

Optimal and Scalable Distribution of Content Updates over a Mobile Social Network

Stratis Ioannidis, Augustin Chaintreau, Laurent Massoulié

Thomson, Paris, France.

stratis@cs.toronto.edu

augustin.chaintreau@thomson.net

laurent.massoulie@thomson.net

Thomson Technical Report
Number: **CR-PRL-2008-08-0001**
Date: August 29th 2008

Abstract: We study the dissemination of dynamic content, such as news or traffic information, over a mobile social network. In this application, mobile users subscribe to a dynamic-content distribution service, offered by their service provider. To improve coverage and increase capacity, we assume that users share any content updates they receive with other users they meet.

We make two contributions. First, we determine how the service provider can allocate its bandwidth optimally to make the content at users as “fresh” as possible. More precisely, we define a global fairness objective (namely, maximizing the aggregate user utility) and prove that the corresponding optimization problem can be solved by gradient descent. Second, we specify a condition under which the system is highly scalable: even if the total bandwidth dedicated by the service provider remains fixed, the expected content age at each user grows slowly (as $\log(n)$) with the number of users n . To the best of our knowledge, our work is the first to address these two aspects (optimality and scalability) of the distribution of dynamic content over a mobile social network.

1. INTRODUCTION

In opportunistic networks, contacts between mobile users can be exploited to exchange data, extending thus the network’s coverage and increasing its capacity. How often and when such contacts occur is dictated by the social interactions and relationships between users. For this reason, recent work [1–6] has focused on how knowledge of the *social network* formed by the mobile users can be used to improve the performance of a variety of applications.

In this paper, we investigate how the social network can assist in the distribution of dynamic content. In this application, users subscribing to a wireless service receive updates on frequently changing content, such as a news-feed, a blog, the price of a stock, traffic congestion information, *etc.* Subscribers of this service share their updates in an opportunistic fashion: Whenever two of them meet, the one whose content is most recent pushes it to the one whose content is older (thereby increasing the number of users that receive fresh information).

One question arising in the above setting is how should the service provider allocate its downlink capacity to ensure that the content at users is as “fresh” as possible. For example, should it allocate its available bandwidth uniformly among subscribers? Alternatively, should it provide more frequent updates to the “most social” subscribers, *i.e.*, the ones that meet other subscribers most often, in the hope that they would spread the content faster? In general, the answer depends on the provider’s downlink bandwidth as well as on how recent the content at users is required to be. Most importantly, it also depends on the users’ social behavior, since the latter determines how they meet. For this reason, answering the above question requires us to understand how the social network formed by mobile users affects the performance of our application.

A second question of importance is how such a system scales as the number of users grows. If more subscribers are added to the system, will the age of content at users increase, thus degrading the service, and, if so, by how much? Again, this depends on the topology of the social network formed by the users. Ideally, one wishes to find a general condition under which the age increases slowly as the network grows.

Our main contribution is providing comprehensive answers to the above two questions. To the best of our knowledge, our work is the first to address these two aspects (optimality and scalability) of the distribution of dynamic content over a mobile social network.

First, we show how the service provider can determine the optimal allocation of its bandwidth. More precisely, the service provider can compute a downlink rate allocation that satisfies a global fairness objective —namely, maximizing the aggregate utility over all users. We

prove that the corresponding optimization problem is convex and therefore can be solved efficiently by gradient descent. Moreover, we give both a centralized and a distributed algorithm for computing the gradient; these can be used by the service provider to compute the optimal allocation, as we illustrate with an empirical study.

Second, we prove that the system described above is scalable, under the condition that the social network formed by the subscribers has a *bounded edge expansion*. In particular, even if the service provider distributes updated content with a fixed total rate, the content ages as seen by users grow slowly (as $\log(n)$) as the number of users n increases. Our second result therefore identifies edge expansion as a key property of the social network that affects scalability. Most importantly, it also implies that the service provider can exploit the social network to offer the service with limited resources, without sustaining a considerable degradation of the service due to system growth.

Our empirical study uses two real-life mobility traces, spanning over different time scales (a few hours and several days, respectively). We compute the optimal downlink allocation and compare it to several simple heuristics, illustrating its dependence on system parameters. An interesting outcome of our study is that the intuition that “most social” users should receive more frequent updates can in fact be wrong: under certain conditions, it is actually optimal to allocate none of the available bandwidth to the most social users in the system.

2. RELATED WORK

Allocating the service provider’s bandwidth among its subscribers has similarities to “spread of influence maximization” problems over traditional social networks. For example, Kempe *et al.* [7] looked at which consumers a product should be marketed to in order to ensure its widespread adoption. The authors show that the objective function exhibits a property called submodularity, and that greedy algorithms find solutions within a constant approximation factor to the optimal. Similarly, Leskovec *et al.* [8] looked at which blogs one should read to quickly detect the outbreak of an important story. Our application and our model are considerably different. Moreover, in our system, the convexity property implies that computing the optimal (rather than an approximate) solution is feasible, as we prove, even in a distributed manner.

In the context of DTNs, algorithms taking advantage of the social behaviour of mobile users have been proposed for publish/subscribe systems [1, 2], routing [3–5] and query propagation [4, 6]. These algorithms exploit concepts from social networks, including node centrality [1, 3, 5], friendship relationships [1, 4], bazaars [6], contact usefulness [2]. However, formally assessing the effect of the social network on the performance of these

algorithms remains largely an open question. Though our focus is on a different application, our work rigorously relates social network properties to the system’s scalability and to algorithms for finding the optimal allocation. As such, it strengthens the foundations of research in the area of mobile social networks.

Measurement studies of human mobility [5,9–11] have mostly focused on inter-contact time statistics and their effect on opportunistic forwarding. Our work illustrates how such human mobility can be used to disseminate content optimally in Section 6.2.

The importance of edge expansion in epidemic dissemination is well known, and has recently been studied in the context of search in unstructured peer-to-peer networks [12], the lifetime of infections by computer viruses [13], and the distributed aggregation of measurements [14]. Our scalability analysis is similar in spirit, as content updates are diffused over the users’ social network. However, the dynamics of the diffusion are much different than in the above works; thus, our work highlights one more application in which edge expansion is of critical importance.

3. SYSTEM DESCRIPTION

3.1 Update Distribution Process

Consider a system of n mobile users that are served by a single service provider (*e.g.*, a cell-phone operator). We denote by $V = \{1, 2, \dots, n\}$ the set of all users. The service provider injects new content updates to the system according to a Poisson process with rate μ , which is bounded by the provider’s downlink capacity. The total injection rate μ is allocated among different users as follows: each new update is pushed to one user chosen from V with a certain probability. As a result, user i receives updates according to a Poisson process with injection rate $x_i \geq 0$, $i \in V$, such that $\sum x_i = \mu$. We denote by \vec{x} the vector of injection rates, which we call the *rate allocation*.

Users share their content with other users they contact (*i.e.* that are within their transmission range) in the following way. First, the content stored at a user is time-stamped with the time at which it was originally downloaded from the service provider. Let $t_i(t)$ be the time-stamp of user i ’s content at time t . The following scheme is then used to share the content:

CONTENT SHARING: A user i will copy user j ’s content when they meet, if the content stored at j is more recent than the content stored at i , *i.e.*, $t_i < t_j$.

Note that, after two users i and j have met, the time-stamp at both becomes $\max(t_i, t_j)$.

We are interested in the *age* Y_i of the content stored

at each user i , defined as:

$$Y_i(t) = t - t_i(t), \quad i \in V.$$

In particular, given a user i we wish to study the distribution of the age Y_i when the system is in steady state or, informally, after the system has operated for a sufficiently long time. Formally, we evaluate $Y_i(T)$ at some time $T > 0$, given that the system has been running in the interval $(-\infty, T]$. Note that Y_i depends both on the rate allocation vector \vec{x} as well as on how contacts between users take place.

