic scenarios (our trace-based evaluation validates this intuition).

The rest of this paper is organized as follows. Sec. II presents a summary of related work and demonstrates the difficulty of introducing heterogeneous mobility in existing Markov models. Then, Sec. III formally defines the problem, and describes our network model. In Sec. IV, we show how to overcome the difficulties raised by the use of heterogeneous pairwise meeting rates and we derive bounds for the delay of epidemic spreading (expectation and distribution), in function of the network size and the *conductance* of the DTN contact graph. Next, in Sec. V, we validate our bounds against simulation results on varied synthetic and real-world DTN traces. Moreover, we empirically show that the delay distribution is concentrated, and that, depending on certain properties of the DTN graph (community structure, scale-freeness), the delay scales very well with network size. In Sec. VI, we discuss applications of our analysis to some DTN algorithms. Finally, we summarize our conclusions and future work in Sec. VII.

II. Related Work and Challenges

Early efforts to evaluate the performance of epidemic spreading in Opportunistic Networks relied on lengthy and complex simulations, using a limited number of simple mobility models (Random Waypoint, Random Direction). Later, analytical studies emerged, using ordinary differential equation models (ODEs) [5], [6], inspired from epidemiology, on the one hand and Markovian models [4], on the other hand. In ODE models, the epidemic spreading process is treated as a fluid flow and the number of copies is approximated as a continuous-valued function of time and the node meeting rate. With Markov models, the spread of a message is modeled with a chain, whose states are the number of copies in the network.

As mentioned above, these studies rely on simple, identical node mobility assumptions, where nodes meet each other at independent identically distributed (IID) time intervals. There is a unique meeting rate λ, describing the contacts of every node pair. As a result, all nodes are equal and can be treated as a group, rather than individually. This is reflected, e.g., in epidemic spreading Markov models, where only the number of message copies is modeled, without regard for the specific nodes carrying those copies, as in Fig. 1. This results in simple chains, with state spaces the size of the network, which can easily be solved for delivery delays and ratios.

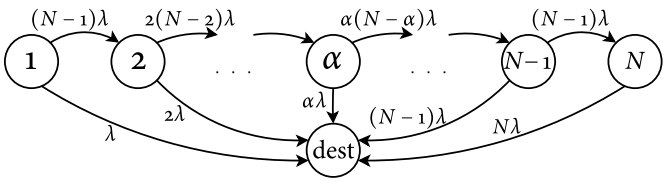

Figure 1. Markov chain from [4]

To incorporate some of the non-trivial structure of mobility and meeting patterns observed in real world experiments [7], [8], newer analyses [3] introduce a small set of different mobility groups: node mobility is different across groups, but inside the same group, nodes are still identical. This results in a relatively complex partial differential equation (PDE) model, with only limited additional diversity in the node mobility.

Here, we use individual pairwise meeting rates λ_{ij} for each node pair (i, j) to model the mobility of nodes. In this case, it is a priori not possible to group nodes (unless we make additional restrictive assumptions about the rates λ_{ij}). Indeed, while in Fig. 1, a state is simply the number of message copies in the network, with our model, the specific nodes carrying those copies need also be modeled. As a result, for each state α of the simple Markov chain, there are $\binom{N}{\alpha}$ states in the new Markov chain (with α the current number of message copies). That is a state space of size $2^N - 1$, resulting in a chain that could only be solved numerically at a rapidly increasing complexity even for modest network sizes. However, this provides no analytical insight into the effect of different contact patterns on performance nor the key mobility characteristics that make the epidemic delay small or large.

One of the contributions of our work is to provide a deeper analytical insight into this complex model, by connecting in a new way the delay of epidemic spreading to a fundamental property, *conductance*, of a well defined graph related to the mobility process in hand. While this approach allows us to only derive performance bounds (a reasonable penalty for tractability), conductance, as we shall see, is closely related to macroscopic mobility patterns such as community-like (clustering) behavior, and is thus directly interpretable. Furthermore, using this approach we are able to show that the delay of epidemic spreading in DTNs has good scaling properties, as the conductance of DTN graphs is relatively unaltered by increasing network size.

Conductance has already been related to the delay of epidemic spreading, in the broader context of traditional computer networking. There, the *push-pull* algorithm has recently received significant attention. This algorithm proceeds in synchronous rounds. In each round, every node in the network randomly chooses a neighbor to communicate with. Then, for every communicating node pair, if one of the two nodes has the information, both will have it at the end of the round. Several bounds on the spreading delay of this algorithm have been proposed and gradually improved [12], [13], [14], [15]. However, the majority of the results are not easily transferable to the DTN context. This is firstly, because the considered model operates synchronously and secondly, because of the implicit connectivity assumption, that nodes have a set of neighbors with whom they can communicate at will.

Only in [12], do the authors consider the asynchronous case and obtain a tail bound which is most relevant to our work. There, epidemic information spreading is analyzed as

an underlying routine of a distributed estimation algorithm. An upper bound is provided for the epidemic spreading delay, in function of the *conductance* of the stochastic matrix that governs communication between node pairs. Here, we derive an upper bound for the entire delay distribution (instead of merely the tail) and we empirically show that our bound is tighter and that the distribution is concentrated.

III. Problem Statement and Model

Let V be our Opportunistic Network, with $|V| = N$ nodes. V is a relatively sparse ad hoc network, where node density is insufficient for establishing and maintaining (end-to-end) multi-hop paths. Instead, data is stored and carried by nodes, and forwarded through intermittent contacts established by node mobility. A *contact* occurs between two nodes who are in range to setup a bi-directional wireless link to each other.

Epidemic spreading in its simplest form is defined as:

Definition 1 (Single-source spreading (SSS)): Given the network V, start from an arbitrary node $i \in V$ and disseminate a message m_i **to all nodes in the network**, via a (mobility-driven) sequence of *local* exchanges between node pairs.

A. Network and Time Model

We describe the network V using a marked point process $(M_n)_{n\in\mathbb{Z}} = \{T_n, \sigma_n\}$, where T_n denotes the starting time of a contact and $\sigma_n = (i, j, \Delta_n)$ denotes the two nodes in contact $i, j \in V$ and the duration of the contact Δ_n [16], [17]. The random variables $J_n = T_n - T_{n-1}$ are the times between the initiations of two successive contacts.

We make the following assumptions:

1) $(T_n)_{n\in\mathbb{Z}}$ are epochs of a Poisson process – the times J_n are IID with intensity[1] λ, dependent on network sparsity.
2) $T_n < T_{n+1}$, $\forall n \in \mathbb{Z}$ – no two contacts start at the same time, i.e., $J_n > 0$.
3) $\Delta_n << J_n$ – the duration of a contact is negligible compared to the time between two contacts, but sufficient for all data transfers to take place. Thus, $\sigma_n = (i, j)$.
4) $(\sigma_n)_{n\in\mathbb{Z}}$ is a discrete IID process, with common distribution $\mathbb{P}[\sigma_n = (i, j)] = p^c_{ij}$. This is the pairwise contact probability, that the next contact in the sequence is between nodes i and j. Then, by item 1), the pairwise contact rate is $\lambda_{ij} = p^c_{ij}\lambda$.

