



HAL
open science

How to Optimally Allocate Your Budget of Attention in Social Networks

Bo Jiang, Nidhi Hegde, Laurent Massoulié, Don Towsley

► **To cite this version:**

Bo Jiang, Nidhi Hegde, Laurent Massoulié, Don Towsley. How to Optimally Allocate Your Budget of Attention in Social Networks. IEEE INFOCOM 2013 - IEEE Conference on Computer Communications, Apr 2013, Torino, Italy. hal-00839599

HAL Id: hal-00839599

<https://inria.hal.science/hal-00839599>

Submitted on 28 Jun 2013

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

How to Optimally Allocate Your Budget of Attention in Social Networks

Bo Jiang*, Nidhi Hegde[†], Laurent Massoulié[‡], Don Towsley*

* University of Massachusetts Amherst, Email: {bjiang, towsley}@cs.umass.edu

[†] Technicolor Paris Research Lab, Email: nidhi.hegde@technicolor.com

[‡] Microsoft Research - INRIA Joint Centre, France, Email: laurent.massoulie@inria.fr

Abstract—We consider the performance of information propagation through social networks in a scenario where each user has a *budget of attention*, that is, a constraint on the frequency with which he pulls content from neighbors. In this context we ask the question “when users make selfish decisions on how to allocate their limited access frequency among neighbors, does information propagate efficiently?” For the metric of average propagation delay, we provide characterizations of the optimal social cost and the social cost under selfish user optimizations for various topologies of interest. Three situations may arise: well-connected topologies where delay is small even under selfish optimization; tree-like topologies where selfish optimization performs poorly while optimal social cost is low; and “stretched” topologies where even optimal social cost is high. We propose a mechanism for incentivizing users to modify their selfish behaviour, and observe its efficiency in the family of tree-like topologies mentioned above.

I. INTRODUCTION

Information dissemination has been transformed by the emergence of online social networks and their enthusiastic adoption by users. Users rely on trust relationships in social networks for accessing information. Relationships form on the basis of the quality of information received, and in turn determine the speed of propagation in the network.

The study of the propagation of information, or a *rumor*, in a social network goes back more than two decades [1]. Previous work has mostly studied the propagation of a rumor originating at a given source, and the spreading mechanism operating in unit-sized “rounds”, where in each round a user chooses a neighbor at random for propagation. The spreading mechanisms considered have been broadly of three types: push mechanisms where the user sends the rumor if he has it, to the chosen neighbor; pull mechanisms where the users pulls the rumor from the chosen neighbor; and a combined mechanism, where the user pushes the rumor if he has it and pulls it if the chosen neighbor has it.

We consider a more realistic scenario of information propagation where several sources of information in the network publish information. We consider an asynchronous pull model, where each node contacts a neighbor after a random delay and pulls any content available at that neighbor. Further, we assume that each user has a limited *budget of attention*, that

is, a limited frequency with which he pulls content from neighbors. Our objective is to identify optimal allocations of this limited frequency among neighbors for each user in the network. In particular, we want to answer the question “when users make selfish decisions on how to allocate their limited access frequency among neighbors, does information propagate efficiently?”

We investigate the efficiency of selfish allocation by considering the metric of average end-to-end delay of content spread. We make the following contributions:

- We study the efficiency of selfish allocation of the budget of attention by characterizing the *price of stability* (PoS) for several network topologies;
- we identify topologies with inefficient propagation under selfish allocations;
- we propose the “plus-one” mechanism, an incentive scheme that coaxes users into mimicking a gradient-descent algorithm, bringing the cost of content propagation closer to optimality;
- we present numerical results that compare the optimal, selfish, and feedback-based allocations.

The rest of the paper is organized as follows. The next subsection presents an overview of related work. We define the social network model in Section II, and present the analysis of optimal and distributed allocation over several topologies in Section III. In Section IV we present our feedback-based mechanism. Numerical results follow in Section V and we conclude in Section VI.

A. Related Work

The literature on information propagation in social networks, or rumor spreading [1] is wide and varied; we mention here only those that are most relevant to our work. The typical model is the randomized broadcast model [2] which is carried out in synchronized rounds. In each round, each user selects a neighbor at random and pushes or pulls, (or both) a rumor. Most previous work focus on the characterization of the delay in spreading a single rumor to all nodes (some examples are [3], [4]). We consider the asynchronous model where each node contacts a neighbor after a random amount of time; nodes then pull content from the selected neighbor. An asynchronous model is considered in [5]; they focus in particular on random regular graphs and derive performance

This work was partially funded by the European Commission under the FIRE SCAMPI (FP7-IST-258414) project and by the French National Research Agency (ANR) under the PROSE project (ANR-09-VERS-007). In addition, it was supported in part by US NSF grant CNS-0721779 and by ARO under MURI W911NF-08-1-0233.

results in terms of optimal delay. In the present paper, we are interested in studying the efficiency of selfish allocation of a limited frequency of neighbor contacts, as the level of attention affects information flow [6]. In particular our objective is to inform the design of algorithms for optimal information spread.

In fact, it turns out that the allocation of attention among friends in a social network can be observed in reality. Backstrom et al. [7] analyze a data set of real measurements on how users split their attention among their friends on Facebook. They consider activities such as communication and viewing and show how the balance of attention varies across activities and other personal characteristics. Our work is complementary in that we are interested in how a user's attention is split when his objective is timely reception of news. We take the approach of conceiving a general model for studying the balance of attention, and an analysis for several network topologies.

II. SOCIAL NETWORK DESCRIPTION

We consider a social network where each user has a set of friends, or *contacts*, that he links to, or *follows*, in order to get content such as news updates, videos, or other messages. Each user then makes all content that he holds available to his followers, such as on the Facebook wall or Twitter stream. Rather than a single source of content, we allow all users to create content, at a given frequency. Further, we assume that users seek to obtain all content circulating in the network. Users consult their contacts for the latest updates of information with the objective of minimizing the average delay for obtaining all information. As in real online social networks, users have a limited *budget of attention*, that is, the total frequency with which they may consult their contacts is limited. As such, this frequency must be allocated among the contacts in a manner that optimizes for delay. We will compare a centralized optimization of consultation frequencies with a distributed one where users optimize the allocation in a selfish manner.

We model the social network as a directed graph $G = (V, E)$, where V is the set of users, and E is the set of directed edges between users, i.e. $(j, i) \in E$ if and only if there is an edge from j to i . Denote by $N(i)$ the set of in-neighbors of user i , i.e. $N(i) = \{j \in V : (j, i) \in E\}$, and let $d_i = |N(i)|$ be the in-degree of i . We assume that G is strongly connected.

