

Co-evolution of Content Popularity and Delivery in Mobile P2P Networks

Srinivasan Venkatramanan and Anurag Kumar
Department of Electrical Communication Engineering
Indian Institute of Science
Bangalore 560012, India
Email: vsrini,anurag@ece.iisc.ernet.in

Abstract—Mobile P2P technology provides a scalable approach for content delivery to a large number of users on their mobile devices. In this work, we study the dissemination of a *single* item of content (e.g., an item of news, a song or a video clip) among a population of mobile nodes. Each node in the population is either a *destination* (interested in the content) or a potential *relay* (not yet interested in the content). There is an interest evolution process by which nodes not yet interested in the content (i.e., relays) can become interested (i.e., become destinations) on learning about the popularity of the content (i.e., the number of already interested nodes). In our work, the interest in the content evolves under the *linear threshold model*. The content is copied between nodes when they make random contact. For this we employ a controlled epidemic spread model. We model the joint evolution of the copying process and the interest evolution process, and derive joint fluid limit ordinary differential equations. We then study the selection of parameters under the content provider's control, for the optimization of various objective functions that aim at maximizing content popularity and efficient content delivery.

I. INTRODUCTION

Peer-Peer (P2P) architectures help relieve file servers from excessive load by enabling clients to communicate among themselves and exchange content, hence providing a scalable approach to content delivery for a large audience. With the proliferation of smart phones and high speed mobile Internet, there is increasing download of media content directly into mobile devices. Mobile P2P technology aims to leverage this trend and several possible architectures have been proposed. In this paper we are concerned with peer-to-peer spread of the content under the mobile opportunistic forwarding paradigm.

While studying content distribution, it is essential to understand the evolution of demand for the content. For modeling demand evolution, we can adopt the point of view that the existing demand for an item influences others to also get interested in it. There has been considerable interest in such models in the area of viral marketing research. Content creators interested in increasing the popularity of the content use data mining techniques to identify optimal viral marketing strategies and to identify the most influential users in a social network [1]. In this paper, we are also concerned with the spread of interest in the item of media content, and for this we adopt an influence spread model introduced earlier in the context of viral marketing [1].

Evidently, content providers need to adopt their content dissemination strategies according to the prevailing interest in

the particular content. Though there have been several works that discuss optimal strategies for content distribution among mobile nodes [2][3] they do not address the issue of jointly modeling the spread of interest in the content and the content itself. In our work, we aim to model the joint evolution of demand and spread of the content in a mobile P2P setting.

A recent work [4] addresses a related problem of demand-aware distribution by employing fluid models for the viral spread of demand and aims at obtaining a hybrid P2P and client-server architecture that can meet the demand. While [4] assumes that the demand spread is uncontrolled and optimizes for content delivery, in this work, we separately discuss optimizing the parameters of the demand evolution process (to increase the popularity the content), as well as the content delivery process (to efficiently meet the content demand).

Our Contributions:

- We develop ordinary differential equation (o.d.e.) models for the co-evolution of popularity of the content and its spread via controlled epidemic copying in a mobile P2P environment
- We then use these o.d.e. models to provide insights into the choice of parameters for the content provider in order to optimize various system performance objectives

In Section II we introduce the combined system model, and obtain the fluid limits for interest evolution and content spread in Sections III and IV. We finally provide numerical results for some practical optimizations (Section V) associated with the evolution of interest and the joint-evolution of interest and spread of the content.

II. THE SYSTEM MODEL

We consider a population of homogeneous mobile *nodes* meeting each other according to a random contact model and we wish to model the spread of a *single* content among these nodes. Pertaining to the content, each node has a current state represented by two bits: the *want* bit indicates whether the node is interested in the content and the *have* bit indicates whether the node has the content. When the *want* bit of a node is 1, we call the node a *destination* (i.e., interested in the content), else we call it a *relay* (i.e., not yet interested in the content). There is a central server that keeps track of the *want* state of all the nodes and broadcasts the total number of destinations (nodes interested in the content), a measure of content popularity at regular intervals via a low-bitrate control

channel. Thus, for the purpose of spread of content popularity, the network is fully connected.