The memoryless (exponential) assumption is key in our DTN model, as well as in all existing analytical studies for DTNs. While this assumption has spurred some controversy [18], Karagiannis et al. established in [19] that after a short power law "characteristic time" of about half a day, the inter-contact time does exhibit exponential decay. Moreover, the authors of [20] recently showed that the observed power law in the aggregate inter-contact time may simply be the result of exponential inter-contacts with different rates drawn from a given (e.g., power law) distribution (a similar fact is known for random graphs with generalized node degree distributions [21]). This further corroborates our approach to treat heterogeneous, general contact rates. Consequently, *approximating* inter-contact times with exponential variables is well-founded, especially when focusing on residual inter-contact times[2], as in analyses of forwarding schemes.

The above defined model discretizes time in the otherwise continuous contact process M_n, by only considering the embedded sequence of contact events σ_n. We can go back to continuous time through the marked point process intensity λ and Wald's equation [22]. While, for small heterogeneous networks, this is only an approximation, as network size increases, the times between events (contact events in our point process) become increasingly smaller and network delays consist of more and more such contact events (discrete time units). Then, by the law of large numbers, the continuous time delays thus computed will converge, allowing us to focus on the embedded discrete contact process.

Under the above model, a given mobility scenario with *heterogeneous node mobility* is described by its pairwise contact probabilities p^c_{ij} forming the ***contact probability matrix***:

$$\mathbf{P}^c = \{p^c_{ij}\}. \tag{1}$$

Probabilities p^c_{ij} can be either measured directly from a given real or synthetic mobility trace (e.g. maximum likelihood estimation of p^c_{ij} for every pair); or they can be calculated using the pair's contact statistics (e.g. frequency), as in [23]. We apply the former approach to all traces we use (Table I).

In the following, we use exclusively the weighted undirected graph $G = (V, E)$, with adjacency matrix $\mathbf{P}^c$.

B. Information Spreading Delay

The quantity we are analysing is the delay for the SSS task in Def. 1 to be completed. In discrete time, the SSS delay is an integer-valued random variable D_{SSS}, defined on the same probability space as our contact process $(M_n)_{n\in\mathbb{Z}} = \{T_n, \sigma_n\}$ from above. Assuming node i starts the message dissemination at time T_0 and denoting by $S(T_n)$ the (sub)set of nodes "infected" with m_i by time T_n:

$$D_{\text{SSS}} = \inf\{n > 0 : S(T_n) = V\}. \tag{2}$$

In the following, we will analyze D_{SSS} by stages. To this end, we define the following two classes of random variables: *partial spreading delays* $D_{\text{SSS}}(\alpha)$ and *transition delays* D_α:

$$D_{\text{SSS}}(\alpha) = \inf\{n > 0 : |S(T_n)| = \alpha\}, \tag{3}$$

$$D_\alpha = D_{\text{SSS}}(\alpha + 1) - D_{\text{SSS}}(\alpha). \tag{4}$$

Thus, $D_{\text{SSS}} = D_{\text{SSS}}(N) = \sum_{\alpha=1}^{N-1} D_\alpha$. Next, we present an absorbing Markov model for the SSS task, which enables a detailed analysis of the D_α's and ultimately of D_{SSS}.

[1] Expected number of points or contact "arrivals" per time unit.

[2] This is the time until the next contact for a node pair from an arbitrary point in time. Intuitively, the residual time reflects how much time a device must wait, before being able to forward a message to another given device.

IV. Theoretical Delay Analysis

We first describe our model for the SSS task, then we analyse the expected spreading delay and finally, we show how the delay distribution can be bounded by a related and much simpler distribution.

A. A Markovian Model for Information Spreading

In the typical Markov chain model for epidemic spreading (shown in Fig. 1), states are the number of copies of m_i in the network (the number of nodes "infected" with m_i). As we point out in Section II, with the introduction of heterogeneous contact probabilities p^c_{ij} instead of the "one-fits-all" probability p, nodes can no longer be treated in groups. As a result, for each state α of that simple Markov chain, there are $\binom{N}{\alpha}$ states in the new Markov chain for the SSS task – with α the current number of copies of m_i. (There is no "dest" state in the new chain, as the SSS task delivers the message to all nodes.)

Denote by $A(\alpha)$ (with $\alpha = 1, \ldots, N$), the set of all α-sized node subsets of our network V (e.g., $A(1) = V$). Then, $|A(\alpha)| = \binom{N}{\alpha}$ and

$$\Omega = \bigcup_{\alpha=1}^{N} A(\alpha) \tag{5}$$

is the state space of the SSS task, where a state $S \in \Omega$ is a subset of infected nodes. While the size of this state space ($|\Omega| = 2^N - 1$) is prohibitive for any quantitative results, its qualitative analysis leads to very interesting findings.

The delay of the epidemic spreading task D_{SSS} defined in Eq. (2) is the absorption time of the Markov chain $(\mathbf{X}_n)_{n\in\mathbb{N}}$ defined on state space Ω by the transition matrix $\mathbf{P}$ in Fig. 2 (non-zero entries are shown in red) and by the initial probability vector $\boldsymbol{\pi}^{(0)}$ (with $\sum_{i=1}^{N} \pi^{(0)}(i) = 1$ and $\pi^{(0)}(i) = 0$ for all $i > N$). The absorbing state is $A(N) = \{V\}$, i.e., the one in which every node has m_i. Transition probabilities are combinatorial functions of pairwise contact probabilities p^c_{ij}. For example, if the chain is in state S_x with nodes $\{a, b, c\}$ infected, a transition to state S_y with nodes $\{a, b, c, d\}$ infected, happens with transition probability $p_{xy} = p^c_{ad} + p^c_{bd} + p^c_{cd}$. With probability $1 - p_{xy}$, the chain stays in state S_x.

Lemma 1: The matrix $\mathbf{P}$ has the following properties:

(i) $\mathbf{P}$ is an **upper triangular matrix**.
(ii) $\mathbf{P}$ is a sparse matrix, specifically a **band(ed) matrix**.
(iii) $\mathbf{P}$ is a **block matrix**.

Proof:

(i) *Upper triangular:* Spreading can only be increased: if the message m_i has currently spread to α nodes (the chain is in state $S_x \in A(\alpha)$), it is impossible to go back to state $S_y \in A(\alpha - \beta)$, with $\alpha > \beta > 0$.

(ii) *Sparse, banded:* When in state $S_x \in A(\alpha)$, the choice of successor states in the Markov chain is limited to states $S_y \in A(\alpha + 1)$. Any state $S_z \in A(\alpha + \beta)$, $\beta \geqslant 2$ is not directly reachable (as no simultaneous contacts occur in our model).