Each user $i \in V$ creates contents according to a Poisson process with rate $\lambda_i > 0$. User i consults his in-neighbor $j \in N(i)$ according to a Poisson process with rate x_{ji} , and pulls all contents available at j . Each user has a limited budget of attention $b_i > 0$, so that the rates of consultation are constrained by $\sum_{j \in N(i)} x_{ji} = b_i$. In terms of the fraction of budget that i devotes to j , $y_{ji} = x_{ji}/b_i$, the budget constraint becomes

$$\sum_{j \in N(i)} y_{ji} = 1. \quad (1)$$

The vector $\mathbf{y}_i = \{y_{ji} : j \in N(i)\}$ represents how user i allocates his budget of attention among his neighbors, which will be referred to as his strategy. The set of all possible strategies of user i is the unit simplex Δ_i in \mathbb{R}^{d_i} . Let $\mathbf{y} = \{\mathbf{y}_i : i \in V\}$ be the strategy profile of the network, and $\Delta = \times_{i \in V} \Delta_i$ is

the set of all possible profiles. Also let \mathbf{y}_{-i} be the strategy profile of all users other than i .

For any $i \neq j$, define $D_{ji}(\mathbf{y})$ to be the delay for user i to receive content originated at user j under profile \mathbf{y} . Define the cost for user i to be

$$C_i(\mathbf{y}) = \frac{1}{\lambda_{-i}} \sum_{j \in V \setminus \{i\}} \lambda_j \mathbb{E} D_{ji}(\mathbf{y}), \quad (2)$$

where $\lambda_{-i} = \sum_{j \in V \setminus \{i\}} \lambda_j$, and define the social cost to be

$$C(\mathbf{y}) = \frac{1}{(n-1)\Lambda} \sum_{i \in V} \sum_{j \in V \setminus \{i\}} \lambda_j \mathbb{E} D_{ji}(\mathbf{y}), \quad (3)$$

where $\Lambda = \sum_{i \in V} \lambda_i$.

Let $C^* = C(\mathbf{y}^*) = \min_{\mathbf{y} \in \Delta} C(\mathbf{y})$ be the optimal social cost, and $\mathbf{y}^* \in \Delta$ any socially optimal profile. When each user minimizes his own cost in a selfish manner, the resulting social cost is in general larger than C^* , and the corresponding profile a Nash equilibrium. In general, there might exist multiple equilibria. We will use $\Gamma \subset \Delta$ to denote the set of profiles that are Nash equilibria.

A standard measure of inefficiency of equilibria is the *price of anarchy* (PoA), which is defined as the ratio of the cost at the worst equilibrium to that of an optimal outcome. In contrast, we are interested in the best equilibria, which would give us a benchmark of what is achievable through distributed means. As such, we focus on the *price of stability* (PoS), defined as the ratio of the cost at the best equilibrium to that of an optimal outcome [8]. Even though the PoS can be seen as a weaker notion of inefficiency, we find it more interesting in a practical sense, since it gives us a target performance for the design of distributed algorithms. More formally, the PoS is defined as follows:

$$\text{PoS} = \min_{\mathbf{y} \in \Gamma} \frac{C(\mathbf{y})}{C(\mathbf{y}^*)} = \frac{\hat{C}}{C^*}, \quad (4)$$

where $\hat{C} = \min_{\mathbf{y} \in \Gamma} C(\mathbf{y})$ is the minimum social cost under a selfish allocation.

In what follows, we will measure inefficiencies in the selfish allocation of the budget of attention in several network topologies. We will show that some topologies lead to large inefficiencies. In Section IV we propose a feedback-based mechanism that results in a distributed allocation that has a cost closer to the optimal cost than does the selfish allocation.

III. EFFICIENCY ANALYSIS

We will now study the efficiency of selfish optimization on several social network topologies. The interest in studying various topologies lies not only in understanding how existing social networks with those topologies perform, but also in identifying efficient structures for information dissemination. The latter may inform smart design for information propagation.

A. Tree Network

We first consider tree topologies. Such structures are interesting as networks since information dissemination can be locally tree-like. Let G be a tree; for all tree networks we study, we will assume G is undirected. Let T_{ji} denote the

component containing j when an edge $(j, i) \in E$ is removed. Let $\lambda_{ji} = \sum_{k \in T_{ji}} \lambda_k$ be the aggregate content creation rate of the nodes in T_{ji} , and $n_{ji} = |T_{ji}|$ the number of nodes in T_{ji} .

For $i \neq j$, let $\mathcal{P}_{j \rightsquigarrow i}$ be the unique shortest path from j to i . The average delay for user i to get contents originated in j is then $\mathbb{E}D_{ji}(\mathbf{y}) = \sum_{e \in \mathcal{P}_{j \rightsquigarrow i}} x_e^{-1}$, where $x_e = x_{uw}$ for an edge $e = (u, w)$. Thus the cost for user i is

$$C_i(\mathbf{y}) = \frac{1}{\lambda_{-i}} \sum_{j: j \neq i} \lambda_j \sum_{e \in \mathcal{P}_{j \rightsquigarrow i}} \frac{1}{x_e} = \frac{1}{\lambda_{-i} b_i} \sum_{k \in N(i)} \frac{\lambda_{ki}}{y_{ki}} + f(\mathbf{y}_{-i}),$$

where $f(\mathbf{y}_{-i})$ represents terms that do not depend on \mathbf{y}_i .

If user i selfishly minimizes C_i , the best rate allocation is given by

$$\hat{y}_{ji} = \frac{\sqrt{\lambda_{ji}}}{\sum_{k \in N(i)} \sqrt{\lambda_{ki}}}, \quad \text{for } j \in N(i). \quad (5)$$

The social cost can be written as follows:

$$\begin{aligned} C &= \frac{1}{(n-1)\Lambda} \sum_{i \in V} \sum_{j: i \neq j} \lambda_j \sum_{e \in \mathcal{P}_{j \rightsquigarrow i}} \frac{1}{x_e} \\ &= \frac{1}{(n-1)\Lambda} \sum_{i \in V} \frac{1}{b_i} \sum_{k \in N(i)} \frac{n_{ik} \lambda_{ki}}{y_{ki}}. \end{aligned}$$

The socially optimal rate allocation is given by:

$$y_{ji}^* = \frac{\sqrt{n_{ij} \lambda_{ji}}}{\sum_{k \in N(i)} \sqrt{n_{ik} \lambda_{ki}}}, \quad \text{for } j \in N(i). \quad (6)$$

The optimal social cost is then as follows:

$$C^* = \frac{1}{(n-1)\Lambda} \sum_{i \in V} \frac{1}{b_i} \left(\sum_{j \in N(i)} \sqrt{n_{ij} \lambda_{ji}} \right)^2, \quad (7)$$

and the cost under selfish allocation is given by:

$$\hat{C} = \frac{1}{(n-1)\Lambda} \sum_{i \in V} \frac{1}{b_i} \left(\sum_{j \in N(i)} n_{ij} \sqrt{\lambda_{ji}} \right) \left(\sum_{j \in N(i)} \sqrt{\lambda_{ji}} \right). \quad (8)$$

We will now study in more detail specific tree structures: the line, the chained star networks, and the k -ary tree.

