Co-evolution of Content Spread and Popularity in Mobile Opportunistic Networks

Srinivasan Venkatramanan, *Student Member, IEEE,* Anurag Kumar, *Fellow, IEEE,*
Department of Electrical Communication Engineering, Indian Institute of Science,
Bangalore - 560012, India.
E-mail: vsrini,anurag@ece.iisc.ernet.in

Abstract—We consider a setting in which a single item of content is disseminated in a population of mobile nodes by opportunistic copying when pairs of nodes come in radio contact. The nodes in the population may either be interested in receiving the content (referred to as destinations) or not yet interested in receiving the content (referred to as relays). We consider a model for the evolution of popularity, the process by which relays get converted into destinations. A key contribution of our work is to model and study the joint evolution of content popularity and its spread in the population. Copying the content to relay nodes is beneficial since they can help spread the content to destinations, and could themselves be converted into destinations. We derive a fluid limit for the joint evolution model and obtain optimal policies for copying to relay nodes in order to deliver content to a desired fraction of destinations, while limiting the fraction of relay nodes that get the content but never turn into destinations. We prove that a time-threshold policy is optimal for controlling the copying to relays, i.e., there is an optimal time-threshold up to which all opportunities for copying to relays are exploited, and after which relays are not copied to. We then utilize simulations and numerical evaluations to provide insights into the effects of various system parameters on the optimally controlled co-evolution model.

Index Terms—Delay tolerant networking, Epidemic spread of information, Influence spread, P2P content spread, Fluid limits

1 Introduction

Due to the ubiquity of cellular networks, there has been a proliferation of hand-held mobile devices. The idea of mobile opportunistic networking is to exploit the mobility of users carrying such devices to transfer content in a device-to-device (D2D) [4] fashion during chance meetings. This is enabled by low-power radio interfaces on these devices (such as Bluetooth or WiFi Direct), and provides the opportunity for creating a multi-hopping communication network completely bypassing the cellular infrastructure. Since such a scheme cannot meet delay guarantees, applications that utilize mobile opportunistic networking must be *delay tolerant*. On the other hand, such a peer-to-peer (P2P) content delivery is scalable [12], as the rate of service scales in proportion to the number of peers in the system with little additional cost to the system planner.

In this paper, we consider such a networking paradigm, and study the problem of dissemination of an item of content among the population of mobile nodes. The content could be a video (news footage, sports highlights, a movie teaser, etc.) or an audio file (a recent song, a popular ringtone, etc.). We assume that the content is neither too big (can be transfered as a single chunk in a pairwise contact) nor too small (there is some advantage to off-loading its transfer directly from the infrastructure wireless network). To begin with, not all nodes in the population are interested in the content. In this paper, we refer to nodes who are *interested* in the content as *destinations* and those who are *not yet interested* in the content as *relays*. In such a situation, the system objective could be to facilitate the spread of content to as many destinations as possible. In doing this, the relay nodes can be used to cache the content, thus making the content more available in the population. This could cost the system planner, due to the need for incentivizing the relay nodes to participate in relaying the content. The overall objective would then be to deliver the content to as many destinations as possible, while minimizing the number of nodes that have the content but never became interested in it.

Literature Survey: There has been considerable work in proposing improved architecture and protocols for an opportunistic network over mobile nodes [7], [11]. In the delay tolerant networking context, there is prior literature (see [14] and [9] and the references therein) on the optimal opportunistic copying of content in order to optimize delivery delay and/or wasteful copying to relays.

An interesting aspect, that is largely unexplored in prior literature, is the evolution of interest in the content. For recently "released" content (such as a new video clip), the set of destinations need not be static. The interest in the content could grow, due to the influence of destinations already present in the population. Such influence could be mediated via a centralized server, which uses a low bit-rate channel

to broadcast the current popularity of the content, or by interactions between the mobile nodes themselves. Thus, it may be necessary to deliver the content to destinations while keeping track of the demand evolution.

Shakkottai and Johari [13] have recently explored this aspect by utilizing an epidemic model to characterize demand evolution and obtained a hybrid of P2P and client-server architecture to efficiently meet the demand. While the work reported in [13] mainly focusses on the tradeoff between centralized and decentralized dissemination, in our work, we are interested in optimizing a purely decentralized system. Also, while in [13] the peer-to-peer file dissemination occurs only among the nodes interested in the content, we consider the notion of relays aiding the spread, thus leading to a more general dynamics.

Our Contributions: In this paper we consider a population of N mobile nodes, and model their pair-wise meetings by independent Poisson point processes. The item of content is provided to a subset of the initially interested set of nodes. Epidemic models are used to model both the spread of popularity and the content. In this framework, we make the following contributions:

1) We model the evolution of popularity and the spread of content by models inspired from epidemiology. While the former is similar to an uncontrolled SIR (Susceptible-Infected-Recovered) process, the latter is modeled as a controlled SI (Susceptible-Infected) process, with the copying probability to a relay (σ) serving as a static control, thus yielding what we call the *SIR-SI* model. The evolutions are modeled as coupled continuous time Markov chains. We scale the population size N, and under certain assumptions on the scaling with N of the model parameters, we use Kurtz's theorem [10] to derive the fluid limit of the joint evolution process. Simulation results are provided to illustrate how large a value of N is required before the fluid limit is a good approximation to the finite population dynamics.
2) Then, permitting the copying probability, σ, to vary with time, we obtain a system of controlled o.d.e.s for which we obtain the optimal control by direct arguments using certain monotonicity properties. This results in a time-threshold structure of the optimal control. We provide an extensive numerical study of this model, thus providing additional interesting insights.

2 Joint Spread of Interest and Content

In this section we describe the system that we wish to model, and motivate our assumptions and simplifications, by taking recourse to an example of a hypothetical mobile application. This application framework that we propose utilises opportunistic links between mobile personal devices and can be used for pre-release publicity of a product. The product could be an upcoming movie, a book soon to be launched, or a new gadget about to be released to the market. The actual item of content under consideration could be a digital discount coupon that could be used to pre-order the product. This may also be bundled with promotional material such as a trailer, a sample chapter, or a demo video; thus, the "coupon" carries a non-trivial quantity of data. The demand for the content could evolve with time and the content creator would wish to spread the digital coupons suitably to meet the demand (as it evolves).

Consider every user in the population carrying a smartphone with a D2D interface turned on. Each of the phones runs the application (app), on which the users can

- Maintain a wishlist of upcoming products (movies, books, gadgets, etc.)
- Share and receive wishlists among other users in the population
- Receive promotional offers either from the central server or via pairwise contacts

By allowing users to declare the interest and enabling the spread of interest, we extend the notion of a personal wishlist (in e-commerce platforms like Amazon) to a *social wishlist*. The application also combine the services of coupon delivery (e.g., Groupon) with such wishlists, thus allowing targeted discount delivery, wherein brands can identify the users who are interested in the coupon and deliver it to them.

For simplicity we will consider a single product to explain the application framework.

Initial seeding: Let an initial subset of users add the product to their wishlist, influenced by extrinsic means such as traditional advertising, mass campaigns, etc. These will form the *initial set of active destinations*. The central app server chooses a subset of these destinations and sends them the promotional offer (digital coupon), thus yielding an *initial set of destinations who have the content*. This concludes the initial seeding phase.

Wishlist exchange: At each pairwise contact, the smartphones exchange their wishlist information. For each product (our analysis is concerned with only one such product) the wishlist information contains the following bits:

- **want** bit : **1** if the node is a destination, **0** otherwise
- **reco** bit : provided the node is a destination, **1** if it is active, **0** if it is inactive. Only active destinations aid in popularity spread.
- **have** bit : **1** if the node has the digital coupon, **0** otherwise

The **want** and **have** bits, as their names suggest, indicate whether the user wants to receive the

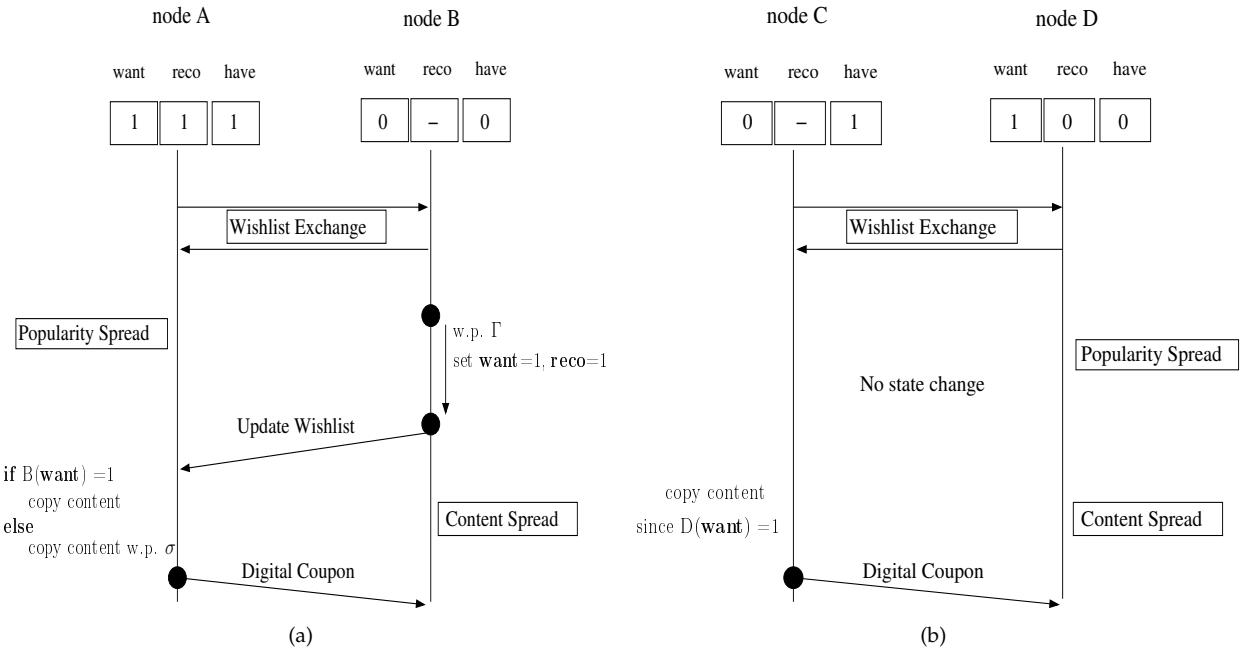

Fig. 1. The above figure shows two possible sample exchanges. (a) Pairwise meeting of an active destination with content (node A) and a relay without content (node B). (b) Pairwise meeting of a relay with content and an inactive destination without content. There is no state change in the popularity spread phase, since neither of the nodes is an active destination.

content, and whether the user possesses the content. The **reco** bit is used to indicate whether the user is actively spreading interest for the content. As long as the reco bit of a user is 1, the app in the user's device will make this recommendation to relays that it meets. Users become inactive by switching the **reco** bit from 1 to 0 (modeled as occurring at rate β_N). Users become inactive when they lose motivation to spread the popularity after a while, due to aging of the content. These users will no longer aid the popularity spread, but still wish to receive the content. This is equivalent to users subscribing to a page on Facebook, but later revoking the rights of the page to *post on the user's behalf*.