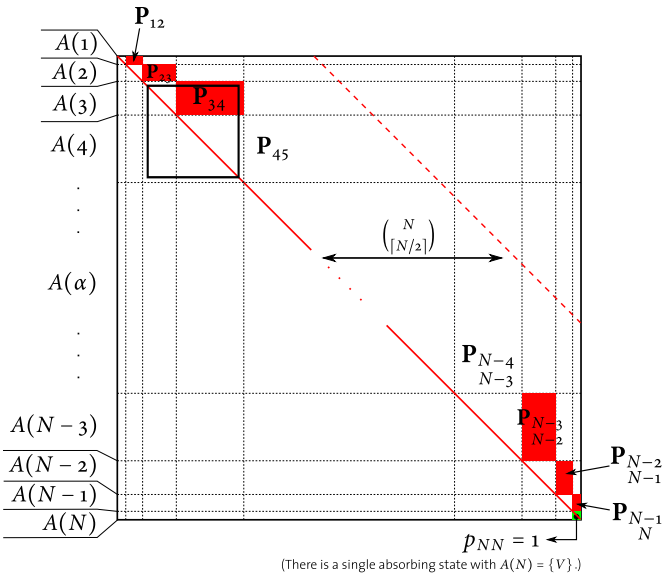

Figure 2. The $(2^N - 1) \times (2^N - 1)$ transition matrix $\mathbf{P}$ for the SSS task

(iii) *Block:* For every $[A(\alpha), A(\alpha + 1)]$ pair, there is an $\binom{N}{\alpha} \times \binom{N}{\alpha+1}$ block, $\mathbf{P}_{\alpha,\alpha+1}$, representing the spread of m_i to a further node, i.e., from α nodes to $\alpha+1$ nodes. Additional blocks are formed by the diagonal entries, corresponding to each $A(\alpha)$. ■

Since all blocks model the same operation (i.e., the spread of the message to one further node), analyzing $\mathbf{P}_{\alpha,\alpha+1}$ will provide significant insight into the properties of the entire matrix and thus, of the SSS task. To analyze a single block, we define a new (absorbing) Markov chain on $\Omega' = A(\alpha) \cup A(\alpha + 1)$ with initial probability vector $\boldsymbol{\pi}^{(0)}_\alpha$ (depending on $\boldsymbol{\pi}^{(0)}$ and on previous blocks) and with transition matrix:

$$\mathbf{P}'(\alpha) = \begin{array}{c} \begin{matrix} A(\alpha) & A(\alpha+1) \end{matrix} \\ \begin{pmatrix} \mathbf{Q}_\alpha & \mathbf{P}_{\alpha,\alpha+1} \\ \mathbf{0} & \mathbf{I}_{\alpha+1} \end{pmatrix} \end{array} \begin{matrix} A(\alpha) \\ A(\alpha+1) \end{matrix} \tag{6}$$

where $\mathbf{I}_{\alpha+1}$ is the $\alpha+1$ identity matrix and $\mathbf{Q}_\alpha$ is a diagonal matrix corresponding to $A(\alpha)$ in the original matrix $\mathbf{P}$. Note that, in the original matrix $\mathbf{P}$, states $A(\alpha+1)$ are not absorbing (there exists a $\mathbf{Q}_{\alpha+1} \neq \mathbf{I}_{\alpha+1}$). However, since we are only interested in the phase $\alpha \to \alpha+1$ of the spreading algorithm, we can safely ignore that for now. As an example, the block forming $\mathbf{P}'(3)$ is denoted with a box in Fig. 2, clearly showing $\mathbf{Q}_3$ and $\mathbf{P}_{34}$ as the upper blocks.

B. Expected Delay Analysis through Absorption

The transition delay D_α defined in Eq. (4) is the absorption time of $\mathbf{P}'(\alpha)$. Its expectation $\mathbb{E}[D_\alpha]$ and its variance $\mathbb{V}[D_\alpha]$ can be obtained from the theory of absorbing Markov chains [24]. Based on this theory, we state the lemma below:

Lemma 2: The expected transition delay $\mathbb{E}[D_\alpha]$ is:

$$\mathbb{E}[D_\alpha] = \sum_{x=1}^{\binom{N}{\alpha}} \frac{\pi^{(0)}_\alpha(x)}{\partial(S_x)}, \tag{7}$$

where $S_x \in A(\alpha)$ and $\partial(S_x) = \sum_{i \in S_x; j \notin S_x} p^c_{ij}$ is the *edge boundary* of the vertex set S_x in the graph G defined by weight matrix $\mathbf{P}^c$. The initial probabilities $\pi^{(0)}_\alpha(x)$ from vector $\boldsymbol{\pi}^{(0)}_\alpha$ sum to 1. A similar expression can be derived for the variance of the delay $\mathbb{V}[D_\alpha]$. Due to space limitations, we refer the interested reader to [25].

Proof: The elements of the diagonal matrix $\mathbf{Q}_\alpha$ from Eq. (6) can be written as:

$$q_{xx} = 1 - \sum_{i \in S_x; j \notin S_x} p^c_{ij}, \quad \forall S_x \in A(\alpha);\ x = 1, \ldots, \binom{N}{\alpha}, \tag{8}$$

with p^c_{ij}, the contact probability between nodes i and j.

By the theory of absorbing Markov chains [24], the expected delay $\mathbb{E}[D_\alpha]$ for leaving the $\alpha \to \alpha+1$ phase can be obtained from the *fundamental matrix* of the Markov chain $\mathbf{P}'(\alpha)$: $\mathbf{N}_\alpha = (\mathbf{I}_\alpha - \mathbf{Q}_\alpha)^{-1}$. Since $\mathbf{Q}_\alpha$ is diagonal, so is $\mathbf{N}_\alpha$, and its elements n_{xx} (with $x = 1, 2, \ldots, \binom{N}{\alpha}$) are easily obtainable as:

$$n_{xx} = \frac{1}{1 - q_{xx}} = \frac{1}{\sum_{i \in S_x; j \notin S_x} p^c_{ij}} = \frac{1}{\partial(S_x)}, \tag{9}$$

where $\partial(S_x)$ is the *edge boundary* of the vertex set S_x in the graph G defined by weight matrix $\mathbf{P}^c$.