Relay nodes, on receiving popularity broadcasts from the central server, might get converted to destinations according to an *influence spread model*. We model this influence process (evolution of interest in the content) using the Homogeneous Influence Linear Threshold (HILT) model, a special case of the Linear Threshold model (LT) introduced in Kempe et al. [1]. For the content copying process, we model the random contacts between pairs of nodes as independent Poisson point processes (a model also used in the context of mobile P2P in [5]), with the copying (when pairs of nodes meet) being controlled probabilistically, in a manner similar to the Susceptible-Infective (SI) epidemic model.

A. Modeling Interest Evolution: The HILT Model

In this section we introduce the HILT model used to model the evolution of interest in the content. In the original LT model [1][6], nodes are part of a weighted directed graph $\mathcal{G} = (\mathcal{N}, \mathcal{E})$, where $\mathcal{E} \subseteq \mathcal{N} \times \mathcal{N}$. With each $i, j \in \mathcal{E}$, there is associated a weight $w_{i,j}$ which gives a measure of *influence* of node i on node j , normalized such that the total weight into any node is at most 1, i.e., $\sum_i w_{i,j} \leq 1$. The Homogeneous Influence Linear Threshold (HILT model) is a special case of the LT model where the network graph is complete and all nodes are homogeneous. Hence, we have a mesh network on the population \mathcal{N} containing $N = |\mathcal{N}|$ nodes with each edge carrying the same influence weight $\gamma_N = \frac{\Gamma}{N-1}$ and $\Gamma \leq 1$.

The evolution of interest in the content is modeled by the following influence process evolving in discrete time. Each node $j \in \mathcal{N}$ independently chooses a random threshold Θ_j from a given distribution F at the beginning. We begin with an initial set of destinations $\mathcal{A}(0)$. In the HILT model, the net influence of a set of destinations on any relay is γ_N times the size of the destination set. The number of destinations (currently interested nodes) is broadcast to all the nodes by the central server after each period, and a relay gets converted into a destination once the net influence exceeds its chosen threshold. In other words, a relay $j \notin \mathcal{A}(k-1)$ gets influenced in step k if,

$$\gamma_N |\mathcal{A}(k-1)| \geq \Theta_j \quad (1)$$

At the end of each period the population will contain three types of nodes (see Figure 1): $\mathcal{A}(k)$, the set of non-infectious destinations, $\mathcal{D}(k) \subseteq \mathcal{A}(k)$, the set of newly infected destinations in that period (infectious for the next period) and the set of relays denoted by $\mathcal{S}(k)$. It is clear that the activation process stops at a random step U when there are no more infectious destinations, i.e., $\mathcal{D}(U) = \emptyset$, and a *terminal set* $\mathcal{A}(U)$ is reached.

B. Modeling Content Copying: SI Model

In order to model the content delivery process, we further classify the nodes depending on whether they have the content (i.e., based on the *have* bit). Let $\mathcal{X}(k) \subseteq \mathcal{A}(k)$ denote the set of destinations that have the content, and $\mathcal{Y}(k) \subseteq \mathcal{S}(k)$ denote the set of relays that have the content. For the evolution




Fig. 1. Spread of Influence in the HILT Model

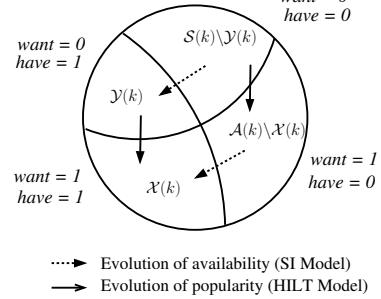


Fig. 2. Joint evolution of content popularity and spread in the HILT-SI Model.

of $(\mathcal{X}(k), \mathcal{Y}(k))$, we model a content copying process based on the Susceptible-Infective (SI model). Between the discrete time steps, pairs of nodes meet at the points of a Poisson process with rate λ_N , and whenever a node that has the content meets a node that does not, content transfer takes place in a probabilistic manner. At such a meeting, we distinguish between the node that does not have the content, being a destination or a relay, by having different copy probabilities α and σ respectively.