After exchanging the wishlists, the following phases may occur depending upon the states of the two nodes in contact.

Popularity spread: If one of the nodes is an active destination (**want=1**, **reco=1**) and the other is a relay (**want=0**), then popularity spread may occur. The app of the relay user will receive a recommendation from the active destination device, and this is notified to the user by the app interface. We assume that the as yet uninterested user (relay) gets influenced with some probability, and model this by the relay switching to (**want=1**, **reco=1**) with probability Γ.

Content spread: If the two nodes that meet have different **have** states, then content (digital coupon) spread may occur. If the node without the content is a destination (**want=1**), then the content is copied with probability 1, while if it is a relay (**want=0**) then the app decides to copy the content with probability σ. Figure 1 shows two possible pairwise interactions that can occur in this system.

The process is terminated when a sufficient fraction of destinations have been given the content (digital coupon). This indicates that the content creator has successfully completed his campaign and the product is removed from the application server.

Remarks:

- Though we might be interested only in delivering the content to the destinations, there are two advantages of copying to a relay. First, a copy to a relay promotes the further spread of the content even to destinations (explored in [14]). Second, the relay we copy to now might later get influenced to become a destination.
- *Incentives and Freeriders:* There is incentive for a node in the population to declare its interest in a particular product, since it is highly likely that it will receive a promotional offer by being in the system. We also assume that the nodes forward the content without engaging in *freeriding* behavior, i.e. receiving the content but not sharing it on future pairwise meetings. The impact of freeriders is an interesting future direction of research and we suspect mechanisms suggested for existing P2P networks [8] can be adopted for this system.
- We also observe that the popularity spread is independent of the possession of content. This

might be an adequate model for the spread of pre-release promotional material among an established fanbase, for example, "fans" of a particular genre of movies, or of a particular author; among fans, having the promotional material may not change the influence level. More complex coupled dynamics will arise if the (want=1, reco=1, have=1) nodes have a different rate of spreading interest in the content, as compared to the (want=1, reco=1, have=0) nodes.

3 CTMC Model of the Joint Evolution

We consider the dissemination of a single item of content among a population $\mathcal{N}$ of mobile nodes of fixed size N . We assume that pairs of nodes meet at the points of a Poisson process with rate $\lambda_N = \frac{\lambda}{N}$. This assumption is justified by earlier studies of several random mobility models in the literature. For instance, in the Random Waypoint Model (RWP), the expected meeting rate between two mobile nodes [5] is given by

$$\lambda = \frac{2\overline{v}r}{A}$$

where $\overline{v}$ is the average velocity of each node, r is the radio range and A is the system area. If we assume that individual node velocities and radio ranges do not scale with N, then the assumption of meeting rate scaling inversely with N, implies the system area scales as N. This can be seen as a result of ensuring a constant node density (number of nodes per unit area).

Define $\mathcal{A}^N(t)$ to be the set of destinations at time t and let $\mathcal{S}^N(t) = \mathcal{N}\backslash\mathcal{A}^N(t)$ be the set of relays. We classify the destinations further as active ($\mathcal{D}^N(t)$) and inactive ($\mathcal{B}^N(t)$). [1] Let $A^N(t)$,$B^N(t)$ and $D^N(t)$ be the sizes of the sets $\mathcal{A}^N(t)$, $\mathcal{B}^N(t)$, and $\mathcal{D}^N(t)$ respectively. We then have $A^N(t) = B^N(t) + D^N(t)$. Only active destinations can convert other relays into destinations, i.e., *infect* them, whereas inactive destinations are still interested in the content, but no longer spread the content's popularity. The spread of popularity is governed by two parameters Γ and β_N. These are similar to the infection and recovery rates in the SIR epidemics. An active destination remains in the active state for an exponentially distributed duration with parameter β_N. Thus the rate at which an active destination gets converted to an inactive destination is β_N, while, the probability with which an active destination infects a relay it meets is given by Γ. From the Poisson process model of meetings, and the exponentially distributed active duration, it follows that the process $(B^N(t), D^N(t))$ is a continuous time Markov chain (CTMC); Figure 2 shows the state transitions in the content popularity process.

1. Due to the similarity with the spread of infection, we will often use terminology from epidemic spread models, such as, susceptible, infected, recovered, etc.

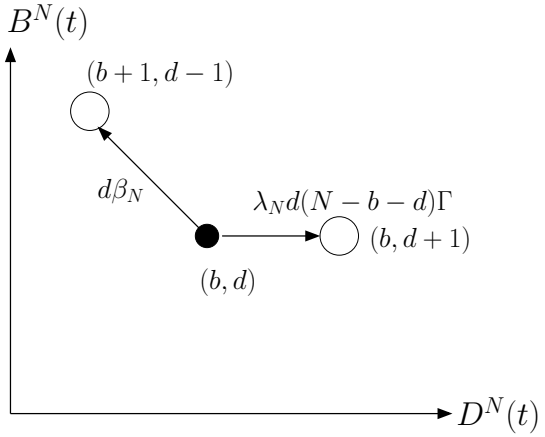

Fig. 2. Possible state transitions for the popularity process $(B^N(t), D^N(t))$.

We aim to model the joint spread of interest in the content and of the content itself. We do this by combining the SIR model for popularity spread with a controlled-epidemic copying process for content spread (probabilistic control, similar to the SI model). We will refer to this joint model as an *SIR-SI* model. In order to model the content delivery process, we further classify the nodes depending on whether they have the content. Let $\mathcal{X}^N(t) \subseteq \mathcal{A}^N(t)$ denote the set of destinations that have the content, and $\mathcal{Y}^N(t) \subseteq \mathcal{S}^N(t)$ denote the set of relays that have the content. Let $\mathcal{X}_b^N(t)$ and $\mathcal{X}_d^N(t)$ respectively be the intersection of $\mathcal{X}^N(t)$ with the sets $\mathcal{B}^N(t)$ and $\mathcal{D}^N(t)$. For the evolution of $(\mathcal{X}^N(t), \mathcal{Y}^N(t))$, we model a content copying process based on the SI model. Whenever a node that has the content meets a node that does not, content transfer takes place based on a controlled copying process. We always copy to a destination that does not have the content, whereas copying to a relay is controlled by a probability $\sigma \in [0, 1]$. We wish to obtain the fluid dynamics of this model. A brief summary of notation is provided in Table 3.1.

3.1 System Evolution

The set of all possible transitions among the different states of nodes is shown in Figure 3. Note that the process evolves at epochs t_k, $k \geq 1$ which are either pairwise meetings (occurring at rate $\lambda_N|\mathcal{N}|(|\mathcal{N}|-1)$) or instances of recovery of an active destination (occurring at rate $\beta_N D^N(t)$). The system state is represented by the tuple $Z^N(t) = (B^N(t), D^N(t), X_b^N(t), X_d^N(t), Y^N(t))$, the sizes of the respective sets. Again, due to the Poisson process model for the meeting instants, and the exponentially distributed active durations for destination nodes, $Z^N(t)$ is a continuous time Markov chain; in Table 2 we show its transition structure.

Notation	Explanation
$\mathcal{A}^N(t)$	set of destinations at time t, of size $A^N(t)$
$\mathcal{S}^N(t)$	set of relays at time t, of size $S^N(t)$
$\mathcal{B}^N(t)$	set of inactive destinations at time t, of size $B^N(t)$
$\mathcal{D}^N(t)$	set of active destinations at time t, of size $D^N(t)$
$\mathcal{X}^N(t)$	set of destinations with the content at time t, of size $X^N(t)$
$\mathcal{X}_b^N(t)$	set of inactive destinations with the content at time t, of size $X_b^N(t)$
$\mathcal{X}_d^N(t)$	set of active destinations with the content at time t, of size $X_d^N(t)$
$\mathcal{Y}^N(t)$	set of relays with the content at time t, of size $Y^N(t)$
$Z^N(t)$	$(B^N(t), D^N(t), X_b^N(t), X_d^N(t), Y^N(t))$
$\tilde{Z}^N(t)$	$\frac{Z^N(t)}{N}$