3.2 Contact Process and Contact Graph

We assume that contacts are symmetric, *i.e.*, user $j \in V$ contacts user $i \in V$ whenever i contacts j , and they last for a time that is negligible compared to the time between two consecutive contacts. Moreover, we assume that the joint contact process, describing contacts among all pairs of users (i, j) , is independent of the content injection process and is stationary ergodic. A simple case for which this holds is when the contact processes between distinct pairs (i, j) are independent renewal processes. However, we do not require these to be independent in our model. For example, i meeting j might increase the chance that i will also meet some $j' \neq j$ within a brief period of time.

In the case where the contact processes are independent renewal processes, we can define the mean contact rate q_{ij} between users i and j , where $q_{ij} \geq 0$. The *contact graph* of the system is a complete, weighted and undirected graph G , whose vertex set is V and each edge (i, j) has weight q_{ij} .

Given a subset of users $A \subseteq V$, let $A^c = V \setminus A$. The *edge expansion* [15, 16] of G is then defined as

$$h_G = \min_{A \subseteq V} \frac{\sum_{i \in A, j \in A^c} q_{ij}}{\min(|A|, |A^c|)}.$$

As we will see in Section 4.2, this property of the contact graph plays an important role in the system’s scalability.

3.3 User Utilities and Optimization Objectives

We would like to choose the rate allocation vector \vec{x} so that a global objective is attained. In general, we assume that a user i is happier when its content is more recent. There are several ways to quantify this notion, one being through a non-increasing utility $u_i : \mathbb{R}_+ \rightarrow \mathbb{R}$ that is a function of the age Y_i . There is no reason to restrict ourselves to positive utility functions —negative utilities can express dissatisfaction or loss of profit.

How much content of a certain age is worth to a user depends on user preferences as well as the nature of the content provided. Some examples of non-increasing utilities are illustrated in Fig. 1. In utilities u_a and u_b , an age threshold value τ exists after which the content is

worthless; this could be the case if, *e.g.*, it is news about sales offers that expire after some time. In utility u_c , even very old content has a vanishing but non-zero value to the user. Finally, the negative utility u_d expresses “dissatisfaction” or “loss” growing linearly with the age.

Denote by $\mathbb{E}_{\vec{x}}[\cdot]$ the expectation of a random variable, given that the rate allocation vector is \vec{x} . A natural goal for the service provider is to maximize the aggregated utility among all users, *i.e.*, the social welfare:

SOCIAL WELFARE MAXIMIZATION

$$\text{Maximize} \quad f(\vec{x}) = \sum_{i=1}^n \mathbb{E}_{\vec{x}}[u_i(Y_i)],$$

$$\text{subject to:} \quad \sum_{i=1}^n x_i \leq \mu \text{ and } x_i \geq 0, 1 \leq i \leq n$$

where $u_i : \mathbb{R}_+ \rightarrow \mathbb{R}$, $1 \leq i \leq n$, are non-increasing, and $\mathbb{E}_{\vec{x}}[u_i(Y_i)]$ the expected utility at user i in steady state under the rate allocation vector \vec{x} .

We note that other global optimization objectives might also be of interest. One example is a weighted version of the above problem, where each expected utility at user i is multiplied by weight $w_i \geq 0$. By setting $w_i = 0$, a server can target only a subset of the users. Another alternative is a “max-min” fair allocation, obtained by replacing the summation in the objective function f by a minimization. Our results (namely, Theorem 1) extend to both of the above cases; the corresponding optimization problems can again be solved with the methods we outline in this paper, as discussed in the end of Section 5.1. For concreteness however, our focus will be on social welfare.

4. MAIN RESULTS

4.1 Optimal Rate Allocation

Our first main result concerns the solution of the social welfare maximization problem. Given a system of mobile users implementing the CONTENT SHARING protocol, we wish to find how the service provider should choose the rate allocation \vec{x} , in order to maximize social welfare. We prove the following theorem, whose generality is surprising: An optimal allocation can be found under all non-increasing utility functions and for general stationary ergodic contact processes. In particular, it is not necessary that contacts between users are independent.

THEOREM 1. *If the user utilities $u_i : \mathbb{R}_+ \rightarrow \mathbb{R}$, $1 \leq i \leq n$, are non-increasing functions, and the joint contact process is stationary ergodic, SOCIAL WELFARE MAXIMIZATION is a convex optimization problem. In particular, the objective function $f(\vec{x}) = \sum_{i=1}^n \mathbb{E}_{\vec{x}}[u_i(Y_i)]$ is concave.*

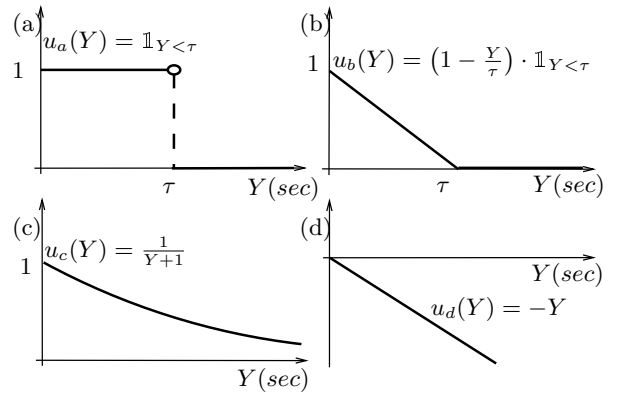


Figure 1: Examples of utility functions. In cases (a) and (b), if the age exceeds a threshold the content has no value to the user, whereas in case (c) even very old content has some positive value. In case (d), the older the content the higher the loss incurred at the user.

Theorem 1 implies that any local maximum of the objective function is a global maximum, and that the maximization can be performed by gradient descent [17], as we describe in Section 6.

In general, to solve the optimization problem with gradient descent, the server needs to know both the user utilities u_i and the steady state c.d.f.’s of the ages Y_i . The latter might be hard to compute in a closed form for a given system, even if the contact processes are independent renewal processes; in Section 6.2, we discuss how they can be estimated in a centralized way by the service provider by gathering simple statistics on the contact processes between users.

However, it is not necessary to follow a centralized approach: In Section 6.3, we present an algorithm with which users can estimate the gradient of the objective function in a fully decentralized manner. Neither user utilities nor traces of user contact processes need to be reported to the service provider using this approach; the users compute and report only their estimates of the gradient, and the service provider can use this information to adjust the injection rates accordingly.

4.2 Scalability

Our second main result addresses the issue of how the CONTENT SHARING protocol scales as the number of mobile users in the system increases. To obtain this result, we assume that the contact processes between users are independent Poisson processes.

THEOREM 2. *Assume that the contact processes between pairs of users (i, j) are independent Poisson processes. If $\vec{x} = [\frac{\mu}{n}]$ (*i.e.*, the service provider chooses uniform rate allocation), the expected age seen by any*

user $i \in V$ in steady state satisfies

$$\mathbb{E}_{\vec{x}}[Y_i] \leq \frac{2}{\mu} \left(2e^{-1/2} + \log(n) \right) + h_G^{-1} \log n. \quad (1)$$

where h_G the edge expansion of the contact graph of the system.

The theorem suggests that sharing content can significantly benefit the system. It is easy to see that, if the users do not share their content, the expected age $\mathbb{E}[Y_i]$ at any user i grows as n/μ , i.e., linearly in n . Theorem 2 states that the ages can grow much slower (as $\log(n)$) when content is shared, if the edge expansion h_G of the contact graph is bounded away from zero. Graphs exhibiting a bounded expansion, also called *expander graphs*, are abundant [16, 18]; in particular, any graph with a sufficiently rich random structure is an expander.

Finally, Theorem 1 can be used to give a lower bound on the social welfare provided that user utilities are convex (as, e.g., utilities u_b, u_c, u_d of Fig. 1).

