

Optimal Channel Choice for Collaborative Ad-Hoc Dissemination

Liang Hu
Tech. University of Denmark
2800 Kgs. Lyngby
Denmark
lhua@fotonik.dtu.dk

Jean-Yves Le Boudec
EPFL
CH 1015 Lausanne
Switzerland
jean-
yves.leboudec@epfl.ch

Milan Vojnović
Microsoft Research
7 JJ Thomson Avenue
CB3 0FB Cambridge, UK
milanv@microsoft.com

July 2009

Technical Report
MSR-TR-2009-26

Microsoft Research
Microsoft Corporation
One Microsoft Way
Redmond, WA 98052
<http://www.research.microsoft.com>

Abstract – Collaborative ad-hoc dissemination of information has been proposed as an efficient means to disseminate information among devices in a wireless ad-hoc network. Devices help in forwarding the information channels to the entire network, by disseminating the channels they subscribe to, plus others. We consider the case where devices have a limited amount of storage that they are willing to devote to the public good, and thus have to decide which channels they are willing to help disseminate. We are interested in finding channel selection strategies which optimize the dissemination time across the channels. We first consider a simple model under the random mixing assumption; we show that channel dissemination time can be characterized in terms of the number of nodes that forward this channel. Then we show that maximizing a social welfare is equivalent to an assignment problem, whose solution can be obtained by a centralized greedy algorithm. We show empirical evidence, based on Zune data, that there is a substantial difference between the utility of the optimal assignment and heuristics that were used in the past. We also show that the optimal assignment can be approximated in a distributed way by a Metropolis-Hastings sampling algorithm. We also give a variant that accounts for battery level. This leads to a practical channel selection and re-selection algorithm that can be implemented without any central control.

1. INTRODUCTION

Several applications relying on opportunistic data transfers between devices have been proposed recently. In [1], the authors propose a wireless ad-hoc podcasting system, where, in addition to downloading content onto devices while docked to a desktop computer, the content is exchanged between devices while users are on the go. They propose several heuristics for content exchange between devices based on the inferred preference of the user owning a device and that of encountered devices. Another related system is CarTorrent [2], a BitTorrent-style content dissemination system designed to exploit the wireless broadcast nature, where the authors propose various solicitation strategies.

We call *channel* an abstraction for various information feeds that generate content recurrently over time. For example, a podcast feed is a channel as well as a profile page of an online social network application (e.g. Facebook or Twitter). While many such services can well be provisioned at mobile devices by accessing the cloud, it is still of interest to speed up information dissemination by augmenting it with device-to-device information transfer. Efficient multi-channel information dissemination through infrastructure and multi-hop wireless transfer would well support various mobile content sharing applications, e.g. Serendipity [?], in particular, in environments where access to the cloud is intermittent, either because of the lack of connectivity or access cost (e.g. during roaming).

We are interested in scenarios where nodes are willing to devote some amount of their resources to help content dissemination. Now the number of information channels can be very large compared to user's interest; for example in the Zune dataset there are 8000+ podcast channels and each user subscribes to 6 channels on average [4, ?]. In such a

setting, we propose to limit the amount of resource that a node devotes to the dissemination of channels other than the ones it subscribes to. This is motivated by the cost for a user in terms of bandwidth usage during meetings, energy consumption, and perhaps also storage. We thus assume that each user device has to decide which channels to help disseminate, in addition to the subscribed ones. We consider a setting where users are cooperative in optimizing the content dissemination, an assumption that underlies the prior work [1]. The cooperation could be induced through various mechanisms like in any other peer-to-peer service. One implicit incentive is indirect reciprocity where users expect that other users would help disseminate the channels subscribed by this user, so the user may well be willing to reciprocate.

We are interested in finding channel selection strategies which optimize channel dissemination times with respect to a system welfare objective. The key assumption that facilitates our framework is a relation between the channel dissemination time and the fraction of the nodes that forward a given channel. Such a relation can be obtained by modeling or empirical analysis, examples of which we show in this paper. However, it is noteworthy that in this paper, we do not advocate any specific function to describe the relation between the dissemination time and the fraction of the forwarding nodes—a thorough analysis of this is left for future work. We cast the problem in the framework of system welfare optimization where the objective is to optimize an aggregate of the utility functions associated with individual channels. We show that for a broad class of utility functions, optimizing the system welfare is equivalent to an assignment problem whose solution can be obtained by a centralized greedy algorithm [3]. We provide empirical evidence, based on real-world data about subscriptions of Zune [4, ?] users to podcasts, that there is a substantial difference between optimal system welfare assignment and some heuristics that were used in the past.

Then we consider the problem of defining a practical, distributed algorithm run by individual nodes to attain a given system objective. We show that the optimal assignment can be approximated in a distributed way by a Metropolis-Hastings sampling algorithm. The algorithm requires knowledge about the fractions of nodes subscribed or forwarding given channel which can be estimated based on local observations by each individual node. We also identify a class of Metropolis-Hastings algorithms that do not require any estimation. We show simulation results that demonstrate that our proposed distributed algorithms concentrate near the optimum system welfare with the rates of convergence of interest in practice.

Our contributions can be summarized as follows:

- We propose a framework for optimizing the dissemination of *multiple* information channels in wireless ad-hoc networks. The optimization is with respect to dissemination times of individual channels subject to the end-user cache capacity constraints. To the best of our knowledge, this is the first proposal for optimizing dissemination of *multiple* information channels in wireless ad-hoc networks with respect to a global system objective.
- The framework enables a *direct engineering* by allowing derivation of the algorithms that decide which channels are helped by which users so as to optimize a given system objective.
- The framework also allows a *reverse engineering* so that

for some given channel selection algorithms deployed by individual nodes, we can determine which underlying global system objective is optimized.

- We show that an optimum system assignment of users to channels for forwarding can be found by a centralized greedy algorithm for a broad class of system objectives identified in this paper.

- Using the data about subscriptions of Zune users to audio podcasts, we demonstrate that there exist scenarios where for given system objective, significant gains can be attained by the optimum system assignment over some heuristics suggested by previous work [1].

- We show that the optimum system objective can be well approximated by a distributed algorithm based on the Metropolis-Hastings sampling run by individual nodes using only local observations.

- We show how to incorporate in our framework and algorithms the objective to optimize the battery expenditure.

- We present extensive simulation results that provide validation and demonstrate practicality of the proposed algorithms.

The paper is structured as follows. Sec. 2 introduces our system model and notation. Sec. 3 presents modeling and empirical analysis about the relation between the channel dissemination time and the fraction of the nodes that forward the channel. In this section, we also define the system objective, the utilities associated to channels, and discuss some of their properties. Sec. 4.1 presents the system welfare problem and the result that the problem can be solved by a centralized greedy algorithm. Sec. 5 presents results on the gain of the optimum system welfare based on the Zune data. Sec. 6 presents our Metropolis-Hastings algorithms. In Sec. 7 we show simulation results. Finally, related work is discussed in Sec. 8 and Sec. 9 concludes the paper. We defer some of our proofs to Appendix.

2. SYSTEM MODEL AND NOTATION

We consider a system of N wireless nodes, or users, participating in the ad-hoc dissemination of J channels. We denote with \mathcal{U} and \mathcal{J} the sets of user and channels, respectively. Every node, say, u has a list $S(u)$ of *subscribed* channels. In the context of this study, we assume that $S(u)$ is fixed for every u . In contrast, every node maintains a variable list of *helped* channels, i.e. channels that this node keeps in its public cache in order to facilitate their dissemination.

When two nodes meet, they update their cache contents. More precisely, if nodes u and u' meet, u gets from u' the content that is newer at u' for the channels that u either subscribes to or helps, and vice-versa. We do not account for the overhead of establishing contacts and negotiating content updates. We assume that when nodes meet the contact duration is large enough for all useful contents to be exchanged, i.e. we assume that the bottlenecks in the system performance are the disconnection times and cache content. In addition, we assume that, once in a while, a node gets direct contact to the Internet and downloads fresh content for the subscribed or helped channels.

At any given point in time, we call x the global system configuration, defined by

$$x_{u,j} = 1 \Leftrightarrow \text{node } u \text{ subscribes to or helps channel } j.$$

Let $H(u, x)$ be the set of channels helped by node u when the configuration is x and let $F(u, x)$ be the set of forwarded

channels, i.e.

$$F(u, x) = H(u, x) \cup S(u), \quad u \in \mathcal{U}.$$

We assume that every node u has a maximum cache capacity C_u , to simplify we count it in the number of channels. We assume that¹ $C_u \geq |S(u)|$, i.e. every node can store all the subscribed channels. The configuration is thus constrained by

$$|F(u, x)| \leq C_u \text{ for all } u \in \mathcal{U}.$$

The problem is then to find a configuration x that satisfies these constraints and maximizes some appropriate performance objective, defined in the next section. Further, we want to find a method to approximate the optimal configuration in a fully distributed way which we do in Sec. 6.

We use the following notation:

$$\begin{aligned} s_j &= \text{proportion of nodes that subscribe to channel } j \\ f_j(x) &= \text{proportion of users that forward channel } j \\ &= \frac{1}{N} \sum_{u \in \mathcal{U}} x_{u,j}. \end{aligned}$$

Without loss of generality and unless indicated otherwise, we assume that channels are labeled in nonincreasing order with respect to their subscription popularity, i.e. $s_1 \geq \dots \geq s_J$. Also $\vec{s} = (s_1, \dots, s_J)$ and $\vec{f} = (f_1, \dots, f_J)$.

3. DISSEMINATION TIME AND UTILITY

To get a better handle on the performance objective we first use an epidemic style analysis, using ordinary differential equations.

3.1 Model-Based Dissemination Time

Consider a channel j and set the time origin to the time at which the most recent version was created by the source. We assume the configuration x is fixed and omit it from the notation in this section. Let $\sigma_j(t)$ be the proportion of j -subscribers that have received the most recent piece at time t , and let $\phi_j(t)$ be the proportion of j -forwarders that have received the most recent piece at time t .

