Mobicom Poster: Distributing Content Updates over a Mobile Social Network

Stratis Ioannidis[a] — *stratis@cs.toronto.edu*

Augustin Chaintreau[b] — *augustin.chaintreau@thomson.net*

Laurent Massoulié[b] — *laurent.massoulie@thomson.net*

[a]CS Department, University of Toronto, Toronto, ON, Canada

[b]Thomson, Paris, France

We study the dissemination of content updates, such as news or traffic information, over a mobile social network. In this application, mobile users receive content updates from their service provider. To improve coverage and increase capacity, we assume that they also share such updates, opportunistically, whenever they meet. We show that the service provider can allocate its downlink bandwidth so that it maximizes the aggregate utility over all users, i.e.*, the social welfare. Moreover, 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.*

I. Introduction

The application considered in this paper is a dynamic-content distribution system over a mobile network. In particular, users subscribe to a content distribution service offered by their wireless service provider. For example, cell-phone users may receive updates from the cell-phone operator regarding a news-feed or a blog of their interest, the price of a stock, *etc.* Users that subscribe to this service share their downloaded content with each other in an opportunistic fashion: whenever two subscribers meet, the one whose content is most recent pushes it to the one whose content is older (thereby extending the network's coverage and increasing its capacity).

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. In particular, assume that the service provider generates new content updates according to a Poisson process with rate μ. Moreover, assume that there are n mobile users and that each update is injected by the service provider to only one user. As a result, each user receives updates according to a Poisson process with rate x_i, such that $\sum_{i=1}^{n} x_i = \mu$.

Let Y_i be the age of the content at user i. Moreover, assume that each user i receives a utility $u_i(Y_i)$ from the content, which only depends on the age Y_i. As a user older content should be less valueable, we consider only non-increasing utilities u_i.

Under these assumptions, a natural goal for the service provider is to maximize the aggregate utility over all users, *i.e.*, the social welfare:

SOCIAL WELFARE MAXIMIZATION

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

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

$$x_i \geq 0,\ 1 \leq i \leq n.$$

where $\mathbb{E}_{\vec{x}}[u_i(Y_i)]$ the expected utility at user i under the rate allocation vector $\vec{x}$.

Note that the optimal allocation will depend on the utilities u_i, as well as the process describing the opportunistic contacts between users. In turn, the latter is governed by the users' social behavior, as this behavior determines who meets whom, and when.

A second important question is how such a system scales as the number of users grows. In particular, if more subscribers are added to the system, while the downlink capacity at the server remains the same, will the age of content at users increase (and, therefore, will the service quality degrade?) and, if so, by how much? This again depends on the social behavior of users; however, it is not a priori clear what aspects of the social behavior have the most important effect on how the age at a user scales.

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 content updates over a mobile social network.

*The full version of this abstract will appear in IEEE-INFOCOM 2009.

II. Main Results

Our first main result concerns the solution of the social welfare maximization problem. We prove the following theorem, whose generality is surprising: obtaining an optimal allocation is feasible when the process describing contacts between users is stationary ergodic. The contact process exhibits this property if, for example, the marginal contact processes between distinct pairs of users i and j, where $i, j = 1, \dots, n$, are independent renewal processes. Another example is the joint contact process resulting from any kind of independent random trip user movements [1].

Theorem 1 *Assume that the process describing contacts between users is stationary ergodic. Then,* Social Welfare Maximization *is a convex optimization problem. In particular, the objective function f, given by Eq.* (1)*, is concave.*

Theorem 1 implies that the optimization problem is feasible and can be solved by gradient descent [2]. Nonetheless, computing the gradient of the objective function f in a closed form is difficult, even if the process describing contacts between users is very simple. In [7], we show that, if the service provider knows the user utilities u_i, $i = 1, \dots, n$, it can compute an unbiased estimator of the gradient of f by collecting simple statistics on contacts between users. Moreover, we also show that the gradient can be computed by the users themselves in a fully distributed fashion, thus suppressing the need for collecting contact statistics.