1) *Line Network*: Suppose G is a line network, with $V = \{1, \dots, n\}$ and $(i, i+1) \in E$ for $i = 1, 2, \dots, n-1$ and $(i, i-1) \in E$ for $i = 2, \dots, n$. Theorem 1 below shows that when the budgets of attention are homogeneous the price of stability is bounded; it is at most five. The proof of this theorem can be found in Appendix A. When the budgets of attention are heterogeneous, the price of stability can be arbitrarily large, as shown in Theorem 2.

Theorem 1. *In a line network of n nodes the PoS is bounded as follows:*

$$\text{PoS} \leq \frac{b_{\max}}{b_{\min}} \min \left\{ 5, 1 + \frac{3\lambda_{\text{ave}}}{2\lambda_{\min}} \right\}, \quad (9)$$

where $\lambda_{\min} = \min_i \lambda_i$, $\lambda_{\max} = \max_i \lambda_i$, $\lambda_{\text{ave}} = \Lambda/n$, $b_{\max} = \max_i b_i$ and $b_{\min} = \min_i b_i$.

Theorem 2. *The price of stability can be arbitrarily large when the budgets of attention are heterogeneous in a line network.*

Proof: Under the assumption of homogeneous content creation rates, set $\lambda_i = 1$. Consider a heterogeneous set of budgets of attention where $b_2 = 1$ and $b_i = \infty$, for all $i \neq 2$. The optimal social cost, under centralized allocation, is given by

$$C^* = \frac{1}{(n-1)n} \left[\sqrt{2(n-2)} + \sqrt{n-1} \right]^2 = \Theta(n^{-1}).$$

The social cost under selfish allocation is

$$\hat{C} = \frac{1}{(n-1)n} (1 + \sqrt{n-2}) (n-1 + 2\sqrt{n-2}) = \Theta(n^{-\frac{1}{2}}).$$

The price of stability, \hat{C}/C^* , is thus $\Theta(\sqrt{n})$. \square

2) *Chained Star Network*: Consider k star networks, each with p nodes. The hubs of the stars are then chained to form a line network, with a total of $n = pk$ nodes. Such topologies are not uncommon in social networks based on communities. Such structure might correspond to communities focused on given topics or interests, that are then connected to the larger social network. Theorem 3 shows that while the socially optimal cost can be large, the PoS is of order 1 in the homogeneous case. The proof is shown in Appendix B.

Theorem 3. *In the chained star network with homogeneous content creation rates and budgets of attention, the price of stability is $\Theta(1)$. Further, the socially optimal cost and cost under selfish allocation are given, respectively, by:*

$$C^* = \Theta(\max\{p, k\}) = \Omega(\sqrt{n}),$$

$$\hat{C} = \Theta(\max\{p, k\}).$$

Remark. As p changes from $\Theta(1)$ to $\Theta(n)$, C^* can have any order between $\Theta(\sqrt{n})$ and $\Theta(n)$. In particular, $C^* = \Theta(n)$ for $p = 1$ and $p = n$, which correspond to the line and star networks, respectively.

3) *k -ary Tree Network*: We now consider rooted trees where each node has k children. Such structures are of interest for social networks with few edges, as sparse random graphs are locally tree-like. Theorem 4 below states the PoS for k -ary trees; its proof can be found in Appendix C. Corollary 5 shows that this PoS can be arbitrarily large for k of constant order and as k scales sublinearly with n .

Theorem 4. *In a k -ary complete tree with n users, the price of stability can be arbitrarily large; it is $\Theta(1 + \frac{\sqrt{n} \log k}{k \log n})$ when content creation rates and budgets of attention are homogeneous.*

Corollary 5. *Let $k = \Theta(n^\alpha)$.*

- 1) *If $\alpha = 0$, i.e. $k = \Theta(1)$, the costs are $C^* = \Theta(\log n)$, $\hat{C} = \Theta(\sqrt{n})$, and $\text{PoS} = \Theta(\sqrt{n}/\log n)$.*
- 2) *If $0 < \alpha < 1/2$, the costs are $C^* = \Theta(n^\alpha)$, $\hat{C} = \Theta(\sqrt{n})$, and $\text{PoS} = \Theta(n^{1/2-\alpha})$.*
- 3) *If $\alpha \geq 1/2$, the costs are $C^* = \Theta(n^\alpha)$, $\hat{C} = \Theta(n^\alpha)$, and $\text{PoS} = \Theta(1)$.*

The above corollary verifies the intuition that long thin networks are less efficient for information propagation than wide networks. The former type of network thus would require an incentive-based mechanism to make them more efficient. We will propose one such mechanism in Section IV.

B. Clique Networks

We now consider the case where G is a clique. This is, in some sense, the best-case scenario, where users have the widest choice possible in allocating their budget of attention. The analysis is more involved since there are many paths between each source-destination pair, creating a more complicated dependency structure between the various links. We consider efficiency for the homogeneous case, where $\lambda_i = 1$ and $b_i = 1$. We will show that there exists a selfish profile that is asymptotically optimal, and thus the price of stability is bounded and approaches one as the network size increases.

Let us consider the *uniform* strategy for user $i \in V$, where the consultation rate is $y_{ji} = u_{ji} = d_i^{-1}$ for all $j \in N(i)$. Let $\mathbf{u}_i = \{u_{ji} : j \in N(i)\}$ and $\mathbf{u}_{-i} = \{u_j : j \in V \setminus \{i\}\}$. The profile $\mathbf{u} = \{\mathbf{u}_i : i \in V\}$ is referred to as the *uniform* profile. For a clique network, $u_{ji} = 1/(n-1)$ for all $i \neq j$.

Theorem 6. *The social cost C^u of the uniform profile is $C^u = H_{n-1} = \sum_{k=1}^{n-1} k^{-1}$, the $(n-1)$ -st harmonic number.*

Proof: We follow the approach in [9]. Let $D_{j(m)}$ be the delay for a content originating from j to reach at least m other users, i.e. $D_{j(m)}$ is the m -th order statistic of $\{D_{ji} : i \in V \setminus \{j\}\}$ for a given j . Let $V_{jk} = D_{j(k)} - D_{j(k-1)}$, with the convention that $D_{j(0)} = 0$. Thus

$$\sum_{i:i \neq j} D_{ji} = \sum_{m=1}^{n-1} D_{j(m)} = \sum_{m=1}^{n-1} \sum_{k=1}^m V_{jk} = \sum_{k=1}^{n-1} (n-k) V_{jk}. \quad (10)$$

Note that V_{jk} is exponentially distributed with parameter $\frac{1}{n-1}k(n-k)$. Taking expectations of (10) we get:

$$\sum_{i:i \neq j} \mathbb{E} D_{ji} = \sum_{k=1}^{n-1} (n-k) \mathbb{E} V_{jk} = \sum_{k=1}^{n-1} \frac{n-1}{k} = (n-1) H_{n-1},$$

where $H_n = \sum_{k=1}^n k^{-1}$ is the harmonic number. Summing over j and dividing by $n(n-1)$, we obtain the social cost $C^u = H_{n-1}$. \square

The next theorem shows that the uniform strategy by a given user is the best response when all other nodes follow the uniform strategy, showing that this is a Nash equilibrium.