TABLE 1
A summary of the notation

Epoch type i	Rate $R_i(Z)$	State update $\delta_i(Z)$
$\mathcal{D}-\mathcal{X}_d$ recovers	$\beta_N(D-X_d)$	(1,-1,0,0,0)
$\mathcal{X}_d$ recovers	$\beta_N X_d$	(1,-1,1,-1,0)
$\mathcal{B}-\mathcal{X}_b$ meets $\mathcal{X}+\mathcal{Y}$	$\lambda_N(B-X_b)(X+Y)$	(0,0,1,0,0)
$\mathcal{D}-\mathcal{X}_d$ meets $\mathcal{Y}$	$\lambda_N(D-X_d)Y$	(0,0,0,1,0) + (0,1,0,1,-1) w.p. Γ
$\mathcal{D}-\mathcal{X}_d$ meets $\mathcal{X}$	$\lambda_N(D-X_d)X$	(0,0,0,1,0)
$\mathcal{X}_d$ meets $\mathcal{Y}$	$\lambda_N X_d Y$	(0,1,0,1,-1) w.p Γ
$\mathcal{S}-\mathcal{Y}$ meets $\mathcal{X}_b+\mathcal{Y}$	$\lambda_N(X_b+Y)(S-Y)$	(0,0,0,0,1) w.p. σ
$\mathcal{S}-\mathcal{Y}$ meets $\mathcal{D}-\mathcal{X}_d$	$\lambda_N(D-X_d)(S-Y)$	(0,1,0,0,0) w.p. Γ
$\mathcal{S}-\mathcal{Y}$ meets $\mathcal{X}_d$	$\lambda_N(X_d)(S-Y)$	(0,1,0,1,0) w.p. Γ (0,0,0,0,1) w.p$(1-\Gamma)\sigma$

TABLE 2
Transition rates and state updates for various possible epochs in the SIR-SI process, when the current state of the CTMC $Z^N(t)$ is given by $Z=(B,D,X_b,X_d,Y)$

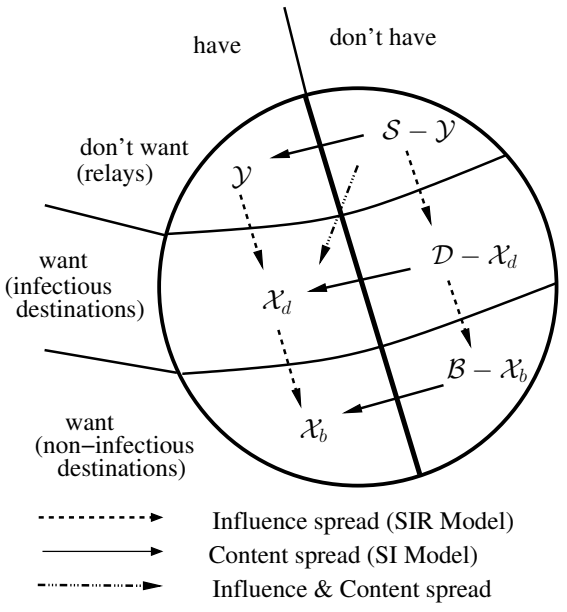

Fig. 3. Possible transitions between the states in the SIR-SI model. Observe that, in addition to the transitions due to SIR model and SI model separately, there are instances where both can occur simultaneously. For instance, when an active destination node $i \in \mathcal{X}_d$ meets a relay node $j \in \mathcal{S}-\mathcal{Y}$, the relay might get converted to a destination and also receive the content, as indicated by the diagonal transition.

3.2 Drift Equations

The drift rate is the expected rate of change of the process out of a given state. We can express the expected drift of the system as follows:

$$F^N(Z) = \sum_{i\in\mathcal{E}} R_i(Z)\delta_i(Z) \tag{1}$$

where $i \in \mathcal{E}$ indexes the epoch type, $R_i(Z)$ and $\delta_i(Z)$ are the respective transition rates and corresponding state changes given the current state of the CTMC $Z^N(t)=Z$, as given in Table 2. The co-evolution process evolves at pairwise meetings or recovery instants. Given the current state of the system, the rates of each possible transition epoch type can be determined (as shown in Table 2), and the state update depends on the type of node(s) involved in the pairwise meeting (or node recovery). The exact expressions for the mean drift rates $F^N(Z)$ are provided in Section 7.1. Define $\tilde{Z}^N(t) = Z^N(t)/N$ and consider the o.d.e.s given below.

$$\dot{b} = \beta d \tag{2}$$

$$\dot{d} = -\beta d + \lambda\Gamma d s \tag{3}$$

$$\dot{x_b} = \beta x_d + \lambda(b-x_b)(x+y) \tag{4}$$

$$\dot{x_d} = \Gamma\lambda(d-x_d)y + \Gamma\lambda x_d s + \lambda(d-x_d)(x+y) - \beta x_d \tag{5}$$

$$\dot{y} = -\Gamma\lambda d y + \lambda\sigma(s-y)(x_b+y+(1-\Gamma)x_d) \tag{6}$$

where $s(t) = 1 - b(t) - d(t)$ and $x(t) = x_b(t) + x_d(t)$. We can then state the following result.

Theorem 1: Given the coevolution Markov process $\tilde{Z}^N(t) = (\tilde{B}^N(t), \tilde{D}^N(t), \tilde{X}_b^N(t), \tilde{X}_d^N(t), \tilde{Y}^N(t))$ with initial conditions $\tilde{Z}^N(0)$ we have for each $T > 0$ and each $\epsilon > 0$,

$$\mathbb{P}\left(\sup_{0 \le u \le T} \left|\left| \tilde{Z}^N(\lfloor Nu \rfloor) - z(u) \right|\right| > \epsilon \right) \overset{N \to \infty}{\to} 0$$

where $z(u) = (b(u), d(u), x_b(u), x_d(u), y(u))$ is the unique solution to the o.d.e. system given by Equations (2)-(6) with $z(0) = \tilde{Z}^N(0)$.

Proof: The proof involves verifying the conditions for applying Kurtz's theorem [10] to the SIR-SI process (see Section 7.1). Since the drifts are Lipschitz, the uniqueness of the solution of the o.d.e. is guaranteed once the initial condition is fixed, by the Cauchy-Lipschitz condition. □

Remark: As is evident from the o.d.e., the interest evolution and the content dissemination processes have independence in one direction, i.e., the evolution of $(b(t), d(t))$ proceeds independently, while the evolution of $(x_b(t), x_d(t), y(t))$ depends on $(b(t), d(t))$.

3.3 Accuracy of the O.D.E. Approximation

Figures 4 and 5 illustrate the convergence of the scaled coevolution process $\tilde{Z}^N(t)$ to the o.d.e. solutions $z(t)$ for increasing values of N for $\sigma = 1$ and $\sigma = 0.3$. We plot $a(t) = b(t) + d(t)$, $x(t) = x_b(t) + x_d(t)$ and $y(t)$ for clarity. Observe that as N increases, the approximation of the original process by the o.d.e. becomes better. Observe that when σ is increased from 0.3 to 1, there is significant increase in $y(t)$ but the contribution to $x(t)$ is not significant. Also for σ a constant, i.e., static control, $x(t)$ asymptotically reaches $a(t)$ irrespective of σ and $y(t)$ asymptotically reaches $s(t) = 1 - a(t)$ provided $\sigma > 0$. Note that from the o.d.e. equations, it is clear that, $x(t)$ is monotonically increasing, whereas in general $y(t)$ need not be monotonic.

3.4 Copying to Relays: Static vs. Dynamic Control

In the above section we considered σ to be a static control, i.e. $\sigma(t) = \sigma$, $\forall t$. We could instead consider a dynamic copying control $\sigma(t) \in [0, 1]$. Then for each N we have a controlled continuous time Markov chain. The optimal control appears difficult to obtain in the finite size problem. We instead replace the constant copying probability in the ODE limit by $\sigma(t)$ which yields a controlled ODE. We can then obtain the optimal (deterministic) control for the controlled ODE (as in 4.1). In certain situations it can be shown that this control is asymptotically optimal for the finite size problem as N increases [3]. The proof given in [3], for discrete time MDPs with finite time horizon, does not directly apply here, and needs to be extended to accommodate our setting. In this paper, we derive the optimal control $\sigma(t)$ for the controlled ODE, and then treat the use of this optimal control for the finite population Markov chain as a heuristic.

In order to justify this approach, the convergence of the co-evolution process to its o.d.e. limit is numerically demonstrated in Figure 6 for a time-threshold type control. In this case, $\sigma(t) = 1$, $t < 4$ and $\sigma(t) = 0$ for $t \ge 4$. Note that when $\sigma(t) = 0$, $\dot{y} \le 0$ from Equation (6). Observe that the trajectory of $x(t)$ is very similar to the one obtained when $\sigma = 0.3$, but the value of $y(t)$ for large t is considerably smaller. This indicates that this threshold policy ($\tau = 4$) will perform better than the static policy $\sigma = 0.3$, for the purpose of copying to a large fraction of the destinations, while keeping in check the number of relays who get the content but do not eventually convert to destinations.

4 The Optimization Problem

In [14], since the fraction of destinations was constant, it was suitable to choose the time of delivery to a fraction α of destinations as the objective to be minimized. In our setup, since the fraction of destinations is time-varying, we define the *target time*,

$$T_\sigma = \inf\{t : x_\sigma(t) \ge \alpha a(\infty)\}$$

where $a(\infty)$ is the terminal fraction of destinations as given by the SIR model for interest evolution $(b(t), d(t))$ and $x_\sigma(t)$ is the fraction of destinations that have the content at time t under the control $\sigma(t)$. Note that since $d(\infty) = 0$ and $a(t) = b(t) + d(t)$, $a(\infty) = b(\infty)$. The cost we are interested in optimizing is of the form

$$C_\sigma = \psi y_\sigma(T_\sigma) + T_\sigma = \psi y_\sigma(T_\sigma) + \int_0^{T_\sigma} 1 dt \qquad (7)$$

Here $y_\sigma(T_\sigma)$ denotes the number of relays that have the content at time T_σ and signifies the number of wasted copies, and ψ is the tradeoff parameter.