C. Co-evolution of Interest and Spread: The HILT-SI Model

In the combined model, underlying the content delivery process (the SI model), the influence process (the HILT model) converts relay nodes into destinations. Thus, in our setup, the fraction of destinations is time-varying (unlike [3]). Also, the content spread is dependent on the interest evolution but not vice versa. Figure 2 shows the possible transitions between the four sets of nodes. An interesting feature of this model is the importance of copying to a relay node. As a content provider, we might be interested in delivering only to the destinations (interested in the content). But, there are two advantages of copying to a relay. First, copy to a relay promotes the further spread of the content even to destinations; this is the aspect explored in a controlled Markov process setting in [3]. Second, the relay we copy to now might later get influenced (by the HILT model) to become a destination, which is a unique feature of the HILT-SI model.

III. EVOLUTION OF INTEREST IN THE CONTENT

Consider the HILT model on N nodes, and with edge weights $\gamma_N = \frac{\Gamma}{N-1}$. Recall Θ_i is the influence threshold of User i , $1 \leq i \leq N$. The $\Theta_i, 1 \leq i \leq N$, are non-negative, i.i.d. random variables chosen according to a continuous distribution $F(\cdot)$, with density function $f(\cdot)$. $(\mathcal{A}(k), \mathcal{D}(k))$, $k \geq 0$ represents the set of destinations and the infectious destinations at the end of period k with $\mathcal{D}(0) = \mathcal{A}(0)$. Define $\mathcal{B}(k) = \mathcal{A}(k-1)$ with $\mathcal{B}(0) = \emptyset$. Thus, for $k \geq 1$,

$\mathcal{B}(k) = \cup_{0 \leq i \leq k-1} \mathcal{D}_i$. Also, since the nodes are homogeneous in the HILT model, it suffices to record the sizes of the respective sets. Let $A(k)$, $D(k)$ and $B(k)$ be the sizes of the sets $\mathcal{A}(k)$, $\mathcal{D}(k)$ and $\mathcal{B}(k)$ respectively.

A. O.D.E. Model for Interest Evolution

We can show that the original HILT process $(B(k), D(k))$ is a Markov chain (see Technical Report [7]). In order to obtain an approximating o.d.e., we work with an appropriately scaled Markov process $(B^N(t), D^N(t))$, which evolves over ‘‘minislots’’ of duration $1/N$ (on a time scale N times faster than that of the original process), with appropriately slowed down dynamics [8]. In each minislot, each infectious destination in $D^N(t)$ is permitted to influence the relays with probability $\frac{1}{N}$ and its influence is deferred with probability $1 - \frac{1}{N}$. In the former case, it contributes its influence of $\frac{\Gamma}{N-1}$ and then moves to the set $B^N(t+1)$, otherwise it stays in the $D^N(t+1)$ set. Define by $C^N(t) \subseteq D^N(t)$ the set of infectious destinations who use their influence at time t . Then,

$$C^N(t) = \frac{D^N(t)}{N} + Z_b^N(t+1)$$

$$B^N(t+1) = B^N(t) + C^N(t)$$

$$\begin{aligned} D^N(t+1) &= D^N(t) - C^N(t) \\ &+ \mathbb{E} \left[\frac{F(\gamma_N(B^N(t) + C^N(t))) - F(\gamma_N B^N(t))}{1 - F(\gamma_N B^N(t))} \right] \times \\ &\left(N - B^N(t) - D^N(t) \right) + Z_d^N(t+1) \end{aligned}$$

where $Z_b^N(t+1)$ and $Z_d^N(t+1)$ are ‘‘noise’’ random variables, with zero mean conditioned on the history, and the expectation in the expression for $D^N(t+1)$ is with respect to $C^N(t)$. Defining $\tilde{B}^N(t)$, $\tilde{C}^N(t)$ and $\tilde{D}^N(t)$ as the fractional processes, we can then state the following theorem.

