



8-12-2009

Quantifying Content Consistency Improvements Through Opportunistic Contacts

Kin-Wah Kwong

University of Pennsylvania, kkw@seas.upenn.edu


Augustin Chaintreau

Thomson

Roch A. Guérin

University of Pennsylvania, guerin@acm.org

Follow this and additional works at: http://repository.upenn.edu/ease_papers

 Part of the [OS and Networks Commons](#), and the [Systems and Communications Commons](#)

Recommended Citation

Kin-Wah Kwong, Augustin Chaintreau, and Roch A. Guérin, "Quantifying Content Consistency Improvements Through Opportunistic Contacts", . August 2009.

Proceedings ACM CHANTS 2009 Workshop Beijing, China, September 2009

This paper is posted at ScholarlyCommons. http://repository.upenn.edu/ease_papers/510

For more information, please contact libraryrepository@pobox.upenn.edu.

Quantifying Content Consistency Improvements Through Opportunistic Contacts

Abstract

The sharing and dissemination of online content is one of the main purposes of social network applications, and the amount of content accessed through them, in particular through portable devices such as smartphones and PDAs, is expected to increase. Consumption of online content, however, does not require a continuous online presence. Content can be downloaded, consumed, modified, and uploaded at different times. An opportunity to improve a user's access to up-to-date information from its own social network is to take advantages of opportunistic contacts between mobile devices, \ie without waiting for connectivity to the network infrastructure. In other words, users of a social network application may receive more fresh content with no extra infrastructure deployment, simply by communicating with mobile devices of other users, in a delay-tolerant manner. Assessing the magnitude of this improvement is, however, challenging. For example, the frequency and patterns of such contacts are partly a function of the social connectivity of users, and so will be the availability of relevant information to share and more importantly the willingness to share that information. All these influence in non-trivial ways the gains that can be realized through opportunistic contacts. The paper's main contribution is in providing a quantitative handle through which these gains can be estimated, while accounting for the above factors.

Keywords

social networks, delay-tolerant, optimization, dissemination

Disciplines

OS and Networks | Systems and Communications

Comments

Proceedings ACM CHANTS 2009 Workshop Beijing, China, September 2009

Quantifying Content Consistency Improvements Through Opportunistic Contacts*

Kin-Wah Kwong
University of Pennsylvania
Philadelphia, U.S.A.
kkw@seas.upenn.edu

Augustin Chaintreau
Thomson
Paris, France
augustin.chaintreau@thomson.net

Roch Guérin
University of Pennsylvania
Philadelphia, U.S.A.
guerin@ee.upenn.edu

ABSTRACT

Contacts between mobile users provide opportunities for data updates that supplement infrastructure-based mechanisms. While the benefits of such opportunistic sharing are intuitive, *quantifying* the capacity increase they give rise to is challenging because both contact rates and contact graphs depend on the structure of the social networks users belong to. Furthermore, social connectivity influences not only users' interests, *i.e.*, the content they own, but also their willingness to share data with others. All these factors can have a significant effect on the capacity gains achievable through opportunistic contacts. This paper's main contribution is in developing a tractable model for estimating such gains in a content update system, where content originates from a server along multiple *channels*, with blocks of information in each channel updated at a certain rate, and users differ in their contact graphs, interests, and willingness to share content, *e.g.*, only to the members of their own social networks. We establish that the added capacity available to improve content consistency through opportunistic sharing can be obtained by solving a convex optimization problem. The resulting optimal policy is evaluated using traces reflecting contact graphs in different social settings and compared to heuristic policies. The evaluation demonstrates the capacity gains achievable through opportunistic sharing, and the impact on those gains of the structure of the underlying social network.

Categories and Subject Descriptors

C.2.1 [Network Architecture and Design]: Wireless communication; C.4 [Performance of Systems]: Modeling techniques

*The work of K.-W. Kwong was performed during an internship at Thomson Research, and its completion was subsequently supported by NSF grant CNS-0721610. The work of R. Guérin was supported in part by NSF grant CNS-0721610.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CHANTS'09, September 25, 2009, Beijing, China.

Copyright 2009 ACM 978-1-60558-741-7/09/09 ...\$10.00.

General Terms

Algorithms, Performance, Theory

Keywords

Social networks, Delay-tolerant networks, Consistency, Dynamic content, Dissemination, Optimization

1. INTRODUCTION

The sharing and dissemination of online content is one of the main purposes of social network applications, and the amount of content accessed through them, in particular through portable devices such as smartphones and PDAs, is expected to increase. Consumption of online content, however, does not require a continuous online presence. Content can be downloaded, consumed, modified, and uploaded at different times. An opportunity to improve a user's access to up-to-date information from its own social network is to take advantages of opportunistic contacts between mobile devices, *i.e.*, without waiting for connectivity to the network infrastructure. In other words, users of a social network application may receive more fresh content with no extra infrastructure deployment, simply by communicating with mobile devices of other users, in a delay-tolerant manner. Assessing the magnitude of this improvement is, however, challenging. For example, the frequency and patterns of such contacts are partly a function of the social connectivity of users, and so will be the availability of relevant information to share and more importantly the willingness to share that information. All these influence in non-trivial ways the gains that can be realized through opportunistic contacts. The paper's main contribution is in providing a quantitative handle through which these gains can be estimated, while accounting for the above factors.