Even though the copying cost is distributed across the nodes, we treat the number of wasted copies at the target time T_σ as part of the system objective. This can be motivated by considering an incentive mechanism for the relay nodes which still hold the content but are not converted into destinations by the target time T_σ. Consider the earlier framework, where the content is a discount coupon for some service/product. The service provider, initially provides these discount coupons to a subset of its customers, who are then encouraged to spread replicas of the coupon. Nodes that are interested in the service, and are in possession of the coupon immediately claim the service. The service provider stops accepting discount coupons once a pre-defined number (here $\alpha a(\infty)$) of coupons have been used. At this time, the relay nodes that are still in possession of the coupon will need to be provided a reimbursement for helping the spread of the coupon. In the example described in Section

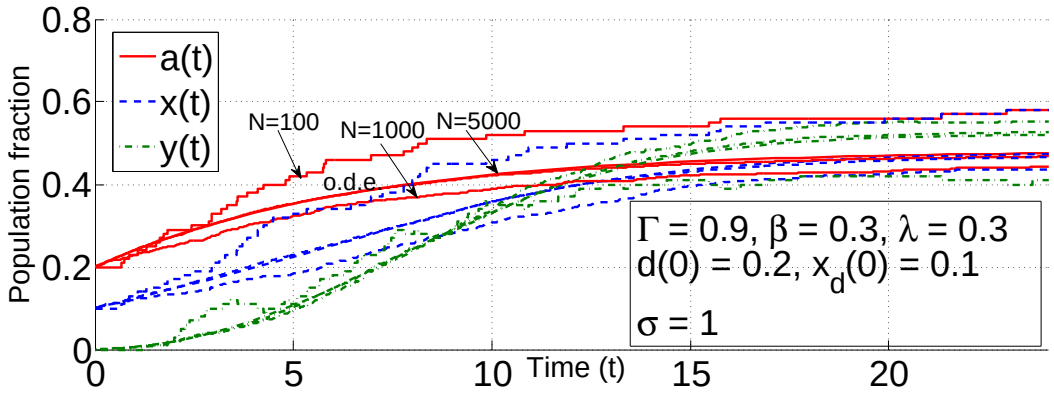

Fig. 4. Convergence of the scaled coevolution process $\tilde{Z}^N(t)$ to the o.d.e. solutions $z(t)$ for increasing values of N with $\sigma = 1$

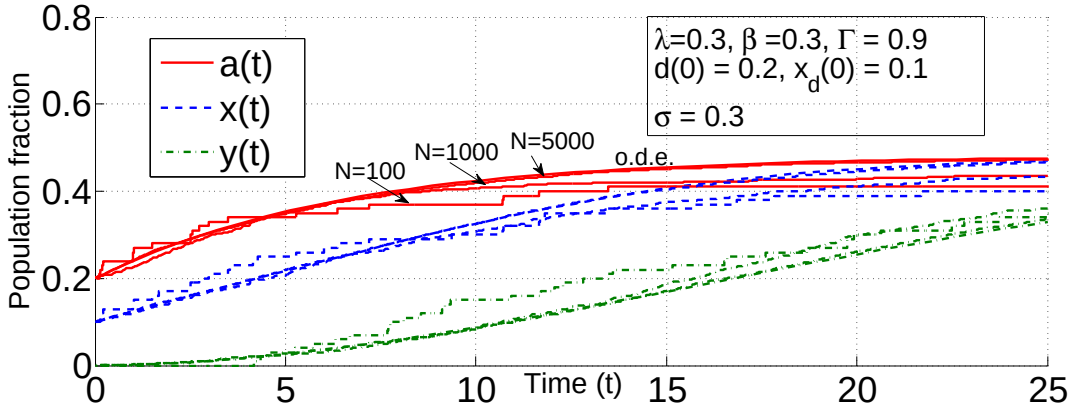

Fig. 5. Convergence of the scaled coevolution process $\tilde{Z}^N(t)$ to the o.d.e. solutions $z(t)$ for increasing values of N with $\sigma = 0.3$

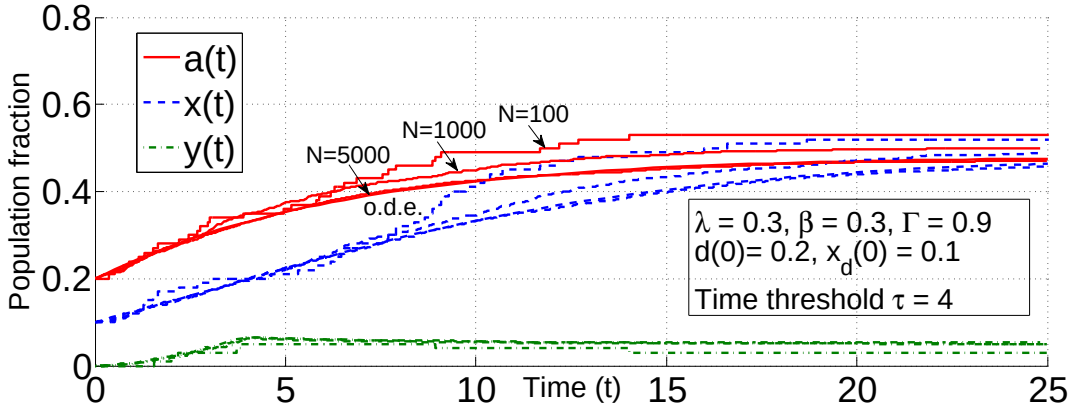

Fig. 6. Convergence of the scaled cofor-evolution process $\tilde{Z}^N(t)$ to the o.d.e. solutions $z(t)$ for increasing values of N with dynamic $\sigma(t)$, a time-threshold function in this case, $\sigma(t) = 1$, $t < 4$ and $\sigma(t) = 0$ for $t \geq 4$.

2, this reimbursement could be in terms of credit on their cellular service, say 100 free SMSs.

4.1 Optimal Control

In this section we will establish the optimality of a time-threshold control for the objective given by the Equation (7). We will use definitions and lemmas provided in Section 7.2 in order to prove the following theorem.

Theorem 2: For the o.d.e system given by Equations (2)-(6) there exists an optimal control of the form,

$$\sigma_\tau(t) = \begin{cases} 1, & 0 < t < \tau \\ 0, & t \geq \tau \end{cases} \tag{8}$$

which optimizes the cost in Equation (7).

Proof: Recall $y_\sigma(T_\sigma)$ is the amount of wasted copies, at the target time T_σ. When $\sigma(t) = \sigma_\tau(t)$, a time-threshold policy (Equation (8)), we will denote this by $y_\tau(T_\tau)$, where τ is the time threshold associated with $\sigma_\tau(t)$. The sketch of the proof is as follows. We first establish that, the set of values taken by $y_\tau(T_\tau)$ form an interval of $[0, \rho_{\max}]$. Then, given any

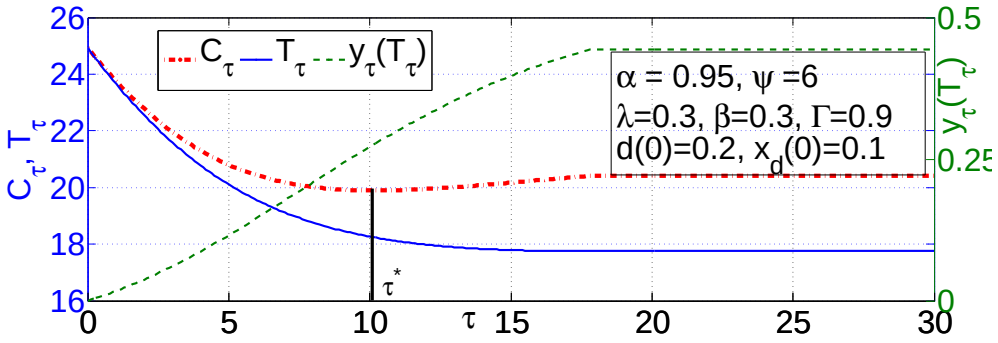

Fig. 7. Evaluation of optimal τ for the time-threshold policy. The plot shows the behaviour of T_τ and $y_\tau(T_\tau)$ for various values of τ, for fixed system parameters.

policy $\sigma(t)$, we show that:

- If $y_\sigma(T_\sigma) = \rho \leq \rho_{\max}$, then $\exists$ a time-threshold policy $\sigma_\tau(t)$ such that $y_\tau(T_\tau) = \rho$ and $T_\tau \leq T_\sigma$.
- If $y_\sigma(T_\sigma) = \rho > \rho_{\max}$, we can find a time-threshold policy $\sigma_\tau(t)$ which has a smaller total cost.

Thus in either case, we have a time-threshold policy, which performs at least as good as the given policy, which proves the optimality of the time-threshold policy. Refer Section 7.2 for the proofs of the above claims. □

Figure 7 shows the variation of T_τ, $y_\tau(T_\tau)$ and C_τ as a function of τ for fixed system parameters. It can be seen that as τ is increased (as we continue to copy to relays for longer) , T_τ decreases monotonically (the target is achieved earlier), and we see an increase in the value of $y_\tau(T_\tau)$ (more wasteful copies). The optimal threshold $\tau^\star$ minimizes the total cost C_τ, by balancing the two component costs, taking the tradeoff parameter ψ into account.

5 Numerical Results

Having observed that the optimal control is of the form 8, we can now numerically compute the optimal time threshold, which we shall refer to as $\tau^\star$. Recall the cost function is of the form

$$C_\sigma = T_\sigma + \psi y_\sigma(T_\sigma) = \psi y_\sigma(T_\sigma) + \int_0^{T_\sigma} 1dt$$

where T_σ is the target time to reach a fraction α of the terminal fraction of destinations under the policy $\sigma(t)$,i.e.,

$$T_\sigma = \inf\{t : x_\sigma(t) \geq \alpha a_\infty\}$$

Since in this case, $\sigma(t) = \sigma_{\tau^\star}(t)$, we shall denote the total cost by $C_{\tau^\star}$, and the two components of the cost by $T_{\tau^\star}$ and $y_{\tau^\star}(T_{\tau^\star})$. We obtain the optimal time-threshold ($\tau^\star$) by numerically sweeping over $[0, T_{\max}]$ to search for the cost minimizing value of τ (as in Figure 7), where $T_{\max}$ is the maximum simulation time. In the following discussion we shall numerically study the effect of various cost parameters, system parameters and initial conditions on $\tau^\star$,$C_{\tau^\star}$, $T_{\tau^\star}$ and $y_{\tau^\star}(T_{\tau^\star})$.