Theorem 7. *The uniform profile is a Nash equilibrium.*

Proof: Suppose users $i = 1, 2, \dots, n-1$ follow the uniform strategy and consider user n . For a given j ,

$$D_{jn}(\mathbf{y}_n, \mathbf{u}_{-n}) = \min_{1 \leq k \leq n-1} \left\{ \frac{1}{y_{kn}} E_{kn} + Y_{jk}^n \right\},$$

where $\{E_{kn}\}_{k=1}^{n-1}$ are i.i.d. exponential random variable with parameter 1, and Y_{jk}^n is the time for an item originating from j to reach k without passing through n . Let $\mathbf{T}_j^n = \{T_{ji}^n\}_{i=1}^{n-1}$ be the order statistics of $\mathbf{Y}_j^n = \{Y_{jk}^n\}_{k=1}^{n-1}$. By symmetry, the

random vectors $(Y_{j1}^n, \dots, Y_{j,n-1}^n)$ and $(T_{j\sigma(1)}^n, \dots, T_{j\sigma(n-1)}^n)$ are identically distributed, where σ is a permutation chosen uniformly randomly from the symmetric group S_{n-1} , independently of \mathbf{T}_j^n . Therefore,

$$\begin{aligned} & \mathbb{P}(D_{jn}(\mathbf{y}_n, \mathbf{u}_{-n}) > x \mid \mathbf{T}_j^n) \\ &= \mathbb{P} \left\{ \bigcap_{k=1}^{n-1} \{E_{kn} > y_{kn}(x - T_{j\sigma(k)}^n)\} \mid \mathbf{T}_j^n \right\} \\ &= \mathbb{E}_\sigma \left\{ \exp \left\{ - \sum_{k=1}^{n-1} y_{kn}(x - T_{j\sigma(k)}^n)^+ \right\} \right\} \\ &\geq \exp \left\{ - \sum_{k=1}^{n-1} y_{kn} \mathbb{E}_\sigma(x - T_{j\sigma(k)}^n)^+ \right\} \\ &= \exp \left\{ - \sum_{k=1}^{n-1} y_{kn} \sum_{l=1}^{n-1} \frac{1}{n-1} (x - T_{jl})^+ \right\} \\ &= \exp \left\{ - \sum_{l=1}^{n-1} \frac{1}{n-1} (x - T_{jl})^+ \right\} \\ &= \mathbb{P}(D_{jn}(\mathbf{u}_n, \mathbf{u}_{-n}) > x \mid \mathbf{T}_j^n), \end{aligned}$$

where the inequality follows from Jensen's inequality. Therefore, for $j = 1, 2, \dots, n-1$,

$$\mathbb{P}\{D_{jn}(\mathbf{y}_n, \mathbf{u}_{-n}) > x\} \geq \mathbb{P}\{D_{jn}(\mathbf{u}_n, \mathbf{u}_{-n}) > x\},$$

and hence $C_n(\mathbf{y}_n, \mathbf{u}_{-n}) \geq C_n(\mathbf{u}_n, \mathbf{u}_{-n})$ for any $\mathbf{y}_n \in \Delta_n$. Therefore, the uniform profile is a Nash equilibrium. \square

The next theorem shows that the cost of the uniform profile is larger than the optimal social cost by at most one. Thus the price of stability approaches one as the network size increases.

Theorem 8. *In a network G of size n where every user publishes at rate $\lambda = 1$ and has a budget of attention $b = 1$, the cost for any user i is lower bounded by*

$$C_i \geq \frac{n}{n-1} H_{n-1} - 1.$$

Proof: The argument uses a *backwards growth process* similar to that used in Section V of [10]. Assume the process is in steady state, i.e. the process started at $-\infty$. Let $B_i(t)$ be the set of users whose states at time $-t$ reach user i by time 0. Note that $B_i(0) = \{i\}$. By stationarity and the independence of the publishing and consulting processes, $\mathbb{P}\{D_{ji} > t\} = \mathbb{P}\{j \notin B_i(t)\}$. Now let $A_{ji} = \inf\{t : j \in B_i(t)\}$. Note that $\{j \notin B_i(t)\} = \{A_{ji} > t\}$. Thus $\mathbb{P}\{D_{ji} > t\} = \mathbb{P}\{A_{ji} > t\}$, i.e. D_{ji} and A_{ji} are identically distributed. Now let $A_{(k)i} = \inf\{t : |B_i(t)| = k + 1\}$. Note that $\{A_{(k)i} : 1 \leq k \leq n-1\}$ are the order statistics of $\{A_{ji} : j \in \{1, \dots, n\} \setminus \{i\}\}$.

If we reverse the arrow of time at time 0, then $B_i(t)$ is the set of infected users at time s in the SI epidemic spreading model where an infected user j contaminates a susceptible user ℓ at rate $y_{\ell j}$. It then follows that $W_{ki} \triangleq A_{(k)i} - A_{(k-1)i}$ is exponentially distributed with parameter μ_{ki} given by

$$\mu_{ki} = \sum_{\substack{j \in B_i(A_{(k-1)i}) \\ \ell \notin B_i(A_{(k-1)i})}} y_{\ell j} \leq |B_i(A_{(k-1)i})| = k,$$

and hence $\mathbb{E}W_{ki} \geq 1/k$. It follows that

$$\begin{aligned} C_i &= \frac{1}{n-1} \sum_{j:j \neq i} \mathbb{E}D_{ji} = \frac{1}{n-1} \sum_{j:j \neq i} \mathbb{E}A_{ji} \\ &= \frac{1}{n-1} \sum_{k=1}^{n-1} \mathbb{E}A_{(k)i} = \frac{1}{n-1} \sum_{k=1}^{n-1} \sum_{\ell=1}^k \mathbb{E}W_{k\ell} \\ &\geq \frac{1}{n-1} \sum_{k=1}^{n-1} \sum_{\ell=1}^k \frac{1}{\ell} = \frac{n}{n-1} H_{n-1} - 1. \end{aligned}$$

□

C. Expander Network

We now consider a network characterized by an expander graph, which might be considered more realistic. An expander graph is a sparse graph with strong connectivity properties. An example is a d -regular graph which is often used in modeling social networks.

Finding a Nash equilibrium when the topology is an expander graph turns out to be quite complex. We thus consider *approximate* Nash equilibria. We define a user's strategy to be an ϵ -approximate NE if the cost to the user under this strategy is no worse than ϵ more than the cost of any other strategy [8]. More formally, the profile $\hat{\mathbf{y}}$ is an ϵ -approximate NE if for any k and any $\tilde{\mathbf{y}}_k \in \Delta_k$,

$$C_k(\hat{\mathbf{y}}) \leq \epsilon + C_k(\tilde{\mathbf{y}}_k, \hat{\mathbf{y}}_{-k}).$$

The ϵ -approximate price of stability is defined by (4) with Γ replaced by Γ_ϵ , the set of ϵ -approximate NE.