There are many possible metrics for quantifying the performance improvement of content delivery achievable through opportunistic contacts. One possible approach is to consider a *generic* network metric that would be relevant to all applications: one may hence focus on the additional flow capacity provided by intermittent links [7], or consider the time needed to exchange data between arbitrary pairs of users [4]. Unfortunately these metrics can be both complex to define and difficult to interpret, as their values vary greatly depending on the pairs of users, the flows or the network load considered. A different approach, which is used in the paper, is to define a *specific* application metric that describes directly the performance of the network to support the requirement of a given application. The improvement gath-

ered by delay-tolerant communication can then directly be interpreted in terms closer to users’ experience, as in a field test. When it is possible to compute such an application metric, one can study under which conditions opportunistic contacts significantly improve services provided by the network.

Here, in contrast to other work, we focus on content consistency as this content is updated over time. We show that it is possible to accurately measure the benefits of opportunistic contacts according to this application-specified metric. Content originates at a server and is structured into different *channels*. Channels can be thought of as information sources of interest to some members of a social network. Channel information is updated according to a stationary renewal process, and the server seeks to keep users in sync with the latest content of each channel, but does so under some capacity limitations. The value of channel information to users is a function of its relative age. Users can select which channel content they are willing to store as well as which users they are willing to share it with during opportunistic contacts¹. The benefits derived from opportunistic contacts is measured through a notion of “capacity”, which measures the number of users with access to recent information on channels they are interested in (*e.g.*, subscribe to).

The paper makes the following contributions:

- It develops a model for quantifying the added capacity available for content updates through opportunistic contacts in mobile networks. The model incorporates the effect of users’ social connectivity and social behavior in sharing content during those opportunistic contacts. These are shown to have a significant impact on the potential gains achievable from opportunistic contacts.
- It demonstrates how a capacity-achieving policy can be explicitly constructed by solving a convex optimization problem, and illustrates how this optimal operating point can be realized using basic information on users’ contacts and interests.
- Using actual mobility traces, the capacity benefits of opportunistic contacts to a content update application operating on cell phones carried by humans are characterized. The experimental results further demonstrate a significant effect that users’ social behavior, *e.g.*, differences in willingness to share with other users, can have on overall performance.

The rest of the paper is organized as follows. The next section gives a brief literature review. Section 3 introduces the specification and model of our content update system. It also presents our notion of capacity and establishes that it can be obtained by solving a convex optimization problem. The content update capacity of some real-life opportunistic mobile networks is explored in Section 4, and the optimal policy is compared to several heuristics. Section 5 concludes the paper.

¹Both storage and transmission of content to other users during contacts have costs, *e.g.*, memory and battery life time. It is, therefore, important to quantify the resulting gains to users, if only to motivate such opportunistic sharing.

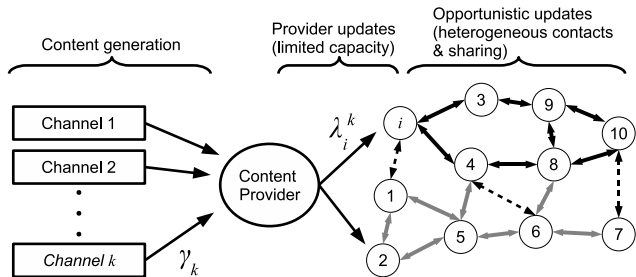


Figure 1: Content update system

2. RELATED WORK

Taking advantage of contacts created by mobility for delay-tolerant applications have recently received much attention, typically for routing (*e.g.*, [10, 6, 2]), and more recently for content dissemination (*e.g.*, [14, 5, 3, 12, 9]) and content updates (*e.g.*, [1, 8, 13]). The use of social behavior to improve the delivery of information, as in viral marketing [11], has been suggested several times [6, 14, 5, 8]. Most of the previous work relies on a utility-based criterion to optimize the dissemination of information, and usually estimates the resulting benefits using generic network metrics such as propagation delay. In particular, [8] showed that content update systems could scale better by leveraging peer-to-peer (opportunistic) sharing of information, and it developed optimal dissemination strategies for different settings [8, 1]. Our work shares this focus on content update systems, but acknowledges the impact that individual connectivity patterns and sharing behavior can have, and explicitly incorporates their effect in devising content update solutions. Furthermore, in contrast to the previous work, we seek to precisely characterize the benefits (here, content consistency) afforded to applications by different opportunistic sharing schemes, including the optimal solution. This calls for a model that accounts for the creation of content, its consumption by different users, as well as users’ behavior when they are asked to opportunistically share content.

3. CONTENT UPDATE SYSTEM

Content is structured into a finite set of channels \mathcal{K} , with a channel identified by its index k . The content of channel $k \in \mathcal{K}$ consists of a sequence of blocks created by a source and updated over time. The set of mobile hosts² is denoted by \mathcal{V} . The structure of the overall content update system is shown in Fig. 1 where the arrows among the mobile hosts denote their sharing behavior, *e.g.*, sharing only between users within their communities or with similar interests.

In order to clarify where the paper differs from the previous work, we first review the propagation of a single block update through the whole network as captured by the previous models, *e.g.*, [8]. Next, we extend significantly this model to define a notion of network capacity when content

²In this paper, the terms “node”, “user” and “mobile host” are synonymous.

spans multiple blocks and is updated according to different statistics, and users follow different behavior with regard to content subscription and sharing.