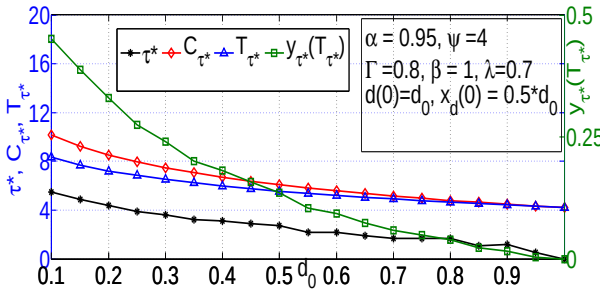

Fig. 8. Decentralized Influence Spread: Effect of d_0 on the optimal cost and optimal τ

5.1 Effect of initial conditions

Figure (8) shows the effect of d_0 on the optimal $\tau^\star$. We see that as we increase d_0 keeping the fraction of $\frac{x_{d_0}}{d_0}$ constant, there are more destinations that have the content. Further, there are fewer relays in the population, and with fixed Γ and β, there is less chance for them to get converted into destinations. This prevents us copying to more relays and hence there is a decrease in the value of $\tau^\star$. Note that when $d_0 = 1$, the entire population consists of only destinations, and hence the optimal $\tau = 0$.

Figure (9) shows the effect of x_{d_0} on the optimal $\tau^\star$. As the fraction $\frac{x_{d_0}}{d_0}$ is increased, since a larger number of destinations have the content, the newly infected relays also obtain the content. Observe that $\frac{x_{d_0}}{d_0} = 1$ implies that all the initial destinations are given the content. This implies that any destination converted in the future, will automatically receive the content. This alleviates the need to copy to any of the relays, since the main purpose of copying to relays is to serve the destinations without the content (either because they were destinations not given the content initially or were converted by other destinations without the content at a later time).

5.2 Effect of system parameters

Figure (10) shows the effect of Γ on $\tau^\star$. We see that as Γ increases, it increases the rate at which relays

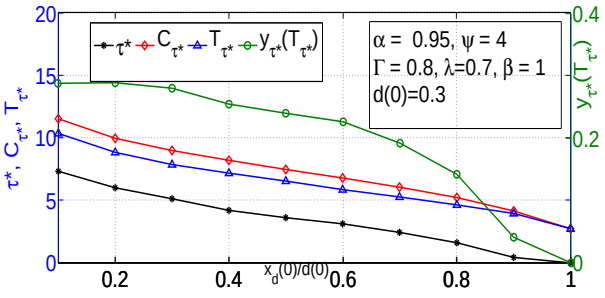

Fig. 9. Decentralized Influence Spread: Effect of $\frac{x_{d_0}}{d_0}$ on the optimal cost and optimal τ

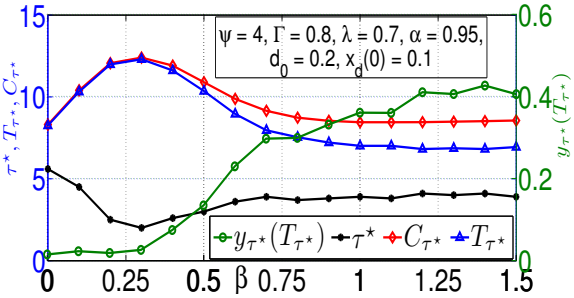

Fig. 11. Decentralized Influence Spread: Effect of β on the optimal cost and optimal τ

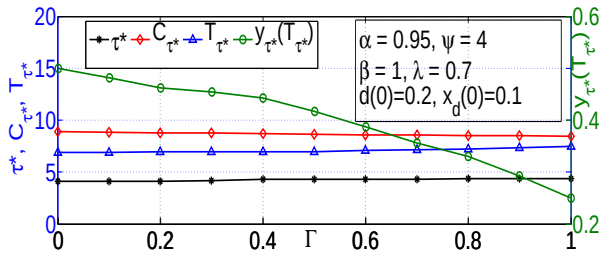

Fig. 10. Decentralized Influence Spread: Effect of Γ on the optimal cost and optimal τ

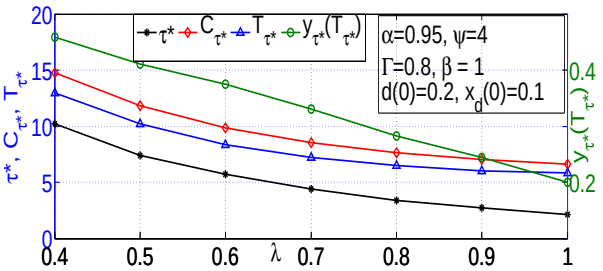

Fig. 12. Decentralized Influence Spread: Effect of λ on the optimal cost and optimal τ

get converted into destinations, without affecting the rate at which content copying occurs (dependent on the meeting rate λ and copying probability $\sigma(t)$). We do not see much effect on $\tau^\star$, and $T_{\tau^\star}$. This is because, though increasing Γ increases the terminal set of destinations (and thus the target $\alpha a(\infty)$), it also helps in achieving the target by the conversion of relays with the contents into destinations with the content. Also due to this conversion, we observe a decrease in $y_\tau(T_\tau)$, further indicated by the fact that $y(t) \leq s(t)$ and the fraction of relays $s(t)$ is decreasing rapidly.

Figure (11) shows the effect of β on $\tau^\star$. We observe that when β is very low, destinations remain active for a longer duration before recovery. This leads to a faster decay in $s(t)$, the fraction of relays, and also an increase in the target $\alpha a(\infty)$. The larger target requires that copying to relays must continue for longer, and the faster decay in the fraction of relays prevents an excessive number of copies being eventually wasted in spite of copying for longer. Thus, a small β results in a large value of τ^*. For large β, since active destinations stay active for a shorter duration, the eventual population of destinations is small. Hence, the population of destinations carrying the content is small, thus requiring the relays to play a greater role in the spread of the content. However, again, since β is large these relays are less likely to get converted into destinations, thus explaining the increase in $y_\tau(T_\tau)$ for large values of β.

Figure (12) shows the effect of λ on the optimal $\tau^\star$. As λ increases, the rate at which content spread increases, and thus a destination in need of the content is highly likely to obtain it from another destination. This results in a more passive policy of copying to relays, leading to a decrease in the optimal $\tau^\star$ and the corresponding costs.

5.3 Effect of cost parameters

Figure (13) shows the effect of α on $\tau^\star$. For the given system parameters, $a_\infty = 0.4352$ and hence for $\alpha < 0.23$, we see that the target is already achieved at time zero (since the initial seeding $x_d(0) = 0.1$). Thus for $\alpha < 0.23$, $\tau^\star = 0$,$T_{\tau^\star} = 0$ and $y_{\tau^\star}(T_{\tau^\star}) = 0$. As α increases, we see that $\tau^\star$ increases and in turn increases the costs, since we need to continue copying to relays for a longer duration to achieve the target. It is also to be noted that as $\alpha \to 1$, T_τ approaches ∞, since by definition $x(t)$ reaches $a(\infty)$ only asymptotically.

Figure (14) shows the effect of ψ on $\tau^\star$. As ψ increases, we see that the emphasis on the wasted copies, $y(T)$, increases, and hence the control needs to be *less aggressive*. Thus we see a decrease in $\tau^\star$ as ψ increases, and this leads to a decrease in the wasted copies ($y_{\tau^\star}(T_{\tau^\star})$) and an increase in the delay of reaching the target ($T_{\tau^\star}$). We see an increase in the total cost $C_{\tau^\star}$ since both the terms in the cost function ($T_{\tau^\star}$ and $\psi y_{\tau^\star}(T_{\tau^\star})$) are increasing.

6 Conclusions

In this paper we studied the joint evolution of popularity and spread of content in a mobile opportunistic

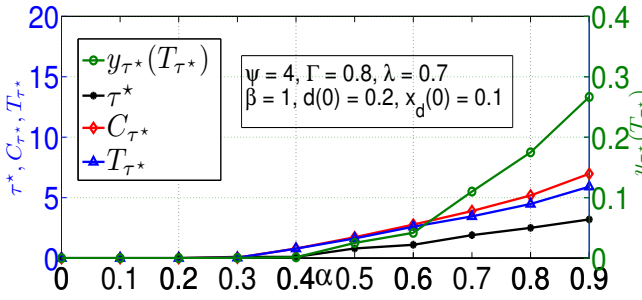

Fig. 13. Decentralized Influence Spread: Effect of α on the optimal cost and optimal τ

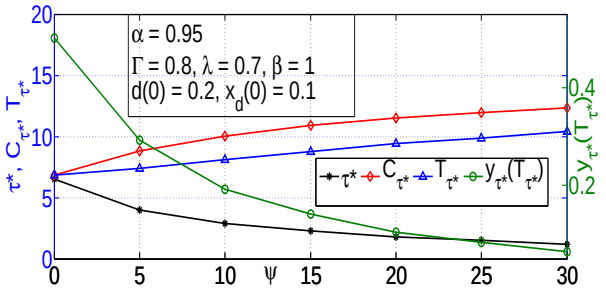

Fig. 14. Decentralized Influence Spread: Effect of ψ on the optimal cost and optimal τ

setting. We proposed a possible application framework for such a system. We developed a continuous time Markov model for the joint evolution and derived its fluid limit. Finally we showed that a time-threshold policy is an optimal copying policy for the joint evolution model, to optimize the combined cost of target time and the amount of wasted copies. We also have reported several interesting insights into the evolution of popularity and the co-evolution of popularity and content delivery, which will help the content producers and distributors understand the interplay of various system parameters. By allowing users to declare the interest and enabling the spread of interest, we extend the notion of a personal wishlist (in e-commerce platforms like Amazon) to a social wishlist. The application also combine the services of coupon delivery (Groupon) with such wishlists, thus allowing targeted discount delivery, wherein brands can identify the users who are interested in the coupon and deliver it to them.