Our second main result addresses system scalability. Assume that the contact processes among pairs of users $i, j = 1, \dots, n$, are independent Poisson processes with rates $q_{ij} \geq 0$. In this case, our system can be represented by a complete, weighted, undirected graph $G(V, E)$ whose vertex set is $V = \{1, \dots, n\}$, and whose edge weights are q_{ij}, $i, j \in V$. Our result, summarized in the following theorem, relates the scalability of our system to the *edge expansion* h_G [4] of the graph G, defined as:

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

Theorem 2 *Assume that the processes describing the contacts between any pair of users i and j, $i, j = 1, \dots, n$, are independent Poisson processes. If the rate allocation is the uniform allocation* (i.e., $x_i = \frac{\mu}{n}$, $i = 1, \dots, n$)*, then the expected age seen by any user i satisfies*

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

where h_G the edge expansion of the graph G representing the system.

Theorem 2 indicates that our system scales very well if the edge expansion h_G of the contact graph is bounded away from zero, as the system size n grows. In such a system, even though the service provider commits only a constant amount of bandwidth (μ) to serve its users, sharing updates guarantees an age that grows slowly with n (as $\log(n)$). Graphs having the constant expansion property are called *expander* graphs, and there is a substantial body of work (see, *e.g.*, [6]) arguing that such graphs are abundant.

III. Empirical Study

III.A. Optimization

We used our algorithm for estimating the gradient of f to find the optimal rate allocation for two human mobility traces collected by other researchers. The Infocom06 data set [3] contains Bluetooth contacts between iMotes distributed to participants of the three-day Infocom '06 conference. We focused on a 10 hour period during the first day of the conference. In the MIT data set, collected by the Reality-Mining project [5], participants carried GSM enabled cell-phones over a period of 9 months. We assume 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. Moreover, due to memory constraints, we limit our analysis of the MIT dataset to an 80 day period.

For simplicity, we assumed that every user has the same utility $u_i(Y_i) = u(Y_i)$, where u is one of the functions shown in Figure 1 (for u_a, we choose the threshold value $\tau = 200$sec).

Figure 2 presents the optimal rate allocation obtained with u_a in the Infocom06 data set, for different values of μ. As shown in Figures 2 (a) and (b), for small μ, the optimal allocation tends to be biased towards the "most social" users in the system, *i.e.*, the

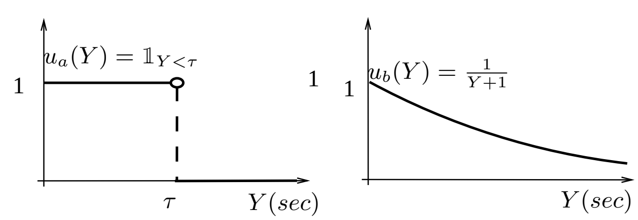

Figure 1: Examples of utilities. For the utility $u_a(Y)$, if the age exceeds the threshold τ, the content has no value to the user. If the utility is $u_b(Y)$, even very old content has some value.

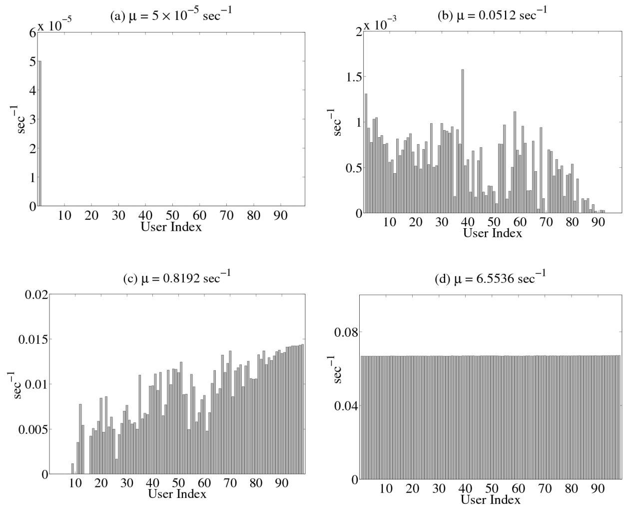

Figure 2: Four optimal rate allocations for the utility function $u_a(Y) = \mathbb{1}_{Y<200}$ under different values of μ (Infocom06). Users are indexed according to their contact rates, in decreasing order. For small μ, all updates are injected to the user with the highest contact rate. However, there is a region of μ for which the 8 most social users receive no updates from the service provider