3.1 Propagation of a single update block

The social network created among users by their opportunistic contacts can be used to reduce the age of dynamic information that all of them are interested to maintain on their devices. We assume that all the data from a block are over-written whenever a newer block of the same channel is received. Consequently nodes maintain only the latest block received on each channel. We assume that blocks are atomic in terms of both content and transmission, *i.e.*, they have a fixed size denoted by b_k (bits) which may depend on the channel. Furthermore, whole blocks can be exchanged during connections with either the infrastructure or other nodes. A more general update model would allow content fragmentation. This adds significant complexity (partial updates need to be considered and tracked) that is beyond the scope of this work.

Content provider.

Let us assume that an updated block is available at the content provider. The provider has a total capacity C (bits per second) available for all channel updates. This constrains its ability to provide all the users with fresh content, in particular when a channel is popular.

The provider has complete freedom in choosing which nodes and blocks to update. We rely on a simple randomized strategy to capture this flexibility. At each time slot (of duration δ seconds), the provider attempts to send the latest block of channel k to node i with probability $\delta \cdot \lambda_i^k$. Following a usual assumption, δ is assumed small enough that the times at which these events occur approach a continuous time Poisson process with rate λ_i^k . The capacity limitation of the provider imposes that λ_i^k 's satisfy

$$\sum_{i \in \mathcal{V}, k \in \mathcal{K}} \lambda_i^k \cdot b_k \leq C$$

where we recall that b_k is the block size of channel k .

Mobile hosts.

In addition to receiving updates from the provider, mobile hosts can also receive updates through opportunistic contacts. We assume that nodes are able to exchange all the blocks for which one of them has a more recent version. However, while those contacts are not bandwidth limited in our model, other factors impact the update capacity that can be realized through them. For example, unlike provider-mobile links, links between mobile hosts are only intermittently available, and their availability is a function of mobility patterns that are unpredictable.

We assume that the process describing opportunistic contacts between all pairs of users (which can be described as a marked point process with marks in $\mathcal{V} \times \mathcal{V}$) is stationary and ergodic. As this process is driven only by users' mobility, it is assumed independent of the block update process at the provider. Contacts between different pairs of nodes are, however, *not* assumed independent.

In order to define capacity and optimal application performance, we assume that the distribution of this opportunistic contact process is known. It is possible to find optimal application performance from an adaptive algorithm even when

this distribution is unknown (see conclusion) but this is beyond the scope of this paper.

3.2 Content generation and sharing

For defining the capacity of an opportunistic network to update content, several important dimensions need to be considered such as content update generation, users' interests and their willingness to share. Those dimensions are discussed in the following.

Content update generation.

The previous work, *e.g.*, [8], assumes that content is continuously updated, so that each new transmission from the content provider is a new and fresher block for this particular channel. The importance of this content (or utility measuring a user's satisfaction) is then assumed to be a function of the propagation delay only. In practice, however, updates are likely to occur at finite intervals of time that vary across channels. Moreover, the satisfaction of a user depends on the *relative* age of the content: a block that is an hour old may be very relevant if content for this channel is updated everyday, but less so if new content is created every minute.

To deal with this issue, we differ from the analysis of [8] in two ways: First, we assume that blocks of a single channel are updated according to a stationary renewal process. We assume that this occurs independently of opportunistic contacts and transmission from content provider. However, we *do not* assume that the processes of updates between different channels are independent. This already allows various scenarios. For example, a deterministic, periodical update process, a Poisson update process, or even update processes with heavy tailed statistics where most updates occur in bursts and there can exist a long period without updates. Second, we assume that the satisfaction of users is not a function of the propagation delay of their content, but rather a function of the consistency of their content with respect to recent updates (see Section 3.3).

Content interest and sharing.

The previous work, *e.g.*, [8], only focused on a single block that all nodes wish to receive and are willing to share with others. In practice, popularity of different channels varies greatly and power, memory, as well as trust limitations affect users' willingness to arbitrarily exchange blocks. Update exchanges typically take place only between users who trust each other. It is important to incorporate such factors when assessing the capacity available from opportunistic contacts. We introduce the $N \times K$ *interest* matrix A as:

$$A_{i,k} = \begin{cases} 1 & \text{if node } i \text{ is interested in channel } k, \\ 0 & \text{otherwise.} \end{cases}$$

The $N \times N$ *sharing* matrix B^k for channel k is:

$$B_{i,j}^k = \begin{cases} 1 & \text{if, whenever a contact } (i,j) \text{ occurs,} \\ & \text{blocks from channel } k \text{ can be sent} \\ & \text{from } i \text{ to } j, \\ 0 & \text{otherwise.} \end{cases}$$

Note that $B_{i,j}^k = 1$ if and only if i and j are interested in channel k (*i.e.*, $A_{i,k} = A_{j,k} = 1$), i agrees to transmit a block to j during an opportunistic contact, and j agrees to receive a block from i .

This allows us to consider arbitrary content sharing patterns, *e.g.*, a network where all nodes agree to store and

forward all blocks, a network with a subset of selfish nodes who may store and not forward blocks to others, a network with cautious nodes who only receive blocks from a subset of nodes they trust.

3.3 Definition of capacity

The goal of this paper is to offer an effective estimate of the capacity improvements that opportunistic contacts can offer, as well as how to realize those improvements. The first step towards realizing this goal is to introduce a precise definition of capacity for such a content update system.