7 Proofs

7.1 Verification of Kurtz's theorem conditions for SIR-SI model

Kurtz's theorem [10] is used to approximate pure jump Markov processes by ordinary differential equations in the limit (usually as system size $N \to \infty$). To do so, we must derive the mean drift rates of an appropriately scaled version of the system, and study the limit of these mean drift rates as the scaling goes to ∞. In our case, for the joint evolution model described in Section 3, we can use the Equation 1 and Table 2 to write down the mean drift rates. Let $Z = (B, D, X_b, X_d, Y)$ denote the current state of the CTMC $Z^N(t)$ and let $\tilde{Z} = \frac{Z}{N}$. The expected drift rate of the CTMC $\tilde{Z}^N(t)$, $F^N(\tilde{Z}) := (F_b^N(\tilde{Z}), F_d^N(\tilde{Z}), F_{x_b}^N(\tilde{Z}), F_{x_d}^N(\tilde{Z}), F_y^N(\tilde{Z}))$ can then be written as:

$$
\begin{aligned}
F_b^N(\tilde{Z}) &= \beta_N \tilde{D} \\
F_d^N(\tilde{Z}) &= -\beta_N \tilde{D} + N\lambda_N \Gamma \tilde{D}\tilde{S} \\
F_{x_b}^N(\tilde{Z}) &= \beta_N \tilde{X}_d + \\
&\quad N\lambda_N(\tilde{B} - \tilde{X}_b)(\tilde{X} + \tilde{Y}) \\
F_{x_d}^N(\tilde{Z}) &= \Gamma N\lambda_N(\tilde{D} - \tilde{X}_d)\tilde{Y} - \\
&\quad \beta_N \tilde{X}_d + \Gamma N\lambda_N \tilde{X}_d \tilde{S} + \\
&\quad N\lambda_N(\tilde{D} - \tilde{X}_d)(\tilde{X} + \tilde{Y}) \\
F_y^N(\tilde{Z}) &= -\Gamma N\lambda_N \tilde{D}\tilde{Y} + \\
&\quad N\lambda_N \sigma(\tilde{S} - \tilde{Y})(\tilde{X}_b + \tilde{Y} + (1-\Gamma)\tilde{X}_d)
\end{aligned}
$$

and the corresponding o.d.e. equations,

$$
\begin{aligned}
\dot{b} &= \beta d && (9)\\
\dot{d} &= -\beta d + \lambda\Gamma ds && (10)\\
\dot{x_b} &= \beta x_d + \lambda(b - x_b)(x + y) && (11)\\
\dot{x_d} &= \Gamma\lambda(d - x_d)y + \Gamma\lambda x_d s + && (12)\\
&\qquad \lambda(d - x_d)(x + y) - \beta x_d \\
\dot{y} &= -\Gamma\lambda dy + \lambda\sigma(s - y)(x_b + y + (1-\Gamma)x_d) && (13)
\end{aligned}
$$

where $s(t) = 1 - b(t) - d(t)$ and $x(t) = x_b(t) + x_d(t)$. Let $f_b, f_d, f_{x_b}, f_{x_d}, f_y$ denote the right hand sides of the above fluid limit equations and let $f := (f_b, f_d, f_{x_b}, f_{x_d}, f_y)$. The following four conditions are shown to be equivalent [2] to the necessary conditions for Kurtz's theorem.

(i) **Lipschitz property** We see that each of the partial derivatives $\frac{\partial f_u}{\partial v}$ where $u, v \in (b, d, x_b, x_d, y)$ is bounded when $(b, d, x_b, x_d, y) \in [0,1] \times [0, 1-b] \times [0, b] \times [0, d] \times [0, 1-b-d]$. Thus the norm of Jacobian is uniformly bounded, and it follows that $f(z)$ is Lipschitz.

(ii) **Uniform Convergence** Taking $N\lambda_N \to \lambda$ and $\beta_N \to \beta$ we see that the uniform convergence of $F^N(z)$ to $f(z)$ is straightforward.

(iii) **Bounded Noise variance** Since in the co-evolution system the maximum jump rate out of a state is bounded $\lambda_N N(N-1) + \beta_N N$, and since $N\lambda_N \to \lambda$ and $\beta_N \to \beta$, the jump rate is $O(N)$; the increments for the scaled process $\tilde{Z}^N(t)$ are $O(N^{-1})$. This is referred to as "hydrodynamic scaling" [2]. This ensures that the noise convergence to zero is ensured.

(iv) **Convergence of initial conditions** The initial conditions are chosen such that

$$(\tilde{B}^N(0), \tilde{D}^N(0), \tilde{X}_b^N(0), \tilde{X}_d^N(0), \tilde{Y}^N(0)) = (b(0), d(0), x_b(0), x_d(0), y(0))$$

Since $F^N(.)$ and $f(.)$ satisfy the above four conditions, by [2], Theorem 1 follows. ■

7.2 Kamke-dominance

Let $\mathbf{w}(t)$ be the solution of the o.d.e $\dot{\mathbf{w}} = f(\mathbf{w}; z)$ with piecewise Lipschitz continuous control $z(t)$, where f is continuously differentiable and Lipschitz in $\mathbf{w}$ and z. Let $\mathbf{w}^{(1)}(t)$ and $\mathbf{w}^{(2)}(t)$ be the trajectories corresponding to two controls $z^{(1)}(t)$ and $z^{(2)}(t)$ respectively, i.e.,

$$\dot{\mathbf{w}}^{(1)} = f(\mathbf{w}^{(1)}; z^{(1)}) \quad (14)$$

$$\dot{\mathbf{w}}^{(2)} = f(\mathbf{w}^{(2)}; z^{(2)}) \quad (15)$$

Motivated by the Kamke condition (see [15]) we define Kamke-dominance between two controls $z^{(1)}(t)$ and $z^{(2)}(t)$ as follows.

Definition 1: We say $z^{(1)}(t)$ *Kamke-dominates* $z^{(2)}(t)$ *in the system* $f(\mathbf{w}; z)$ if for each t_0 where $\mathbf{w}^{(1)}(t_0) \geq \mathbf{w}^{(2)}(t_0)$ and $\exists i$ with $w_i^{(1)}(t_0) = w_i^{(2)}(t_0)$ then

$$f_i(\mathbf{w}^{(1)}(t_0); z^{(1)}(t_0)) \geq f_i(\mathbf{w}^{(2)}(t_0); z^{(2)}(t_0))$$

Lemma 1: Let $\mathbf{w}^{(1)}(t)$ and $\mathbf{w}^{(2)}(t)$ be as in Equations (14) and (15), and $z^{(1)}(t)$ Kamke-dominate $z^{(2)}(t)$ in the system $f(\mathbf{w}; z)$. If $\mathbf{w}^{(1)}(0) \geq \mathbf{w}^{(2)}(0)$ then $\forall t$, $\mathbf{w}^{(1)}(t) \geq \mathbf{w}^{(2)}(t)$.

Proof: The proof of Proposition 1.1 in [15] can be adopted directly in writing this proof. Since f is continuously differentiable and Lipschitz in $\mathbf{w}$ and z, by the Cauchy-Lipschitz condition, we have a unique solution given the initial conditions. Let $\phi_t^{(1)}(\mathbf{w}^{(1)}(0))$ denote the solution of (14) with initial condition $\mathbf{w}^{(1)}(0)$, and $\phi_t^{(2)}(\mathbf{w}^{(2)}(0))$ denote the solution of (15) with initial condition $\mathbf{w}^{(2)}(0)$. We wish to show that $\phi_t^{(1)}(\mathbf{w}^{(1)}(0)) \geq \phi_t^{(2)}(\mathbf{w}^{(2)}(0))$. Consider m an integer, and let $\phi_t^{(1),m}(\mathbf{w}^{(1)}(0))$ be the solution corresponding to

$$\dot{\mathbf{w}}^{(1)} = f(\mathbf{w}^{(1)}; z^{(1)}) + (\frac{1}{m})\mathbf{e}$$

Then, by continuity of o.d.e. solutions with respect to drift and initial conditions [6, Chap.1, Lemma 3.1], $\phi_s^{(1),m}(\mathbf{w}^{(1)}(0) + \frac{\mathbf{e}}{m})$ defined on $0 \leq s \leq t$ for all large m, say $m > M$ and

$$\phi_s^{(1),m}(\mathbf{w}^{(1)}(0) + \frac{\mathbf{e}}{m}) \to \phi_s^{(1)}(\mathbf{w}^{(1)}(0)) \quad (16)$$

as $m \to \infty$, uniformly in $s \in [0, t]$, where $\mathbf{e} = (1, 1, \cdots, 1)$. We claim that for $0 \leq s \leq t$, for all $m > M$,

$$\phi_s^{(1)m}(\mathbf{w}^{(1)}(0) + \frac{\mathbf{e}}{m}) >> \phi_s^{(2)}(\mathbf{w}^{(2)}(0)) \quad (17)$$

where $x >> y$ implies $x_i > y_i$, $\forall i$.