The fundamental matrix $\mathbf{N}_\alpha$ is a $\binom{N}{\alpha} \times \binom{N}{\alpha}$ diagonal matrix, where each $n_{xx} = \partial(S_x)^{-1}$ represents the expected time the Markov chain $\mathbf{P}'(\alpha)$ will spend in state $S_x \in A(\alpha)$ before being absorbed, i.e., moving on to phase $\alpha + 1$. Therefore, the expected absorption time of the Markov chain $\mathbf{P}'(\alpha)$ (or expected transition delay $\mathbb{E}[D_\alpha]$) is a weighted sum of the n_{xx}'s with the initial probabilities $\pi^{(0)}_\alpha(x)$ as weights.

$$\mathbb{E}[D_\alpha] = \sum_{x=1}^{\binom{N}{\alpha}} \pi^{(0)}_\alpha(x) \cdot n_{xx} \tag{10}$$

and, with Eq. (9), this completes the proof. ∎

Note that the edge boundary of a vertex set S_x is equal to the edge boundary of the set's complement $\overline{S_x}$, i.e., $\partial(S_x) = \partial(\overline{S_x})$. This means that the SSS task from Def. 1 exhibits some symmetry: the expected sojourn times in states $S_x \in A(\alpha)$ and $\overline{S_x} \in A(N-\alpha)$ are equal, i.e., $\mathbf{N}_\alpha = \mathbf{N}_{N-\alpha}$.

However, this does not mean that the respective expected absorption delays $\mathbb{E}[D_\alpha]$ and $\mathbb{E}[D_{N-\alpha}]$ are also equal. This is because the expected absorption time of a Markov chain depends not only on its fundamental matrix, but also on its initial probability distribution. Since the initial probability distributions $\boldsymbol{\pi}^{(0)}_\alpha$ for each phase of the SSS task depend on the previous phases, this vector is not easily obtainable and it is very likely that $\boldsymbol{\pi}^{(0)}_\alpha \neq \boldsymbol{\pi}^{(0)}_{N-\alpha}$.

While we cannot obtain a closed form expression for the expected transition delays $\mathbb{E}[D_\alpha]$, we *can* use the fundamental matrices $\mathbf{N}_\alpha$ to bound them:

Theorem 3 (Moments of the delay.): The expected delay of the SSS task is bounded as follows:

$$\mathbb{E}[D_{\text{SSS}}] < \frac{2\ln(N-1)}{N\Phi}, \tag{11}$$

where Φ is the *conductance* of the DTN graph G, defined as $\Phi = \min_{S \in \Omega} \phi(S)$. The conductance of a cut $\phi(S)$ is[3]

$$\phi(S) = \frac{\partial(S)}{|S| \cdot |\overline{S}|} = \frac{\partial(\overline{S})}{|S| \cdot |\overline{S}|} = \phi(\overline{S}), \tag{12}$$

where $|S|$ is the cardinality of the set S.

Once again, a similar bound can be calculated for the variance of the delay, which we omit for space reasons. The interested reader is referred to [25].

Proof: Using Eq. (7) and the fact that $\pi^{(0)}_\alpha(x) < 1$ for all x, we bound the expected transition delay as follows:

$$\mathbb{E}[D_\alpha] \leqslant \frac{1}{\min\limits_x \partial(S_x)} = \frac{1}{\alpha(N-\alpha) \min\limits_x \phi(S_x)}. \tag{13}$$

Let the $\alpha-$conductance of our DTN graph G be the minimum conductance among partitions of size α, i.e., $\Phi_\alpha = \Phi_{N-\alpha} = \min\limits_x \phi(S_x)$, for $S_x \in A(\alpha)$ and $x = 1, \ldots, \binom{N}{\alpha}$. Then, for all $\alpha = 1, \ldots, N-1$,

$$\mathbb{E}[D_\alpha] \leqslant \frac{1}{\alpha(N-\alpha)\Phi_\alpha}. \tag{14}$$

Since $D_{\text{SSS}} = \sum_{\alpha=1}^{N-1} D_\alpha$, using the linearity of expectation and the two above inequalities, we obtain:

$$\mathbb{E}[D_{\text{SSS}}] \leqslant \sum_{\alpha=1}^{N-1} \frac{1}{\alpha(N-\alpha)\Phi_\alpha}. \tag{15}$$

This is a relatively tight bound, however, it is not easily interpretable. To this end, we observe that $\Phi_\alpha \leqslant \Phi$ for all $\alpha = 1, \ldots, N-1$ and obtain:

$$\mathbb{E}[D_{\text{SSS}}] \leqslant \sum_{\alpha=1}^{N-1} \frac{1}{\alpha(N-\alpha)\Phi} < \frac{2\ln(N-1)}{N\Phi}. \tag{16}$$

and this completes the proof. ∎

Naturally, the bound in Thm. 3 can also be calculated for the partial spreading delays $D_{\text{SSS}}(\alpha)$, defined in Eq. (3). This can be done by simply stopping the sum in Eq. (16) at the desired stage. This results in:

$$\mathbb{E}[D_{\text{SSS}}(\alpha)] \leqslant \sum_{\beta=1}^{\alpha} \frac{1}{\beta(N-\beta)\Phi} < \frac{1}{N\Phi}\left(\ln \frac{\alpha(N-1)}{N-\alpha}\right). \tag{17}$$

Note that, to obtain these bounds, we used the exact average delay formula and essentially replaced all cut volumes $\partial(S_x)$ for partitions of a certain size α, by the minimum cut volume for that size, $\alpha(N-\alpha)\Phi_\alpha$. Consequently, in scenarios where the contact graph does not have any well-defined communities or if those communities are relatively small, the bound is close to the actual delay value. Conversely, if the contact graph exhibits larger and more defined communities (i.e., well connected vertex subsets, with weaker links towards the outside), the bound is somewhat looser. In Section V, we will in fact observe this subtle distinction on our various traces.

[3]This is one of several definitions of conductance to be found in the literature. We use it as it provides the best bound.

	MIT	INFO	HCMM	SLAW
Scale and context	92 campus students & staff	41 conference attendees	100 nodes	100 nodes
Structure (from [26])	6 communities	no strong communities	10 communities	Hurst param. $h = 0.75$
Period	9 months	3 days	2.5 months	4 months
Scanning Interval	300s (Bluetooth)	120s (Bluetooth)	N/A	N/A
# Contacts total	81 961	22 459	1 129 242	192 245

Table I
MOBILITY TRACES CHARACTERISTICS.

C. Bounding the Delay's Distribution

In addition to the mean epidemic spreading delay, it is important to know the probability that the spreading is completed within a certain timeframe. This is practically relevant for most algorithms based on epidemic spreading. For example, in routing, this translates to the delivery probability/ratio. In this section, we show how to construct a new and much simpler Markovian model, whose absorption time distribution upper bounds the absorption time distribution of $\mathbf{P}$, that is D_{SSS}.

Theorem 4 (Bounding Markov Chain): The random variable D_{SSS}, the absorption time of $\mathbf{P}$, is upper bounded (in the usual stochastic order) by $\hat{D}_{\mathrm{SSS}}$, the absorption time of the discrete-time pure-birth Markov chain $\hat{\mathbf{P}}$ (Fig. 3) [27] with initial probability distribution $\hat{\boldsymbol{\pi}}^{(0)} = (1, 0, \ldots, 0)$.