Consistency-based utility.

We do so by way of a *utility function* that expresses users' satisfaction. For simplicity, we assume that for every channel $k \in \mathcal{K}$ the utility of the content associated with channel k at node i is given by the following binary variable:

$$U_i^k(t) = \begin{cases} 1 & \text{if node } i \text{ is interested in channel } k, \\ & \text{and it has the latest block at time } t, \\ 0 & \text{otherwise.} \end{cases}$$

$U_i^k(t)$ measures the satisfaction of a user in the sense that how much time the user has the latest content for channel k . There are many possible extensions to this basic utility function, *e.g.*, utility could be a non-increasing function of the block versions (score 1 if the node has the latest block, score 1/2 if it has the second latest, etc.). These different extensions can still be handled using the convex optimization formulation developed in the paper, but it becomes much harder to extract insight from their behavior. For this reason, we concentrate on the above simple binary consistency function for the rest of the paper.

Capacity region.

We can now formulate an intuitive definition of the capacity region of a content update system. Let $u_i^k = \mathbb{E}[U_i^k(t)] \in [0, 1]$ which measures the fraction of time node i has the latest block created for channel k (or, equivalently, the probability that this occurs in steady state). One can say that a vector $(f_k)_{k \in \mathcal{K}}$ in $[0, 1]^{\mathcal{K}}$ is inside the "capacity region" of the system if there exist parameters $(\lambda_i^k)_{i \in \mathcal{V}, k \in \mathcal{K}}$ such that for any i and k , the expected fraction of interested nodes that hold the latest blocks on channel k is at least f_k . Of course, a difficult part is to figure out how λ_i^k can be chosen to test this assumption. The measure u_i^k corresponds to the utility of the system from the user's point of view (for channel k).

The boundary of this network capacity region can be generally expressed by considering any non-decreasing concave function $\phi: \mathbb{R}^{\mathcal{V} \times \mathcal{K}} \rightarrow \mathbb{R}$ and then solving the following optimization problem **CAP**:

$$\begin{aligned} & \underset{\lambda_i^k, i \in \mathcal{V}, k \in \mathcal{K}}{\text{maximize}} && \phi \left(\left(u_i^k \right)_{i \in \mathcal{V}, k \in \mathcal{K}} \right), \\ & \text{subject to} && \sum_{i \in \mathcal{V}, k \in \mathcal{K}} \lambda_i^k \cdot b_k \leq C \end{aligned}$$

The function ϕ of all users' utilities denotes the utility of the system not from a user's point of view, but as a whole.

THEOREM 3.1. *For ϕ non-decreasing and concave,*

$$(\lambda_i^k)_{i \in \mathcal{V}, k \in \mathcal{K}} \mapsto \phi \left(\left(u_i^k \right)_{i \in \mathcal{V}, k \in \mathcal{K}} \right)$$

is a concave function.

Theorem 3.1 establishes that the optimization problem **CAP** can be efficiently solved due to its structural properties. This allows us to compute the optimal performance of the network in terms of maximizing ϕ , and hence characterize the corresponding capacity region. A consequence of this theorem is that any locally optimum choice of $(\lambda_i^k)_{i \in \mathcal{V}, k \in \mathcal{K}}$ is a global optimum, which can hence be found through a simple gradient search. Note that u_i^k 's are the functions of λ_i^k 's, the mobility patterns of the users, the content generation processes of the channels, as well as the interest and sharing behavior of users (captured through matrices A and $(B^k)_{k \in \mathcal{K}}$). Because the users' mobility patterns and the content generation processes of the channels are independent of the provider transmission process, they do not affect the concavity result of Theorem 3.1. The proof of Theorem 3.1 can be found in Appendix A.

Fairness.

The above definition of ϕ is very general and can be used to incorporate fairness. For any $\alpha \geq 0$, let the α -fairness function h_α be

$$h_\alpha(x) = \begin{cases} \frac{x^{1-\alpha}}{1-\alpha}, & \text{if } \alpha \neq 1 \\ \ln(x), & \text{if } \alpha = 1. \end{cases}$$

For channel k , let

$$f_k = \frac{1}{\sum_{i \in \mathcal{V}} A_{i,k}} \sum_{i \in \mathcal{V}} u_i^k$$

denote the average fraction of nodes for channel k with the latest block. We define the per-block α -fairness as

$$\sum_{k \in \mathcal{K}} h_\alpha(f_k). \quad (2)$$

Note that when $\alpha = 1$ maximizing this function corresponds to ensuring proportional fairness: at this equilibrium point, varying parameters to improve the fraction of nodes for one channel will necessarily result in a proportional decrease in same proportion of the fraction of nodes in another channel.

4. EVALUATION WITH ACTUAL TRACES

This section evaluates the capacity benefits of opportunistic contacts in settings that exhibit different contact statistics and user's sharing behavior.