The dynamics of the system can be described by the system of ordinary differential equations:

$$\frac{d}{dt} \sigma_j(t) = (\lambda_j + \eta \phi_j(t)) (\sigma_j - \sigma_j(t)) \quad (1)$$

$$\frac{d}{dt} \phi_j(t) = (\lambda_j + \eta \phi_j(t)) (\phi_j - \phi_j(t)) \quad (2)$$

where λ_j is the contact rate between a node and an infrastructure able to deliver channel j , and η is the contact rate between nodes. These equations correspond to the ‘‘random node mixing’’ assumption and are asymptotically valid when N is large. It follows that $d\sigma_j/d\phi_j = (s_j - \sigma_j)/(f_j - \phi_j)$, hence

$$\sigma_j(t) = \frac{f_j \sigma_j(0) - s_j \phi_j(0)}{f_j - \phi_j(0)} + \frac{s_j - \sigma_j(0)}{f_j - \phi_j(0)} \phi_j(t). \quad (3)$$

We can solve Eq. (2) explicitly. Note that

$$\frac{1}{(\lambda + \eta \phi)(f - \phi)} = \frac{1}{\lambda + f} \left(\frac{1}{\lambda + \eta \phi} + \frac{1}{f - \phi} \right) \quad (4)$$

¹ $|A|$ denotes the cardinality of a finite set A .

from which we get

$$\phi_j(t) = \frac{1}{\eta} \left(-\lambda_j + \frac{(\lambda_j + \eta\phi_j(0))(\lambda_j + f_j\eta)}{\lambda_j + \eta\phi_j(0) + \eta(f_j - \phi_j(0))e^{-(\eta f_j + \lambda_j)t}} \right).$$

By Eq. (3) we obtain

$$\sigma_j(t) = \sigma_j(0) + (s_j - \sigma_j(0)) \times \frac{(\lambda_j + \eta\phi_j(0))(1 - e^{-(\eta f_j + \lambda_j)t})}{\lambda_j + \eta\phi_j(0) + \eta(f_j - \phi_j(0))e^{-(\eta f_j + \lambda_j)t}}. \quad (5)$$

Dissemination Time

Suppose that at time T_0 a piece is issued by the source. Let T_1 be the time at which a proportion α of the subscribers have received this or a more recent piece. We call $t_j = T_1 - T_0$ the dissemination time and take it as metric for channel j .

We compute t_j as follows. First note, from Eq. (5):

$$e^{-(\eta f_j + \lambda_j)t_j} = \frac{(\lambda_j + \eta\phi_j(0)(1 - K_j))}{\eta f_j K_j + \lambda_j + \phi_j(0)\eta(1 - K_j)}$$

where

$$K_j = \frac{\alpha - \frac{\sigma_j(0)}{s_j}}{1 - \frac{\sigma_j(0)}{s_j}}.$$

It follows

$$t_j = \frac{1}{\lambda_j + f_j\eta} \ln \frac{(f_j - \phi_j(0))\eta K_j + \lambda_j + \eta\phi_j(0)}{(\lambda_j + \eta\phi_j(0))(1 - K_j)}. \quad (6)$$

PROPOSITION 3.1. *The dissemination time t_j is a monotonic nonincreasing, strictly convex function of f_j .*

Proof is in Appendix A.

Of particular interest is the small injection rate regime, where the dissemination is dominated by epidemic content. In this case, we have

$$\begin{aligned} \sigma_j(0) &\ll \frac{\lambda_j}{\eta} \ll s_j \\ \phi_j(0) &\ll \frac{\lambda_j}{\eta} \ll f_j \end{aligned}$$

and Eq.(6) becomes

$$t_j \approx \frac{1}{\eta f_j} \left(\ln \frac{\alpha}{1 - \alpha} + \ln \frac{\eta f_j}{\lambda_j} \right). \quad (7)$$

3.2 Empirical Dissemination Time

We consider the dissemination time evaluated from real mobility traces. In particular, we consider (CAM) a data trace of human mobility in the area of Cambridge, UK [5] and (SF-TAXI) a data trace of the routes of taxis in the area of San Francisco [6]. The CAM dataset contains information about the contacts between 36 human-carried, Bluetooth-equipped devices over slightly more than 10 days. SF-TAXI contains GPS coordinates for each of about 500 taxis over a month period. For the latter trace, we define a contact between two nodes if the distance between two nodes is smaller or equal to 500 meters [7].

We infer the dissemination times by conducting the following experiment. For given data trace (either CAM or SF-TAXI), we fix a portion of nodes and then pick uniformly at random the given portion of nodes from the set of all nodes and designate them as forwarders. We then inject a message

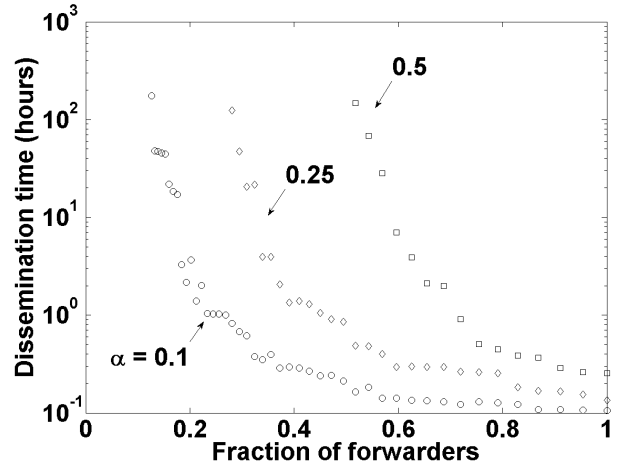


Figure 1: Dissemination time versus the fraction of forwarding nodes in CAM data. Each mark shows the median value of the dissemination time obtained by taking each node as a source and repeating for 10 random selections of the forwarding nodes.

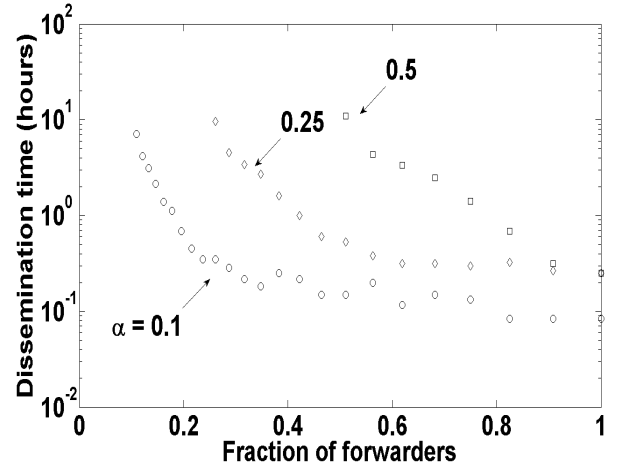


Figure 2: Same as in Fig. 1 but for SF-TAXI data.

to one of the forwarders at an instance of time and then pass through the trace forward in time, recording the instances at which a forwarder first received the message by encountering a forwarder that already received the message. For the CAM data, we repeat the experiment for each source and 10 random samples of the set of forwarders. Finally, for each given portion of the forwarding nodes, we compute the median dissemination time.

Fig. 1 and Fig. 2 show the empirical dissemination time versus the portion of the forwarding nodes for CAM and SF-TAXI traces, respectively. In both cases, they confirm that the dissemination time is well fitted by a curve that exhibits diminishing returns with increasing the number of forwarders.

3.3 Utility Function

We assume that for each channel j there is an underlying utility function $U_j(t_j)$ that specifies the satisfaction of a subscriber with the channel j given that the dissemination




Figure 3: Utility of the dissemination time. (Left) A concave decreasing utility with respect to the dissemination time capturing the increasing rate of the user’s unhappiness. (Right) Finite utility for dissemination time up to some time t_0 and $-\infty$ for larger values.

time is t_j . It is natural to assume $U_j(t_j)$ is a nonincreasing function of t_j . We will discuss later in this section some additional conditions for a channel utility function.

We denote with $V_j(f_j) = U_j(t_j(f_j))$ the utility function for channel j with respect to the fraction of the users who forward channel j . It is natural to assume that $V_j(f_j)$ is a monotonic nondecreasing function of f_j . This indeed follows if both $U_j(t_j)$ and $t_j(f_j)$ are nonincreasing functions which are rather natural assumptions.

It remains to discuss what the system welfare utility is, i.e. when considering all channels together. We admit standard definition that the system welfare is a weighted sum of the utilities over all channels, i.e. for given positive weights $\vec{w} = (w_1, \dots, w_J)$,

$$V(\vec{f}) = \sum_{j \in \mathcal{J}} w_j V_j(f_j).$$

Two special cases may be of interest, which correspond to different fairness objectives. The former is channel centric, in that it considers each channel as one entity, regardless of the number of subscribers. This utility is obtained by setting all the weights w_j to 1, hence we have

$$V_{CH}(\vec{f}) = \sum_{j \in \mathcal{J}} V_j(f_j) \quad (8)$$

where V_j is a per-channel metric, for example as in Eq.(6) or Eq.(7). The latter is user centric and has the weights such that w_j is proportional to the proportion of j -subscribers, s_j , hence we consider

$$V_{US}(\vec{f}) = \sum_{j \in \mathcal{J}} s_j V_j(f_j) \quad (9)$$

with V_j as before.

In Sec. 6 we will show that this utility framework can easily be extended to battery saving.

Sufficient Conditions for a Concave Utility

We discuss a set of *sufficient* conditions that ensure the utility $V_j(f_j)$ is a concave function of f_j . This class of utility functions will be of interest in Section 4.1.

PROPOSITION 3.2. *Suppose (C1) $U_j(t_j)$ is a nonincreasing, concave function of t_j and (C2) $t_j(f_j)$ is a convex function of f_j . Then, $V_j(f_j)$ is a concave function of f_j .*

PROOF. By simple differential calculus,

$$\begin{aligned} V_j'(f_j) &= U_j'(t_j) t_j'(f_j) \\ V_j''(f_j) &= U_j''(t_j) (t_j'(f_j))^2 + U_j'(t_j) t_j''(f_j). \end{aligned}$$

From the last equation, (C1) $U_j'(t_j) \leq 0$, $U_j''(t_j) \leq 0$, and (C2) $t_j''(f_j) \geq 0$, it follows $V_j''(f_j) \leq 0$, i.e. $V_j(f_j)$ is a concave function of f_j . \square

Condition (C1) says that the utility function $U_j(t_j)$ captures the increasing dissatisfaction of a subscriber of channel j with the dissemination time t_j . See Figure 3–left for an illustration. Such a utility function could be seen as a smooth version of a step function (see Figure 3–right) where the utility is finite up to some time and then becomes $-\infty$ for larger dissemination times. This captures scenarios where a subscriber values the information only if received within some time.