Theorem 1: Given the interest evolution Markov process $(\tilde{B}^N(t), \tilde{D}^N(t))$, for the threshold distribution with density $f(\cdot)$, with bounded $f'(\cdot)$ and hazard function $h_F(x) = \frac{f(x)}{1-F(x)}$, we have for each $T > 0$ and each $\epsilon > 0$,

$$\mathbb{P} \left(\sup_{0 \leq u \leq T} \left\| (\tilde{B}^N(\lfloor Nu \rfloor), \tilde{D}^N(\lfloor Nu \rfloor)) - (b(u), d(u)) \right\| > \epsilon \right) \xrightarrow{N \rightarrow \infty} 0$$

where $(b(u), d(u))$ is the unique solution to the o.d.e.,

$$\dot{b} = d \quad (2)$$

$$\dot{d} = h_F(\Gamma b) \Gamma d (1 - b - d) - d \quad (3)$$

with initial conditions $(b(0) = 0, d(0) = a(0))$.

Proof: This is essentially an instance of Kurtz’s Theorem [9]. A detailed proof is provided in our technical report [7]. ■

B. Effect of the Threshold distribution

In the HILT model, while Γ is indicative of the total level of influence each individual can receive from the others, the threshold distribution $F(\cdot)$ captures the variation among the individuals’ susceptance levels for getting interested in the

content. An empirical analysis on the effect of threshold distributions on collective behavior is available in [10]. For the HILT model of interest evolution under the uniform threshold distribution, we see that the fluid limit becomes,

$$b(t) = \frac{d_0}{r} - \frac{d_0}{r} e^{-rt} \quad (4)$$

$$d(t) = d_0 e^{-rt} \quad (5)$$

We see that this matches well with the discrete expression given in [6]. When the influence threshold is exponentially distributed with parameter Λ , we have $h_F(x) = \Lambda$, implying a *memoryless* property for the influence process, i.e., the relay nodes are equally likely to get influenced at a given time instant, irrespective of the net accumulated influence in the past. Under exponential distribution, the fluid limit is the solution to the o.d.e.,

$$\begin{aligned} \dot{b} &= d \\ \dot{d} &= \Lambda \Gamma d (1 - b - d) - d \end{aligned}$$

This is equivalent to the SIR epidemic model with infection rate $\Lambda \Gamma$ and recovery rate of 1. Note that $b(t)$ is then equivalent to the recovered set, and $d(t)$ is equivalent to the infected set.

IV. JOINT EVOLUTION OF INTEREST IN AND SPREAD OF THE CONTENT

In this section, we shall adopt techniques similar to Section III to analyse the joint evolution of interest and spread of the content (HILT-SI model). Recall that, while the HILT model evolves independently, the evolution of the SI model depends upon the HILT model. For this section, we shall assume for simplicity that the thresholds in the HILT model are uniformly distributed; the analysis can be easily extended to HILT with general distribution F .