4.1 Simulation setting

4.1.1 Data sets

We use mobility traces³ from two different environments, *Infocom05* and *MIT*. The *Infocom05* data set logs Bluetooth contacts between 41 devices carried by participants of the INFOCOM'05 conference. The *MIT* data set was built using GSM cell-tower associations of 100 cell phones carried by students and faculty during a semester. For the latter data set, we assume that two phones are in contact whenever connected to the same GSM base station, and we remove isolated nodes (87 cell phones were included). To speed-up the evaluation, a simplification was introduced that preserves the heterogeneity in contact rates present in the traces. First, we extracted a contact graph from each

³The data sets are available at <http://www.crowdad.org>.

trace that explicitly identified the average contact rates of each individual pair of nodes. These rates were then used to generate independent memoryless contact processes for all pairs of nodes. This is an approximation of the real traces as it removes dependencies between contact processes of different pairs of nodes, as well as their detailed statistics. However, it still captures heterogeneity in contact rates between different pairs of nodes. The computational procedure for solving the optimization problem **CAP** can be found in Appendix B, and the computation relies on 50 samples of contact latencies between nodes⁴. Matlab 7.6 was used to solve the optimization problem **CAP**. In future work, we plan to carry out investigation using actual contact traces.

4.1.2 Provider update policies

We assume that blocks have all equal size and hence the provider update capacity C is simply given in updates per minute. Moreover, content updates are generated for each channel k according to a Poisson process with rate γ_k . The following policies for allocating $(\lambda_i^k)_{k \in \mathcal{K}, i \in \mathcal{V}}$ are studied:

- **Uniform no sharing:** The provider update capacity is uniformly shared across blocks, *i.e.*,

$$\lambda_i^k = A_{i,k} \frac{C}{\sum_{i \in \mathcal{V}, k \in \mathcal{K}} A_{i,k}}.$$

We further assume that there is no sharing between nodes, so that updates are only from the server.

- **Uniform:** This is the same as **Uniform no sharing**, except that sharing between nodes is now enabled during opportunistic contacts, and follows a sharing behavior specified through $(B_{i,j}^k)_{k \in \mathcal{K}, i, j \in \mathcal{V}}$.
- **Optimal no sharing:** The server allocates capacity across channels and nodes in a manner that is optimal given its knowledge of nodes' interests and channel block generation rates $(\gamma_k)_{k \in \mathcal{K}}$. However, sharing between nodes is not allowed and not taken into account in the server capacity allocation.
- **Optimal oblivious:** This is the same as **Optimal no sharing** but with sharing now enabled. The important aspect is that the server's update policy remains oblivious to the presence and structure of opportunistic contacts.
- **Optimal:** The provider chooses the optimal capacity allocation $(\lambda_i^k)_{i \in \mathcal{V}, k \in \mathcal{K}}$ by solving the convex optimization problem **CAP**.

Among these policies, only **Optimal** requires to know or estimate the contact patterns between the nodes and their sharing behavior. Note that another possible policy is to consider the server capacity allocation obtained from **CAP** but the sharing among nodes is actually not allowed in reality. However, we decide not to consider it as, by definition, **Optimal no sharing** outperforms this policy.

The main purpose for comparing these different policies is to develop a better understanding of how social factors, *e.g.*, heterogeneous contact rates and users' sharing behavior, affect the network's ability to keep content up-to-date

⁴Samples of contact latencies are denoted by $\hat{s}_{i,j}^{k,(l)}$'s in Appendix B.

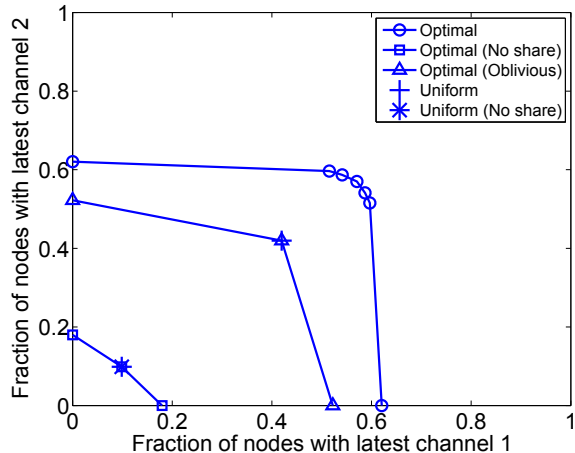


Figure 2: MIT: Capacity regions. Fraction of up-to-date nodes for channel 1 (x -axis) and channel 2 (y -axis).

across users. As expected, **Optimal** always outperforms all other policies because it is cognizant of both user and channel characteristics, and the effect of social factors on opportunistic updates. **Optimal oblivious** that takes differences in block generation rates into account generally outperforms **Uniform**, but not always. This is because it ignores the possibility of sharing among nodes, which can occasionally result in an inefficient allocation decision.

4.2 Evaluating capacity

We start by comparing the capacity *regions* realized by different policies.

A combination (f_1, f_2) is deemed inside the “capacity region” if we can find parameters $(\lambda_i^k)_{i \in \mathcal{V}, k \in \mathcal{K}}$ such that for any channel k , the expected fraction of nodes with the latest blocks on this channel is at least f_k . In order to explore the space of feasible combinations, we introduce a simple weighted average objective function:

$$\text{maximize } w \cdot f_1 + (1 - w) \cdot f_2 \text{ for } w \in [0, 1].$$

$$\lambda_i^k, i \in \mathcal{V}, k \in \mathcal{K}$$

By varying the value of w , a different objective is defined and the maximum value attainable for f_2 can be obtained as a function of f_1 , from which the boundary of the capacity region can be found. Note that, for a policy that does not incorporate an objective (*i.e.*, **Uniform no sharing**, and **Uniform**), the value of w has no effect, hence the capacity region is essentially “rectangular”, based on a single data point given by the specific (f_1, f_2) combination that this policy achieves.