Condition (C2) says that the dissemination time $t_j(f_j)$ exhibits diminishing returns with increasing the portion of forwarders f_j . We have already demonstrated cases in Section 3.1 and Section 3.2 that support this assumption.

4. SYSTEM WELFARE PROBLEM

4.1 The Greedy Algorithm

We pose a system welfare problem where the objective is to optimize the aggregate utility of channel dissemination times subject to the end-user capacity constraints. Solving the system welfare problem amounts to finding an *assignment* of users to channels that solves the following problem:

SYSTEM	
maximize	$\sum_{j=1}^J w_j V_j \left(\frac{1}{N} \sum_{u=1}^N x_{u,j} \right)$
over	$x_{u,j} \in \{0, 1\}$
subject to	$\sum_{j=1}^J x_{u,j} \leq C_u$
	$x_{u,j} = 1, (u, j) : j \in S(u).$

Defining the system welfare utility as a sum of individual utilities is rather standard in the microeconomics framework of the resource allocation and was successfully applied, for example, in the contexts of wireline Internet [8] and wireless networks [9]. Note that in SYSTEM, w_j are positive constants that can be arbitrarily fixed. In particular, it is of interest to set w_j proportional the portion of users subscribed to channel j (i.e. s_j). In this case, the utility $V_j(\cdot)$ can be interpreted as the utility for channel j for a typical j -subscriber.

We rephrase SYSTEM as an optimization over the number of helper users per channel. Consider $\vec{H} = (H_1, \dots, H_J)$ where H_j is the number of helper users for channel j . Let us define $v(A)$ for $A \subseteq \mathcal{J}$, by

$$v(A) = \sum_{u \in \mathcal{U}} \min \left(\sum_{j \in A} 1_{j \in \mathcal{J} \setminus S(u)}, C_u - |S(u)| \right). \quad (10)$$

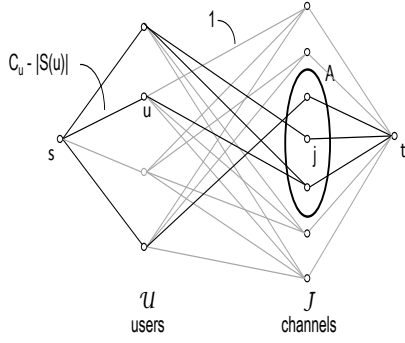


Figure 4: Feasible assignment of users to channels as a max-flow problem. \mathcal{U} is the set of users and \mathcal{J} is the set of channels. s is a source node connected to a user u with an edge of capacity $C_u - |S(u)|$. There is an edge of capacity 1 between any user u and channel j if and only if user u is not subscribed to channel j , i.e. $j \in \mathcal{J} \setminus S(u)$.

Let $P(v)$ be the polyhedron defined by²

$$P(v) = \{x \in \mathbb{N}_0^J : x(A) \leq v(A), A \subseteq \mathcal{J}\}.$$

We consider the following problem.

SYSTEM-H

maximize $\sum_{j=1}^J w_j V_j \left(s_j + \frac{1}{N} H_j \right)$

over $\vec{H} \in P(v)$.

PROPOSITION 4.1. *The optimal value of the solution of SYSTEM is equal to that of SYSTEM-H.*

PROOF. By the definition of \vec{s} and \vec{H} , the objective functions of SYSTEM and SYSTEM-H are the same.

(1) We now claim that the constraints of SYSTEM imply the constraints of SYSTEM-H. To this end, for any set $A \subseteq \mathcal{J}$, let $H(A) := \sum_{j \in A} H_j$. Note that $H(A) = \sum_{u \in \mathcal{U}, j \in A} x_{u,j}$ can be interpreted as a flow into the set A in the following graph. Let $G' = (V', E')$ with the set of vertices V' defined by $V' = \mathcal{U} \cup \mathcal{J} \cup \{s\} \cup \{t\}$ where s and t are interpreted as a source and a terminal vertex, respectively. The set of edges E' is defined as follows. The source s is connected to a node u with an edge of capacity $c(s, u) = C_u - |S(u)|$. A user u and a channel j are connected by an edge of capacity $c(u, j) = 1$ if and only if the user u is not subscribed to channel j , i.e. $j \in \mathcal{J} \setminus S(u)$. Finally, each vertex $j \in A$ is connected to the vertex t with an edge of sufficiently large capacity ($|A||\mathcal{J}|$ suffices). See Fig. 4 for an illustration.

By the max-flow-min-cut theorem [10, Theorem 8.6], the max value of an s - t flow equals the minimum capacity of an s - t cut. Graph G' is a collection of edge-disjoint trees and thus finding an s - t min-cut amounts to deciding for each user u whether to cut an edge between the vertices s and u or all the edges between the vertex u and vertices $j \in A$ (if there exist any). In the former case, the capacity of the cut is $C_u - |S(u)|$ while in the latter case it is $\sum_{j \in A} 1_{j \in \mathcal{J} \setminus S(u)}$.

²We use the notation $x(A) = \sum_{j \in A} x_j$.

It follows that the max-flow is given by $v(A)$ in (10). The announced claim is thus proven. It follows that the optimum of SYSTEM is \leq that of SYSTEM-H.

(2) Conversely, let \vec{H}^* be a point that achieves the optimum of SYSTEM-H. It exists since $P(v)$ is finite. Consider the graph G' derived from Figure 4 where we define the capacity of the edges (j, t) by setting $c(j, t) = H_j^*$. Following [11] (Proof of Lemma 4.1, p.103, text after Eq. (4.4), $(\{s\} \cup \mathcal{U} \cup \mathcal{J}, \{t\})$ is a min cut separating s and t , equal to $\min(v(\mathcal{J}), H^*(\mathcal{J})) = H^*(\mathcal{J})$ where $H^*(\mathcal{J}) = \sum_{j \in \mathcal{J}} H_j^*$ by definition.

The maximal flow s - t in this graph is thus $H^*(\mathcal{J})$. By the Integral Flow Theorem ([10], Corollary 8.7) there exists an integer flow on this graph such that the flow through j - t is H_j^* . It follows that there exists a solution x^* to the constraints of SYSTEM such that $\sum_{u \in \mathcal{U}} x_{u,j}^* = N s_j + H_j^*$, thus the optimum of SYSTEM is \geq that of SYSTEM-H. \square

We denote with $\Delta_j V(\vec{s} + \vec{H}/N)$ the increment of the aggregate utility function by assigning a user to channel j , i.e.

$$\begin{aligned} & \Delta_j V(\vec{s} + \vec{H}/N) \\ &= V(\vec{s} + (\vec{H} + e_j)/N) - V(\vec{s} + \vec{H}/N) \\ &= w_j [V_j(s_j + (H_j + 1)/N) - V_j(s_j + H_j/N)] \end{aligned}$$

where e_j is a vector of dimension $|\mathcal{J}|$ with all the coordinates equal to 0 but the j th coordinate equal to 1.

Algorithm 1 Centralized GREEDY Algorithm for Allocation of Helped Channels.

```

1:  $H = 0$ 
2: while 1 do
3:   Find  $I \in \mathcal{J}$  such that  $\vec{H} + e_I \in P(v)$ 
4:   and  $\Delta_I V(\vec{s} + \vec{H}/N) \geq \Delta_j V(\vec{s} + \vec{H}/N)$  for all  $j \in \mathcal{J}$ 
5:   such that  $\vec{H} + e_j \in P(v)$ 
6:
7:   if there exists no such  $I$  then break
8:   end if
9:    $H_I \leftarrow H_I + 1$ 
10: end while

```

THEOREM 4.1. *Assume that for each $j \in \mathcal{J}$, $V_j(x)$ is a concave function of x . Then, a solution of SYSTEM is obtained by GREEDY.*

PROOF. Under the assumption that $V_j(x)$ is a concave function with respect to x we have that $V_j(s_j + x)$ is a concave function with respect to x . Showing in addition that $P(v)$ is a submodular polyhedron, we verify the assumptions of Corollary 1 in Feedergruen and Groenevelt [3] from which the asserted result follows.

A polyhedron $P(v)$ is submodular if and only if $v(\cdot)$ is a submodular function, i.e.

$$v(A \cup B) + v(A \cap B) \leq v(A) + v(B), \quad A, B \subseteq \mathcal{J}. \quad (11)$$

But this follows from the fact that $v(\cdot)$ is the characteristic function of the graph in Figure 4 and [11, Lemma 3.2]. \square

The interested reader is referred to Appendix B where we consider a relaxed version of SYSTEM which allows providing some characterization of the solution to this relaxed version.

4.2 Particular Channel Selection Strategies

In this section, we introduce three particular channel selection strategies. Under the assumption of random mixing, the first two strategies correspond to uniform and most-solicited strategies in [1]. The third strategy is new and arises from the Metropolis sampling in Sec. 6.

4.2.1 Uniform

Under the uniform channel selection, each user u picks a subset of $C_u - |S(u)|$ channels by sampling uniformly at random without replacement from the set of channels that user u is not subscribed to, i.e. from the set of channels $\mathcal{J} \setminus S(u)$.

The uniform channel selection biases to forwarding less popular channels. This is quite intuitive as by the channel selection process the users select channels to which they are not subscribed to. The interested reader is referred to [?] where more discussion is provided along with making a connection to an underlying system welfare problem of the uniform channel selection.

4.2.2 Top Popular

Under this scheme, each user u picks channels from the set of channels $\mathcal{J} \setminus S(u)$ without replacement in decreasing order of the channel subscription popularity and random tie break until $C_u - |S(u)|$ channels are picked or there are no channels left. This is a greedy scheme that favours popular channels. We consider this scheme in numerical evaluations in Sec. 5.

4.2.3 Pick from Neighbour

We consider channel selection strategies under which each user u upon encountering another user u' picks a candidate channel from the user u' and then based on some decision process decides whether to replace a channel to which user u currently helps with the candidate channel. The decision process is assumed to be local, independent of the current assignment of users to channels, which makes these strategies of quite practical interest.

We will construct one such a scheme, in Sec. 6, based on the Metropolis-Hastings sampling. We will see that such a scheme is associated with a system welfare problem with the following objective function:

$$V_{PFN}(\vec{f}) = \sum_{j \in \mathcal{J}} V_j^{PFN}(f_j)$$

with

$$V_j^{PFN}(f_j) = (\alpha_j + C)f_j + Df_j \ln f_j \quad (12)$$

where C and D are system constants and $\alpha_j \geq 0$ is a constant for channel j , which expresses its relative importance (the higher the α_j , the more important the channel j).