COROLLARY 1. *If the user utilities $u_i, i \in V$, are convex, then the aggregate expected utility under the optimal rate allocation is at least*

$$\sum_{i=1}^n u_i \left(\frac{2}{\mu} \left(2e^{-1/2} + \log(n) \right) + h_G^{-1} \log n \right).$$

5. ANALYSIS

In this section, we give the proofs of Theorems 1 and 2. We first give a simple characterization of the content age in terms of a simple message propagation scheme.

We will say that a message originating at user $i \in V$ at some time T is flooded over the system if it is propagated as follows: every user having a copy of the message forwards it to every other user it contacts.

For a given time T , we define the process $B_i^T(t) \subset V$ at user i as follows: A user j is in $B_i^T(t)$ if a message placed at a user j at time $T-t$ can reach user i through flooding by time T . In other words, if j forwards the message to every user it contacts, and every other user that receives it also relays it to all other users it meets, the message will reach i before T .

Alternatively, $B_i^T(t)$ can be defined through a flooding that starts from user i and is propagated over the “backwards” contact process: Suppose that, at time T , we “reverse the arrow of time” and look at the process describing the contacts between users “going backwards”. If a message originating at i is flooded over this backwards process, then $B_i^T(t)$ is precisely the set of users that will have the message after time t . For this reason, we call $B_i^T(t)$ the *backwards growth process* at i .

We define s_{ij}^T as

$$s_{ij}^T = \inf_{t \geq 0} \{t \mid j \in B_i^T(t)\}.$$

Looking at the contact process with the arrow of time reversed, s_{ij}^T is the time it takes until a message originating at i reaches j (over the backwards process). For this reason, we call s_{ij}^T the *backwards latency* from i to j . Note that, by definition, $B_i^T(t) = \{j, \text{ s.t. } s_{ij}^T \leq t\}$ and that, if there is no $t > 0$ such that $j \in B_i^T(t)$, then $s_{ij}^T = \infty$.

Recall that $Y_i(T)$ is the age of user i 's content at time T . We can succinctly express $Y_i(T)$ in terms of the latencies s_{ij}^T :

LEMMA 1. *Let $Z_i(T)$ be the elapsed time since user i downloaded content directly from the service provider. Then, for all $T \geq 0$ and all $i \in V$,*

$$Y_i(T) = \min_{j \in V} \{s_{ij}^T + Z_j(T - s_{ij}^T)\}. \quad (2)$$

The proof can be found in Appendix A.1.

5.1 Proof of Theorem 1

Lemma 1 relates the age at user i at time T to the backwards latencies s_{ij}^T . The following Lemma, whose proof is in Appendix A.2, uses this relationship to express the distribution of the age of a user in terms of the latencies in steady state. We denote by $\mathbb{P}_{\vec{x}}(\cdot)$ the probability of an event given that the rate allocation vector is \vec{x} .

LEMMA 2. *Let Y_i be user i 's the steady-state content age and $s_{ij}, j \in V$, the steady-state backwards latencies from i . Then*

$$\mathbb{P}_{\vec{x}}(Y_i > t) = \mathbb{P}_{\vec{x}}(Y_i \geq t) = \mathbb{E} \left[e^{-\sum_{j=1}^n x_j \cdot (t - s_{ij})_+} \right] \quad (3)$$

where the expectation is over the latencies $s_{ij}, j \in V$, in steady state and $(\cdot)_+ \equiv \max(\cdot, 0)$.

An immediate implication of this lemma is that, for every user i , the c.d.f. of Y_i is a concave function of \vec{x} .

COROLLARY 2. *For all $i = 1, \dots, n$, and for any fixed $t > 0$, $\mathbb{P}_{\vec{x}}(Y_i \leq t)$ is concave in \vec{x} .*

To see this, observe that the function $e^{-\sum_{j=1}^n x_j (t - s_{ij})_+}$ is convex in \vec{x} , as the composition of a convex and a linear function. Moreover, if every element in a family of functions $g(\vec{x}, u), u \in \Omega$, is convex in \vec{x} and ν is a positive measure in Ω , the integral $\int_{\Omega} g(\vec{x}, u) d\nu$ is also convex [17]. Hence, the expectation of the above functions over $s_{ij}, j \in V$, is also convex, and the corollary follows from Lemma 2.

As $\mathbb{E}_{\vec{x}}[\mathbb{1}_{Y_i \leq \tau}] = \mathbb{P}_{\vec{x}}(Y_i \leq \tau)$, the above corollary effectively says that if the utility is a step function (like u_a in Fig. 1(a)) its expectation in steady state is concave. It is straightforward to extend this result to general non-increasing utilities.

LEMMA 3. *If $u : \mathbb{R}_+ \rightarrow \mathbb{R}$ be a non-increasing function, then for every $i \in V$, $\mathbb{E}_{\vec{x}}[u(Y_i)]$ is concave.*

The proof can be found in Appendix A.3.

Theorem 1 therefore follows from the fact that the sum of concave functions is concave. A consequence of Lemma 3 is that Theorem 1 extends to any function of the expected utilities $\mathbb{E}_{\vec{x}}[u_i(Y_i)]$ that preserves concavity. For this reason, Theorem 1 holds, *e.g.*, for weighted sums of the expected utilities as well as for “max-min fair” allocations, as noted in Section 3.3.

5.2 Proof of Theorem 2

Assume that the contact processes between pairs (i, j) , $i, j \in V$ are independent Poisson processes. Suppose that at time T a message is placed at user i and flooded over the *forwards* process. We define the *forwards growth process* $A_i^T(t)$ as the set of users reached by the message by time $T + t$. The reversibility and stationarity of the Poisson process imply that, in steady state, the backwards growth process is indistinguishable from the forwards growth process.

LEMMA 4. *In steady state, $\{A_i^T(t); t \geq 0\}$ is identically distributed as $\{B_i^T(t); t \geq 0\}$.*

The steady state behavior of the backwards growth process $B_i^T(t)$ can thus be understood by simply looking at the behavior of the forwards growth process. The latter relates to the edge expansion h_G of our system’s contact graph through the following lemma, whose proof is in Appendix A.4.

LEMMA 5. $\mathbb{P}(|A_i^T(t)| \geq k) \geq (1 - e^{-h_G t})^{k-1}$ when $k \leq n/2$, for every $i \in V$.

Lemma 5, along with Lemma 1, allows us to obtain the following upper bound on the steady-state expected age at i . Intuitively, this bound is derived by observing that, for any $t > 0$, the age at some time T will be no more than t plus the minimum age among all users in the set $B_i^T(t)$.

LEMMA 6. *Suppose that $x_j = \frac{\mu}{n}$, for all $j \in V$. Then, in steady state, for any $t > 0$, and any $i \in V$,*

$$\mathbb{E}[Y_i] \leq t + \frac{n}{\mu} \left[\left[\frac{n}{2} \right]^{-1} (1 - e^{-h_G t})^{\lfloor \frac{n}{2} \rfloor - 1} + \frac{h_G t e^{-h_G t}}{1 - e^{-h_G t}} \right].$$

The proof can be found in Appendix A.5. Using Lemma 6, we can bound $\mathbb{E}[Y_i]$ by appropriately choosing t . Theorem 2 follows as a corollary by setting $t = h_G^{-1} \log n$. The proof of Corollary 1 is an application of Jensen’s inequality.

6. CENTRALIZED AND DISTRIBUTED OPTIMIZATION

In this section, we discuss how the service provider can compute the optimal rate allocation. To do so, it needs to implement gradient descent [17], which requires computing the gradient vector $\nabla f = \left[\frac{\partial f}{\partial x_i} \right]_{i \in V}$ of

the objective function $f(\vec{x}) = \sum \mathbb{E}_{\vec{x}}[u_i(Y_i)]$. As noted in Section 4.1, knowledge of the users’ utilities as well as the c.d.f.’s of the ages Y_i in steady state is required to compute the gradient. In general, it is not always possible to obtain the latter in a closed form, even for simple contact processes. For this reason, we settle for estimating ∇f through an unbiased estimator, which we denote with $\widehat{\nabla} f$.