Proof of claim: Fix $m > M$. We know that $\mathbf{w}^{(1)}(0) + \frac{\mathbf{e}}{m} = \phi_0^{(1),m}(\mathbf{w}^{(1)}(0) + \frac{\mathbf{e}}{m}) >> \mathbf{w}^{(2)}(0) = \phi_0^{(2)}(\mathbf{w}^{(2)}(0))$. By continuity of $\phi^{(1)}$ and $\phi^{(2)}$, the claim (17) holds for small s. If the claim (17) were false, $\exists\ 0 < t_0 \leq t$ such that (17) holds for $0 \leq s < t_0$, and an index i such that,

$$\left(\phi_{t_0}^{(1),m}(\mathbf{w}^{(1)}(0) + \frac{\mathbf{e}}{m})\right)_i = \left(\phi_{t_0}^{(2)}(\mathbf{w}^{(2)}(0))\right)_i$$

and

$$\left.\frac{\mathrm{d}}{\mathrm{d}s}\right|_{s=t_0} \left(\phi_s^{(1),m}(\mathbf{w}^{(1)}(0)+\frac{\mathbf{e}}{m})\right)_i \leq \left.\frac{\mathrm{d}}{\mathrm{d}s}\right|_{s=t_0} \left(\phi_s^{(2)}(\mathbf{w}^{(2)}(0))\right)_i \quad (18)$$

But since $\forall j \neq i$

$$\left(\phi_{t_0}^{(1)m}(\mathbf{w}^{(1)}(0) + \frac{\mathbf{e}}{m})\right)_j \geq \left(\phi_{t_0}^{(2)}(\mathbf{w}^{(2)}(0))\right)_j$$

and $z^{(1)}(t)$ Kamke-dominates $z^{(2)}(t)$ in the system $f(\mathbf{w}, z)$, we have

$$\begin{aligned} f_i(\phi_{t_0}^{(1),m}(\mathbf{w}^{(1)}(0) + \frac{\mathbf{e}}{m}); z^{(1)}(t_0)) + \frac{1}{m} \\ > f_i(\phi_{t_0}^{(1),m}(\mathbf{w}^{(1)}(0) + \frac{\mathbf{e}}{m}); z^{(1)}(t_0)) \\ \geq f_i(\phi_{t_0}^{(2)}(\mathbf{w}^{(2)}(0)); z^{(2)}(t_0)) \end{aligned}$$

which implies

$$\left.\frac{\mathrm{d}}{\mathrm{d}s}\right|_{s=t_0} \left(\phi_s^{(1),m}(\mathbf{w}^{(1)}(0)+\frac{\mathbf{e}}{m})\right)_i > \left.\frac{\mathrm{d}}{\mathrm{d}s}\right|_{s=t_0} \left(\phi_s^{(2)}(\mathbf{w}^{(2)}(0))\right)_i$$

This contradicts (18) and hence proves the claim (17). Applying (16) to this claim completes the proof of the lemma. □

Consider the system of equations representing the evolution of $(x_b(t), x_d(t), y(t))$ in the SIR-SI process.

$$\dot{x_b} = \beta x_d + \lambda(b - x_b)(x + y) \quad (19)$$

$$\dot{x_d} = \Gamma\lambda(d - x_d)y + \Gamma\lambda x_d s + \lambda(d - x_d)(x + y) - \beta x_d \quad (20)$$

$$\dot{y} = -\Gamma\lambda dy + \lambda\sigma(s - y)(x_b + y + (1 - \Gamma)x_d) \quad (21)$$

For identifying the optimal control $\sigma(t)$ it is sufficient to restrict our attention to Equations (19)-(21), since by the definition of the SIR-SI process, $(b(t), d(t))$ are independent of the control, i.e., the popularity evolution is fixed and not influenced by the content dissemination. Let $w^{(1)} = (x_b^{(1)}, x_d^{(1)}, y^{(1)})$ and $w^{(2)} = (x_b^{(2)}, x_d^{(2)}, y^{(2)})$ be the solutions respectively for the controls $\sigma^{(1)}(t)$ and $\sigma^{(2)}(t)$ for the above system with $\sigma^{(1)}(t) \leq \sigma^{(2)}(t)$. Define

$$(u_b, u_d, v) := (x_b^{(1)}, x_d^{(1)}, y^{(1)}) - (x_b^{(2)}, x_d^{(2)}, y^{(2)})$$

and $\Delta\sigma(t) = \sigma^{(1)}(t) - \sigma^{(2)}(t)$. Then we can write,

$$\dot{u}_b|_{u_b=0} = \beta u_d + \lambda(b - x_b^{(1)})(u_d + v) \quad (22)$$

$$\dot{u}_d|_{u_d=0} = \Gamma\lambda(d - x_d^{(1)})v + \lambda(d - x_d^{(1)})(u_b + v) \quad (23)$$

$$\dot{v}|_{v=0} = \lambda(s - y^{(1)})[\Delta\sigma y^{(1)} + \sigma^{(1)}u_b + \Delta\sigma x_b^{(2)} + (1 - \Gamma)(\sigma^{(1)}u_d + \Delta\sigma x_d^{(2)})] \quad (24)$$

Lemma 2: Control Domination

(a) If $\forall t$, $\sigma^{(1)}(t) \geq \sigma^{(2)}(t)$, then $\sigma^{(1)}(t)$ Kamke-dominates $\sigma^{(2)}(t)$ in the system given by equations (19)-(21).

(b) If $\forall t$, $y^{(1)}(t) \geq y^{(2)}(t)$, then $y^{(1)}(t)$ Kamke-dominates $y^{(2)}(t)$ in the system given by equations (19)-(20).

Proof:

(a) From equations (22)-(24). Since $\Delta\sigma \geq 0$, we find that,
- if $u_d, v \geq 0$ and $u_b = 0$, then $\dot{u}_b \geq 0$
- if $u_b, v \geq 0$ and $u_d = 0$, then $\dot{u}_d \geq 0$
- if $u_b, u_d \geq 0$ and $v = 0$ $\dot{v} \geq 0$

This verifies the conditions of Definition 1, and thus proves (a).

(b) From equations (22)-(23). Since $v \geq 0$, we find that,
- if $u_d \geq 0$ and $u_b = 0$, then $\dot{u}_b \geq 0$
- if $u_b \geq 0$ and $u_d = 0$, then $\dot{u}_d \geq 0$

This verifies the conditions of Definition 1, and thus proves (b).

□

Lemma 3: (a) If $\forall t$, $\sigma^{(1)}(t) \geq \sigma^{(2)}(t)$, and $w^{(1)}(0) \geq w^{(2)}(0)$, then $\forall t$, $w^{(1)}(t) \geq w^{(2)}(t)$

(b) If $\forall t$, $y^{(1)}(t) \geq y^{(2)}(t)$, and $(x_b^{(1)}(0), x_d^{(1)}(0)) \geq (x_b^{(2)}(0), x_d^{(2)}(0))$, then $\forall t$,
$$(x_b^{(1)}(t), x_d^{(1)}(t)) \geq (x_b^{(2)}(t), x_d^{(2)}(t))$$

Proof: The proof follows by applying Lemma 1 to the systems in Lemma 2 (a) and (b). □

Define $\mathbf{w}(t) = (x_b(t), x_d(t), y(t))$ We will now use the above lemmas to prove the claims in Section 4.1, in order to establish the optimality of a time-threshold policy. Recall that α has been fixed earlier. Recall the definitions of T_σ and $y_\sigma(T_\sigma)$. We shall replace them with T_τ and $y_\tau(T_\tau)$ whenever $\sigma(t) = \sigma_\tau(t)$, a time-threshold policy.

Consider τ such that $\tau \geq T_\tau =: \tau'$. Evidently, $T_{\tau'} = \tau'$, and $y_\tau(T_\tau) = y_{\tau'}(T_{\tau'})$, $\forall \tau \geq \tau'$.

Define $\mathcal{T} = \{\tau : \tau \geq 0\}$ and $\underline{\mathcal{T}} = \{\tau : \tau \leq T_\tau\} \subset \mathcal{T}$. Observe that $\{\rho : \exists\ \tau \in \mathcal{T}, y_\tau(T_\tau) = \rho\} = \{\rho : \exists\ \tau \in \underline{\mathcal{T}}, y_\tau(T_\tau) = \rho\}$. Hence we limit our discussion to $\tau \in \underline{\mathcal{T}}$.

Lemma 4: The set $\{\rho : \exists\ \tau \in \underline{\mathcal{T}}, y_\tau(T_\tau) = \rho\}$ forms an interval of the form $[0, \rho_{\max}]$.

Proof: Setting $\tau = 0$, i.e., $\sigma(t) = 0$, $\forall t$, we see from Equation (21), since $y(0) = 0$, we have $y(t) = 0$, $\forall t$. This establishes that $\rho = 0$ belongs to the set. It is also easy to observe that $\rho_{\max}$ is achieved by τ such that $\tau = T_\tau$. By the intermediate value theorem, it now suffices to show that $y_\tau(T_\tau)$ is a continuous function of τ.

Consider $\tau, \tau + \delta \in \underline{\mathcal{T}}$, i.e., $\tau \leq T_\tau$, $\tau + \delta \leq T_{\tau+\delta}$. Observe from Lemma 3 , for $\delta > 0$, $T_{\tau+\delta} \leq T_\tau$. With the above constraints, evidently the only case to be considered is
$$\tau < \tau + \delta \leq T_{\tau+\delta} \leq T_\tau$$

Let $\mathbf{w}_\tau(.)$ and $\mathbf{w}_{\tau+\delta}(.)$ be the trajectories corresponding to $\sigma_\tau(t)$ and $\sigma_{\tau+\delta}(t)$, with $\mathbf{w}_\tau(0) = \mathbf{w}_{\tau+\delta}(0)$. Since $\sigma_\tau(t) = \sigma_{\tau+\delta}(t)$, $0 \leq t \leq \tau$, it follows that $\mathbf{w}_\tau(\tau) = \mathbf{w}_{\tau+\delta}(\tau)$. Further, from Equations (19)-(21) we observe that,

$||\mathbf{w}_\tau(\tau + \delta) - \mathbf{w}_{\tau+\delta}(\tau + \delta)||_2 \leq C_1\delta$, where
$C_1 = \sqrt{(\beta + \lambda)^2 + (2\Gamma\lambda + \lambda + \beta)^2 + (\lambda + \Gamma\lambda)^2}$

$\forall t > \tau + \delta$, $\sigma_\tau(t) = \sigma_{\tau+\delta}(t)$ and hence by continuity of o.d.e. system w.r.t initial conditions at $\tau + \delta$ implied by the Lipschitz nature of f w.r.t $\mathbf{w}$ (see [1]), $\mathbf{w}_\tau(.)$ is continuous w.r.t τ.