Pairwise meetings of the nodes constitute points of a Poisson process with rate $\lambda_N = \frac{\lambda}{N}$, and let α and σ be the copy probabilities to destinations and relays respectively (Section II-C (Figure 2)). We can then express their sizes X_k and Y_k as a sum of contributions $P_{(\cdot)}(k)$, due to the content copying process (SI model) and $Q_{xy}(k)$, due to the interest evolution process (HILT model). By the contact model, the probability that a destination node without the content, receives the content in one time slot is $1 - e^{-\lambda_N \alpha (X(k) + Y(k))}$ and similarly for a relay. Also, by the HILT model, the probability of a transition from \mathcal{Y}_k may transition to \mathcal{X}_k for each of the relay nodes would be $\frac{\gamma_N D(k)}{1 - \gamma_N B(k)}$. For large N , we can then write,

$$P_x(k) \stackrel{\text{dist.}}{=} \text{Bin} \left(A(k) - X(k), \lambda_N \alpha (X(k) + Y(k)) \right)$$

$$P_y(k) \stackrel{\text{dist.}}{=} \text{Bin} \left(N - A(k) - Y(k), \lambda_N \sigma (X(k) + Y(k)) \right)$$

$$Q_{xy}(k) \stackrel{\text{dist.}}{=} \text{Bin} \left(Y(k), \frac{\gamma_N D(k)}{1 - \gamma_N B(k)} \right)$$

Hence we can write the evolution of the processes $X(k)$ and $Y(k)$ as follows.

$$\begin{aligned}
X(k+1) - X(k) &= \lambda_N \alpha (X(k) + Y(k))(A(k) - X(k)) \\
&\quad + \frac{\gamma_N D(k)}{1 - \gamma_N B(k)} Y(k) + Z_x(k) \\
Y(k+1) - Y(k) &= \lambda_N \sigma (X(k) + Y(k))(N - A(k) - Y(k)) \\
&\quad - \frac{\gamma_N D(k)}{1 - \gamma_N B(k)} Y(k) + Z_y(k)
\end{aligned}$$

where $Z_x(k)$ and $Z_y(k)$ are noise terms with zero mean conditional on the history of the joint evolution process until period k . In order to obtain the fluid limit, we again work with the scaled process which evolves over mini-slots of width $\frac{1}{N}$.

Theorem 2: Given the joint evolution Markov process $(\tilde{B}^N(t), \tilde{D}^N(t), \tilde{X}^N(t), \tilde{Y}^N(t))$, we have for each $T > 0$ and each $\epsilon > 0$,

$$\mathbb{P} \left(\sup_{0 \leq u \leq T} \left\| \left(\tilde{B}^N(\lfloor Nu \rfloor), \tilde{D}^N(\lfloor Nu \rfloor), \tilde{X}^N(\lfloor Nu \rfloor), \tilde{Y}^N(\lfloor Nu \rfloor) \right) - (b(u), d(u), x(u), y(u)) \right\| > \epsilon \right) \xrightarrow{N \rightarrow \infty} 0$$

where $(b(u), d(u), x(u), y(u))$ is the unique solution to

$$\dot{b} = d \quad (6)$$

$$\dot{d} = \frac{\Gamma d}{1 - \Gamma b} (1 - b - d) - d \quad (7)$$

$$\dot{x} = \lambda \alpha (x + y)(a - x) + \frac{\Gamma d}{1 - \Gamma b} y \quad (8)$$

$$\dot{y} = \lambda \sigma (x + y)(1 - a - y) - \frac{\Gamma d}{1 - \Gamma b} y \quad (9)$$

with initial conditions $(b(0) = 0, d(0) = d_0, x(0) = x_0, y(0) = y_0)$.

Proof: A detailed proof verifying the conditions for Kurtz's theorem to hold, are presented in our technical report [7]. ■

A. Accuracy of the Fluid limit

We know from Kurtz's theorem that the scaled process will converge to the fluid limit ODEs, for large N (see technical report [7] for the accuracy plots). We wish to compare the original HILT-SI process $(B(k), D(k), X(k), Y(k))$ with the solutions of the o.d.e. by superimposing multiple sample paths of the original discrete time HILT-SI process on the o.d.e solutions. (see Figure 3). We find that the o.d.e solution approximates the original process really well, and permits using the fluid limit approximation for sufficiently large N .