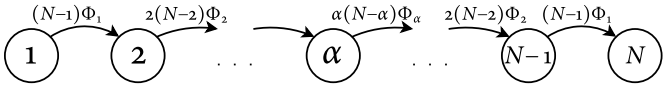

Figure 3. The bounding pure-birth Markov chain

Proof: The distribution of D_{SSS} is a discrete phase-type distribution (a mixture of geometric random variables). We note that $D_{\mathrm{SSS}} = \sum_{\alpha=1}^{N-1} D_\alpha$ and we will once more use the stages of the spreading process to analyze this distribution. The random variables D_α are weighted sums of geometric random variables. More precisely, note that, for each state $S \in \Omega$, the time spent in this state is geometrically distributed with success probability $\partial(S)$ (success means leaving the state). Therefore, for all $\alpha = 1, \ldots N-1$ and all $x = 1, \ldots, \binom{N}{\alpha}$ the variable $D_{x(\alpha)} \sim \mathcal{G}(\partial(S_x))$ is the delay of leaving state S_x.

Further, we recall that we denoted by $\boldsymbol{\pi}_\alpha^{(0)}$ the initial probability vector for stage α. The probability that stage α starts in state S_x is $\pi_\alpha^{(0)}(x)$ and $\sum_x \pi_\alpha^{(0)}(x) = 1$.

Let $\hat{D}_\alpha \sim \mathcal{G}(\alpha(N-\alpha)\Phi_\alpha)$ be the delay of stage α in the Markov chain $\hat{\mathbf{P}}$. Note that this delay is equal to $D_{x^\star(\alpha)}$, where $x^\star$ is the index of the cut achieving Φ_α, the minimum α−conductance. Then, the delay of the $\alpha \to \alpha+1$ stage in the original Markov chain $\mathbf{P}$ is

$$D_\alpha = \sum_{x=1}^{\binom{N}{\alpha}} \boldsymbol{\pi}_\alpha^{(0)}(x) \cdot D_{x(\alpha)} \preceq \sum_{x=1}^{\binom{N}{\alpha}} \boldsymbol{\pi}_\alpha^{(0)}(x) \cdot \hat{D}_\alpha \tag{18}$$

$$= \hat{D}_\alpha \sum_{x=1}^{\binom{N}{\alpha}} \boldsymbol{\pi}_\alpha^{(0)}(x) = \hat{D}_\alpha, \tag{19}$$

where "$\preceq$" denotes the usual stochastic order (inequality in complementary cumulative distribution function (CCDF)). Consequently,

$$D_{\mathrm{SSS}} = \sum_{\alpha=1}^{N-1} D_\alpha \preceq \sum_{\alpha=1}^{N-1} \hat{D}_\alpha = \hat{D}_{\mathrm{SSS}} \tag{20}$$

and this completes the proof. ∎

The bounding delay $\hat{D}_{\mathrm{SSS}}$ is a simple (non-weighted) convolution of geometric distributions. Denoting the success probabilities $p_\alpha = \alpha(N-\alpha)\Phi_\alpha$, the probability mass function of this variable is given by [28]:

$$\mathbb{P}[\hat{D}_{\mathrm{SSS}} = 2k] = \sum_{\alpha=1}^{\lfloor N/2 \rfloor} w_\alpha p_\alpha (1-p_\alpha)^{k-1}, \tag{21}$$

where the weights w_α are defined as:

$$w_\alpha = \prod_{\beta=1;\, \beta\neq\alpha}^{\lfloor N/2 \rfloor} \frac{p_\beta}{p_\beta - p_\alpha} \tag{22}$$

and $\sum_{\alpha=1}^{\lfloor N/2 \rfloor} w_\alpha = 1$. This holds for all $k \geqslant \lfloor N/2 \rfloor$, provided that $p_\alpha \neq p_\beta$ for all $\alpha \neq \beta$.

Thm. 4 also holds if we replace all Φ_α with Φ in Fig. 3. In this case, the weights w_α simplify to:

$$w_\alpha = \begin{cases} (-1)^{\alpha-1}\binom{N}{\alpha}, & \text{if } N \text{ is even,} \\ (-1)^{\alpha-1}\left(1-\frac{2\alpha}{N}\right)\binom{N}{\alpha}, & \text{if } N \text{ is odd.} \end{cases} \tag{23}$$

The probability mass function of the partial spreading delay $D_{\mathrm{SSS}}(\alpha)$ can be calculated in the same way for $\alpha < \lfloor N/2 \rfloor$, by simply stopping the summation in Eq. (21) at the desired stage α. A formula can also be derived for $\lfloor N/2 \rfloor \leqslant \alpha < N-1$, but it is more involved.

Similarly to the expected delay bound, the above distribution bound is more or less tight depending on the strength and size of node communities in the analyzed contact graph.

V. EMPIRICAL ANALYSIS

In this section, we validate the results of the analysis against simulation results for the SSS task, using well known real world traces as well as synthetic mobility models. Then, we show that the delay distribution of the SSS task is concentrated and that the delay has good scaling properties, as the conductance of considered DTN graphs is relatively unaltered by increasing network size.

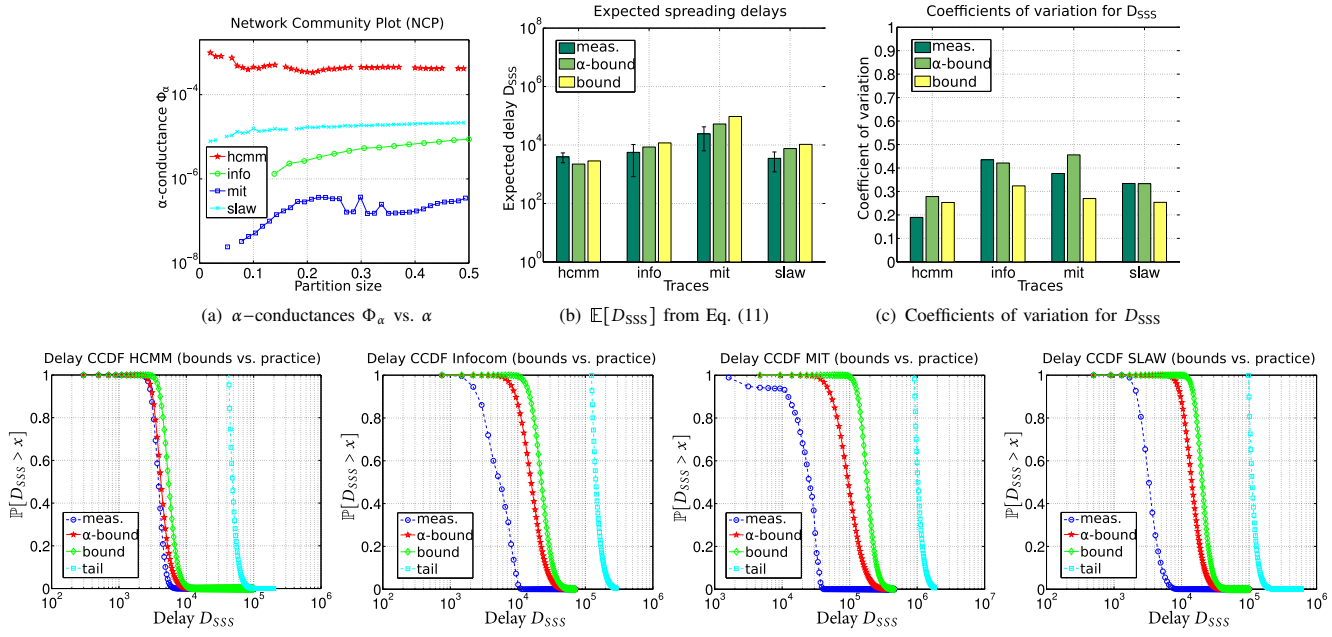

(a) $\alpha-$conductances Φ_α vs. α (b) $\mathbb{E}[D_{\text{SSS}}]$ from Eq. (11) (c) Coefficients of variation for D_{SSS}