The function $V_j^{PFN}(f_j)$ in Eq. (12) is a monotonic non-decreasing function of f_j . Note, however, that $V_j^{PFN}(f_j)$ is a *convex* function of f_j . It is thus not concave and hence does not validate the condition discussed in Sec. 3.3, which ensures optimality of the greedy assignment in Sec. 4.1.

Proof is in Appendix C.

5. SYSTEM OPTIMUM VS. HEURISTICS

In this section, we demonstrate:

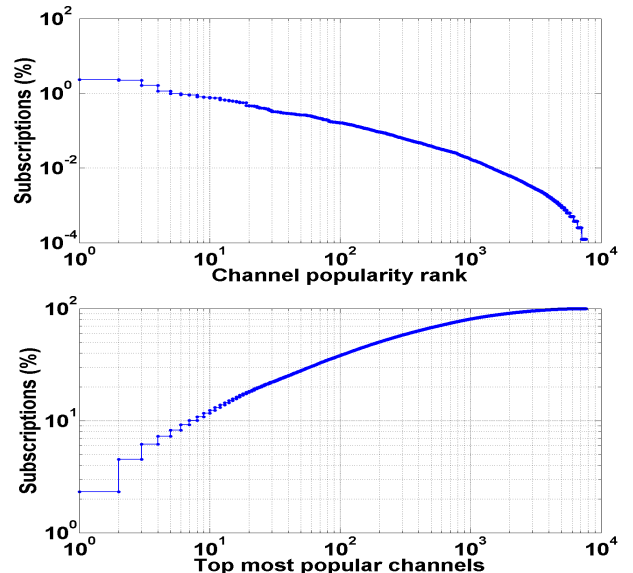


Figure 5: Channel subscription popularity from the Zune podcasts data. (Top) Fraction of the subscriptions per channel. (Bottom) Fraction of the subscriptions covered by a set of most popular channels.

A system optimal assignment of channels can yield significantly larger system welfare than some heuristics suggested by prior work.

In particular, we compare with the Uniform and Top Popular assignments defined in the preceding section.

We use the subscription assignments of users to channels that we derive from the subscriptions of the users of Zune to audio podcasts. This dataset consists of 8,000+ distinct podcast feeds and more than a million of users. The data provides us with complete information about subscriptions of users to podcasts. In Fig. 5-top, we show the fraction of subscriptions covered by individual channels. This metric corresponds to our definition of \vec{f} . We note that the distribution is quite skewed with a few channels with many subscriptions and many with a few. The median number of the fraction of subscriptions per channel is as small as about $2 \cdot 10^{-5}$. Moreover, only about 1% of all the channels have the fraction of subscriptions at least the factor 1/10 of that of the most popular channel. The body of the distribution in Fig. 5-top is well approximated by a line (power-law) with the slope of about $-2/3$. In Fig. 5-bottom, we re-plot the same data but show the fraction of the subscriptions covered by a set of most popular channels. From this figure, we note that about half of the subscriptions are covered by as few as 2.5% of the most popular channels.

We consider the user-centric system welfare with the channel utility functions $V_j(f_j) = -t_j(f_j)$ where $t_j(f_j)$ is the dissemination time given by Eq. (6). For each user u , we set $C_u = |S(u)| + C$ where $|S(u)|$ is specified by the input data and C is a parameter. We compute optimum assignment by using the algorithm GREEDY (Sec. 4.1). Uniform and Top Popular assignments are computed as prescribed by their respective definitions.

In Fig. 6, we show the dissemination time per subscription versus the per node capacity C . The rate of the access to the infrastructure is fixed to 1 access per day by each user.




Figure 6: Dissemination time per subscription versus the size of the public cache C , $C_u = |S(u)| + C$.

The rate at which each user encounters other users is fixed to 100 users per day. If the dissemination is solely by direct access to the infrastructure, then the dissemination time is about 13.5 hours. We note that the dissemination time under the system optimum assignment can be reduced for the order of several hours if the dissemination is augmented with the peer-to-peer dissemination. Perhaps even more interestingly, we observe that the gap between the system optimum and that of the Uniform and Top Popular assignments can be significant.

In Fig. 7, we present the results under the same setting as in Fig. 6 but varying the encounter rate and holding the cache size C fixed to 5 (Top) and 20 (Bottom). These results show a lack of order for the Uniform and Top Popular assignments – for some cases one is better than the other one and vice-versa in other cases. In any case, system optimum indeed provides best performance.

6. A DISTRIBUTED METROPOLIS HASTINGS ALGORITHM

We now consider the problem of designing a distributed algorithm. The goal is for each node to control its set of helped channels so that the resulting global configuration x maximizes a system welfare of the form

$$V(x) = \sum_{j \in \mathcal{J}} w_j V_j(f_j(x)) \quad (13)$$

as discussed in Sec. 3 (note that, unlike in Sec. 3, we make the dependence on the global configuration x explicit).

6.1 Metropolis-Hastings

We propose to use a Metropolis-Hastings algorithm [12], as it lends itself well to distributed optimization, and was successfully used in distributed control problems in wireless networks [13]. Before giving our distributed algorithm, we first give a short description of a centralized version of the Metropolis-Hastings algorithm:

At every time step, the algorithm picks a tentative configuration x' , with probability $Q(x, x')$, where x is the current configuration. We assume that the matrix $Q(x, x')$ has the

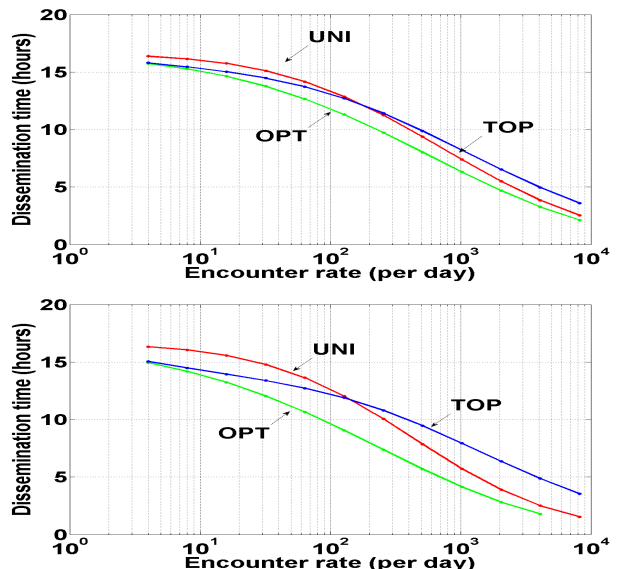


Figure 7: Dissemination time per subscription versus the rate of encounters η . The cache for user u set as $C_u = |S(u)| + C$ with (Top) $C = 5$ and (Bottom) $C = 20$.

weak symmetry property:

$$Q(x, x') > 0 \Rightarrow Q(x', x) > 0$$

for all $x \neq x'$. The tentative configuration is accepted (i.e. becomes the new configuration) with probability $p = \min(1, q)$ with

$$q = \frac{\pi(x')Q(x', x)}{\pi(x)Q(x, x')} \quad (14)$$

where $\pi(\cdot)$ is a probability distribution on the set of possible configurations. The algorithm does not converge stricto sensu, however, after a large number of iterations, the probability distribution of the configuration x converges to the a priori distribution $\pi(\cdot)$. Typically, one uses for $\pi(\cdot)$ a Gibbs distribution, given by

$$\pi(x) = \frac{1}{Z} e^{-\frac{V(x)}{T}} \quad (15)$$

where T is a system parameter (the “temperature”) and Z is the normalizing constant. If T is small, the distribution $\pi(\cdot)$ is very much concentrated on the large values of $V(x)$, so that the algorithm produces random configurations that tend to maximize $V(x)$.

6.2 A Distributed Rewiring Algorithm

We use Metropolis-Hastings as follows. We use a Gibbs distribution, as in Eq.(15) with $V(\cdot)$ the utility function in Eq. (13). We consider every meeting between two nodes as one step of the algorithm. When two nodes meet, they opportunistically exchange content updates; then one of them, say u is selected as leader and attempts to replace one of its helped channels by one of the channels forwarded from the set held by the other node, say v , as described in Algorithm 2. We now turn to the computation of the acceptance probability (line 5 of the algorithm), as given by Eq.(14). First we compute $Q(x, x')$ where $x' = x - 1^{u,j} + 1^{u,j'}$ is the

Algorithm 2 Distributed Algorithm for Allocation of Helped Channels

```

1: if  $F(u, x) \subset F(v, x)$  then do nothing
2: else
3:    $u$  selects one channel  $j$  uniformly at random in the
   set  $H(u, x)$ 
4:    $u$  selects one channel  $j'$  uniformly at random in the
   set  $F(v, x) \setminus F(u, x)$ 
5:   compute the acceptance probability  $p = \min(1, q)$ 
   with  $q$  given by Eq.(18)
6:   draw a random number  $U$  uniformly in  $[0, 1]$ ;
7:   if  $U < p$  then drop channel  $j$  and adopt channel  $j'$ 
   as a helped channel
8:   end if
9: end if

```

new configuration ($1^{u,j}$ is the configuration vector defined by $1_{u',j'}^{u,j} = 1$ if $u = u'$ and $j = j'$, 0 otherwise):

PROPOSITION 6.1. *The following holds*

$$\frac{Q(x', x)}{Q(x, x')} = \frac{\sum_{v \neq u} \frac{1_{j \in F(v, x)}}{|F(v, x) \setminus F(u, x)| + 1_{j' \notin F(v, x)}}}{\sum_{v \neq u} \frac{1_{j' \in F(v, x)}}{|F(v, x) \setminus F(u, x)|}}. \quad (16)$$

Proof is in Appendix D.

We will make use of the following approximation, derived in Appendix E,

$$\frac{Q(x', x)}{Q(x, x')} \approx \frac{f_j(x)}{f_{j'}(x)}. \quad (17)$$

We also note the following result (proof in Appendix F.)