ones with the highest aggregate contact rates. In Infocom06, for both utility functions, whenever μ is less than 6×10^{-3} updates per second, the optimal allocation concentrates on the "most central" user, from which content is disseminated to all users the fastest. In Infocom06, this node also has the highest aggregate contact rate —in other words, it is also the "most social" user. The optimal allocation also concentrates on a single "most central" user for MIT, when μ is less than 4×10^{-4} sec^{-1}, though this user does not have the highest contact rate in this data set. This indicates that the "most social" user need not also be the "most central". For μ very large, the optimal rate allocation becomes uniform. The improvement provided by content sharing becomes negligible and the system behaves as if nodes were isolated and no sharing takes place. It can be verified that, by Theorem 1, if all users are isolated, the uniform allocation is optimal.

We observe an interesting phenomenon for utility u_a when μ is between 0.2 and 0.8192 sec^{-1} in the Infocom06 data set. In this region, the rate at the 8 "most social" users becomes zero. This indicates that the behavior of the "most central" user is fully reversed: while for low values of μ this user acts as a global source of all incoming information for all users, in this region of μ it receives no updates from the service provider. A similar phenomenon was observed in the MIT dataset, for a value of μ around 0.2sec^{-1}.

In Fig. 3 and 4 we plot the ratio between the social welfare achieved under three heuristics and the optimal. We consider (a) the uniform allocation, (b) a skewed allocation, in which all the injection rate is concentrated at the "most social" user (*i.e.*, the one with highest aggregate contact rate) 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, the fraction of users with content age below the threshold for u_a is given in Table 1 and compared with the case of isolated nodes (no sharing).

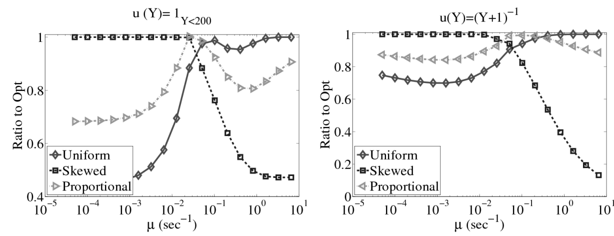

Figure 3: Ratio of social welfare of heuristic rate allocations to the optimal (Infocom06 data set).

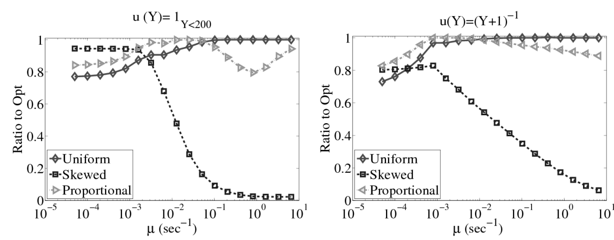

Figure 4: Ratio of social welfare of heuristic rate allocations to the optimal (MIT data set).

μ (sec^{-1})	no-sharing	skewed	prop.	unif.	optim.
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.

The comparison of these allocations confirms our earlier observations. The skewed allocation performs well for small values of μ, but not always optimally as it may not select the most central node (see Fig. 4). Uniform is always optimal for large values of μ. Proportional is the best among the three for intermediate values of μ.

In conclusion, our results highlight a transition depending on μ for the optimal rate allocation from concentrated to uniform. The topology of the users' social network plays an important role (*e.g.*, in selecting the "most central" nodes) which becomes more prominent when μ takes intermediate values.

III.B. Edge Expansion and Content Age

To illustrate the relevance of Theorem 2, we looked at how foregoing certain opportunities for sharing content would affect the system. In particular, although the actual contact rate between two users i and j may be q_{ij}, a user may choose to utilize only a fraction $q'_{ij} \leq q_{ij}$ for sharing updates. *E.g.*, users may reduce their bandwidth consumption according to the following strategies:

Proportional(p). Every contact rate is scaled by $0 < p < 1$, *i.e.*, $q'_{ij} = pq_{ij}$. This can be implemented by ignoring contact events with probability $1 - p$ and using them for content sharing with probability p.