V. OPTIMIZATION PROBLEMS

A. Interest Evolution

Content creators are often interested in understanding the evolution of popularity of their content, and would wish to maximize the level of popularity achieved. In most cases, the content creator does not have control over influence weight Γ or the threshold distribution F of the population, and the only parameter under control is d_0 , the initial fraction of destinations in the population. We might be interested in




Fig. 3. Comparison of the unscaled HILT-SI process ($N = 1000$) with the corresponding fluid limit. The solid lines indicate the evolution of o.d.e. solutions. Shown in dots alongside are multiple runs of the original discrete process.

choosing the right d_0 which can give us the required b_∞ , and we see that by letting $t \rightarrow \infty$ in Equation (4) we get,

$$d_0 = \frac{b_\infty(1 - \Gamma)}{1 - b_\infty \Gamma} \quad (10)$$

Also, given the initial fraction of destinations d_0 in an HILT network with parameter Γ , the time we have to wait to get the final fraction of destinations to be at least β ($\beta < \frac{d_0}{r}$) is given by,

$$T(\beta, d_0, \Gamma) = \frac{1}{r} \ln \left(\frac{1 - r}{1 - \frac{\beta}{d_0} r} \right) \quad (11)$$

with r as defined earlier. Since we are observing the process at a finite time T , $d(T)$ is not zero. Hence, we consider $a(T) = b(T) + d(T)$ and setting $a(T) = \beta$ we get the above result. We can also determine the initial fraction d_0 to be chosen so that by time T we will have at least β fraction of the nodes in the destination set. (see Technical Report [7]).

B. Joint evolution of interest and spread

In this section we discuss the optimization problems that might be of interest for the joint evolution process. With the interest in the content evolving independently, the main motive of the content provider would then be to ensure that the content is delivered to as many destinations as possible. Since we intend to deliver to the destinations, it would be optimal to set $\alpha = 1$. Hence the only control parameter is the copy probability to a relay, σ . This may be controlled in practice by incentivizing copies to relay nodes. Having a high value of σ accelerates the $x(t)$ process, but it also leads to higher value of $y(t)$, the number of *wasted* copies at time t . Here we discuss two possible optimization problems for σ that can be posed, keeping in mind this tradeoff between $x(t)$ and $y(t)$.

1) *Maximize target spread:* As content providers, we might be interested to deliver the content to as many destinations as possible by some fixed time. This might be the case, when the content is time-dependent and its usefulness expires by that fixed time. The problem can be formally posed as,

$$\max_{\{\sigma: y(\tau) \leq \zeta\}} x(\tau)$$

We compute the feasible set $\{\sigma : y(\tau) \leq \zeta\}$ numerically obtain the optimal solution σ^* that maximizes $x(\tau)$. From Figure 4 we observe that, for τ small, we can afford a high value of σ since $y(t)$ may not exceed ζ in such a short time. Also, when τ is large, we can have a high σ , provided the interest evolution process is strong enough (high Γ). This will




Fig. 4. Maximize target spread: The optimal solution plotted for a fixed value of ζ and varying τ .




Fig. 5. Maximize target spread: The optimal solution plotted for a fixed value of τ and varying ζ .

ensure that the number of relays $s(t)$ will be low enough by τ , and since $y(t) \leq s(t)$, $\forall t$, $y(t)$ cannot exceed ζ . Figure 5 shows that as ζ , the constraint on $y(\tau)$, is relaxed, we see that the optimal σ^* increases. It can be seen that $y(\tau)$ stabilizes at the value dictated by the number of relay nodes $s(\tau)$ for high value of σ . Also note that, if the fixed time of interest τ is high, then there is negligible contribution due to increase in σ since $x(t)$ increases to $a(t)$ eventually irrespective of σ .