In the following, we first outline how the service provider can compute the optimal rate allocation given an unbiased estimator $\widehat{\nabla} f$. We then show two ways to derive such an estimator. The first is centralized: the service provider needs to know the user utilities and collect contact statistics to apply it. The second is distributed: the gradient is computed directly by the users and reported to the service provider.

6.1 Implementing Gradient Descent with a Gradient Estimator

Given an unbiased estimator $\widehat{\nabla} f$ of ∇f , the service provider can use the following projected gradient descent algorithm to compute the optimal rate allocation vector:

$$\vec{x}_{k+1} = \Pi \left(\vec{x}_k + \gamma_k \widehat{\nabla} f(\vec{x}_k) \right), \quad (4)$$

where γ_k is some positive gain parameter such that $\sum_{k=0}^{\infty} \gamma_k = \infty$, $\lim_{k \rightarrow \infty} \gamma_k = 0$ and Π is the projection on the set $\{\vec{x} \in \mathbb{R}_+^n : \sum_i x_i \leq \mu\}$.

The study of such algorithms constitutes the field of stochastic approximation, and there are known technical conditions on the sequences of gradient estimates $\widehat{\nabla} f(\vec{x}_k)$ and gain parameters γ_k which guarantee convergence of \vec{x}_k to a maximiser of the objective function f . One example is the following lemma:

LEMMA 7 (BENAÏM, [19]). *Suppose that for some $q \geq 2$ $\sup_{x_k} \mathbb{E}[|\nabla f(x_k) - \widehat{\nabla} f(x_k)|^q] < \infty$ and $\sum_k \gamma_k^{1+q/2} < \infty$. Then the sequence (4) converges to a maximizer of f a.s.*

In the following, we will use the above lemma to guarantee the convergence of our algorithms to an optimal allocation vector \vec{x} under certain conditions. These conditions are sufficient but not necessary; both our centralized and our distributed algorithms may converge even if these conditions do not hold, as we illustrate in Section 7.

6.2 A Centralized Implementation

We first assume the service provider knows the utility functions u_i and collects traces of user contacts. For example, the mobile devices may log contacts and upload their logs to their service provider; alternatively, the positions and collocations of users can be monitored (*e.g.*, by triangulation). Both assumptions are removed in the next section.

As noted above, the service provider needs to know the c.d.f. of the ages Y_i in steady state to compute the gradient. Lemma 2 suggests a way to estimate these distributions from samples of the backwards latencies s_{ij}^T , $i, j \in V$. To begin with, by eq. (3), for $i, j \in V$,

$$\frac{\partial \mathbb{P}_{\bar{x}}(Y_i < t)}{\partial x_j} = \mathbb{E} \left[(t - s_{ij})_+ e^{-\sum_{j=1}^n x_j (t - s_{ij})_+} \right] \quad (5)$$

If the service provider has traces of user contacts, for each $i \in V$, it can generate samples of s_{ij}^T , $j \in V$, at different times T , by flooding messages from i over the backwards contact process (e.g., by “running” the contact traces backwards). In steady state (i.e., for large T), the empirical mean of the r.h.s. of (5) is an unbiased estimator of $\frac{\partial \mathbb{P}_{\bar{x}}(Y_i < t)}{\partial x_j}$.

The gradient can then be estimated as follows:

$$\widehat{\frac{\partial f}{\partial x_j}}(\bar{x}) = \int_0^\infty \sum_{i=1}^n \frac{\partial \mathbb{P}_{\bar{x}}(\widehat{Y_i < u_i^{-1}(t)})}{\partial x_j} dt$$

where $\widehat{\frac{\partial \mathbb{P}_{\bar{x}}(Y_i < t)}{\partial x_j}}$ are the estimators of $\frac{\partial \mathbb{P}_{\bar{x}}(Y_i < t)}{\partial x_j}$, $i, j \in V$.

As the service provider collects more traces of user contacts and generates more samples of backwards latencies, it can adapt its allocation vector using (4). The following lemma states that convergence to an optimal solution can be guaranteed if the user utilities are bounded and integrable (as, e.g., u_a and u_b of Fig. 1). The proof can be found in Appendix A.6.

LEMMA 8. *Assume that the user utilities are bounded and integrable (i.e., $\int_0^\infty |u_i(t)| dt < \infty$). Then, the assumptions of Lemma 7 hold for $q = 2$ and $\gamma_k = 1/k$.*

6.3 A Distributed Implementation

The approach described in Section 6.2 relies on the service provider collecting information on both the user contacts and the utility functions u_i , where $i \in V$. While there are scenarios in which this can be done, it is clearly appealing to avoid this requirement. We now describe an alternative way of optimizing the rates x_i , which suppresses the need to learn the above quantities.

The main step consists of obtaining an unbiased estimate of the derivative

$$\frac{\partial}{\partial x_i} \mathbb{E}_{\bar{x}} [u_j(Y_j(t))] \quad (6)$$

of the expected utility of some user j , with respect to the injection rate x_i at a user i . To this end, we rely on the following result of [20]. Let $\{Z_t\}_{t \geq 0}$ be a stationary process driven by a Poisson process N with intensity $y > 0$. To stress its dependency on process N , we write $Z_t(N)$. Then, under suitable integrability assumptions,

$$\frac{\partial}{\partial y} \mathbb{E}_y [Z_t(N)] = \mathbb{E}_y \left[\int_0^\infty [Z_t(N + \delta_0) - Z_t(N)] dt \right]. \quad (7)$$

In the above expression, \mathbb{E}_y denotes the expectation when the intensity of N is y , and $N + \delta_0$ denotes the process consisting of the events in process N , plus an *additional event* occurring at time 0.

In other words, an unbiased estimate of the derivative in the left-hand side is given by the integral

$$\int_0^\infty [Z_t(N + \delta_0) - Z_t(N)] dt.$$

In our context, the process Z_t is the instant system utility, $\sum_j u_j(Y_j(t))$. Thus, the above expression can be interpreted as the overall additional utility brought to user j by an *extra* content injection at user i at time 0.

Let us see how this estimate can be computed in the present context. To estimate the derivative of $\mathbb{E}_{\bar{x}}[u_j(Y_j)]$ with respect to x_i , we proceed as follows. At some arbitrary time instant, say at 0, user i generates a *dummy* event, pretending to have received fresh information from the service provider. From there on, it maintains both its true age process $Y_i(t)$ and a dummy age process $\tilde{Y}_i(t)$, which has been artificially set to zero at time 0, but otherwise evolves as process $Y_i(t)$.

When two users k, ℓ meet, if k currently maintains a dummy age process, it then communicates to ℓ both its true and its dummy age; from there on, ℓ also runs both a true and a dummy age process, Y_ℓ and \tilde{Y}_ℓ . Note that, for any user j , the two processes Y_j and \tilde{Y}_j will eventually coincide (this is clearly enforced when the service provider injects new content at j). Provided that user j kept track of both its actual and its dummy age process (from the time t_{start} when it first received a dummy age till the time t_{end} when the two processes coincide), it can then locally compute the quantity

$$\begin{aligned} \Delta_{\bar{x},i}(j) &:= \int_{t_{start}}^{t_{end}} [u_j(\tilde{Y}_j(t)) - u_j(Y_j(t))] dt \\ &= \int_0^\infty [u_j(\tilde{Y}_j(t)) - u_j(Y_j(t))] dt. \end{aligned}$$

By the result (7) of [20], the quantity $\Delta_{\bar{x},i}(j)$ is an unbiased estimate of (6). This quantity, being the overall increase in utility to user j due to the addition of the dummy update at user i , is related to the notion of a *sample path shadow price* introduced in Kelly and Gibbens [21]. Indeed, they define this quantity as the pathwise cost increase due to a particular packet. Thus, the estimate $\Delta_{\bar{x},i}(j)$ above can be seen as the sample path shadow utility at user j of the dummy update at user i . Note that the creation of one dummy process by user i allows all users j to evaluate the corresponding estimate.