PROPOSITION 6.2. *Suppose that for a finite constant $D > 0$, $\lim_{N \rightarrow +\infty} NT = D$. Then,*

$$\lim_{N \rightarrow +\infty} \frac{V(x') - V(x)}{T} = \frac{1}{D} (w_{j'} V_{j'}'(f_{j'}(x)) - w_j V_j'(f_j(x))).$$

In view of the last proposition, we have

$$\begin{aligned} q &= \frac{Q(x', x)}{Q(x, x')} e^{\frac{1}{T}(V(x') - V(x))} \\ &\approx \frac{Q(x', x)}{Q(x, x')} e^{\frac{1}{NT}(w_{j'} V_{j'}'(f_{j'}(x)) - w_j V_j'(f_j(x)))}. \end{aligned}$$

Combining with (17) we obtain for q the value

$$q = \frac{f_j(x)}{f_{j'}(x)} e^{\frac{1}{D}(w_{j'} V_{j'}'(f_{j'}(x)) - w_j V_j'(f_j(x)))} \quad (18)$$

where $D = NT$ is a global system parameter.

Algorithm 2 requires node u to estimate f_j and $f_{j'}$. This can be done by having the nodes exchange, when they meet, updates of channel popularity for all channels that they know of, and then performing exponential smoothing. A simple, but memory hungry scheme, is as follows. Every node u maintains for every channel j an estimate \hat{f}_j . When node u meets node u' , for all channels that u' helps or subscribes to, node u does $\hat{f}_j \leftarrow a + (1 - a)\hat{f}_j$ and for all other channels $\hat{f}_j \leftarrow (1 - a)\hat{f}_j$ where $0 < a < 1$.

A less memory hungry scheme can be obtained by implementing a lazy evaluation scheme; it avoids the scalability

problem that would arise if all nodes would carry popularity information for all channels, while achieving approximately the same goal. We keep f_j in memory, for a total of at most W records, but only for the forwarded channels and otherwise for the most recently seen channels. When a node u meets a node u' and has to push out an \hat{f}_j record due to the W limit, the record with the oldest modification date is pushed out. The record for the new channel at this node is set to $\hat{f}_j \leftarrow a \frac{1}{1+W} + (1 - a)\hat{f}_j(u')$, where $\hat{f}_j(u')$ is the record at u' (which must be present since u' forwards j).

Further, all nodes need to share the global system variable D , and know the utility function of each channel (the latter can be contained as meta-information in the channel data). In Section 6.3, we give a simplified algorithm, which does not require such estimations.

6.3 A Simplified Algorithm

It is possible to entirely avoid the estimation of the f_j quantities, albeit at the expense of imposing a family of utility functions. The idea is to pick a set of utility functions $V_j(\cdot)$ such that f_j and $f_{j'}$ cancel out in Eq.(18). This results in a scheme that belongs to the class of schemes pick from neighbour that was introduced in Section 4.2.3.

THEOREM 6.1. *If for each channel j , the utility function is $V_j^{PFN}(\cdot)$ in Eq.(12) then q in Eq.(18) is given by*

$$q = \frac{\beta_{j'}}{\beta_j} \quad (19)$$

with $\beta_j = e^{\frac{\alpha_j}{D}}$ and $\beta_{j'} = e^{\frac{\alpha_{j'}}{D}}$. In particular, q is thus independent of $f_j(x)$, $f_{j'}(x)$ and more generally of the configuration x .

PROOF. Follows from Eq. (12) and Eq. (18). \square

With this simplified algorithm, nodes need to know the static parameters $\beta_j > 0$ associated with each channel. There is no global constant, nor it is necessary to evaluate $f_j(x)$. Higher values of j mean that we give more value to disseminating channel j more quickly. Note that only the relative values of β_j matter, as Eq.(19) uses only ratios, and β_j can thus be interpreted as the priority level for channel j . The resulting algorithm is as follows.

Algorithm 3 Distributed Algorithm for Allocation of Helped Channels when Utility is Given by Eq.(12). Every channel j has a static priority level $\beta_j > 0$.

```

1: if  $F(u, x) \subset F(v, x)$  then do nothing
2: else
3:    $u$  selects one channel  $j$  uniformly at random in the
   set  $H(u, x)$ 
4:    $u$  selects one channel  $j'$  uniformly at random in the
   set  $F(v, x) \setminus F(u, x)$ 
5:   if  $\beta_{j'} \geq \beta_j$  then drop channel  $j$  and adopt channel
    $j'$  as a helped channel
6:   else
7:     draw a random number  $U$  uniformly in  $[0, 1]$ ;
8:     if  $U < \frac{\beta_{j'}}{\beta_j}$  then drop channel  $j$  and adopt chan-
     nel  $j'$  as a helped channel
9:   end if
10:  end if
11: end if

```

If we set $\beta_j = 1$ for all channels, i.e. we give all channels the same utility function, then Algorithm 3 always accepts the proposed change. Note however that, even in this case, the resulting allocation is, in general, not uniform, as the optimal allocation depends on the proportion of subscribers s_j for each channel; indeed, the algorithm will tend to give more help to channels that have few subscribers. Note also that, in general, the scheme is different from that in Sec. 4.2.1 as under the scheme therein, each user picks from the set of all distinct channels for which this user is not a subscriber, while for the algorithm in the present section, the picking is from the forwarding channels of an encountered user.

6.4 A Battery Saving Algorithm

The previous algorithm may be improved to account for battery saving. The motivation is that a node may be reluctant to help disseminate channels if its battery level is low. We address this issue as follows. Assume that every node u knows its battery level $b_u \geq 0$. The battery is empty when $b_u = 0$. Assume to simplify that all nodes measure b_u in the same scale, for example, number of remaining hours of operation at full activity. We can replace the global utility in Eq.(13) by

$$\sum_{j \in \mathcal{J}} w_j V_j(f_j) - \sum_{u \in \mathcal{U}} W_u(b_u)$$

where $W_u()$ is a convex, decreasing function of its argument (for example $W_u(b) = \frac{1}{b^m}$), such that $W_u(b)$ expresses the penalty perceived by user u when its battery level is b . We can apply the Metropolis-Hastings algorithm with this new function. The only difference is in the computation of the acceptance probability. This can be applied to Algorithms 2 or 3 in the same way, we give the details only for Algorithm 3. The computation of q in Eq.(19) is replaced by

$$q = \frac{\beta_{j'}}{\beta_j} e^{-[h_u(b_u) - h_{u'}(b_{u'})]} \quad (20)$$

where u and u' are the two nodes involved in the interaction and $h_u(b) > 0$ is the marginal cost of exchanging a channel when two nodes meet, divided by the temperature T (an increasing function of b).

The resulting algorithm is the same as Algorithm 2 with Eq.(18) on line 5 replaced by Eq.(20). The required configuration is (1) every channel j has a static priority level $\beta_j > 0$ and (2) every node u knows its own function $h_u(b)$ for the cost of exchanging one channel with a neighbour when this node's battery level is b .

7. SIMULATION RESULTS

In this section, we present simulation results that address the following goals: (i) demonstrate the concentration of the distributed Metropolis-Hastings algorithm to the optimum system welfare and (ii) demonstrate that optimizing a system welfare under real-world mobility produces good forwarding assignments of channels to users.

In order to cover a broad set of parameters, we conducted simulations by varying the parameters along the following dimensions: (i) node mobility either random mixing or using a real mobility trace, (ii) small and large system scale with respect to the number of users and the number of channels, (iii) different distributions for the subscriptions per channel,

(iv) the fractions of nodes forwarding or subscribed to a channel either known or locally estimated, and (v) a range of the temperatures for the Metropolis-Hastings algorithm. We consider random mixing mobility in order to provide results for scenarios for which we have a good understanding of the relation between the channel dissemination time and the fraction of the forwarding nodes. We used our own discrete-event simulator.

7.1 Random Mixing Mobility

We simulate random mixing mobility where each user encounters other users uniformly at random. In such a system, we indeed have that the dissemination time for any channel depends only on the portion of the nodes that forward the channel (Section 3.1).

We consider a small- and a large-scale system where for the former the number of users and the number of channels are both set to 20 while for the latter the number of users is 200 and the number of channels is 100. For the fractions of subscribers per channel (\vec{s}), we assume a Zipf distribution with the scale parameter equal to either 2/3 or 1. The former value is motivated by the empirical distribution derived from the Zune data (Fig. 5 discussed in Section 5) while the latter value was used in previous work [1]. For the objective of the system welfare, we consider both the channel- and user-centric cases with the utility function $V_j(f_j) = -t_j(f_j)$ for channel j where $t_j(f_j)$ is the dissemination time and f_j is the fraction of forwarding nodes. In particular, we admit Eq. (6). In cases when \vec{f} or \vec{s} are locally estimated, each node uses an exponential weighted averaging with the smoothing constant (weight of a sample) set as follows. For the estimation of \vec{f} , the constant is set to 0.9. For the estimation of \vec{s} , the constant is equal to 0.1 and 0.02 for the channel- and user-centric case, respectively.

In Fig. 8, we present the results obtained for the channel-centric case. The graphs show the mean dissemination time per channel, i.e. $(\sum_{j \in \mathcal{J}} t_j(f_j))/J$, versus the number of encounters per node. We show the results for the Metropolis-Hastings with \vec{f} assumed to be either known or locally estimated by individual nodes. We observe that the system welfare under the Metropolis-Hastings algorithm concentrates near the optimum system welfare. The results in Fig. 8 indicate a faster concentration in cases when \vec{f} is globally known. In Fig. 9, we present analogous results for the user-centric case. In this case, we show the mean dissemination time per user, i.e. $(\sum_{j \in \mathcal{J}} s_j t_j(f_j))/\sum_{j \in \mathcal{J}} s_j$, with \vec{f} and \vec{s} either globally known or locally estimated by individual nodes. In summary, the presented results in either channel- or user-centric case support the following claim:

The system welfare under the Metropolis-Hastings algorithm concentrates near the optimum system welfare with \vec{f} (and \vec{s} in the user-centric case) either globally known or locally estimated.

7.2 Real Trace Mobility

We compare the system performance under the assignment of channels to users that optimizes a system welfare (OPT) with that of heuristics Uniform (UNI) and Top Popular (TOP), respectively introduced in Sec. 4.2.1 and Sec. 4.2.2. Our goal is to demonstrate that OPT can do a better job compared to the heuristics UNI and TOP.