We use the *MIT* data set to illustrate the capacity regions of the different policies in Fig. 2. For simplicity, the two channels were chosen to have the same block generation rate $\gamma_1 = \gamma_2 = 1/360\text{min}$, and all nodes in the system were interested in receiving both, *i.e.*, the interest matrix A is full. The sharing matrix B was, however, chosen with 20% of non-zero entries, *i.e.*, only 20% of the nodes are willing to share, and they were the 20% of nodes with the highest contact rates. The provider capacity is $C = 0.05$ update per minute. Fig. 2 shows the performance seen by channels

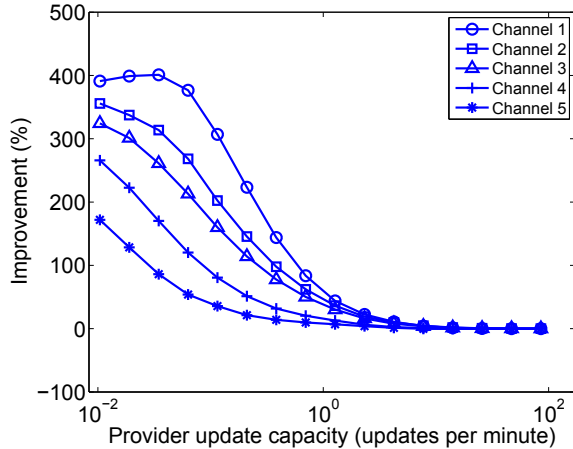


Figure 3: *Infocom05*: Capacity improvement of **Optimal** over **Optimal oblivious**.

1 and 2 in terms of the average fraction of nodes with the most up-to-date block for the channel, *i.e.*, f_1 and f_2 . Each point corresponds to a different value of the weighing factor $w = \{0, 0.2, 0.4, 0.5, 0.6, 0.8, 1\}$ in the objective function. Note that when $w = 0.5$, **Optimal no sharing** overlaps with **Uniform no sharing**, and **Optimal oblivious** with **Uniform**. In **Optimal no sharing** and **Optimal oblivious**, the results are overlapped under some values of w and hence the number of data points in a corresponding policy shown in the figure is fewer.

We see that opportunistic sharing alone yields a significant improvement in the fraction of nodes with consistent, *i.e.*, up-to-date, versions of both blocks, and that **Optimal** yields an additional improvement of about a third. This demonstrates the benefits by incorporating specific knowledge about opportunistic contacts and users' behavior in allocating the provider capacity.

4.3 Heterogeneous channels

This section investigates the benefits of the optimal server policy when channels are heterogeneous in their update rates. We consider a scenario with 5 channels using the *Infocom05* contact traces. All nodes are interested in receiving all channels, but only 20% of nodes accept to exchanges blocks during opportunistic contacts. We assume that the 20% of nodes with the highest contact rates accept to transmit blocks from channel 1 to other nodes. For channel 2, we choose the 20% of nodes with the next highest contact rates, and so on for the remaining channels so that each channel is assigned to a distinct group of nodes with decreasing contact rates. This creates significant heterogeneity in how channels with identical popularity are able to benefit from opportunistic contacts. We assume that block generation rates across channels are:

$$\begin{aligned}
 & (\gamma_1, \gamma_2, \gamma_3, \gamma_4, \gamma_5) \\
 & = \left(\frac{1}{20\text{min}}, \frac{1}{40\text{min}}, \frac{1}{60\text{min}}, \frac{1}{120\text{min}}, \frac{1}{240\text{min}} \right).
 \end{aligned}$$

Specifically, the generation rate for channel 1 is the largest so that it is hardest to maintain the latest copy for this channel, but at the same time it is associated with the sharing nodes

with the highest contacts rates. In order to allow trading efficiency for fairness across blocks in the network, we chose to maximize the block-based objective function with $\alpha = 1$ (Eq. (2)) that corresponds to the per-block proportional fairness.

Fig. 3 plots the relative improvement ratio of f_k , $k = 1, 2, \dots, 5$, under **Optimal** over **Optimal oblivious** across channels as a function of the provider update capacity. We observe that the improvement achieved by **Optimal** is substantial for all channels, and sometimes multiplies the fraction of nodes that have the latest content by a factor of up to four. This is because **Optimal** exploits the users' sharing behavior and opportunistic contacts to efficiently disseminate up-to-date content. On the other hand, **Optimal oblivious** ignores these factors and simply relies on relative differences in content generation rates when making transmission choices. This often results in shortsighted choices that prevent it from leveraging opportunistic transmissions.

5. CONCLUSION

Recent work has advocated using opportunistic contacts between mobile nodes to improve users access to content. Although promising, this has left open a difficult issue: Accurately measuring the improvements this affords applications as a function of user connectivity patterns and social behavior (*e.g.*, interest in content and willingness to share it). This paper makes a novel and important contribution to this problem for a time-sensitive content update application, for which it provides a complete characterization of the capacity available through opportunistic updates. Realizing this capacity is shown to critically depend on how the content provider allocates its own updates to nodes. In particular, this allocation depends on content generation rates, node contact rates, as well as nodes' interests and sharing behavior. Surprisingly, despite these complex dependencies, it is actually possible to compute an exact optimal policy that realizes capacity. Our results further establish that in the presence of heterogeneous contact rates and sharing behavior among nodes, simple heuristics that are oblivious to those parameters can translate into subpar performance.

