Delay-cost Optimal Coupon Delivery in Mobile Opportunistic Networks

Srinivasan Venkatramanan
Joint work with Prof. Anurag Kumar

Department of ECE, IISc.

January 9, 2014

Mobile Opportunistic Networks (MON)

- Proliferation of smart mobile devices
  - Smartphones, tablets, wearable computing devices, etc.
- Mobile content delivery - a challenge
  - Direct delivery - not scalable
  - Alternative: p2p delivery (via Bluetooth, WiFi-Direct, etc.)
- A single item of content may be of **interest** to several co-located users
  - e.g., slides of a conference keynote/ course lecture
- The item can be forwarded between devices when their p2p radio interfaces make contact
  - Akin to **epidemic** spread
  - A multi-hop opportunistic mobile network
  - Provides an approach to *delay tolerant* networking

Mobile Coupon Delivery Ecosystem

- **Scenario:** Pre-release promotions of a product (movie, book, etc.)
- **Content:** discount coupon for pre-ordering the product
- Mobile users can
  - maintain a wishlist of products of interest
  - share wishlists with peers
  - receive coupons from peers or central server
- wishlists + peer recommendation + coupon delivery

A Possible Application Framework

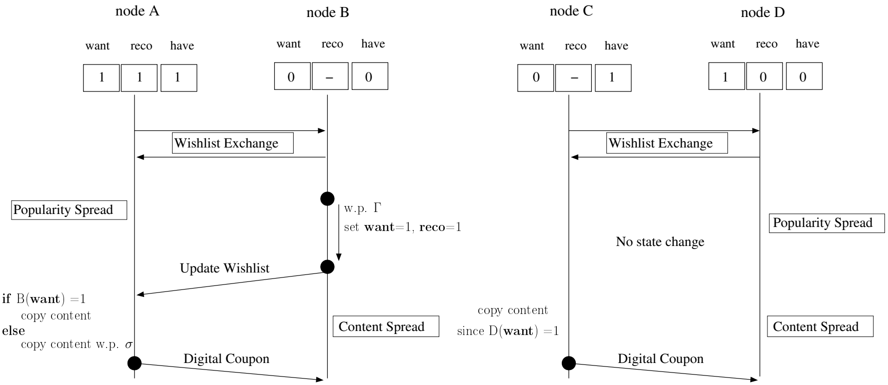

Content Popularity and Dissemination: SIR-SI model

- Population size: N (fixed)
- Pairwise meetings at points of independent Poisson processes
- The coupon needs to reach certain *destination* nodes
- Some destination nodes are given the coupon initially
- The set of destination nodes grows, as more nodes express their interest
- *Do-not-want* nodes could help in forwarding (relays)

Objective: Quickly spread content to a large fraction of destinations nodes while minimizing the residual number of relays that have the content

SIR-SI States and Evolution

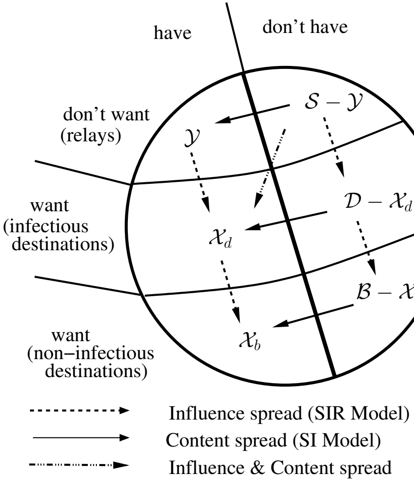

- λ_N: Poisson meeting rate for each pair of nodes; $\lambda_N = \frac{\lambda}{N}$
- β_N: Recovery rate of infectious destinations; $\beta_N = \beta$
- Γ: Influence probability
- σ: Copying probability to a relay (control)

Fluid Limits - Kurtz Theorem

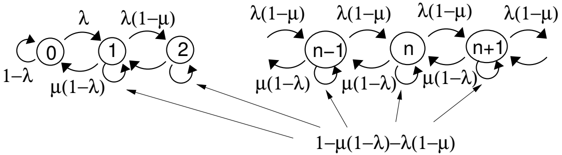

- Queue length process, $X(k), k \geq 0$, $\bar{X}^{(N)}(t) := \frac{1}{N} X(\lfloor Nt \rfloor)$
- ODE: $\dot{x}(t) = (\lambda - \mu) I_{\{x(t)>0\}}$ with $x(0) = \bar{X}^{(N)}(0)$, for each N