We define the system welfare using the dissemination func-

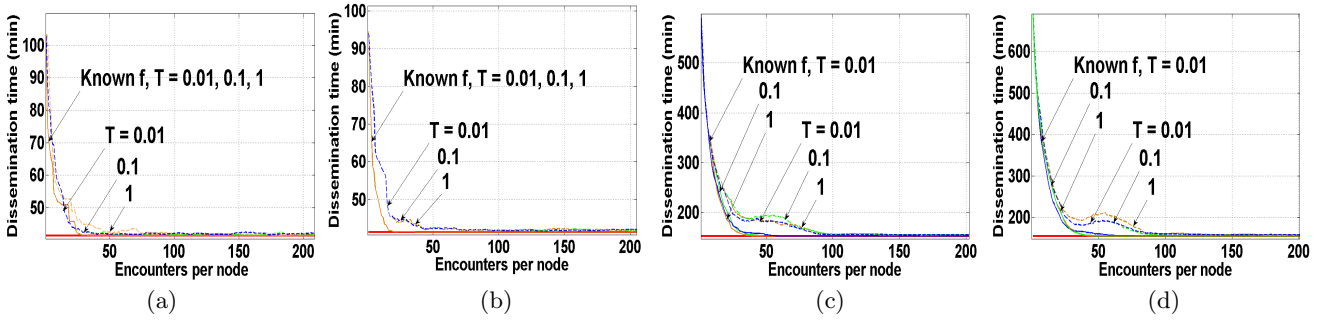


Figure 8: Convergence of the Metropolis-Hastings algorithm under channel-centric system welfare: (a) small-scale, Zipf-2/3, (b) small-scale, Zipf-1, (c) large-scale, Zipf-2/3, (d) large-scale, Zipf-1. Small-scale refers to $(N, J) = (20, 20)$ and the large-scale refers to $(N, J) = (200, 100)$. The y-axis is the mean dissemination time over all channels. The thick horizontal line denotes the system optimum mean dissemination time. Other solid curves are for the Metropolis-Hastings algorithm with the portions of nodes that forward any given channel known (\vec{f}). The dashed lines denote the same but with \vec{f} locally estimated.

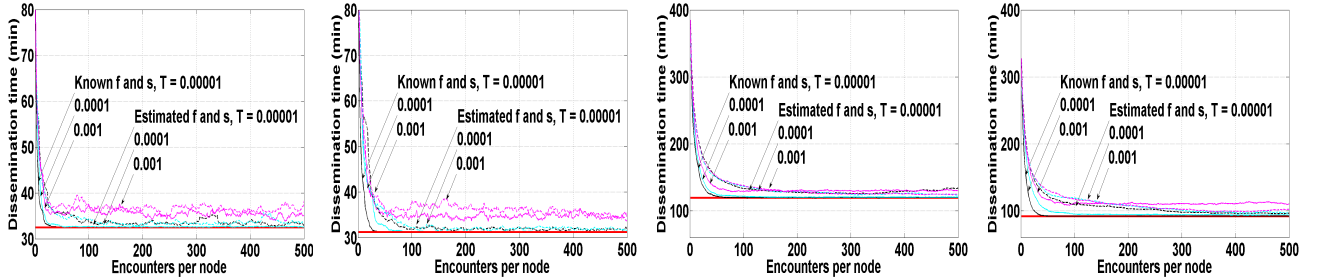


Figure 9: Same as in Fig. 8 but for the user-centric case.




Figure 10: Empirical dissemination curve for the target fraction of nodes $\alpha = 0.25$ from the CAM mobility trace.

tion $t_j(f_j)$ inferred from the mobility trace CAM and letting $V_j(f_j) = -t_j(f_j)$ as in the preceding section. Specifically, we define the logarithm of $t_j(f_j)$ by a concatenation of linear segments that closely follow the empirical data as showed in Fig. 10. While different methods could be used to infer a dissemination curve like that in Fig. 10, we relied on hand-picking which suffices for our purpose. We consider a scenario with $J = 7$ channels, two subscriptions per each user, and two channels helped by each user. We assume that the channel subscription rates follow a Zipf distribution with the scale parameter equal to 1. For each setting of the simulation parameters, we repeat the experiment five times, each time injecting a message of a channel to a user picked uniformly at random from the users who are either subscribers or helpers for the channel at the beginning of the

Table 1: Per-channel and per-user dissemination times in minutes for CAM trace.

Channel-centric	UNI	TOP	OPT
Median	70.2500	133.1000	52.1429
Mean	70.4700	137.1250	57.2000
User-centric	UNI	TOP	OPT
Median	70.4028	97.4528	56.9333
Mean	70.0578	102.7284	59.4089

trace. Note that there are 36 distinct users in the CAM data and that the encounter rate η is equal to 0.001 per second, i.e. 1.2 users every two minutes.

In Table 1 we present the median and mean dissemination time per channel, and per user, for the channel- and user-centric cases, respectively. For both mean and median dissemination time, OPT substantially outperforms UNI and TOP for either channel-centric or user-centric case. In particular, in the channel-centric case, OPT achieves over 70 minutes less dissemination time than TOP and over 10 minutes less dissemination time than UNI for both mean and median dissemination time. In the user-centric case, OPT achieves over 40 minutes less dissemination time than TOP and over 10 minutes less dissemination time than UNI for both mean and median dissemination time. Furthermore, in Fig. 11, we show the mean dissemination time for each channel for channel-centric case (top) and user-centric case (bottom). We discuss the channel-centric case (qualitatively similar conclusions hold for the user-centric case). First, under the channel assignment UNI, some intermediate popular channels may be penalized with a high dissemination time.

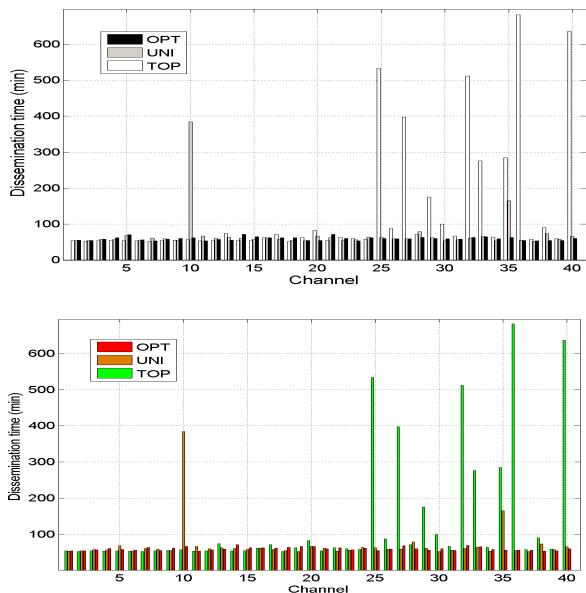


Figure 11: Mean channel dissemination time under CAM mobility with channel-centric system welfare (top) and user-centric system welfare (bottom). Channels are enumerated in decreasing popularity (channel 1 is most popular, etc).

In particular, in Fig. 11(top), we note that the tenth most popular channel gets as much as five hours larger dissemination time than under other channel assignments. Second, same can happen under TOP where the results conform to the expected bias against less popular channels. To be specific, many less popular channels get as much as several hours larger dissemination time than under other channel assignment. The results demonstrate cases where assigning channels by optimizing a system welfare avoids penalizing some channels, which can occur under the heuristics such as UNI or TOP.

8. RELATED WORK

[1] proposes several heuristics for content exchange between devices based on the inferred preference of the user owning a device and that of encountered devices. Each device is assumed to forward an unlimited number of feeds and prioritizes the download of pieces of the content feeds from encountered devices. Feeds subscribed by a device are prioritized over other feeds. In addition, each device uses a solicitation strategy to decide which pieces to fetch from the encountered devices. Specifically, the solicitation strategies considered in [1] include the *most solicited* and *uniform* which essentially correspond to the top popular and uniform channel assignments considered in this paper. The approach in [1] was to evaluate the system performance for a set of solicitation strategies. In this paper, our approach is different—we start with a system welfare objective from which a channel selection strategy follows.

Another related system is CarTorrent [2], a peer-to-peer file sharing tailored for vehicular network scenarios by using an epidemic-style content dissemination. Our work is distinct from that on epidemic-style dissemination in that unlike to previous work, our focus is on efficient dissemination of *multiple* content streams.

A related line of research is that of peer-to-peer storage. [14] modeled a peer-to-peer data sharing system, originally proposed in [15] to enable access to content under limited access to the Internet. [14] studied the performance of various cache policies with limited cache size at individual devices. Several content replication strategies were investigated in [16]. In these systems, nodes query for the content through multiple hops that is supported by the system. Our work has some similarity with that on peer-to-peer storage in that our system welfare amounts to deciding what portion of nodes should “cache” a given channel. Note, however, that our objective is different – our goal is to optimize caching of channels with respect to channel dissemination times that derive from mobility of devices.

Another system welfare problem was recently considered in [9] but for a different problem of optimizing the access rates of mobile devices to a server. The age of single epidemic was recently considered in [17].

Last but not least, we mention the work on characterization of real-world mobility. An early analysis of human mobility was presented in [18] where it was found that the distribution of the inter-contact time between mobile devices decays as a power-law over a range from minutes to a portion of a day. In [7], it was found that this distribution, in fact, is well characterized by a power-law decay with an exponential cut-off. Finally, the authors in [19] studied the diameter of random temporal networks; on the basis of analytical and empirical results, they found that such networks are characterized by small diameter.

9. CONCLUSION

We proposed a framework for optimizing the dissemination of *multiple* information channels in wireless ad-hoc networks. The problem amounts to finding an assignment of users to channels for forwarding the content of channels that optimizes a given system welfare objective. We showed that a system-optimum assignment can be found by a centralized greedy algorithm. Moreover, we proposed a distributed algorithm using the Metropolis-Hasting sampling that stabilizes around the system optimum. We also discussed how to incorporate the battery expenditure of devices into the optimization framework.

There are several interesting directions for future investigation. First, it is of interest to examine the relation between the dissemination time and the fraction of the forwarding nodes across a large set mobility traces. Second, our distributed algorithm involves control over two timescales, a slow timescale for the assignment of users to channels and a fast timescale for the online estimation of the parameters – it is of interest to examine the rates of convergence of the two controls. Third, it may be worth exploring other Metropolis-Hastings samplings for speeding up the convergence and alternative online estimators that are both fast and robust. Forth, it would be important to examine which particular system welfare objective would be of special interest in practice. Fifth, one may analyze the gap between the problems SYSTEM and SYSTEM-R (Appendix). Last but not least, it is of interest to consider the system welfare problem proposed in this paper in scenarios where the dissemination time of a channel depends not only on the number of the nodes that forward the channel, but also on which nodes in particular are the forwarding nodes.