Our result points to several important research challenges that remain to be addressed:

- Although our model deals with general statistics of content creation and contacts between nodes, we have studied quantitative performance improvement in a simple case: content and meeting generated according to memoryless statistics. It would be important to understand analytically as well as empirically how other statistics impact the improvement provided by opportunistic contacts.
- Computing optimal server allocation policy requires knowledge about distribution of latencies between users' contact times in the network. However, since this problem follows a convex optimization, it is possible to come up with a distributed scheme which targets the optimal allocation with limited information about processes of contacts (see [8] for an example).
- We assume that the content provider follows a simple randomized memoryless strategy for updates (essentially choosing to update blocks in nodes independently with different probabilities). A more complex

strategy may take into account time of content creation and the current age of content in nodes, it would be important to understand how that impacts the role of opportunistic contacts.

- Similarly, we assume that nodes decide to always share with a subset of other nodes, as captured in the sharing matrix. A more complex strategy may be developed to reflect other criteria such as minimizing energy consumption in opportunistic sharing while providing users with sufficiently fresh content.

We hope that the concrete evidence of these quantitative benefits presented in the paper will foster other contributions aimed at improving application performance in opportunistic mobile networks.

6. REFERENCES

- [1] E. Altman, P. Nain, and J. Bermond. Distributed storage management of evolving files in delay tolerant ad hoc networks. In *Proc. IEEE INFOCOM*, 2009.
- [2] A. Balasubramanian, B. Levine, and A. Venkataramani. DTN routing as a resource allocation problem. In *Proc. ACM SIGCOMM*, 2007.
- [3] N. Banerjee, M. D. Corner, D. Towsley, and B. N. Levine. Relays, base stations, and meshes: enhancing mobile networks with infrastructure. In *Proc. ACM MobiCom*, 2008.
- [4] A. Chaintreau, A. Mtibaa, L. Massoulie, and C. Diot. The diameter of opportunistic mobile networks. In *Proc. ACM CoNEXT*, 2007.
- [5] P. Costa, C. Mascolo, M. Musolesi, and G. P. Picco. Socially-aware routing for publish-subscribe in delay-tolerant mobile ad hoc networks. *IEEE JSAC*, 26(5), 2008.
- [6] E. M. Daly and M. Haahr. Social network analysis for routing in disconnected delay-tolerant manets. In *Proc. ACM MobiHoc*, 2007.
- [7] D. Hay and P. Giaccone. Optimal routing and scheduling for deterministic delay tolerant networks. In *Proc. IEEE WONS*, 2009.
- [8] S. Ioannidis, A. Chaintreau, and L. Massoulie. Optimal and scalable distribution of content updates over a mobile social network. In *Proc. IEEE INFOCOM*, 2009.
- [9] E. Jaho and I. Stavarakakis. Joint interest- and locality-aware content dissemination in social networks. In *Proc. IEEE WONS*, 2009.
- [10] S. Jain, K. Fall, and R. Patra. Routing in a delay tolerant network. In *Proc. ACM SIGCOMM*, 2004.
- [11] D. Kempe, J. Kleinberg, and E. Tardos. Maximizing the spread of influence in a social network. In *Proc. KDD*, 2003.
- [12] J. Leguay, A. Lindgren, J. Scott, T. Friedman, and J. Crowcroft. Opportunistic content distribution in an urban setting. In *Proc. ACM SIGCOMM workshop on Challenged networks*, 2006.
- [13] V. Lenders, G. Karlsson, and M. May. Wireless ad hoc podcasting. In *Proc. IEEE SECON*, 2007.
- [14] E. Yoneki, P. Hui, S. Chan, and J. Crowcroft. A socio-aware overlay for publish/subscribe communication in delay tolerant networks. In *Proc. ACM/IEEE MSWiM*, 2007.

APPENDIX

A. PROOF OF THEOREM 3.1

We first introduce the concept of block propagation delay $Y_i^k(t)$ which measures the latency for node i to receive the latest block of channel k . There are two ways that node i can receive the latest block for channel k . One way is to directly receive it from the provider, and the another way is through opportunistic sharing with other users. It can be shown from [8] (Lemma 2) that for any $y \geq 0$

$$\mathbb{P} \left[Y_i^k > y \right] = \mathbb{E} \left[e^{-\sum_{j \in \mathcal{V}} \lambda_j^k \cdot (y - s_{i,j}^k)_+} \right] \quad (3)$$

where $(\cdot)_+$ denotes $\max(\cdot, 0)$, $s_{i,j}^k$ is defined as the minimum value s such that a message created at time $t - s$ in j can reach i before time t , using any opportunistic contacts (i', j') such that $B_{i',j'}^k = 1$, and the expectation on the right-hand side is taken over all the values of $s_{i,j}^k$. Note that all $s_{i,j}^k$'s do not depend on $(\lambda_i^k)_{i \in \mathcal{V}, k \in \mathcal{K}}$, or the content generation process.