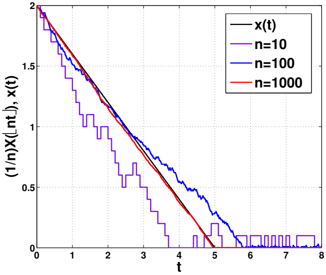

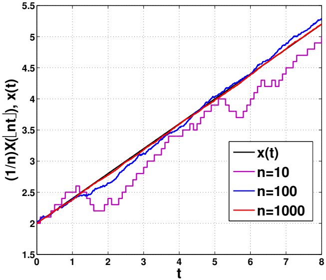

SIR-SI CTMC Markov Chain: Transitions at Meeting Epochs

- Let $k = 0, 1, 2, \cdots$ index the meeting/recovery epochs at times t_k
- System state: $Z(k) = (B(k), D(k), X_b(k), X_d(k), Y(k))$

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

$$
\begin{aligned}
\dot{b} &= \beta d \\
\dot{d} &= -\beta d + \lambda \Gamma d s \\
\dot{x}_b &= \beta x_d + \lambda(b - x_b)(x + y) \\
\dot{x}_d &= \lambda(d - x_d)(x + y) + \lambda \Gamma d y + \Gamma \lambda x_d(s - y) - \beta x_d \\
\dot{y} &= -\Gamma \lambda d y + \lambda \sigma (s - y)(x_b + y + (1 - \Gamma)x_d)
\end{aligned}
$$

where $s(t) = 1 - b(t) - d(t)$

Convergence of the CTMC to O.D.E. Limit

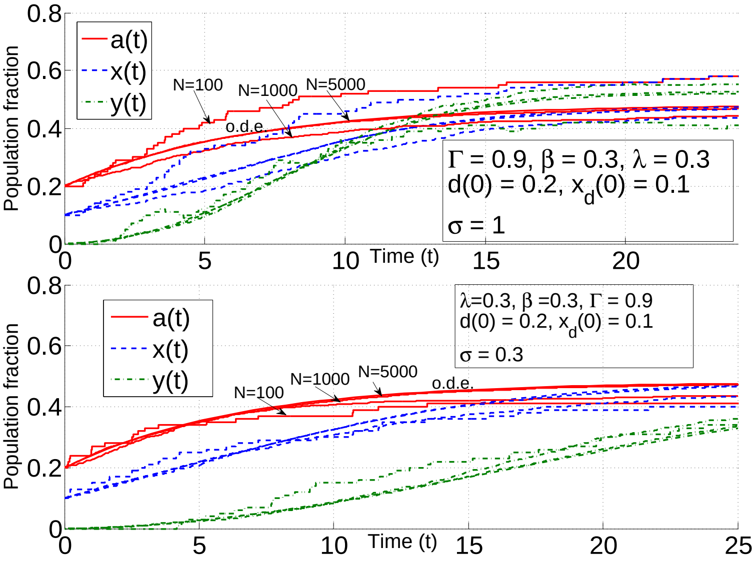

SIR-SI Model: Dynamic Control of Copying

- Dynamic control $\sigma : Z(k) \rightarrow [0, 1]$
- CTMDP for each N: Obtaining optimal control is difficult
- Replace probabilistic control σ in the ODE by $\sigma(t)$ (controlled ODE)
- Optimal (deterministic, open-loop) control for the controlled ODE
- Can be shown to be asymptotically optimal for the finite size problem

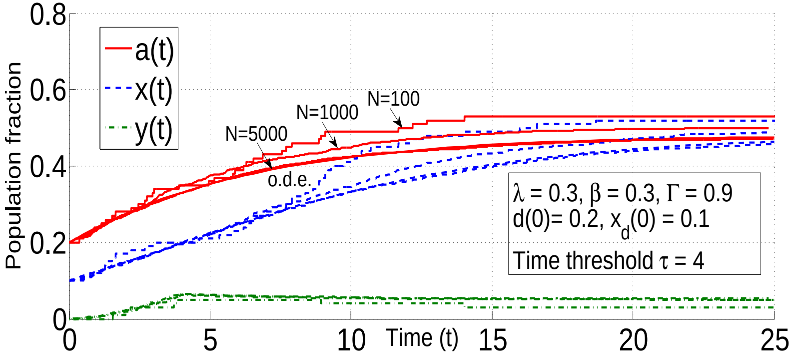

Optimal Control

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

- $a(\infty) = b(\infty) + d(\infty)$: terminal fraction of destinations
- $x_\sigma(t)$: fraction of destinations that have the coupon at time t