10. REFERENCES

- [1] V. Lenders, G. Karlsson, and M. May. Wireless Ad Hoc Podcasting. In *IEEE SECON*, 2007.
- [2] K. C. Lee, S. H. Lee, R. R. Cheung, U. Lee, and M. Gerla. First Experience with CarTorrent in a Real Vehicular Ad Hoc Network Testbed. *Mobile Networking for Vehicular Env.*, pages 109–114, 2007.
- [3] A. Federgruen and H. Groenevelt. The greedy procedure for resource allocation problems: Necessary and sufficient conditions for optimality. *Operations Research*, 34(6):909–918, 1986.
- [4] Microsoft. Zune social. www.zune.net, December 2008.
- [5] J. Scott, R. Gass, J. Crowcroft, P. Hui, C. Diot, and A. Chaintreau. CRAWDAD data set cambridge/haggle (v. 2006-09-15). Downloaded from <http://crawdad.cs.dartmouth.edu/cambridge/haggle>, September 2006.
- [6] M. Piorkowski, N. Sarafijanovic-Djukic, and M. Grossglauser. CRAWDAD data set epfl/mobility (v. 2009-02-24). Downloaded from <http://crawdad.cs.dartmouth.edu/epfl/mobility>, February 2009.
- [7] T. Karagiannis, J.-Y. Le Boudec, and M. Vojnović. Power-Law and Exponential-Decay of Inter-Contact Time Between Mobile Devices. In *ACM MOBICOM*, 2007.
- [8] F. P. Kelly, A. K. Maulloo, and D. K. H. Tan. Rate Control for Communication Networks: Shadow Prices, Proportional Fairness and Stability. *Journal of the Operational Research Society*, 49:237–252, 1998.
- [9] S. Ioannidis, A. Chaintreau, and L. Massoulié. Optimal and Scalable Distribution of Content Updates over a Mobile Social Network. In *IEEE INFOCOM*, 2009.
- [10] B. Korte and J. Vygen. *Combinatorial Optimization: Theory and Algorithms*. Springer, 3 edition, 2000.
- [11] N. Megiddo. Optimal flows in networks with multiple sources and sinks. *Mathematical Programming*, 7(1):97–107, 1974.
- [12] P. Brémaud. *Markov Chains: Gibbs Fields, Monte Carlo Simulation and Queues*. Springer, 1999.
- [13] B. Kauffmann, F. Baccelli, A. Chaintreau, V. Mhatre, K. Papagiannaki, and C. Diot. Measurement-Based Self Organization of Interfering 802.11 Wireless Access Networks. In *IEEE INFOCOM*, 2007.
- [14] C. Lindemann and O. P. Waldhorst. Modeling epidemic information dissemination on mobile devices with finite buffer. In *ACM Sigmetrics*, 2005.
- [15] M. Papadopouli and H. Schulzrinne. Effects of power conservation, wireless coverage and cooperation on data dissemination among mobile devices. In *ACM Mobihoc*, 2001.
- [16] E. Cohen and S. Shenker. Replication strategies in unstructured peer-to-peer networks. In *ACM Sigcomm*, 2002.
- [17] A. Chaintreau, J.-Y. Le Boudec, and N. Ristanović. The age of gossip: Spatial mean-field regime. In *ACM Sigmetrics*, 2009.
- [18] A. Chaintreau, P. Hui, J. Scott, R. Gass, J. Crowcroft, and C. Diot. Impact of human mobility on opportunistic forwarding algorithms. *IEEE Trans. on Mobile Computing*, 6(6):606–620, 2007.
- [19] A. Chaintreau, A. Mtibaa, L. Massoulié, and C. Diot. The diameter of opportunistic mobile networks. In *ACM CoNEXT*, 2007.
- [20] Hal R. Varian. *Microeconomic Analysis*. W. W. Norton & Company Ltd., 3 edition, 1992.

APPENDIX

A. PROOF OF PROPOSITION 3.1

Let $t(f) = \eta t_j(f)$, $\lambda = \lambda_j/\eta$, $K = K_j$, $\phi(0) = \phi_j(0)$. From Eq. (6), we have

$$t(f) = \frac{1}{\lambda + f} \log(af + b)$$

where

$$\begin{aligned} a &= \frac{K}{(1-K)(\lambda + \phi(0))} \\ b &= \frac{\lambda + \phi(0)(1-K)}{(1-K)(\lambda + \phi(0))}. \end{aligned} \quad (21)$$

Now, note

$$\begin{aligned} \frac{d^2}{df^2} t(f) &= \frac{2}{(\lambda + f)^3} \log(af + b) \\ &\quad - \frac{2a}{(\lambda + f)^2} \frac{1}{af + b} - \frac{a^2}{\lambda + f} \frac{1}{(af + b)^2}. \end{aligned}$$

It follows that $(d^2/df^2)t(f) > 0$, i.e. $t(f)$ is strictly convex, if and only if

$$\log(af + b) - a \frac{\lambda + f}{af + b} - \frac{1}{2} \left(a \frac{\lambda + f}{af + b} \right) > 0. \quad (22)$$

Let

$$x = a \frac{\lambda + f}{af + b}. \quad (23)$$

From the last identity, we have

$$af + b = \frac{b - a\lambda}{1 - x}.$$

Note that it holds

$$b - a\lambda > 1 \quad (24)$$

$$x < 1 \quad (25)$$

Indeed, from Eq. (21) we have

$$b - a\lambda = 1 + \frac{K\lambda}{(1-K)(\lambda + \phi(0))} > 1.$$

Thus, Eq. (24) follows. From Eq. (23) and Eq. (21),

$$x = \frac{\lambda + f}{\frac{\lambda + (1-K)\phi(0)}{K} + f}.$$

Combining with the fact $K < 1$ (as $\alpha < 1$), Eq. (25) follows.

Eq. (22) can be rewritten as

$$\log \frac{b - a\lambda}{1 - x} - x - \frac{1}{2}x^2 > 0$$

i.e.

$$\frac{1 - x}{b - a\lambda} < e^{-x(1 + \frac{1}{2}x)}.$$

Under Eq. (24), the last inequality is implied by

$$1 - x < e^{-x(1 + \frac{1}{2}x)}$$

but this indeed holds as by the mean-value theorem, $h(x) = e^{-x(1+\frac{1}{2}x)}$ satisfies

$$h(x) = 1 - x + \frac{1}{2}h''(x^*)$$

for some $0 < x^* < x$ and the readily checked property that $h''(x) > 0$, for any $x > 0$.

B. RELAXED SYSTEM

We get some insight in how the system forwarding capacity is assigned over channels, for given channels subscription rates $\vec{s} = (s_1, \dots, s_J)$, by considering the following relaxation of SYSTEM-H.

<p>SYSTEM-R</p> <p style="margin-left: 40px;">maximize $\sum_{j=1}^J w_j V_j(s_j + h_j)$</p> <p style="margin-left: 40px;">over $h_j \in [0, 1 - s_j]$</p> <p style="margin-left: 40px;">subject to $\sum_{j=1}^J h_j \leq c - s$</p>
--

Here c and s are the per node capacity and per node number of subscriptions, respectively,

$$c = \frac{1}{N} \sum_{u=1}^N C_u \text{ and } s = \frac{1}{N} \sum_{u=1}^N |S(u)|.$$

SYSTEM-R is obtained by removing all the constraints in SYSTEM-H but those for which $|A| = 1$ or $|A| = J$ in the definition of the polyhedron $P(v)$. Moreover, the portion of user that help any given channel is relaxed to take fractional values – this is a good approximation for systems with a large number of users, i.e. large N . For each user u , we only retain (i) the constraint that u can be assigned to a channel j only if u is not subscribed to channel j , i.e. $j \in \mathcal{J} \setminus S(u)$ and (ii) the user capacity constraint, i.e. that a user u can help only $C_u - |S(u)|$ channels. We have the following result.

PROPOSITION B.1. *The optimum value of SYSTEM is less than or equal to that of SYSTEM-R.*

While we believe that in many cases the system welfare under SYSTEM-R would be equal to that under SYSTEM, we leave for future study to examine the gap between the two problems. In the sequel, we consider the solution to SYSTEM-R.

PROPOSITION B.2. *There is a unique solution to SYSTEM-R that satisfies the following:*

$$h_j(\mu) = \begin{cases} 0 & D_j(\mu/w_j) \leq s_j \\ D_j(\mu/w_j) - s_j & s_j < D_j(\mu/w_j) < 1 \\ 1 - s_j & 1 \leq D_j(\mu/w_j) \end{cases} \quad (26)$$

where $D_j(\cdot)$ is the inverse of $V_j'(\cdot)$ and μ is the solution of

$$\sum_{j=1}^J h_j(\mu) = c - s. \quad (27)$$

The solution can be interpreted as follows. The function $D_j(x)$ can be interpreted as the *demand* for the forwarding capacity for channel j , given a shadow price x .³ The shadow price for a channel j is equal to μ/w_j . If the demand by a channel is less than or equal to the forwarding capacity provided by the subscribers of this channel, then the channel gets no extra forwarding capacity. If the demand of a channel is greater than what can be supported by the system (all users forward the channel), then all users forward this channel. Otherwise, the portion of users that forward the channel is equal to the demand for this channel. See Figure 12 for an illustration.