Suppose G is an expander network with bounded degree. We now show that the profile where users implement a uniform allocation of their budget of attention is an approximate NE.

Theorem 9. *The uniform profile is a d -approximate Nash equilibrium, where $d = \max_k d_k$ is the maximum degree of the graph.*

Proof: Consider user ℓ . For a given j ,

$$D_{j\ell}(\mathbf{y}_\ell, \mathbf{y}_{-\ell}) = \min_{k \in N(\ell)} \left\{ \frac{1}{y_{k\ell}} E_{k\ell} + Y_{jk}^\ell \right\},$$

where Y_{jk}^ℓ is the time for an item originating from j to reach k without passing through ℓ . Define a random variable K by $K = \min\{k^* : Y_{jk^*}^\ell = \min_k Y_{jk}^\ell\}$. Since $\{Y_{jk}^\ell\}$ are independent of $\{E_{k\ell}\}$, so is K . Thus

$$\mathbb{E}\{E_{K\ell}\} = \sum_k \mathbb{E}\{E_{k\ell}\} \mathbb{P}\{K = k\} = 1.$$

We then have the following;

$$\begin{aligned} \mathbb{E}D_{j\ell}(\mathbf{u}_\ell, \mathbf{y}_{-\ell}) &= \mathbb{E} \left\{ \min_{k \in N(\ell)} \{d_\ell E_{k\ell} + Y_{jk}^\ell\} \right\} \\ &\leq \mathbb{E} \{d_\ell E_{K\ell} + Y_{jK}^\ell\} = d_\ell + \mathbb{E} \left\{ \min_{k \in N(\ell)} Y_{jk}^\ell \right\} \\ &\leq d_\ell + \mathbb{E}D_{j\ell}(\mathbf{y}_\ell, \mathbf{y}_{-\ell}). \end{aligned}$$

Summing over j and then minimizing over \mathbf{y}_ℓ , we obtain

$$C_\ell(\mathbf{u}_\ell, \mathbf{y}_{-\ell}) \leq d_\ell + \min_{\mathbf{y}_\ell \in \Delta_\ell} C_\ell(\mathbf{y}_\ell, \mathbf{y}_{-\ell}),$$

for any $\mathbf{y}_{-\ell}$. □

In view of Theorem 8, the next theorem shows that, for an expander network with edge expansion bounded away from zero, the uniform profile is order optimal, and hence the d -approximate price of stability is bounded.

Theorem 10. *The social cost C^u of the uniform profile is bounded by*

$$C^u \leq \frac{2d}{h_G} H_{\lfloor n/2 \rfloor},$$

where $d = \max_i d_i$, and h_G is the edge expansion of G defined by

$$h_G = \min_{|A| \leq \lfloor n/2 \rfloor} \frac{|\partial A|}{\min\{|A|, |A^c|\}},$$

with $\partial A = \{(u, v) \in E : u \in A, v \in A^c\}$.

Proof: The proof is essentially the same as that for Theorem 6. The difference is that the exponential random variable V_{jk} in (10) now has parameter μ_{jk} given by

$$\mu_{jk} = \sum_{u \in F_{jk}, v \in F_{jk}^c} y_{uv} \geq \frac{1}{d} |\partial F_{jk}| \geq \frac{h_G}{d} \min\{k, n-k\},$$

where F_{jk} is the set consisting of the first k users that has got the content originating from j , with $F_{j1} = \{j\}$. Hence

$$\sum_{i:i \neq j} \mathbb{E}D_{ji} \leq \sum_{k=1}^{n-1} \frac{(n-1)d}{h_G \min\{k, n-k\}} \leq \frac{2d(n-1)}{h_G} H_{\lfloor n/2 \rfloor}.$$

The desired result follows by summing over j and dividing by $n(n-1)$. □

IV. INCENTIVIZING EFFICIENT BEHAVIOR

As we have just seen, topologies can be classified according to their performance under optimal and selfish allocations as

- A. **Efficient:** These topologies have bounded PoS and socially optimal cost of order $\log n$; they do not require additional mechanisms. Examples are expanders and cliques.
- B. **Inefficient amenable:** These topologies have high PoS yet low (logarithmic) optimal social cost. As we shall show, incentive mechanisms can be introduced to change users' behavior and reduce their otherwise inefficient performance. Examples are the k -ary tree with bounded degree ($k = \Theta(1)$), and with low-scaling degree ($k = \Theta(n^\alpha)$, $\alpha \ll 1/2$).
- C. **Inefficient suboptimal:** These topologies show inefficient content spread even under socially optimal allocations. No mechanism that preserves the topology and the budgets of interest can therefore lead to good performance. Examples are line and star networks, and k -ary trees with high-scaling degree ($k = \Theta(n^\alpha)$, α close to $1/2$ and $\alpha > 1/2$).

We now propose an incentive mechanism that will prove particularly appealing for inefficient amenable graphs.

A. The Plus-One mechanism

Our mechanism relies on using incentives as a form of feedback for reallocating attention. Each receiver k , upon receiving useful information, sends a reward of 1, that we call a +1, to each node involved in relaying this information from its source s . Note that by useful information, we mean that the piece of information that arrived earliest at r , among all copies of this information at r . Therefore, a node sends a +1 to the neighbor through which the earliest copy was received. We now provide details of this mechanism.

- Each receiver node k , upon reception of useful information from source s , sends a +1 $_k$ to each node along the path to s that was involved in relaying that piece of information. Each node i along this path then keeps a score $O_j^i = \sum_{s,k:(j,i) \in \mathcal{P}_{s \rightsquigarrow k}^*} +1_k$, where $\mathcal{P}_{s \rightsquigarrow k}^*$ is the quickest (shortest) path from s to k .
 - The receiver is not required to know the topology of the network nor the path to each source. A completely distributed implementation consists in each receiver k sending a +1 $_k$ to the neighbor through which it received the useful information from s . Each node along the path would then aggregate the +1s it receives along with its own +1 before sending the sum up to its neighbor.
- At time intervals that are much longer than the slots over which +1s are sent, each user i updates his allocation rates as follows:

$$y_{ji}(t+1) = y_{ji}(t) - \gamma_t \left(\delta_{ji}(t) - \frac{\sum_k \delta_{ki}(t)}{d_i} \right), \quad (11)$$

where

$$\delta_{ji}(t) = -\frac{1}{n(n-1)y_{ji}(t)^2} O_j^i,$$

and with γ_t such that $\sum_{t=0}^{\infty} \gamma_t = \infty$, $\lim_{t \rightarrow \infty} \gamma_t = 0$.