(d) D_{SSS} CCDF for HCMM (e) D_{SSS} CCDF for Infocom (f) D_{SSS} CCDF for MIT (g) D_{SSS} CCDF for SLAW

Figure 4. Prediction accuracy and concentration for the delay of the SSS task

A. *The Accuracy of the Bounds*

To cover a broad range, we use two real-world connectivity traces and two synthetic mobility traces for validation: (i) the *Reality Mining* trace (MIT) [7], (ii) the Infocom 2005 trace (INFO) [29], (iii) a synthetic scenario created with the HCMM mobility model (HCMM) [30] and (iv) a synthetic scenario created with the SLAW mobility model (SLAW) [31]. Their characteristics are shown in Table I.

For all traces, we obtained the contact probability matrix $\mathbf{P}^c$, using a maximum likelihood estimator for the meeting probability of each node pair. We then calculated approximations for the $\alpha-$conductances Φ_α – the minimum conductance among partitions of size α, defined just before Eq. (14) – of the weighted graph with adjacency matrix $\mathbf{P}^c$, using the graph partitioning package Metis [32]. Metis uses a very fast partitioning algorithm (linear running time) and, as shown in [33], it produces very high quality partitions, with conductance values close to theoretical lower bounds. In fact, within the partitions produced by Metis there are sometimes even disconnections, resulting in conservative conductance values and thus quite conservative bounds in our analysis.

For each graph and each cluster size, we take the minimum $\alpha-$conductance Φ_α over 10 000 Metis runs. We plot the resulting $\alpha-$conductances of our traces for all α values in Fig. 4(a)[4]. This plot is called the *Network Community Plot (NCP)* [34] and it offers significant insight into the graph's community structures at various scales (α values).

Increasing NCP curves mean that the edge boundaries, $\partial(S) = \sum_{i\in S; j\notin S} p^c_{ij}$, consistently increase with $\alpha = |S|$. That is, no subset of size α distinguishes itself from the rest of the network, suggesting an absence of community structure. This is the case of the Infocom and SLAW graphs, also recognized in [26]. For these traces the delay bounds will be quite close to the empirical values for these traces. This is also true for the HCMM trace, where the NCP curve is almost constant, due to the very small built-in communities (10 nodes each). On the other hand, the MIT trace exhibits relatively sudden dips in conductance values (around $\alpha = 0.3N$) as the partition size α increases. This attests the existence of strong and relatively large ($\sim 0.3N$ nodes) communities in this trace, as reported also in [26]. Consequently, the delay bound will be somewhat less tight for this trace than for the previous three.

Using the $\alpha-$conductances Φ_α or alternatively the graph conductance Φ, we can now calculate the upper bounds for both the expected delay of epidemic spreading (Eq. (11)) as well as its complementary cumulative distribution function (CCDF), easily obtained from Eq. (21). We compare these analytically predicted quantities with results obtained from simulating epidemic spreading on our traces. For each trace, we measure the time it takes until *all* nodes in the network receive a message started at a uniformly chosen node. The message generation process produces a sample of at least 1000 observations per source node for shorter traces and up to 50 000 observations per source node for longer ones. For all measured expected delays, we compute the 95$^{\text{th}}$ percentile using the normal distribution.

Fig. 4(b) shows the expected delays obtained from the simulations ("meas.") compared to the bounds, calculated using the $\alpha-$conductance values Φ_α from Metis ("α-bound") and respectively the graph conductance Φ ("bound"). (Note the logarithmic scale of the $y-$axis.) The expected epidemic

[4]For better visualization, the $\alpha-$conductances are normalized by the duration of the traces. This avoids curve overlaps, while keeping the trends.

spreading delays obtained in simulations are indeed smaller than the predicted upper bounds. Only for the HCMM model is the prediction slightly off, presumably on account of a non-optimal approximation of the conductance by Metis.

In Fig. 4(c), we plot the coefficients of variation for the distribution of the delay D_{SSS} of each trace. The coefficient of variation, defined as $\sqrt{\mathbb{V}[D_{\mathrm{SSS}}]}/\mathbb{E}[D_{\mathrm{SSS}}]$, is a measure of the dispersion of a probability distribution. Values less than 1 (the coefficient of variation for the exponential distribution) are considered small and indicate a concentrated distribution. As it can be seen in Fig. 4(c), all our traces have very small and similar coefficients of variation, indicating that most delay values fall within a small interval of the expected delay. A similar more rigorous statement about concentration can be made using Chebyshev's inequality, as shown in [25]. This has positive implications, e.g., for the predictability of the time until termination of epidemic-based distributed estimation (often tricky to detect), despite the considerably heterogeneity in the mobility environments considered.

Figs. 4(d)–4(g) show the empirical CCDFs for the spreading delays in comparison with our theoretical CCDFs, obtained from Eq. (21) by using the α–conductance values Φ_α ("a-bound") and respectively the graph conductance Φ ("bound"). We also show the tail bound obtained by Shah in [12] (Theorem 3.1.). Once again, the empirical results are well within all the predicted bounds. Note that the tail bound is at least one order of magnitude larger than the bounds we obtained here. Note also that the concentration property discussed above is apparent in these plots.

As expected, a relatively loose bound, as in the case of MIT attests the presence of at least two strong communities connected through a relatively small set of edges (i.e., small conductance), as observed in [26]. On the other hand, for the community-free traces Infocom and SLAW [26], as well as for the HCMM trace (which has built-in very small communities of 10 nodes each) the bounds are, as expected, (much) closer to the empirical results.

B. The Scaling Behavior of the Conductance

As we have seen previously, the derived bounds are dependent on the conductance profile (or NCP) of the graph under analysis, as well as on the number of nodes in the network. This dependence is crucial for two reasons, as follows.