Fig. 2 illustrates how this estimate is effectively computed if the utility at j is $u_a = \mathbb{1}_{Y < \tau}$ or $u_d = -Y$, of Fig. 1. For u_a , the integral is simply the length of the period during which \tilde{Y}_j is below τ while Y_j is above τ , consisting of at most one non-empty interval ($t_2 - t_1$ in

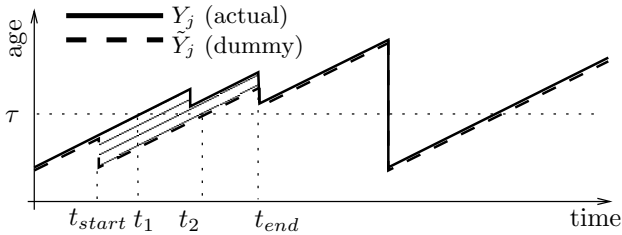


Figure 2: The original and dummy processes at user j . If the utility at j is $u_a(Y) = \mathbb{1}_{Y < \tau}$, then $\Delta_{x,i}(j) = t_2 - t_1$. If the utility is $u_d(Y) = -Y$, then $\Delta_{x,i}(j)$ is equal to the shaded region.

Fig. 2). For u_d , this integral is the area of one or more parallelograms (shaded in Fig. 2).

It remains to communicate such estimates to the service provider. For instance, this could happen whenever the service provider injects new content. Given some current choice \vec{x} of rates, the service provider can compute an unbiased estimate $\widehat{\nabla} f(\vec{x})$ of the gradient of the objective function f by taking its i -th coordinate to be

$$\frac{\partial f}{\partial x_i}(\vec{x}) = \sum_j \Delta_{\vec{x},i}(j).$$

We can again show the convergence of the sequence (4) provided that the utilities are again bounded and integrable. In this case, we also place an additional requirement on the time it takes to send a message from user i to j . The proof can be found in Appendix A.7.

LEMMA 9. *Assume that the user utilities are bounded and integrable. Moreover, assume that $\mathbb{E}[\ell_{ij}^T]^2 < \infty$ for all T and for all $i, j \in V$, where ℓ_{ij}^T the time it takes to send a message from user i to j through flooding over the forwards process. Then, the assumptions of Lemma 7 hold for $q = 2$ and $\gamma_k = 1/k$.*

7. EMPIRICAL STUDY

We implemented the centralized algorithm of Section 6.2 and used it to compute the optimal rate allocation for two real-world data sets of human mobility traces. The Infocom06 data set [22] contains opportunistic Bluetooth contacts between 98 iMotes, 78 of which distributed to Infocom 06 participants and 20 of which were static. We focused on a 10 hour period during the first day of conference. The MIT data set, collected by the Reality-Mining project [11], comprises 95 participants carrying GSM enabled cell-phones over a period of 9 months. We consider, as in [22], that two phones are in contact when they share the same GSM base station. We exclude 12 users from our analysis, as they were isolated. Due to memory size limitations, we limited our analysis of the MIT data set to an 80 day period.

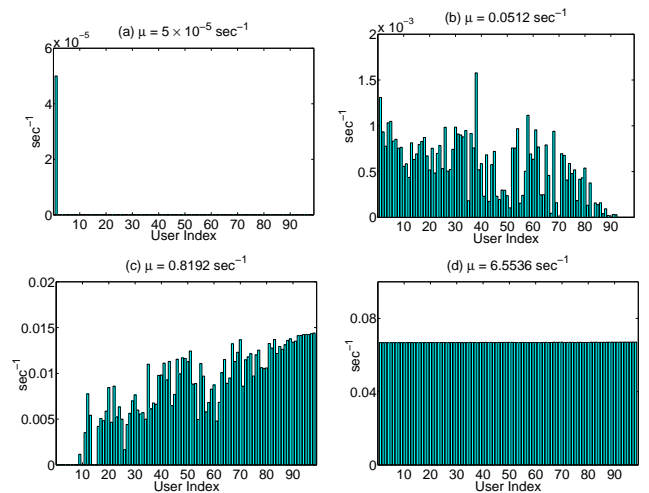


Figure 3: Four optimal allocations for the utility function $u_a(Y) = \mathbb{1}_{Y < 200}$ (Infocom06). Users are indexed according to their contact rates, in decreasing order. Although for small μ all updates are injected at the user with the highest contact rate, in (c) no updates are injected to the highest contact rate users.

We assume that every user has the same utility (*i.e.*, $u_i(Y_i) = u(Y_i)$), where u is one of the functions u_a and u_c shown in Fig. 1. For u_a , we chose the threshold value $\tau = 200\text{sec}$. Although utilities u_c are not integrable (and, thus, Lemma 8 does not apply), our algorithm converged for both utilities in both datasets. The allocation found outperformed all other allocations we obtained heuristically, as discussed below.

Fig. 3 presents the optimal rate allocation under u_a in the Infocom06 data set, for different values of μ . For small μ , the optimal allocation tends to be skewed towards users with high contact rates, as shown in Figures 3 (a) and (b). For both utility functions, the optimal allocation concentrates on a single user whenever μ is less than $6.4 \times 10^{-3} \text{sec}^{-1}$ (*i.e.*, one update every 2.6 minutes). We also observed this on the MIT data set when μ is less than $4 \times 10^{-4} \text{sec}^{-1}$ (*i.e.*, one update every 41 minutes). Intuitively, the injection rate is concentrated on the most “central” user, *i.e.* the one from which content is disseminated to all users the fastest. In the Infocom06 data set, the most central user is also the “most social” user, *i.e.* the one with the highest contact rate. This is not the case however for the MIT data set; the most central user had the third highest contact rate.

For higher values of μ , more users are allocated positive rates. In the Infocom 06 dataset, we observe an interesting phenomenon for utility u_a when μ is between 0.2 and 0.8 sec^{-1} . In this region, the injected rate at the users with the top 8 contact rates is zero

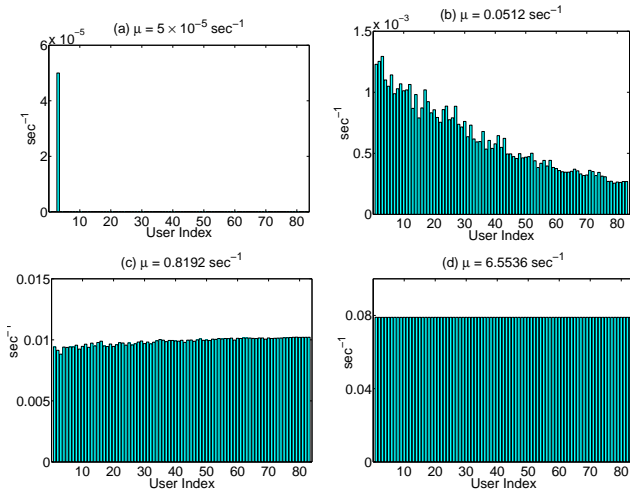


Figure 4: Four optimal allocations for the utility function $u_a = \mathbb{1}_{Y < 200}$ (MIT). Users are indexed according to their contact rates, in decreasing order. For small μ , all updates are injected to the user with the third highest contact rate. The bias towards nodes with high contact rates flips on the range $[0.2 \text{ sec}^{-1}, 0.8 \text{ sec}^{-1}]$: in (c), users with high contact rates receive updates less frequently than users with low contact rates, though the difference is not as pronounced as in Infocom 06.