We refer to the +1s as incentives because, as feedback, they represent the importance of a link, thus the value of the incentive to provide to bring about a favorable change in that link's allocation. In the present paper we keep the form of these incentives quite general, but they may be regarded as monetary or non-monetary. Non-monetary incentives might include a form of reputation or recognition, such as in networks like Klout [11]. In such networks, users receive votes that count towards their reputation or expertise, in return for some service (like answering questions) they provide to other users. A gain in reputation incites users to respond favorably when there is a possibility of receiving such votes. Note that such a method incentivizes a user to "serve" other users, thus going beyond a selfish allocation of attention.

The behavior induced by the Plus-One mechanism turns out to perform a stochastic gradient descent. The +1s collected by a node that correspond to a link e indeed serve to estimate the gradient of the cost with respect to the allocation on link e , x_e . We now show how δ_{ji} is an estimate of the gradient of the cost: $\frac{\partial C}{\partial y_{ji}}$. For ease of exposition we will assume homogeneous budgets of attention ($b_i = 1$ for all i) and content creation rates

($\lambda_i = 1$ for all i). The expected delay from source s to user k can be written as:

$$\mathbb{E}D_{sk}(\mathbf{y}) = \mathbb{E} \left[\min_{\mathcal{P}_{s \rightsquigarrow k}} \sum_{e \in \mathcal{P}_{s \rightsquigarrow k}} \frac{1}{y_e} E_e \right].$$

From (3), the social cost is $C(\mathbf{y}) = \frac{1}{n(n-1)} \sum_{k \in V} \sum_{s: s \neq k} \mathbb{E}D_{sk}$. The gradient of the average delay from s to k with respect to edge (j, i) can be estimated as follows:

$$\frac{\partial \mathbb{E}D_{sk}}{\partial y_{ji}} \approx \mathbb{E} \left[\mathbb{1}_{(j,i) \in \mathcal{P}_{s \rightsquigarrow k}^*} \frac{-1}{y_{ji}^2} E_{ji} \right],$$

where $\mathbb{1}_A$ is the indicator of A , and the gradient of the overall cost is then given by:

$$\frac{\partial C}{\partial y_{ji}} \approx -\frac{1}{n(n-1)} \mathbb{E} \left[\frac{E_{ji}}{y_{ji}^2} \sum_{s,k} \mathbb{1}_{(j,i) \in \mathcal{P}_{s \rightsquigarrow k}^*} \right].$$

An estimator $\tilde{\delta}_{ji}$ of $\frac{\partial C}{\partial y_{ji}}$ can then be written as follows:

$$\hat{\delta}_{ji} = -\frac{1}{n(n-1)y_{ji}^2} \sum_{s,k} \mathbb{1}_{(j,i) \in \mathcal{P}_{s \rightsquigarrow k}^*}.$$

Note that $\hat{\delta}_{ji}$ corresponds exactly to δ_{ji} , with $O_j^i = \sum_{s,k} \mathbb{1}_{(j,i) \in \mathcal{P}_{s \rightsquigarrow k}^*}$.

A study of the convergence properties of the Plus-One mechanism requires a careful analysis of the interchange of expectation and differentiation, which is beyond the scope of the present paper. Indeed results from simulations presented in Section V show convergence for all graphs that we consider.

B. Inefficient suboptimal graphs

The Plus-One mechanism performs well for the inefficient amenable graphs: the PoS is reduced and the cost under this distributed mechanism is very close to optimal, as we will see in Section V. For inefficient suboptimal graphs, however, regardless of the PoS, the socially optimal cost is still quite high. Our results from Section III show that line and star networks, and k -ary trees with $k = \Theta(n^\alpha)$, $\alpha > 0$, fall in this category of graphs. For such topologies, incentive mechanisms that promote only shifting of attention are not sufficient. More complex mechanisms that change the graph structure, or modify the budgets of attention of some nodes would seem more suitable. We leave the study of such mechanisms for future work.

V. SIMULATION RESULTS

We now perform simulations to verify our efficiency results and validate the *Plus-One* (PO) mechanism. For each network topology, we ran a discrete-event simulation for three different scenarios. In the first scenario, the PO mechanism is implemented, so that the users are incentivized to minimize the social cost. In the second scenario, the users behave selfishly, so they optimize only for their own cost. The tuning of their allocation of attention is similar to the PO mechanism except that since there is no reward from the downstream users, there is only a local optimization. In the first two scenarios, the initial allocation of attention is uniform for all users, which is

reasonable without prior knowledge of the network. In the third scenario, the users do not tune their allocation, and maintain the uniform strategy.

We set the each user’s publishing rate λ_i to $\lambda = 1$ and the budget of attention b_i to $b = 1$. The users update their allocation every 100 time units. We ran the simulation for a length of time long enough so that the network reaches steady state as illustrated in Figure 1, which shows the average delay over time in a complete ternary tree with 1093 nodes. The average is taken within a window of 100 time units. We observe convergence around a small set of values in all scenarios for all topologies we consider.




Fig. 1. Average delay over time in a complete ternary tree with 1093 nodes.

We now study the steady-state average delay for various network topologies. For almost all cases we will plot the average delay derived from both theoretical and simulation results. Figure 2 plots the average delay over increasing network size for a line network. The PO mechanism indeed improves upon the selfish allocation, achieving a cost close to the theoretical socially optimal cost. Note however that the socially optimal cost scales linearly with the network size. This line network is an example of inefficient suboptimal graphs that needs additional mechanisms beyond PO to improve the linear cost.




Fig. 2. Average delay against network size for a line network.

Figure 3 plots the average delay against network size for complete k -ary trees with $k = 4$. The PO mechanism achieves the theoretical socially optimal cost, which scales logarithmically with the network size. The social cost of the uniform strategy is only slightly higher than the socially optimal cost, though the gap is increasing as the network size increases. In contrast, the cost for the selfish optimization is significantly higher, and increases much faster ($\sim n^{1/2}$) with the network size. This topology is an example of an inefficient

amenable graph shows the power of the PO mechanism, in bringing the social cost down to one very close to the optimal cost.




Fig. 3. Average delay against network size for complete quaternary tree.

Figure 4 plots the average delay in random 3-regular networks. Since random d -regular graphs have good expansion property with high probability, we know from Theorems 9 and 10 that the uniform strategy achieves a social cost that scales logarithmically, and that it is a 3-approximate NE. Figure 4 shows that the costs associated with PO, selfish optimization and the uniform strategy actually coincide for the network scales used in the simulation. This expander graph is an example of the efficient graphs, where the diversity of paths leads to socially optimal costs, but also bounded PoS.




Fig. 4. Average delay against network size for random 3-regular network.

VI. CONCLUSION

We have shown that social network topologies can be categorized into three classes according to efficiency of information spreading: efficient, inefficient amenable, and inefficient suboptimal. We have proposed the *Plus-One* mechanism, an incentive-based mechanism that brings the costs in inefficient amenable graphs close to optimal. Inefficient suboptimal graphs, on the other hand, are resilient to our mechanism, in that the cost under distributed attention allocation is reduced close to the optimal social cost, but the optimal social cost itself is quite high. For such graphs, mechanisms that go beyond incentives for attention shifting, those that induce a change in graph structure or in the budgets of attention are needed.