2) *Minimize reach time*: Another problem of interest is to deliver the content to a fraction of the destinations as early as possible. Defining $\tau_\eta = \inf\{t : x(t) \geq \eta\}$, we can formally pose this problem as,

$$\min_{\{\sigma: y(\tau_\eta) \leq \zeta\}} \tau_\eta$$

We numerically compute the feasible set by computing τ_η for each σ and verifying if $y(\tau_\eta) \leq \zeta$. In Figure 6, it is clear that a higher value of ζ allows us to use a high value of σ so that we can accelerate the content delivery to destinations. As a result, we also see that the time taken to deliver the content decreases with increasing ζ . In Figure 7, note that since we begin with $x_0 = 0.2$, for $\eta \leq 0.2$, $\tau_\eta = 0$ and σ can be any value. When η increases, as long as τ_η is low enough, we can afford to have a high σ and still keep $y(\tau_\eta)$ under control, since the $y(t)$ process may not have enough time to reach ζ . But as η increases further, σ needs to be reduced, since an increased value of τ_η implies, the process will run longer and $y(\tau_\eta)$ might then exceed ζ .

VI. CONCLUSIONS

In this paper we studied the co-evolution of popularity and spread of content in a mobile P2P environment using a joint model (HILT-SI) and derived their fluid limits. We derive explicit solutions for uniform threshold distribution, and show that for exponential distribution of threshold the HILT fluid limit is equivalent to the SIR model. We then used fluid limits to address several optimization problems that might be of practical interest. This work can be extended in several possible directions. One could generalize the problem




Fig. 6. Minimize reach time: The optimal solution plotted for a fixed value of η and varying ζ .




Fig. 7. Minimize reach time: The optimal solution plotted for a fixed value of ζ and varying η .

for analysing evolution of popularity and spread for multiple P2P content. And finally one could consider other models for both interest evolution and for the copying process, including those where the processes of interest evolution and content delivery are dependent on each other.

VII. ACKNOWLEDGEMENTS

The authors would like to acknowledge the Department of Science and Technology (DST) and the Indo-French Center for the Promotion of Advanced Research (CEFIPRA) for funding this work.

REFERENCES

- [1] D. Kempe, J. Kleinberg, and É. Tardos, "Maximizing the spread of influence through a social network," in *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 2003, pp. 137–146.
- [2] M. Khouzani, S. Sarkar, and E. Altman, "Dispatch then stop: Optimal dissemination of security patches in mobile wireless networks," in *49th IEEE Conference on Decision and Control (CDC)*, 2010. IEEE, 2010, pp. 2354–2359.
- [3] C. Singh, A. Kumar, R. Sundaresan, and E. Altman, "Optimal forwarding in delay tolerant networks with multiple destinations," in *9th Intl. Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOpt)*, 2011.
- [4] S. Shakkottai and R. Johari, "Demand-aware content distribution on the internet," in *IEEE/ACM Transactions on Networking (TON)*, vol. 18, no. 2. IEEE Press, 2010, pp. 476–489.
- [5] R. Nuez-Queija and B. Prabhu, "Scaling laws for file dissemination in p2p networks with random contacts," in *16th International Workshop on Quality of Service (IWQoS) 2008*. IEEE, 2008, pp. 75–79.
- [6] S. Venkatraman and A. Kumar, "Information dissemination in socially aware networks under the linear threshold model," in *2011 National Conference on Communications (NCC)*, jan. 2011, pp. 1–5.
- [7] V. Srinivasan and A. Kumar, "Co-evolution of content popularity and delivery in mobile p2p networks," *CoRR*, vol. abs/1107.5851, 2011.
- [8] M. Benaïm and J. Le Boudec, "A class of mean field interaction models for computer and communication systems," in *Performance Evaluation*, vol. 65, no. 11–12. Elsevier, 2008, pp. 823–838.
- [9] T. Kurtz, "Solutions of ordinary differential equations as limits of pure jump markov processes," in *Journal of Applied Probability*. JSTOR, 1970, pp. 49–58.
- [10] M. Granovetter, "Threshold models of collective behavior," in *American journal of sociology*. JSTOR, 1978, pp. 1420–1443.