First, let us examine the expected spreading delay in Eq. (11) more closely. In particular, we transpose this equation to continuous time (by Wald's equation, as mentioned in Sec. III):

$$\mathbb{E}[T_{\mathrm{SSS}}] < \frac{2\ln(N-1)}{\lambda N \Phi}, \tag{24}$$

with T_{SSS}, the continuous-time counterpart of D_{SSS}. As the network size N increases, so does the intensity of the contact process λ and therefore, the term λN will roughly tend to 1 as $N \to \infty$. Therefore, the scaling behavior of the spreading delay with increasing network size is almost entirely determined by the conductance Φ – defined in Section IV-B – since the term $\ln(N-1)$ increases very slowly as the network grows.

The second reason why understanding the evolution of the conductance as the network size increases is crucial is that, as we have seen above, the tightness of the bound is also conductance-dependent.

The evolution of conductance has already been studied for a special category of graphs, i.e., graphs built using the preferential attachment model. Mihail et al. analytically showed in [35] that the conductance of these graphs remains constant as the network size increases. The main feature of the preferential attachment model is that it generates scale-free graphs (i.e., graph with power-law degree distributions). Since our graph model is rather generic, it is hardly possible to analytically investigate the behavior of its conductance. However, the scale-free property has been consistently observed on many DTN contact graphs, obtained both from real-world traces and from mobility models. It is thus reasonable to hypothesize that these graphs do as well conserve their conductance value as their size increases. In the following, we will empirically verify the validity of this hypothesis.

We use the same traces and models as previously and employ two simple methods to obtain networks of various sizes. In the first method, we start from a large network and randomly remove entire communities (if they exist). In the second method, we start from a large network and randomly remove single nodes. We then calculate the conductances Φ of the newly obtained smaller networks using Metis.

For SLAW and HCMM, we start with networks of maximum size 500 nodes and the same parameters as previously (default 0.75 Hurst parameter for SLAW, and 10 communities for HCMM). For MIT, the maximum network size is obviously fixed to the 92 nodes participating in the experiment. We do not use the Infocom trace for this part, as it is already quite small at 41 nodes.

Fig. 5 shows the evolution of the graph conductances Φ in function of the network size.

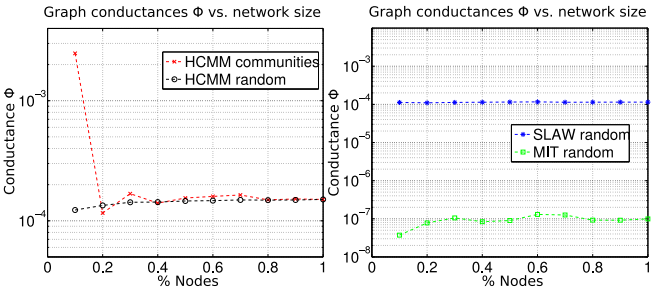

(a) HCMM (max. 500 nodes) (b) SLAW (500 nodes), MIT (92 nodes)

Figure 5. Evolution of the graph conductance Φ with network size

It is apparent from Fig. 5 that the conductance of typical DTN scenarios is relatively constant with increasing network size. This is also consistent with observations from [34]. There, the authors empirically analyze a series of large real-world networks (online social networks, peer-to-peer

networks, citation networks etc) and find that, when communities exist, their sizes are limited to about 100 nodes, regardless of the size of the network. When partitions larger than 100 nodes are sought for, the resulting $\alpha-$conductances Φ_α are consistently increasing with partition size.

In other words, the minimum conductance Φ is already reached for small network sizes. This is very important from a practical viewpoint, as shown in the beginning of this section. In particular, it means that the increase in spreading delay is very slow (logarithmic) as the network size increases. Incidentally, it also means that the derived bounds can easily be calculated even for larger networks.

VI. Applications

In this section, we discuss the possible usage of our derived bounds in practice, by showing how they apply to epidemic-based DTN algorithms. Some of these algorithms (e.g., for distributed estimation) rely on *multi-source spreading* (MSS) instead of the basic single-source spreading (SSS) task we defined in Sec. III. There, each i of an arbitrary subset of nodes $A \subseteq V$ has a distinct message to distribute to all nodes $V \setminus \{i\}$. For the MSS task, similar delay bounds can be obtained, e.g., by applying Boole's inequality (the union bound) to the results for the SSS task. We defer a more thorough analysis of the MSS task to future work.

Epidemic routing and its variants. Epidemic routing (ER) was the first DTN routing protocol [1] and it is essentially a brute-force aproach to routing, as it is trying all possible space-time paths. Obviously, epidemic routing is highly inefficient. To reduce the overhead, randomized forwarding (RF) was introduced. In this case, for a given message m, any node carrying the message transfers it to any other node it encounters *with a fixed probability* p. Then, the delivery delay of m can be analyzed similarly to the previous section, as shown below.

Theorem 5 (Delay of Randomized Forwarding.): The expected delay of Randomized Forwarding is, similarly to Eq. (11), given by:

$$\mathbb{E}[D_{\mathrm{RF}}] < \frac{2\ln(N-1)}{Np\Phi}. \tag{25}$$

A similar equivalent can be obtained for Eq. (21), simply by replacing Φ with $p\Phi$ in the equation.

Proof: With simple epidemic spreading, the meeting probability p^c_{ij} is also the infection probability (i.e., the probability that i infects j, given that i is infected and j is not infected). By contrast, with randomized forwarding, both the meeting and the forwarding probabilities must be considered, resulting in new infection probabilities: $p \cdot p^c_{ij}$. Thus, the contact process of our DTN is effectively *thinned*. It is known that thinning by ε transforms a geometric distribution with success probability p_{succ} into a geometric distribution with probability $\varepsilon \cdot p_{\mathrm{succ}}$. Therefore, randomized forwarding can be analyzed, by analyzing a new DTN contact graph, defined by the contact probability matrix $\mathbf{P}^c_{\mathrm{RF}} = p \cdot \mathbf{P}^c$.

By constructing the corresponding Markov chain as in Section IV and applying the subsequent analysis, one can obtain an equivalent for Eq. (11) as shown in this theorem. The same goes for Eq. (21), simply by replacing Φ with $p\Phi$ in the equation. ■

For simple epidemic routing, as well as for time-limited epidemic routing, the bounds derived in the previous section can be used to determine an appropriate Time-To-Live (TTL) and respectively, the time-threshold parameter. For probabilistic forwarding, in addition to the TTL, the fixed forwarding probability p can be derived, so as to achieve a certain delivery probability.

Distributed estimation. Epidemic spreading also underlies a number of distributed estimation algorithms. While the proposed analysis may not directly and easily apply to all of them, we give here an example where it does. In [9], the authors propose a simple epidemic-based algorithm for estimating separable functions of individual values present at nodes. The algorithm relies on a well-known property of the extrema (minimum or maximum) of n random variables of certain probability distributions. For example, the minimum of exponential variables $X_1 \sim \mathcal{E}(\lambda_1)$ and $X_2 \sim \mathcal{E}(\lambda_2)$ is exponentially distributed with rate $\lambda_1+\lambda_2$ (also valid for more than two variables). Thus, to estimate an additive function of local nodes values v_i (e.g., network size), it suffices that each node i generate K samples from $X_i \sim \mathcal{E}(v_i)$. Then, on each encounter, the two participating nodes calculate the pointwise minimum of their two samples. After enough encounters, the global parameter is estimated by maximum likelihood from the local vector of minimum values.