Outlook: A natural extension of our work is the design of distributed mechanisms for networks with inefficient suboptimal graph structure. We have assumed in the present

paper that all users are interested in receiving information from all sources. We may consider a more interesting and realistic case where users have differing sets of interests. Mechanism design for such scenarios is decidedly more complex, with a more intricate contact structure.

REFERENCES

- [1] B. Pittel, "On spreading a rumor," *SIAM J. Appl. Math.*, vol. 47, no. 1, pp. 213–223, Mar. 1987.
- [2] R. Karp, C. Schindelhauer, S. Shenker, and B. Vocking, "Randomized rumor spreading," in *Proc. IEEE FOCS*. IEEE, 2000, pp. 565–574.
- [3] F. Chierichetti, S. Lattanzi, and A. Panconesi, "Rumor spreading in social networks," *Automata, Languages and Programming*, pp. 375–386, 2009.
- [4] K. Censor-Hillel and H. Shachnai, "Fast information spreading in graphs with large weak conductance," in *Proc. ACM-SIAM SODA*, 2011.
- [5] H. Amini, M. Draief, and M. Lelarge, "Flooding in weighted random graphs," in *Proc. SIAM ANALCO*, 2011, pp. 1–15.
- [6] G. Kossinets, J. Kleinberg, and D. Watts, "The structure of information pathways in a social communication network," in *Proc. 14th ACM SIGKDD*. ACM, 2008, pp. 435–443.
- [7] L. Backstrom, E. Bakshy, J. Kleinberg, T. M. Lento, and I. Rosenn, "Center of attention: How facebook users allocate attention across friends," in *Proc. 5th ICWSM*, 2011.
- [8] N. Nisan, T. Roughgarden, É. Tardos, and V. V. Vazirani, *Algorithmic Game Theory*, N. Nisan, T. Roughgarden, É. Tardos, and V. V. Vazirani, Eds. Cambridge Univ Press, 2007.
- [9] S. Janson, "One, two and three times $\log n/n$ for paths in a complete graph with random weights," *Comb. Probab. Comput.*, vol. 8, no. 4, pp. 347–361, Jul. 1999.
- [10] S. Ioannidis, A. Chaintreau, and L. Massoulié, "Optimal and scalable distribution of content updates over a mobile social network," in *IEEE INFOCOM'09*, April 2009, pp. 1422–1430.
- [11] "Klout," <http://klout.com>.

APPENDIX

A. Proof of Theorem 1

Note that $n_{i+1,i} = n - i$ and $n_{i-1,i} = i - 1$. The socially optimal cost and cost under selfish allocation given in (7) and (8) respectively, can be bounded as follows:

$$C^* \geq \frac{1}{(n-1)\Lambda b_{\max}} \sum_{i=1}^n \left(\sum_{j=i\pm 1} \sqrt{n_{ij}\lambda_{ji}} \right)^2 \geq \frac{C^o}{b_{\max}},$$

and

$$\begin{aligned} \widehat{C} &\leq \frac{1}{(n-1)\Lambda b_{\min}} \sum_{i=1}^n \left[\sum_{j=i\pm 1} n_{ij}\sqrt{\lambda_{ji}} \right] \left(\sum_{k=i\pm 1} \sqrt{\lambda_{ki}} \right) \\ &= \frac{C^o + C^d}{b_{\min}}, \end{aligned}$$

where

$$C^o = \frac{1}{(n-1)\Lambda} \sum_{i=1}^n (n_{i,i+1}\lambda_{i+1,i} + n_{i,i-1}\lambda_{i-1,i}),$$

and

$$C^d = \frac{n+1}{(n-1)\Lambda} \sum_{i=1}^n \sqrt{\lambda_{i+1,i}\lambda_{i-1,i}}.$$

When $\lambda_i = 1$ for all i , C^o is equal to $(n+1)/3$. Thus for general λ_i ,

$$C^o \geq \frac{\lambda_{\min}(n+1)}{3\lambda_{\text{ave}}} \quad (12)$$

Next we will show bounds on C^o that are independent of λ_i . Let $p_i = \Lambda^{-1} \sum_{j=1}^i \lambda_j$. Note that p_i is increasing in i . Let $i^* = \min\{i : p_i \geq 1/2\}$. Then $1 - p_i \geq 1/2$ for all $i < i^*$ and $p_i \geq 1/2$ for all $i \geq i^*$. Now we have the following:

$$\begin{aligned} C^o &= \frac{1}{n-1} \sum_{i=1}^n [p_{i-1}(n-i+1) + i(1-p_i)] \\ &\geq \frac{1}{n-1} \left[\sum_{i=i^*+1}^n \frac{1}{2}(n-i+1) + \sum_{i=1}^{i^*-1} \frac{1}{2}i \right] \\ &= \frac{1}{2(n-1)} \left[\left(i^* - \frac{n+1}{2} \right)^2 + \frac{n^2-1}{4} \right] \geq \frac{n+1}{8}. \end{aligned} \quad (13)$$

On the other hand, C^d can be bounded by

$$C^d = \frac{n+1}{n-1} \sum_{i=2}^{n-1} \sqrt{p_{i-1}(1-p_i)} \leq \frac{(n+1)(n-2)}{2(n-1)} \leq \frac{n+1}{2}. \quad (14)$$

Now, since we have

$$\text{PoS} = \frac{\widehat{C}}{C^*} \leq \frac{b_{\max}}{b_{\min}} \left(1 + \frac{C^d}{C^o} \right),$$

(9) follows from (12), (13) and (14). \square

B. Proof of Theorem 3

We set $b_j = 1$ and $\lambda_j = 1$ for all nodes j . Let i be the hub of the i -th star, and L_i the leaf nodes connected to it. Then $n_{il} = n - 1$ and $n_{li} = 1$ for any $l \in L_i$. For adjacent hubs, $n_{i-1,i} = (i-1)p$ and $n_{i+1,i} = (k-i)p$. Then the socially optimal cost is given by:

$$C^* = \frac{1}{n(n-1)} \left(\sum_{i=1}^k \sum_{l \in L_i} (n-1) + S_1 \right) = 1 - \frac{k}{n} + \frac{S_1}{n(n-1)},$$

where

$$S_1 = \sum_{i=1}^k \left[\sum_{l \in L_i} \sqrt{n-1} + \sum_{j=i\pm 1} \sqrt{(n-n_{ji})n_{ji}} \right]^2.$$

Using the inequalities $\sum_{i=1}^3 a_i^2 \leq \left(\sum_{i=1}^3 a_i \right)^2 \leq 3 \sum_{i=1}^3 a_i^2$, we obtain $S_1 = \Theta(S_2)$, where

$$\begin{aligned} S_2 &= k(p-1)^2(n-1) + \sum_{i=1}^k \sum_{j=i-1}^i p^2 j(k-j) \\ &= (n-k)(p-1)(n-1) + \frac{1}{3}n(nk-p) = \Theta(n^2 \max\{p, k\}). \end{aligned}$$

Thus $C^* = \Theta(\max\{p, k\}) = \Omega(\sqrt{n})$.