Keep Highest(k). Every user i restricts its contacts to only k users. These k users are the ones with the highest contact rates with this node.

Rate Limit (q). For a given q, each user i with $q_i > q$ equalizes its highest contact rates (in other words, trims them), so that its aggregate rate is $q'_i = q$.

The last two strategies may yield asymmetric contacts between two users. We therefore consider two variants of them, one in which a contact between i and j is used if *both* i and j choose to use it, and one where a contact is used if *any* of the two chooses to use it. In the top row of Fig. 5 we see the average content for different strategies in Infocom06 and MIT. We see that Rate Limit is the best, given that it reduces the average bandwidth consumption per user without increasing the age considerably. We also plot in the second row of Fig. 5 an approximation of h_G^{-1} in terms of the spectral gap λ of the graph G, as a function of the average contact rate. We observe that the qualitative comparison between the strategies in terms of h_G^{-1} agrees with the one derived through the maximum age for both data sets. This is quite surprising, given that Theorem 2 merely establishes an upper bound and that our approximation overestimates h_G^{-1}. Moreover, it suggests that the maximum age is inherently linked to the edge expansion of the graph G.

IV. Conclusions

Our results show that content updates can be distributed over a mobile social network in a scalable way. Most importantly, our work shows that the social network can be successfully exploited to improve the system's performance —in particular, by maximizing the social welfare. Investigating different applications, such as cases in which the content is directly generated by users, is interesting future work.

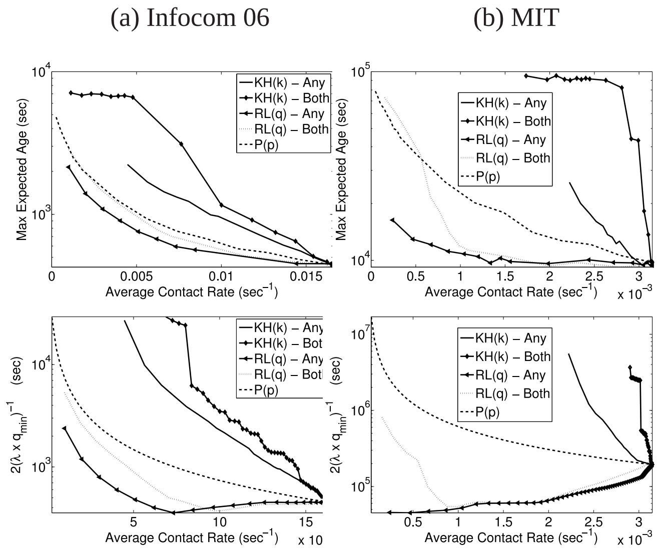

Figure 5: Maximum age and inverse of expansion for different rate reduction strategies. The expansion h_G is approximated by $0.5\lambda \times q_{\min}$, where λ the spectral gap of the graph and $q_{\min}$ the minimum contact rate of all users. The qualitative comparison between different strategies for reducing bandwidth based on h_G aggrees with the one derived through the maximum age.

References

[1] J.-Y. Le Boudec and M. Vojnovic. The random trip model: Stability, stationary regime, and perfect simulation. *IEEE/ACM Trans. Netw.*, 14(6), 2006.

[2] Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.

[3] A. Chaintreau, A. Mtibaa, L. Massoulié, and C. Diot. The diameter of opportunistic mobile networks. In *Proc. of ACM CoNEXT*, 2007.

[4] Fan Rong K. Chung. *Spectral Graph Theory*. American Mathematical Society, 1997.

[5] N. Eagle and A. Pentland. Reality mining: Sensing complex social systems. *Journal of Personal and Ubiquitous Computing*, 10(4), 2005.

[6] Shlomo Hoory, Nathan Linial, and Avi Wigderson. Expander graphs and their applications. *Bulletin of the AMS*, 43(4), 2006.

[7] Stratis Ioannidis, Augustin Chaintreau, and Laurent Massoulié. Optimal and scalable distribution of content updates over a mobile social network. In *IEEE INFOCOM*, 2009.