Let us consider how to compute $\mathbb{E} [U_i^k(t)]$ for node i and channel k . If node i is not interested in channel k (i.e., $A_{i,k} = 0$), then this expected value is null. Otherwise, assuming $A_{i,k} = 1$, we have

$$U_i^k(t) = \mathbb{I} \left\{ Y_i^k(t) \leq \Gamma^k(t) \right\},$$

where $\mathbb{I} \{ \cdot \}$ is the indicator function and $\Gamma^k(t)$ is the time elapsed since the creation of the latest blocks on channel k . Since the process of block creation is assumed independent of the provider update process and the process of opportunistic contacts between mobile hosts, thus we have

$$\mathbb{E} \left[U_i^k(t) \mid \Gamma^k(t) \right] = F^{Y_i^k} \left(\Gamma^k(t) \right)$$

where $F^{Y_i^k}$ denotes the cumulative distribution function of Y_i^k , i.e., $F^{Y_i^k}(y) = \mathbb{P} [Y_i^k \leq y]$.

Note that, as a consequence of Eq.(3), one can easily deduce (see [8]) that the function $(\lambda_i^k)_{i \in \mathcal{V}, k \in \mathcal{K}} \mapsto F^{Y_i^k}(y)$, for any $y \geq 0$, is concave, and

$$\begin{aligned} \mathbb{E} \left[U_i^k(t) \right] &= \mathbb{E} \left[\mathbb{E} \left[U_i^k(t) \mid \Gamma^k(t) \right] \right] \\ &= \int_0^\infty \mathbb{P} \left[\Gamma^k(t) = y \right] \mathbb{E} \left[U_i^k(t) \mid \Gamma^k(t) \right] dy. \end{aligned} \quad (4)$$

Since the integral of a family of concave functions with respect to a positive measure is a concave function, it proves that $(\lambda_i^k)_{i \in \mathcal{V}, k \in \mathcal{K}} \mapsto \mathbb{E} [U_i^k(t)]$ is a concave function. The theorem then follows from the fact that a composition of a non-decreasing concave function ϕ with a concave function is concave and non-decreasing.

Note that Eq. (4) can be further simplified when the process of block creation follows simple statistics, e.g.,

- If blocks are created according to a Poisson process with rate γ_k , then

$$\begin{aligned} \mathbb{E} \left[U_i^k(t) \right] &= \int_0^\infty \gamma_k e^{-\gamma_k \cdot y} F^{Y_i^k}(y) dy \\ &= \mathbb{E} \left[\exp(-\gamma_k \cdot Y_i^k) \right]. \end{aligned} \quad (5)$$

- If blocks are created according to a deterministic period of $1/\gamma_k$, then

$$\begin{aligned}\mathbb{E} [U_i^k(t)] &= \int_0^{1/\gamma_k} \gamma_k F^{Y_i^k}(y) dy \\ &= \mathbb{E} [\max(0, 1 - \gamma_k Y_i^k)].\end{aligned}\quad (6)$$

B. COMPUTATIONAL PROCEDURE

We describe how the service provider can solve the optimization problem **CAP** and compute the optimal values of $(\lambda_i^k)_{i \in \mathcal{V}, k \in \mathcal{K}}$.

First, the provider needs to be aware of the rate $(\gamma_k)_{k \in \mathcal{K}}$ and the statistics of the content generation process for each channel. It also needs to know the interest of nodes for each channel via the matrix A . Last, in order to compute the expected utilities seen by each node on each channel, we assume that it knows a certain number of samples for $(s_{i,j}^k)_{k \in \mathcal{K}, i, j \in \mathcal{V}}$. Samples are indexed by $l = 1, \dots, L$, and denoted by $(\hat{s}_{i,j}^{k,(l)})_{k \in \mathcal{K}, i, j \in \mathcal{V}, l=1, \dots, L}$. Note that the provider does not need to know the sharing matrices $(B^k)_{k \in \mathcal{K}}$ explicitly, because they are implicitly contained in the samples of $\hat{s}_{i,j}^{k,(l)}$'s which are sufficient to run the procedure.

Second, the provider needs to use these samples to estimate expected utility. Note first that for any y , i , and k the provider can estimate $\mathbb{P}[Y_i^k > y]$, seen as a function of $(\lambda_i^k)_{i \in \mathcal{V}, k \in \mathcal{K}}$, based on Eq. (3):

$$\mathbb{P}[Y_i^k > y] \approx \frac{1}{L} \sum_{l=1}^L e^{-\sum_{j \in \mathcal{V}} \lambda_j^k \cdot (y - \hat{s}_{i,j}^{k,(l)})_+}.\quad (7)$$

According to this expression, the utility for a given node and channel can be expressed as a function of $(\lambda_i^k)_{i \in \mathcal{V}, k \in \mathcal{K}}$. As an example, if blocks are created according to a Poisson process, we have by Eq. (5),

$$\mathbb{E}[U_i^k(t)] \approx \int_0^\infty \gamma_k e^{-\gamma_k \cdot y} \left(1 - \frac{1}{L} \sum_{l=1}^L e^{-\sum_{j \in \mathcal{V}} \lambda_j^k \cdot (y - \hat{s}_{i,j}^{k,(l)})_+}\right) dy.$$

This allows us to derive an estimator for any objective ϕ as a function of $(\lambda_i^k)_{i \in \mathcal{V}, k \in \mathcal{K}}$. As this function is concave, a maximum can be found using convex optimization techniques. The estimator becomes closer to the expectation as the number of samples gets large.