The cost under selfish allocation can be written as follows:

$$\widehat{C} = \frac{1}{n(n-1)} \left(\sum_{i=1}^k \sum_{l \in L_i} (n-1) + S_3 \right) = 1 - \frac{k}{n} + \frac{S_2}{n(n-1)},$$

where

$$S_3 = \sum_{i=1}^k \left(\sum_{l \in L_i} 1 + \sum_{j=i \pm 1} \sqrt{n_{ji}} \right) \times \left[\sum_{l \in L_i} (n-1) + \sum_{j=i \pm 1} (n-n_{ji}) \sqrt{n_{ji}} \right].$$

Using the facts that $|L_i| = p-1$, $n_{ji} \leq n$ and the inequality $(n-x)\sqrt{x} \leq \frac{2}{3\sqrt{3}}n^{3/2}$ for $x \in [0, n]$, we obtain

$$S_3 \leq \sum_{i=1}^k (p+2\sqrt{n}) \left(pn + \frac{4}{3\sqrt{3}}n^{3/2} \right) = O\left((n+k\sqrt{n})(pn+n^{3/2})\right) = O(n^2 \max\{k, p\}).$$

Thus $\widehat{C} = O(\max\{p, k\})$. Since $\widehat{C} \geq C^* = \Theta(\max\{p, k\})$, we have $\widehat{C} = \Theta(\max\{p, k\})$ and the result follows. \square

C. Proof of Theorem 4

We will consider a homogeneous case, where all budgets of attention are equal, $b_i = 1$. Further, we will assume that λ_i is constant for all i , so we have $\lambda_{ji} \propto n_{ji}$. The socially optimal cost (7) and cost under selfish allocation (8) can now be written as follows:

$$C^* = \frac{1}{n(n-1)} \sum_{i \in V} \left(\sum_{j \in N(i)} \sqrt{(n-n_{ji})n_{ji}} \right)^2, \quad (15)$$

and

$$\widehat{C} = \frac{1}{n(n-1)} \sum_{i \in V} \left[\sum_{j \in N(i)} (n-n_{ji}) \sqrt{n_{ji}} \right] \left(\sum_{j \in N(i)} \sqrt{n_{ji}} \right). \quad (16)$$

Let the depth of the tree be h . Then $n = \frac{k^{h+1}-1}{k-1}$. Label the nodes in such a way that for $i = 1, 2, \dots, n$, $p(i) = \lceil (i-1)/k \rceil$ is the parent of i , $c_j(i) = k(i-1) + j + 1$ is the j -th child of i , for $j = 1, 2, \dots, k$. Then $n_{p(i),i} = \frac{k^{h+1}-k^{h-h(i)+1}}{k-1}$ and $n_{c_j(i),i} = \frac{k^{h-h(i)-1}}{k-1}$, where $h(i) = \lfloor \log_k(ki-i) \rfloor$ is the depth of node i .

By (15), the socially optimal cost is

$$C^* = \frac{1}{n(n-1)} \sum_{i=1}^n \left[\sqrt{(n-n_{p(i),i})n_{p(i),i}} + \sum_{j=1}^k \sqrt{(n-n_{c_j(i),i})n_{c_j(i),i}} \right]^2,$$

which is bounded by

$$\frac{1}{n(n-1)} (S_1 + k^2 S_2) \leq C^* \leq \frac{2}{n(n-1)} (S_1 + k^2 S_2),$$

where

$$S_1 = \sum_{i=1}^n (n - n_{p(i),i}) n_{p(i),i},$$

and

$$S_2 = \sum_{i=1}^n (n - n_{c_1(i),i}) n_{c_1(i),i}.$$

Using the facts that $n = \frac{k^{h+1}-1}{k-1}$, $n_{p(i),i} = \frac{k^{h+1}-k^{h-h(i)+1}}{k-1}$ and $n_{c_1(i),i} = \frac{k^{h-h(i)-1}}{k-1}$, we obtain

$$\begin{aligned} S_1 &= \frac{1}{(k-1)^2} \sum_{i=1}^n (k^{h-h(i)+1} - 1)(k^{h+1} - k^{h-h(i)+1}) \\ &= \frac{1}{(k-1)^2} \sum_{h'=0}^h k^{h'} (k^{h-h'+1} - 1)(k^{h+1} - k^{h-h'+1}) \\ &= \frac{hk - h - 2}{(k-1)^3} k^{2h+2} + \frac{hk - h + 2k}{(k-1)^3} k^{h+1} \\ &= \Theta(hk^{2h}) = \Theta(n^2 \log_k n), \end{aligned}$$

and

$$\begin{aligned} S_2 &= \frac{1}{(k-1)^2} \sum_{i=1}^n (k^{h+1} - k^{h-h(i)})(k^{h-h(i)} - 1) \\ &= \frac{1}{(k-1)^2} \sum_{h'=0}^h k^{h'} (k^{h+1} - k^{h-h'}) (k^{h-h'} - 1) = \frac{1}{k} S_1. \end{aligned}$$

Therefore,

$$\frac{1+k}{n(n-1)} S_1 \leq C^* \leq \frac{2(1+k)}{n(n-1)} S_1,$$

and hence $C^* = \Theta(k \log_k n)$.

By (16), the social cost under selfish allocation is

$$\begin{aligned} \widehat{C} &= \frac{1}{n(n-1)} \sum_{i=1}^n \left[(n - n_{p(i),i}) \sqrt{n_{p(i),i}} \right. \\ &\quad \left. + k(n - n_{c_1(i),i}) \sqrt{n_{c_1(i),i}} \right] \times (\sqrt{n_{p(i),i}} + 2\sqrt{n_{c_1(i),i}}) \\ &= \frac{1}{n(n-1)} (S_1 + k^2 S_2 + k S_3), \end{aligned}$$

where

$$S_3 = \sum_{i=1}^n (2n - n_{c_1(i),i} - n_{p(i),i}) \sqrt{n_{c_1(i),i} n_{p(i),i}}.$$

Note that $2n \geq 2n - n_{c_1(i),i} - n_{p(i),i} \geq n$, and $n \geq n_{p(i),i} \geq n/2$ for $i \neq 1$, so

$$S_3 = \Theta\left(n^{3/2} \sum_{i=2}^n \sqrt{n_{c_1(i),i}}\right).$$

We also have

$$\sum_{i=2}^n \sqrt{n_{c_1(i),i}} = \frac{1}{\sqrt{k-1}} \sum_{h'=1}^{h-1} k^{h'} \sqrt{k^{h-h'} - 1} = \Theta(n/k),$$

yielding $S_3 = \Theta(kn^2 \log_k n + n^{5/2})$ and hence $\widehat{C} = \Theta(S_3/n^2) = \Theta(k \log_k n + \sqrt{n})$. Taking the ratio \widehat{C}/C^* , the result follows. \square