PROOF. SYSTEM-R is a convex optimization problem and thus has a unique global optima. The dual problem is $\min_{\mu \geq 0} F(\mu)$, where

$$F(\mu) = \max_{\vec{h} \in \mathcal{H}} F(\vec{h}, \mu) \quad (28)$$

$$F(\vec{h}, \mu) = \sum_{j=1}^J w_j V_j(s_j + h_j) - \mu \left(\sum_{j=1}^J h_j - (c - s) \right)$$

$$\mathcal{H} = \{x \in [0, 1]^J : x_j \leq 1 - s_j, j = 1, \dots, J\}.$$

The problem (28) separates into the following optimization problems, for $j = 1, \dots, J$,

CHANNEL- j

$$\begin{aligned} &\text{maximize} && w_j V_j(s_j + h_j) - \mu h_j \\ &\text{over} && h_j \in [0, 1 - s_j] \end{aligned}$$

For each j , CHANNEL- j is a convex optimization problem so there is a unique optimum solution h_j . The objective is a concave function with h_j . Let us first consider the optimization without the constraint $h_j \in [0, 1 - s_j]$. The necessary and sufficient optimality condition is

$$w_j V_j'(s_j + h_j) = \mu.$$

Hence, $s_j + h_j = V_j'^{-1}(\mu/w_j) := D_j(\mu/w_j)$. In case, $h_j \in (0, 1 - s_j)$ we have that the solution to SYSTEM-R satisfies $h_j = D_j(\mu/w_j) - s_j$. The condition $h_j \in (0, 1 - s_j)$ can be rewritten as $s_j < D_j(\mu/w_j) < 1$. In other cases for the unconstrained problem we either have $h_j \leq 0$ or $1 - s_j \leq h_j$ and hence the solution to SYSTEM-R is 0 and $1 - s_j$, respectively. This concludes the proof. \square

B.0.1 Properties of the Solution

We consider how the subscription popularity of channels relates to the allocated forwarding capacity for the solution to SYSTEM-R. Specifically, we identify cases when the bias is to forwarding either most popular or least popular channels.

Let us introduce an enumeration of channels so that the following holds

$$w_1 V_1'(s_1) \geq w_2 V_2'(s_2) \geq \dots \geq w_J V_J'(s_J). \quad (29)$$

We argue that in the following sense this order of channels corresponds to a *priority order* – if a set of k channels is allocated zero helper capacity, then this must be the channels $J - k + 1, \dots, J$. To see this, note from Proposition B.2 that the optimum allocation \vec{h} can be obtained by the following water filling algorithm:

³The concepts of a demand function and a shadow price are rather intuitive and are standard in the microeconomic analysis [20].




Figure 12: The helper capacity allocated to channel j versus the shadow price under SYSTEM-R.

FIND SYSTEM-R OPT: Initialize $\mu = w_1 V_1'(s_1)$.
Continuously decrease μ with unit rate until equality in (27) is attained.

Applying this algorithm, we note that by (29) and the definition of the functions in Eq. (26), we have the asserted property of the priority order.

We discuss two special cases. First, consider symmetric channel-centric system welfare with the weights and the channel utility functions equal for all channels. Under the assumption that the channel utility functions are concave, we have that the priority order is in nonincreasing channel subscription order, i.e. $s_1 \leq s_2 \leq \dots \leq s_J$. This can be seen as a consequence of the concavity of the channel utility functions, i.e. the diminishing returns with the allocated forwarding capacity. Second, we consider symmetric user-centric system welfare, where all the channel utilities are assumed to be identical. Note that in (29) we have $\bar{w} = \bar{s}$ and $s_j V_j'(s_j) = s_j V_1'(s_j)$ for each channel j . We will show that channels are prioritized in a way that depends on the elasticity of the function $xV_1'(x)$, which we discuss in the following.

PROPOSITION B.3. *Suppose $xV_1'(x)$ is a nonincreasing (resp. nondecreasing) function of x , then the priority order is a nondecreasing (resp. nonincreasing) channel subscription order i.e. $s_1 \leq \dots \leq s_J$ (resp. $s_1 \geq \dots \geq s_J$).*

PROOF. Indeed, if $xV_1'(x)$ is nonincreasing (resp. nondecreasing) with x , then for any $i, j \in \mathcal{J}$ such that $s_i V_1'(s_i) \geq s_j V_1'(s_j)$, we have $s_i \leq s_j$ (resp. $s_i \geq s_j$). \square

Let us define the elasticity of the function $xV_1'(x)$ by

$$e(x) = \frac{xV_1''(x)}{V_1'(x)} = \frac{\frac{dV_1'(x)}{V_1'(x)}}{\frac{dx}{x}}, \quad x \geq 0.$$

This is a standard measure that indicates the relative change of the function with the relative change of x .

Saying that $xV_1'(x)$ is a nonincreasing function of x is equivalent to saying that $e(x) \leq -1$ for all $x \geq 0$. Likewise, for $xV_1'(x)$ nondecreasing with x we have the correspondence $e(x) \geq -1$ for all $x \geq 0$. We thus have that if the function $xV_1'(x)$ is sufficiently elastic then the less popular channels are prioritized. In contrast, if the function $xV_1'(x)$ is sufficiently flat then the more popular channels are prioritized.

C. PROOF OF PROPOSITION 4.5

Proof follows from the derivation formulae

$$\begin{aligned} \frac{d}{dt_j} U_j^{PFN} &= \frac{dV_j^{PFN}}{df_j} \frac{df_j}{dt_j} \\ \frac{d^2}{dt_j^2} U_j^{PFN} &= \frac{d^2 V_j^{PFN}}{df_j^2} \left(\frac{df_j}{dt_j} \right)^2 + \frac{dV_j^{PFN}}{df_j} \frac{d^2 f_j}{dt_j^2} \\ &= \frac{d^2 V_j^{PFN}}{df_j^2} \left(\frac{df_j}{dt_j} \right)^2 - \frac{dV_j^{PFN}}{df_j} \frac{d^2 t_j}{df_j^2} \left(\frac{df_j}{dt_j} \right)^3. \end{aligned}$$

The result follows under the assumptions that $t_j(f_j)$ is a nonincreasing convex function of f_j and that $V_j^{PFN}(f_j)$ is a nondecreasing, convex function of f_j .

D. PROOF OF PROPOSITION 6.1

A node u at encounter of a node v replaces a channel $j \in H(u, x)$, selected uniformly at random, with a channel j' picked uniformly at random from $F(v, x) \setminus F(u, x)$. This amounts to the following transition probabilities Q for the candidate configuration change (we consider that no step of the algorithm is performed if the condition $F(u, x) \subset F(v, x)$ is fulfilled):

$Q(x, x') = 0$ if $x' \neq x - 1^{u,j} + 1^{u,j'}$ for any u, j and $j' \neq j$.
Else if $x' = x - 1^{u,j} + 1^{u,j'}$ for some u, j and $j' \neq j$ then

$$\begin{aligned} Q(x, x') &= \frac{1}{N} \frac{1_{j \in H(u, x)}}{|H(u, x)|} \frac{1}{N} \sum_{v \neq u} \frac{1_{j' \in F(v, x)}}{|F(v, x) \setminus F(u, x)|} \\ Q(x', x) &= \frac{1}{N} \frac{1_{j' \in H(u, x')}}{|H(u, x')|} \frac{1}{N} \sum_{v \neq u} \frac{1_{j \in F(v, x')}}{|F(v, x') \setminus F(u, x')|} \end{aligned}$$

Note the following relations

$$F(u, x') = [F(u, x) \setminus \{j\}] \cup \{j'\} \quad (30)$$

$$F(v, x') = F(v, x), \quad v \neq u \quad (31)$$

For $v \neq u$ and $j \in F(v, x)$, we have

$$|F(v, x') \setminus F(u, x')| = \begin{cases} |F(v, x) \setminus F(u, x)| & \text{if } j' \in F(v, x) \\ |F(v, x) \setminus F(u, x)| + 1 & \text{else.} \end{cases} \quad (32)$$

We thus have,

$$Q(x', x) = \frac{1}{N} \frac{1_{j' \in H(u, x')}}{|H(u, x')|} \frac{1}{N} \sum_{v \neq u} \frac{1_{j \in F(v, x)}}{|F(v, x) \setminus F(u, x)| + 1_{j' \notin F(v, x)}}.$$

which shows Eq. (16).

E. APPROXIMATION IN EQUATION (18)

We can re-write Eq.(16) as $\frac{Q(x', x)}{Q(x, x')} = \frac{a}{b}$ with

$$a = \frac{1}{N-1} \sum_{v \neq u} \frac{1_{j \in F(v, x')}}{|F(v, x') \setminus F(u, x')|} \quad (33)$$

$$b = \frac{1}{N-1} \sum_{v \neq u} \frac{1_{j' \in F(v, x)}}{|F(v, x) \setminus F(u, x)|} \quad (34)$$

Now let V be a random variable equal to one user drawn uniformly at random in $\mathcal{U} \setminus \{u\}$:

$$\begin{aligned} a &= \mathbb{E} \left(\frac{1}{|F(V, x') \setminus F(u, x')|} 1_{j \in F(V, x')} \right) \\ &= \mathbb{E} \left(\frac{1}{|F(V, x') \setminus F(u, x')|} \Big| j \in F(V, x') \right) \mathbb{P}(j \in F(V, x')) \\ b &= \mathbb{E} \left(\frac{1}{|F(V, x) \setminus F(u, x)|} \Big| j' \in F(V, x) \right) \mathbb{P}(j' \in F(V, x)) \end{aligned}$$

Note that for large N

$$\mathbb{P}(j \in F(V, x')) = \frac{f_j N - 1}{N - 1} \approx f_j \quad (35)$$

$$\mathbb{P}(j' \in F(V, x)) = \frac{f_{j'} N}{N - 1} \approx f_{j'}, \quad (36)$$

N . Further, under enough mixing, we conjecture that

$$\begin{aligned} & \mathbb{E} \left(\frac{1}{|F(V, x') \setminus F(u, x')|} \middle| j \in F(V, x') \right) \\ & \approx \mathbb{E} \left(\frac{1}{|F(V, x) \setminus F(u, x)|} \middle| j' \in F(V, x) \right) \end{aligned}$$

which gives Eq. (17).

F. PROOF OF PROPOSITION 6.2

First, note

$$\begin{aligned} V(x') - V(x) &= w_{j'} [V_{j'}(f_{j'}(x')) - V_{j'}(f_{j'}(x))] \\ &\quad - w_j [V_j(f_j(x')) - V_j(f_j(x))]. \end{aligned}$$

Second, we have $f_{j'}(x') = f_{j'}(x) + \frac{1}{N}$ and $f_j(x') = f_j(x) - \frac{1}{N}$. Thus, we have

$$\begin{aligned} N[V_{j'}(f_{j'}(x')) - V_{j'}(f_{j'}(x))] &\rightarrow V_{j'}'(f_{j'}(x)) \\ N[V_j(f_j(x')) - V_j(f_j(x))] &\rightarrow V_j'(f_j(x)) \end{aligned}$$

as N tends to infinity. The result follows.