Recall that $\tau < \tau + \delta \leq T_{\tau+\delta} \leq T_\tau$. By the just observed continuity, for $\epsilon > 0$, we can obtain $\delta > 0$ such that, $|x_{\tau+\delta}(T_{\tau+\delta}) - x_\tau(T_{\tau+\delta})| \leq \epsilon$, i.e.,

$$|\alpha a(\infty) - x_\tau(T_{\tau+\delta})| \leq \epsilon$$

However over the interval $(T_{\tau+\delta}, T_\tau]$, the rate of increase of $x_\tau(.)$ is bounded below by $\lambda a(t)^2(1-\alpha)\alpha$ (from Equations (19)-(21)). Hence,

$$|T_{\tau+\delta} - T_\tau| \leq \frac{\epsilon}{\lambda a(t)^2(1-\alpha)\alpha}$$

which can be made small by choosing an appropriate $\delta > 0$. We have thus shown that T_τ continuous w.r.t τ. $\tau \in \underline{\mathcal{T}}$.

To show the continuity of $y_\tau(T_\tau)$, consider,

$|y_\tau(T_\tau) - y_{\tau+\delta}(T_{\tau+\delta})| \leq$
$$|y_\tau(T_\tau) - y_\tau(T_{\tau+\delta})| + |y_\tau(T_{\tau+\delta}) - y_{\tau+\delta}(T_{\tau+\delta})|$$

In the above equation, the first term on the right hand side can be made arbitrarily small by the continuity of T_τ w.r.t τ and the fact that $y_\tau(.)$ is a continuous trajectory, and the second term can be made arbitrarily small by the continuity of $\mathbf{w}_\tau(.)$ w.r.t τ. Thus $y_\tau(T_\tau)$ is a continuous function of τ, in the space of threshold policies. And since $y_\tau(T_\tau) : \tau \to \rho$ and $\tau \in [0, \infty)$, we see that $\rho \in [0, \rho_{max}]$. □

Lemma 5: Let $\sigma(t)$ be a policy such that $y_\sigma(T_\sigma) < \rho_{\max}$. Then there exists a threshold policy $\sigma_\tau(t)$ whose cost is no worse than that of $\sigma(t)$.

Proof: Since $\rho := y_\sigma(T_\sigma) < \rho_{\max}$, there exists $\tau \geq 0$ such that $y_\tau(T_\tau) = \rho$. We will argue that $\sigma_\tau(t)$ is such that $T_\tau \leq T_\sigma$.

Let $\mathbf{w}_\tau(t)$ and $\mathbf{w}_\sigma(t)$ be the system trajectories for the controls $\sigma_\tau(t)$ and $\sigma(t)$ respectively with $\mathbf{w}_\tau(0) = \mathbf{w}_\sigma(0)$. By definition, for $t \leq \tau$, $\sigma_\tau(t) \geq \sigma(t)$ and hence, by Lemma 3, $\mathbf{w}_\tau(\tau) \geq \mathbf{w}_\sigma(\tau)$.

Suppose, contrary to the claim, $T_\tau > T_\sigma$. Evidently, this cannot happen if for all t, $0 \leq t \leq T_\sigma$, $y_\tau(t) \geq y_\sigma(t)$. This is because, by Lemma 3, we will have $x_\tau(t) \geq x_\sigma(t)$ and hence $T_\tau \leq T_\sigma$.

Hence, there exists t_0, $\tau \leq t_0 < T_\sigma$, such that $y_\tau(t_0) = y_\sigma(t_0)$. Then we have $y_\tau(t_0) = y_\sigma(t_0) > \rho$ (since $y_\tau(T_\tau) = \rho$ and $\tau \leq t_0 < T_\tau$, and for $t \geq \tau$, $\dot{y}_\tau(t) < 0$). Also, for $t \geq t_0$, $\sigma_\tau(t) = 0 \leq \sigma(t)$. Thus from (21), $\forall s \in \{s : s > t_0, y_\tau(s) = y_\sigma(s)\}$, $\dot{y}_\tau(s) \leq \dot{y}_\sigma(s)$. Thus for $s \geq t_0$, $y_\tau(s) \leq y_\sigma(s)$. Since $y_\tau(T_\tau) = y_\sigma(T_\sigma) = \rho$, this implies $T_\tau \leq T_\sigma$.

□

Lemma 6: Consider a non-threshold policy $\sigma(t)$ which achieves $\rho > \rho_{max}$. Consider the time-threshold policy $\hat{\sigma}(t)$ of the form,

$$\hat{\sigma}(t) = \begin{cases} 1, & 0 \leq t < \sup\{t : \sigma(t) > 0\} \\ 0, & \text{otherwise} \end{cases}$$

Then $\hat{\sigma}(t)$ policy achieves a smaller total cost (given by Equation 7) than $\sigma(t)$.

Proof: Let $\sigma(t)$ and $\sigma_{\hat{\tau}}(t)$ be as above. Let $\mathbf{w}_{\hat{\tau}}(t)$ and $T_{\hat{\tau}}$ be as defined earlier for time-threshold policies, corresponding to $\sigma_{\hat{\tau}}(t)$. Then by Lemma 4, $y_{\hat{\tau}}(T_{\hat{\tau}}) \leq \rho_{max} < \rho = y_{\sigma}(T_{\sigma})$, and by Lemma 3, $T_{\hat{\tau}} \leq T_{\sigma}$, and hence the $\sigma_{\hat{\tau}}(t)$ policy achieves a smaller total cost than $\sigma(t)$. □

8 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] V. Borkar, *Stochastic approximation: a dynamical systems viewpoint*. Cambridge Univ Pr, 2008.

[2] R. Darling, "Fluid limits of pure jump markov processes: A practical guide," *ArXiv preprint arXiv:math/0210109*, 2002.

[3] N. Gast, B. Gaujal, and J. Boudec, "Mean field for markov decision processes: from discrete to continuous optimization," *Arxiv preprint arXiv:1004.2342*, 2010.

[4] N. Golrezaei, K. Shanmugam, A. G. Dimakis, A. F. Molisch, and G. Caire, "Femtocaching: Wireless video content delivery through distributed caching helpers," in *INFOCOM, 2012 Proceedings IEEE*. IEEE, 2012, pp. 1107–1115.

[5] R. Groenevelt, P. Nain, and G. Koole, "The message delay in mobile ad hoc networks," *Performance Evaluation*, vol. 62, no. 1, pp. 210–228, 2005.

[6] J. Hale, *Ordinary Differential Equations*. RE Krieger, Malabar, FL, 1980.

[7] Ó. R. Helgason, E. A. Yavuz, S. T. Kouyoumdjieva, L. Pajevic, and G. Karlsson, "A mobile peer-to-peer system for opportunistic content-centric networking," in *Proceedings of the second ACM SIGCOMM workshop on Networking, systems, and applications on mobile handhelds*. ACM, 2010, pp. 21–26.

[8] S. D. Kamvar, M. T. Schlosser, and H. Garcia-Molina, "Incentives for combatting freeriding on p2p networks," in *Euro-Par 2003 Parallel Processing*. Springer, 2003, pp. 1273–1279.

[9] 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.

[10] T. Kurtz, "Solutions of ordinary differential equations as limits of pure jump markov processes," in *Journal of Applied Probability*. JSTOR, 1970, pp. 49–58.

[11] R. Ramanathan, R. Hansen, P. Basu, R. Rosales-Hain, and R. Krishnan, "Prioritized epidemic routing for opportunistic networks," in *Proceedings of the 1st international MobiSys workshop on Mobile opportunistic networking*. ACM, 2007, pp. 62–66.

[12] T. Repantis and V. Kalogeraki, "Data dissemination in mobile peer-to-peer networks," in *Proceedings of the 6th international conference on Mobile data management*. ACM, 2005, pp. 211–219.

[13] 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.

[14] C. Singh, E. Altman, A. Kumar, and R. Sundaresan, "Optimal forwarding in delay-tolerant networks with multiple destinations," *Networking, IEEE/ACM Transactions on*, vol. PP, no. 99, pp. 1–1, 2013.

[15] H. Smith, *Monotone dynamical systems: An introduction to the theory of competitive and cooperative systems*. American Mathematical Soc., 2008.

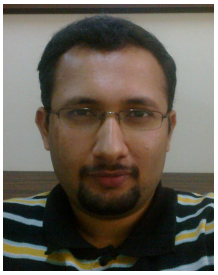

Srinivasan Venkatramanan obtained his B.E. degree in Electronics and Communication Engineering from the College of Engineering Guindy, Anna University in 2008. He is currently a PhD student under Prof. Anurag Kumar at the Department of Electrical Communication Engineering, Indian Institute of Science. His area of research is social information networks, specifically, modeling, analysis and control of information dissemination on social information networks and mobile peer-peer networks. He is also interested in characterizing collective behavior on real-world networks such as social networks, societal networks (transportation, power grids) and biological networks.

Anurag Kumar obtained his B.Tech. degree from the Indian Institute of Technology at Kanpur, and the PhD degree from Cornell University, both in Electrical Engineering. He was then with Bell Laboratories, Holmdel, N.J., for over 6 years. Since 1988 he has been with the Indian Institute of Science (IISc), Bangalore, in the Department of Electrical Communication Engineering, where he is now a Professor. He is currently also the Chair of the Electrical Sciences Division at IISc. From 1988 to 2003 he was the Coordinator at IISc of the Education and Research Network Project (ERNET), India's first wide-area packet switching network. His area of research is communication networking, specifically, modeling, analysis, control and optimisation problems arising in communication networks and distributed systems. Recently his research has focused primarily on wireless networking. He is a Fellow of the IEEE, of the Indian National Science Academy (INSA), of the Indian Academy of Science (IASc), and of the Indian National Academy of Engineering (INAE). He is a coauthor of the postgraduate text-books "Communication Networking: An Analytical Approach," and "Wireless Networking" both by Kumar, Manjunath and Kuri, published by Elsevier.