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

- $y_\sigma(T_\sigma)$: fraction of relays that have the coupon at time T_σ

Theorem

For the above o.d.e. system, for the cost function displayed earlier, there exists an optimal control of the form,

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

- Distributed implementation: Time stamping of the coupon

Optimality of a Time Threshold Control: Sketch of Proof

- Define Kamke dominance: extension of Kamke condition for o.d.e.s
- Here Kamke dominance holds, hence $\forall t$, $\sigma^{(1)}(t) \geq \sigma^{(2)}(t)$, and

$$(x_b^{(1)}(0), x_d^{(1)}(0), y^{(1)}(0)) \geq (x_b^{(2)}(0), x_d^{(2)}(0), y^{(2)}(0)),$$

  implies that, $\forall t$,

$$(x_b^{(1)}(t), x_d^{(1)}(t), y^{(1)}(t)) \geq (x_b^{(2)}(t), x_d^{(2)}(t), y^{(2)}(t))$$

- Consider $\sigma(t)$ (any action function) and $\sigma_\tau(t)$ (a time threshold action function) such that $y_\sigma(T_\sigma) = y_{\sigma_\tau}(T_{\sigma_\tau}) = \rho$
  - We can show that $T_{\sigma_\tau} \leq T_\sigma$ and hence the time threshold action function is optimal in $\{\sigma(\cdot) : y_\sigma(T_\sigma) = \rho\}$
- Define $\rho_{max} = \max\{y_{\sigma_\tau}(T_{\sigma_\tau}) : \tau \geq 0\}$
  - Consider an action function $\sigma(t)$ with $y_\sigma(T_\sigma) > \rho_{max}$
  - The threshold policy with $\tau := \sup\{t : \sigma(t) > 0\}$ has lower cost

Optimal Control for the Running Example

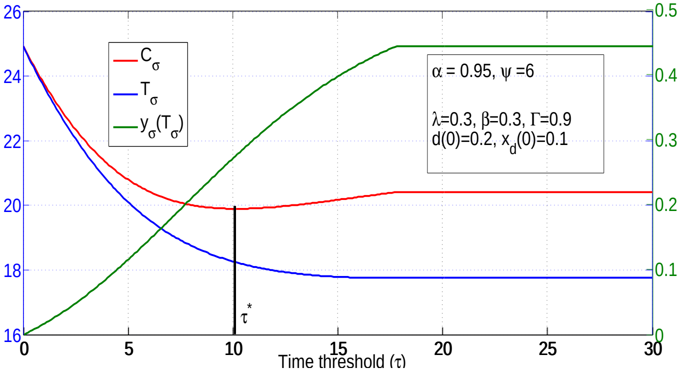

- $\alpha = 0.95$, $C_\sigma = 6y_\sigma(T_\sigma) + T_\sigma$
- The optimal control is to copy until $\tau = 10.1$ and then stop copying

Conclusion and Future Work

- A possible application framework for coupon delivery
- Delay-cost optimal forwarding
  - Joint evolution of content delivery and popularity
  - Modeled as CTMC and obtained the fluid limits
- Existence of Time-threshold control which is delay-cost optimal
- Performance of optimal fluid policy for the finite N case
- **Possible extensions:**
  - Multiple items of content; communities of interest
  - Large content: divided into several chunks
  - Service pricing, and incentive mechanisms for relays

References

- Chandramani Singh, Anurag Kumar, Rajesh Sundaresan and Eitan Altman, "Optimal Forwarding in Delay Tolerant Networks with Multiple Destinations", IEEE/ACM Transactions on Networking (TON) 2013.
- **SV** and Anurag Kumar, "Coevolution of Content Popularity and Delivery in Mobile P2P Networks", IEEE Infocom'12 (mini-conference)
- Shakkottai, S. and Johari, R., "Demand-aware content distribution on the internet", IEEE/ACM Transactions on Networking (TON) 2010
- Thomas G. Kurtz, "Solutions of Ordinary Differential Equations as Limits of Pure Jump Markov Processes", J. Appl. Prob. 7, 49-58 (1970)
- N. Gast, B. Gaujal, and J. Boudec, "Mean field for Markov decision processes: from discrete to continuous optimization," Arxiv preprint arXiv:1004.2342, 2010.
- H. Smith, "Monotone dynamical systems: An introduction to the theory of competitive and cooperative systems." American Mathematical Soc., 1995.