(as in Fig. 3(c)), contradicting the intuition that very “social” users should receive higher injection rates. In fact, while for low values of μ the “most social” user accumulates all the injected rate, thus acting as a global hub of all incoming information, in this region of μ it receives all its updates from its neighbors. A similar behavior is observed on the MIT dataset (see Fig. 4), as well as for utilities u_b .

Last, when μ is very large, the optimal rate allocation becomes uniform among all users. Intuitively, the improvement provided by content sharing becomes negligible, as any user receives updates from its neighbors at a rate much smaller than its injection rate. Thus, the system behaves as if users were isolated (no sharing); the concavity of the expected utility $\mathbb{E}_{\bar{x}}[u(Y_i)]$, implies that, in this case, the optimal allocation is indeed the uniform.

In Figures 5 and 6, we plot the ratio between the social welfare achieved under several simple heuristics and the optimal. The heuristics considered are (a) the uniform allocation, (b) a skewed allocation, in which all the injection rate is concentrated at the “most social” user and (c) an allocation in which each user receives an injection rate that is proportional to its aggregate contact rate. For the Infocom06 data set, we also show the fraction of users with content age below the threshold

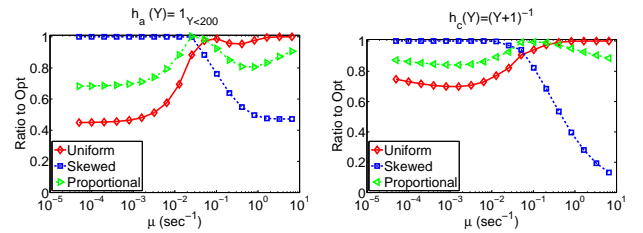


Figure 5: Ratio of social welfare of heuristic rate allocations to the optimal in the Infocom06 data set. Skewed is optimal for small μ , while for large μ uniform becomes optimal.

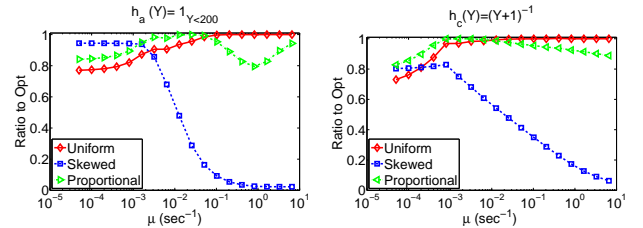


Figure 6: Ratio of social welfare of heuristic rate allocations to the optimal in the MIT data set. Skewed is not optimal for low μ , as the “most central” user is not also the “most social”.

for u_a in Table 1.

The comparison of the heuristic allocations confirms the above observations. The skewed allocation performs well for small values of μ , but not always optimally as it may not select the most central user (see Figure 6). Uniform is always optimal for large values of μ . Proportional is sometimes the best among the three for intermediate values of μ . Moreover, from Table 1, we see that the improvement under content sharing is significant: when $\mu = 0.0128$ (an update is injected every 78 sec), the expected number of users below the age threshold grows from 2.5% to about a third of the network as content sharing is used and service provider targets the most central user. In contrast, when an update is injected every 10 sec, 60% of the users on average receive the content on time (instead of 19% without sharing), and this is achieved when rates are allocated uniformly.

Our results highlight a transition depending on μ for the optimal rate allocation from skewed to uniform. These two simple heuristics perform well, but the social network plays an important role in selecting the most central users, as well as when μ takes intermediate values.

8. CONCLUSIONS

Our results show that content updates can be distributed over a mobile social network in a scalable way. Moreover, the social network can be successfully ex-

μ (sec ⁻¹)	no-sharing	skewed	prop.	uniform	optimal
0.0128	2.5%	34%	24%	30%	34%
0.0256	5.1%	42%	42%	37%	42%
0.1024	19%	46%	56%	60%	60%

Table 1: Expected fraction of users with age below 200sec in Infocom06.

ploited to obtain an optimal allocation of the service provider’s aggregate injection rate.

We see several other applications that could be explored with our model. For instance, content updates may actually be generated by the users, as opposed to being injected by a service provider; such an application is very appealing from a social networking perspective. Our distributed method for computing the gradient implies that such a system may support a pricing scheme. This is because it essentially outlines how to compute a user’s sensitivity to the injection rates of other users, in a distributed manner.

9. REFERENCES

- [1] E. Yoneki, P. Hui, S. Chan, and J. Crowcroft, “A socio-aware overlay for publish/subscribe communication in delay tolerant networks,” in *Proc. of ACM MSWiM*, 2007.
- [2] P. Costa, C. Mascolo, M. Musolesi, and G.-P. Picco, “Socially-aware routing for publish-subscribe in delay-tolerant mobile ad hoc networks,” *to appear in IEEE JSAC Special issue on Delay-Tolerant Networks*, 2008.
- [3] E. Daly and M. Hahr, “Social network analysis for routing in disconnected delay tolerant manets,” in *Proc. of ACM MobiHoc*, 2007.
- [4] A. G. Miklas, K. K. Gollu, K. K. W. Chan, S. Saroiu, P. K. Gummadi, and E. de Lara, “Exploiting social interactions in mobile systems,” in *Proc. of Ubicomp*, 2007.
- [5] V. Srinivasan, M. Motani, and W. T. Ooi, “Analysis and implications of student contact patterns derived from campus schedules,” in *Proc. of ACM MobiCom*, 2006.
- [6] M. Motani, V. Srinivasan, and P. S. Nuggehalli, “Peoplenet: engineering a wireless virtual social network,” in *Proc. of ACM MobiCom*, 2005.
- [7] D. Kempe, J. Kleinberg, and Éva Tardos, “Maximizing the spread of influence through a social network,” in *Proc. of ACM KDD*, 2003.
- [8] J. Leskovec, A. Krause, C. Guestrin, C. Faloutsos, J. VanBriesen, and N. Glance, “Cost-effective outbreak detection in networks,” in *Proc. of ACM KDD*, 2007.
- [9] A. Chaintreau, P. Hui, J. Crowcroft, C. Diot, R. Gass, and J. Scott, “Impact of human mobility on opportunistic forwarding algorithms,” *IEEE Trans. on Mob. Comp.*, vol. 6, no. 6, 2007.
- [10] T. Karagiannis, J.-Y. L. Boudec, and M. Vojnović, “Power law and exponential decay of inter contact times between mobile devices,” in *Proc. of ACM MobiCom*, 2007.
- [11] N. Eagle and A. Pentland, “Reality mining: Sensing complex social systems,” *Journal of Personal and Ubiquitous Computing*, vol. 10, no. 4, 2005.
- [12] C. Gkantsidis, M. Mihail, and A. Saberi, “Random walks in peer-to-peer networks,” in *Proc. of IEEE INFOCOM*, 2004.
- [13] A. Ganesh, L. Massoulié, and D. Towsley, “The effect of network topology on the spread of epidemics,” in *Proc. of IEEE INFOCOM*, 2005.
- [14] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah, “Randomized gossip algorithms,” *IEEE/ACM Trans. Netw.*, vol. 14, no. SI, 2006.
- [15] F. R. K. Chung, *Spectral Graph Theory*. American Mathematical Society, 1997.
- [16] S. Hoory, N. Linial, and A. Wigderson, “Expander graphs and their applications,” *Bulletin of the AMS*, vol. 43, no. 4, 2006.
- [17] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2004.
- [18] J. Friedman, “A proof of Alon’s second eigenvalue conjecture,” in *Proc. of ACM STOC*, 2003.
- [19] M. Benaïm, “Dynamics of stochastic approximation algorithms,” in *Séminaire de Probabilités XXXIII*, ser. Lect. Notes in Math., vol. 1709. Springer, 1999.
- [20] P. Brémaud and F. J. Vázquez-Abad, “On the pathwise computation of derivatives with respect to the rate of a point process: the phantom RPA method,” *Queueing Syst. Theory Appl.*, vol. 10, no. 3, 1992.
- [21] R. Gibbens and F. P. Kelly, “Resource pricing and the evolution of congestion control,” *Automatica*, no. 35, 1999.
- [22] A. Chaintreau, A. Mtibaa, L. Massoulié, and C. Diot, “The diameter of opportunistic mobile networks,” in *Proc. of ACM CoNEXT*, 2007.
- [23] R. G. Gallager, *Discrete Stochastic Processes*. Kluwer, 1996.