It is proven that the estimation delay of the above algorithm is a simple function of the delay of the underlying epidemic spreading algorithm. As a consequence, our analysis also applies to this distributed estimation algorithm. Mosk-Aoyama and co-authors also provide an upper bound for their algorithm, equivalent to the one in [12], which we have shown to be much looser than the one we propose.

In the case of distributed estimation, the proposed results can be used to, e.g., easily obtain an approximation of the time until termination, i.e., the time when the estimate at each node has reached a certain precision. Detecting termination is a common concern with distributed estimation, which may be alleviated through the use of these bounds.

VII. Conclusion and Future Work

In this paper, we have used a much more realistic DTN mobility model than previously to analyze the delay of epidemic information spreading in Opportunistic Networks. Unlike earlier analytical research work, our network and mobility model captures the full heterogeneity of node mobility, which has been observed in real scenarios [7], [8].

We have derived upper bounds for the expectation and the distribution of the delay of epidemic spreading. We have also shown that the distribution is concentrated and that, depending on graph properties (communities, scale-freeness), the delay scales very well with network size. This

delay is very important, as epidemic spreading underlies a series of DTN algorithms, e.g., for routing, and for distributed estimation. Moreover, we have shown that our bounds are tighter than previously derived ones [12].

In the future, we plan to further relax the assumptions of our model, as well as to extend our analysis to more sophisticated epidemic-based routing algorithms and to further distributed estimation algorithms.

REFERENCES

[1] Vahdat A and Becker D. *Epidemic routing for partially connected ad hoc networks*. Tech. Rep. CS-200006, Duke University, 2000.

[2] Guerrieri A, Carreras I et al. *Distributed estimation of global parameters in delay-tolerant networks*. Comput Commun, 33:1472–1482, 2010.

[3] Chaintreau A, Le Boudec JY et al. *The age of gossip: spatial mean field regime*. SIGMETRICS 2009, 109–120.

[4] Groenevelt R, Nain P et al. *The message delay in mobile ad hoc networks*. Perform Eval, 62:210–228, 2005.

[5] Zhang X, Neglia G et al. *Performance modeling of epidemic routing*. Comput Netw, 51:2867–2891, 2007.

[6] Haas ZJ and Small T. *A new networking model for biological applications of ad hoc sensor networks*. IEEE/ACM Trans Netw, 2006.

[7] Eagle N and Pentland A. *Reality mining: sensing complex social systems*. Pers Ubiq Comput, 10(4):255–268, 2006. ISSN 1617-4909.

[8] Lenders V, Wagner J et al. *Measurements from an 802.11b mobile ad hoc network*. WOWMOM 2006.

[9] Mosk-Aoyama D and Shah D. *Fast distributed algorithms for computing separable functions*. IEEE Trans Inf Theory, 2008.

[10] Krifa A, Baraka C et al. *Optimal buffer management policies for delay tolerant networks*. IEEE SECON 2008, 260–268.

[11] El Fawal A, Le Boudec JY et al. *Self-Limiting Epidemic Forwarding*. Tech. rep., 2006.

[12] Shah D. *Gossip Algorithms*. Foundations and Trends in Networking, 3(1):1–125, 2009.

[13] Chierichetti F, Lattanzi S et al. *Almost tight bounds for rumour spreading with conductance*. STOC 2010, 399–408.

[14] Giakkoupis G. *Tight bounds for rumor spreading in graphs of a given conductance*. STACS 2011, vol. 9, 57–68.

[15] Censor-Hillel K, Haeupler B et al. *Global computation in a poorly connected world: Fast rumor spreading with no dependence on conductance*. CoRR, abs/1104.2944, 2011.

[16] Baccelli F and Brémaud P. Elements of Queueing Theory. 2003.

[17] Carreras I, Miorandi D et al. *A simple model of contact patterns in delay-tolerant networks*. Wireless Networks, 16, 2010.

[18] Chaintreau A, Hui P et al. *Impact of human mobility on the design of opportunistic forwarding algorithms*. INFOCOM 2006.

[19] Karagiannis T, Le Boudec JY et al. *Power law and exponential decay of inter contact times between mobile devices*. MobiCom 2007, 183–194.

[20] Passarella A and Conti M. *Characterising aggregate inter-contact times in heterogeneous opportunistic networks*. IFIP NETWORKING 2011.

[21] van der Hofstad R. *Random graphs and complex networks*, 2011. URL http://www.win.tue.nl/~rhofstad/NotesRGCN.pdf.

[22] Picu A and Spyropoulos T. *Distributed optimization in DTNs: Towards understanding greedy and stoch. algos*. Tech. Rep. 326, ETHZ, 2010.

[23] Hossmann T, Spyropoulos T et al. *A complex network analysis of human mobility*. NetSciCom 2011.

[24] Kemeny J and Snell L. Finite Markov Chains. 1960.

[25] Picu A, Spyropoulos T et al. *An analysis of the information spreading delay in heterogeneous mobility DTNs*. Tech. Rep. 341, ETHZ, 2011.

[26] Hossmann T, Spyropoulos T et al. *Putting contacts into context: Mobility modeling beyond inter-contact times*. ACM MobiHoc 2011.

[27] Brémaud P. Markov chains : Gibbs fields, Monte Carlo simulation, and queues. Springer, Berlin, Germany, 2001. ISBN 0-38798-509-3.

[28] Sen A and Balakrishnan N. *Convolution of geometrics and a reliability problem*. Statistics & Probability Letters, 43(4):421–426, 1999.

[29] Hui P, Chaintreau A et al. *Pocket switched networks and human mobility in conference environments*. WDTN 2005.

[30] Boldrini C and Passarella A. *Hcmm: Modelling spatial and temporal properties of human mobility driven by users' social relationships*. Comput Commun, 33:1056–1074, 2010.

[31] Lee K, Hong S et al. *Slaw: A new mobility model for human walks*. IEEE INFOCOM 2009, 855–863.

[32] Karypis G. *Metis – family of graph and hypergraph partitioning software*. URL http://glaros.dtc.umn.edu/gkhome/views/metis.

[33] Leskovec J, Lang KJ et al. *Empirical comparison of algorithms for network community detection*. ACM WWW 2010, 631–640.

[34] Leskovec J, Lang KJ et al. *Statistical properties of community structure in large social and information networks*. ACM WWW 2008, 695–704.

[35] Mihail M, Papadimitriou C et al. *On certain connectivity properties of the internet topology*. Journal of Computer and System Sciences, 72(2):239 – 251, 2006.