APPENDIX

A. PROOF OF TECHNICAL LEMMAS

A.1 Proof of Lemma 1

For a given $T > 0$, we define a *backwards contact sequence from i to j* as a sequence of pairs $(a_0, t_0), \dots, (a_k, t_k)$ in $V \times \mathbb{R}_+$ such that (a) $a_0 = i$ and $a_k = j$, (b) $0 = t_0 \leq t_1 \leq \dots \leq t_k$ and (c) for $1 \leq \ell \leq k$, user a_ℓ contacts $a_{\ell-1}$ on time $T - t_\ell$. Intuitively, a backwards contact sequence indicates the path followed by a packet originating at i until it reaches j , when packets are flooded over the backwards process. Note that if a backwards contact sequence $(a_0, t_0), \dots, (a_k, t_k)$ from i to j exists then $j \in B_i^T(t_k)$.

Consider w.l.o.g. a user j with $s_{ij}^T \neq \infty$. The definition of $B_i^T(t)$ implies that there exists a contact sequence $(a_0, t_0), \dots, (a_k, t_k)$ from i to j such that $t_\ell = s_{ia_\ell}^T$ for $0 \leq \ell \leq k$. Intuitively, this sequence is the route followed by the first message reaching j (in the backwards process). At each contact, the age of $a_{\ell-1}$ is set to the minimum of its current age and the age of a_ℓ . These minimizations imply (it can formally be shown by induction over the backwards contact sequence) that

$$Y_i(T) \leq s_{ij}^T + Y_j(T - s_{ij}^T) \leq s_{ij}^T + Z_j(T - s_{ij}^T)$$

as $Y_j(t) \leq Z_j(t)$ for all t . Since the above is true for any j in V ,

$$Y_i(T) \leq \min_{j \in V} \{s_{ij}^T + Z_j(T - s_{ij}^T)\}.$$

On the other hand, there must be a user, say j' , such that the content at user i at time T was originally downloaded by j' at some time $t \leq T$ and reached i through CONTENT SHARING, so that $Y_i(T)$ is $T - t$. Then, there exists a backwards contact sequence $(b_0, t'_0), \dots, (b_m, t'_m)$ from i to j' in so that (a) $t'_m < T - t$ and (b) user j' does not download the content from the service provider in the interval $(t, T - t'_m]$. We then have that $Y_i(T) = T - t = t'_m + Z_{j'}(T - t'_m)$. Observe that, by the existence of the above contact sequence, $j' \in B_i^T(t'_m)$ and, hence, $s_{ij'}^T \leq t'_m$. On the other hand, observe that the function $s - Z_{j'}(T - s)$ is increasing in s . For this reason we have that

$$Y_i(T) \geq s_{ij'}^T + Z_{j'}(T - s_{ij'}^T) \geq \min_{j \in V} \{s_{ij}^T + Z_j(T - s_{ij}^T)\}$$

and the lemma follows.

A.2 Proof of Lemma 2

From Lemma 1, we have that

$$\mathbb{P}_{\bar{x}}(Y_i(T) > t) = \mathbb{P}_{\bar{x}}\left(\min_{j \in B_i^T(t)} \{s_{ij}^T + Z_j(T - s_{ij}^T)\} > t\right),$$

as, by definition, $s_{ij}^T > t$ for any $j \notin B_i^T(t)$. Hence,

$$\mathbb{P}_{\bar{x}}(Y_i(T) > t) = \mathbb{P}_{\bar{x}}\left(\bigcap_{j \in B_i^T(t)} (Z_j(T - s_{ij}^T) > t - s_{ij}^T)\right).$$

Assume that the process is in steady state, *i.e.*, the contact and injection processes started at $-\infty$. Recall that, as the aggregate injection process is a Poisson process with rate μ , the injection processes at each user j are Poisson processes with rates x_j . It is a fundamental property of the Poisson process that these processes are independent (see, *e.g.* [23]). As they are also independent of the contact process, the random variables $Z_j(T - s_{ij}^T)$, $j \in V$, are independent and exponentially distributed with parameters x_j , by the memoryless property of the exponential distribution. Therefore,

$$\begin{aligned} \mathbb{P}_{\bar{x}}(Y_i(T) > t) &= \mathbb{E}\left[\prod_{j \in B_i^T(t)} e^{-x_j(t - s_{ij}^T)}\right] \\ &= \mathbb{E}\left[e^{-\sum_{j \in B_i^T(t)} x_j(t - s_{ij}^T)}\right] \end{aligned}$$

and the lemma follows as $j \in B_i^T(t)$ iff $s_{ij}^T \leq t$. The above derivation can also be repeated for $\mathbb{P}_{\bar{x}}(Y_i \geq t)$ and the equality is due to the continuity of the exponential density.

A.3 Proof of Lemma 3

Suppose first that $u \geq 0$. We then have that:

$$\mathbb{E}_{\bar{x}}[u(Y_i)] = \int_0^\infty \mathbb{P}_{\bar{x}}(u(Y_i) > z) dz = \int_0^\infty \mathbb{P}_{\bar{x}}(Y_i \in u^{-1}(I_z)) dz$$

where u^{-1} the inverse mapping of h and $I_z = (z, \infty)$. Since u is non-increasing, $u^{-1}(I_z)$ is either \emptyset, \mathbb{R}_+ , an interval of the form $[0, y)$ or an interval of the form $[0, y]$. In all four cases, $\mathbb{P}_{\bar{x}}(Y_i \in u^{-1}(I_z))$ is concave. Hence $\mathbb{E}_{\bar{x}}[u_i(Y_i)]$ is concave as the integral of a family of parametrized concave functions over a positive measure. The above argument can be extended to real, non-increasing functions $u : \mathbb{R}_+ \rightarrow \mathbb{R}$ by noting that,

$$\mathbb{E}_{\bar{x}}[u(Y_i)] = \int_0^\infty \mathbb{P}_{\bar{x}}(u(Y_i) > z) dz - \int_{-\infty}^0 \mathbb{P}_{\bar{x}}(u(Y_i) \leq z) dz.$$

A.4 Proof of Lemma 5

Suppose that a message is placed in i at time T and is propagated through flooding. For some $j \geq 0$, let $K_j = \inf\{t \text{ s.t. } |A_i^T(t)| \geq j\}$, be the first time for which at least j users have the message. Then, for $1 < j \leq n$,

$$\mathbb{P}(|A_i^T(t)| < k) = \mathbb{P}(K_j > t) \leq \mathbb{P}\left(\sum_{k=1}^{j-1} \mathcal{B}_k > t\right), \quad t \geq 0 \quad (8)$$

where \mathcal{B}_k , $1 \leq k < j$ are independent, exponentially distributed random variables with parameters β_k given by $\beta_k = kh_G$, for $1 \leq k \leq n/2$, and $\beta_k = (n - k)h_G$

for $n/2 < k < n$, and h_G is the edge expansion of the contact graph G . We prove this statement by induction on j .

PROOF OF (8). Let $T_j = K_{j+1} - K_j$, for $1 \leq j < n$, be the time between two consecutive increases of $|A_i^T(t)|$. Then, by definition $K_j = \sum_{k=1}^{j-1} T_k$, $1 < j \leq n$. For $j = 2$ the statement holds by the definition of h_G . Suppose that the statement is true for $j = k$, where $1 < k \leq n$. Then

$$\begin{aligned} \mathbb{P}(K_{k+1} \geq t) &= \mathbb{P}\left(\sum_{j=1}^k T_j > t\right) = \mathbb{P}(T_k + K_k \geq t) \\ &= \int_0^\infty \mathbb{P}(T_k \geq t - s \mid K_k = s) f_{K_k}(s) ds \quad (9) \end{aligned}$$

Conditioned on $A_i^T(K_k)$, T_k does not depend on K_k , so where

$$\begin{aligned} \mathbb{P}(T_k \geq t - s \mid K_k = s) &= \sum_{A \in V, |A|=k} \mathbb{P}(T_k \geq t - s \mid A_i^T(K_k) = A) \\ &\quad \cdot \mathbb{P}(A_i^T(K_k) = A \mid K_k = s) \end{aligned}$$

On the other hand, conditioned on $A_i^T(K_k) = A \subset V$ (where, by definition of K_k , $|A| = k$), T_k is the time until a user within A contacts a user in $A^c = V \setminus A$. This is exponentially distributed with rate $\text{vol}(\partial A) = \sum_{i \in A, j \in A^c} q_{ij}$, hence

$$\mathbb{P}(T_k > t \mid A_i^T(K_k) = A) = e^{-\text{vol}(\partial A)t}.$$

We have $\text{vol}(\partial A) = \text{vol}(\partial A^c)$, as $q_{ij} = q_{ji}$ for all $i, j \in V$. Furthermore, by the definition of h_G we have

$$\text{vol}(\partial A) \geq h_G \min(|A|, |A^c|) \geq h_G \min(k, n - k).$$

We thus get that $\mathbb{P}(T_k > t \mid A_i(K_k) = A) \leq e^{-\beta_k t}$ where β_k as in (8). As β_k only depends on k , not on A , we get that $\mathbb{P}(T_k \geq t - s \mid K_k = s) \leq e^{-\beta_k(t-s)}$ for $s \leq t$. Using the above bound in (9) and applying Fubini's Theorem yields the statement. \square

By induction over k we can show that $\mathbb{P}(\sum_{j=1}^{k-1} \mathcal{B}_j > t) = 1 - (1 - e^{-h_G t})^{k-1}$, for $1 \leq k \leq n/2$, and the lemma follows.

A.5 Proof of Lemma 6.

$B_i^T(t) \subseteq V$ and therefore, by Lemma 1,

$$Y_i(T) \leq \min_{j \in B_i^T(t)} \{s_{ij}^T + Z_j(T - s_{ij}^T)\} \leq t + \min_{j \in B_i^T(t)} Z_j(T - t)$$

where the last inequality is true because $s_{ij}^T \leq t$ for $j \in B_i^T(t)$ and $s \mapsto s + Z_j(T - s)$ is an increasing function. We therefore have that, for every $t > 0$,

$$\mathbb{E}[Y_i(T)] \leq t + \mathbb{E}\left[\min_{j \in B_i^T(t)} Z_j(T - t)\right]. \quad (10)$$

Condition on $B_i^T(t) = B \subseteq V$. The r.v. $\min_{j \in B} Z_j(T - t)$ is the elapsed time at $T - t$ since a user j in B last

downloaded content from the service provider. Since each user downloads new content independently according to a Poisson process with rate μ/n , we have that

$$\mathbb{P}\left(\min_{j \in B_i^T(t)} Z_j(T - t) > \tau \mid B_i^T(t) = B\right) = e^{-\mu|B|\tau/n},$$

$$\text{hence } \mathbb{E}\left[\min_{j \in B_i^T(t)} Z_j(T - t)\right] = \mathbb{E}[n/(\mu|B_i^T(t)|)]. \quad (11)$$

In steady state, by Lemma 4, $B_i^T(t)$ is distributed as $A_i^T(t)$. From Lemma 5, the cardinality of the latter is stochastically bounded from below by a truncated geometric random variable, which implies that $\mathbb{E}[|B_i^T(t)|^{-1}]$ is upperbounded by

$$\begin{aligned} &\left\lfloor \frac{n}{2} \right\rfloor^{-1} (1 - e^{-h_G t})^{\lfloor \frac{n}{2} \rfloor - 1} + \sum_{k=1}^{\lfloor \frac{n}{2} \rfloor - 1} \frac{(1 - e^{-h_G t})^{k-1} e^{-h_G t}}{k} \\ &\leq \left\lfloor \frac{n}{2} \right\rfloor^{-1} (1 - e^{-h_G t})^{\lfloor \frac{n}{2} \rfloor - 1} + \frac{h_G t e^{-h_G t}}{1 - e^{-h_G t}} \end{aligned}$$

as $\sum_{k=1}^n \frac{x^k}{k} \leq \sum_{k=1}^\infty \frac{x^k}{k} = -\log(1 - x)$ for $0 < x < 1$. The lemma follows by replacing $\mathbb{E}[|B_i^T(t)|^{-1}]$ with the above bound in (11) and using (10) to bound $\mathbb{E}[Y_i]$.

A.6 Proof of Lemma 8

The integrability assumption, along with the fact that the utilities are decreasing, imply that the utilities are positive. Observe that

$$|(t - s_{ij})_+ e^{-\sum_{j=1}^n x_j(t - s_{ij})_+}| \leq t. \quad (12)$$

Hence, if $u_i(Y) \leq B$ for all $Y \in \mathbb{R}^+$, $i \in V$,

$$\begin{aligned} \left| \frac{\partial f}{\partial x_j} \right|^2 &= \left| \int_0^B \sum_{i=1}^n \frac{\partial \mathbb{P}_{\vec{x}}(Y_i < u_i^{-1}(t))}{\partial x_j} dt \right|^2 \\ &= \left(\sum_{i=1}^n \int_0^B \left| \frac{\partial \mathbb{P}_{\vec{x}}(Y_i < u_i^{-1}(t))}{\partial x_j} \right| dt \right)^2 \\ &\stackrel{(12)}{\leq} \left(\sum_{i=1}^n \int_0^B |u_i^{-1}(t)| dt \right)^2 \\ &\leq (nC)^2 \end{aligned}$$

for some C , as $|u_i^{-1}(t)|$ are integrable. Hence, $\|\nabla f(\vec{x})\|^2$ is bounded by $n^3 C^2$ for all \vec{x} . The same bound holds for every evaluation of the estimator as well, so it also holds for the variance in Lemma 7.

A.7 Proof of Lemma 9

Formally, we can define ℓ_{ij}^T as

$$\ell_{ij}^T = \inf\{t \mid j \in A_i^T(t)\},$$

where $A_i^T(t)$ the forwards growth process of in Section 5.2.

The conditions on the utilities guarantee that $\|\nabla f(\vec{x})\|^2$ is bounded for all \vec{x} (see proof of Lemma 8). On the

other hand

$$\begin{aligned}\mathbb{E}[|\Delta_{\bar{x},i}|^2] &\leq \mathbb{E}\left[\left(\int_{t_{start}}^{t_{end}} \left|u_j(\tilde{Y}_j(t)) - u_j(Y_j(t))\right| dt\right)^2\right] \\ &\leq 4B^2\mathbb{E}[|t_{end} - t_{start}|^2]\end{aligned}$$

where B the bound on the utilities. Let k be a user such that $x_k > \mu/n$ (there exists at least one as $\sum x_i = \mu$). Let t_* be the first time after t_{start} that k downloads new content. Then

$$\begin{aligned}\mathbb{E}[|t_{end} - t_{start}|^2] &\leq \mathbb{E}[|t_* - t_{start}|^2] + \mathbb{E}[|t_{end} - t_*|^2] \\ &\leq 2n/\mu + \mathbb{E}[|\ell_{kj}^{t_*}|^2]\end{aligned}$$

which is bounded.