Delay-cost Optimal Coupon Delivery in Mobile Opportunistic Networks

Srinivasan Venkatraman [1]

Department of ECE, IISc.

19 December 2013

[1] Joint work with Prof. Anurag Kumar

Outline

Outline

Mobile Opportunistic Networks (MON)

- Proliferation of smart mobile devices
  - Smartphones, tablets, wearable computing devices, etc.

Mobile Opportunistic Networks (MON)

- Proliferation of smart mobile devices
  - Smartphones, tablets, wearable computing devices, etc.
- Mobile content delivery - a challenge
  - direct delivery from infrastructure is not scalable
  - must optimally combine infrastructure with peer-to-peer content delivery (via Bluetooth, WiFi-Direct, etc.)

Mobile Opportunistic Networks (MON)

- Proliferation of smart mobile devices
  - Smartphones, tablets, wearable computing devices, etc.
- Mobile content delivery - a challenge
  - direct delivery from infrastructure is not scalable
  - must optimally combine infrastructure with peer-to-peer content 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
  - e.g., a just released popular video song spreading in a campus setting

Mobile Opportunistic Networks (MON)

- Proliferation of smart mobile devices
  - Smartphones, tablets, wearable computing devices, etc.
- Mobile content delivery - a challenge
  - direct delivery from infrastructure is not scalable
  - must optimally combine infrastructure with peer-to-peer content 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
  - e.g., a just released popular video song spreading in a campus setting
- The item can be forwarded between devices when their p2p radio interfaces make contact
  - Akin to **epidemic** spread
  - A multi-hop opportunistic mobile network
- Can provide an approach for *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 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

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

Related Literature

- Combining infrastructure and p2p delivery [2]
- Optimal forwarding to a fixed set of destinations [3]
- Equivalence between influence spread models (viral marketing) and epidemic spread models [4]

[2] Shakkottai, S. and Johari, R., "Demand-aware content distribution on the internet", IEEE/ACM Transactions on Networking (TON) 2010

[3] 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.

[4] **SV** and Anurag Kumar, "Coevolution of Content Popularity and Delivery in Mobile P2P Networks", IEEE Infocom'12 (mini-conference).

Outline

Content Popularity and Dissemination 1/2

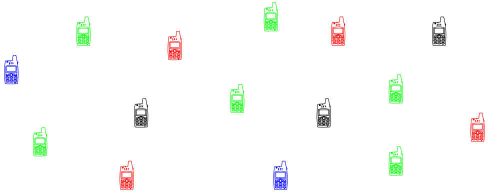

want:have nodes
do−not−want:have−not nodes
do−not−want:have nodes
want:have−not nodes

- A fixed population of nodes; size N
- A single item of content needs to reach certain *destination* nodes
  - Also called *want* nodes
  - Some want nodes are given the coupon initially

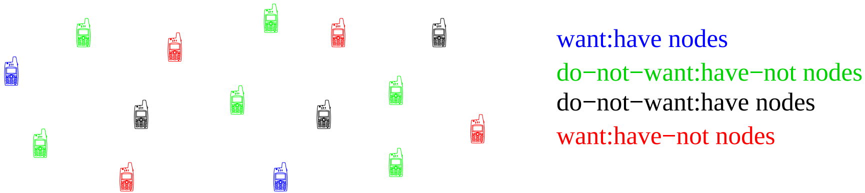

- A fixed population of nodes; size N
- A single item of content needs to reach certain *destination* nodes
  - Also called *want* nodes
  - Some want nodes are given the coupon initially
- The set of want nodes grows, as more nodes express their interest in the content
  - Infrastructure might broadcast the number of nodes who want the content
  - Meetings with want nodes can help to spread the popularity (in this talk)

Content Popularity and Dissemination 1/2

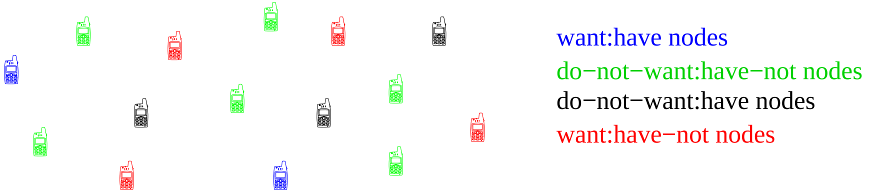

- A fixed population of nodes; size N
- A single item of content needs to reach certain *destination* nodes
  - Also called *want* nodes
  - Some want nodes are given the coupon initially
- The set of want nodes grows, as more nodes express their interest in the content
  - Infrastructure might broadcast the number of nodes who want the content
  - Meetings with want nodes can help to spread the popularity (in this talk)
- *Do-not-want* nodes could help in forwarding
  - Also called *relay* nodes

Content Popularity and Dissemination 2/2

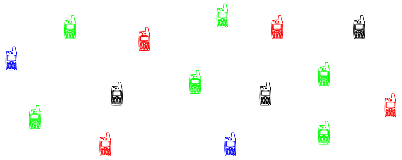

want:have nodes
do−not−want:have−not nodes
do−not−want:have nodes
want:have−not nodes

- Random meeting instants between pairs of nodes
  - Modeled as independent Poisson point processes
  - Instantaneous copying is assumed at a meeting instant

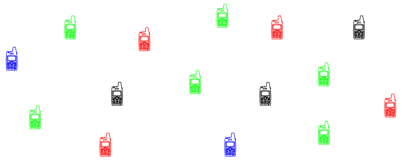

- Random meeting instants between pairs of nodes
  - Modeled as independent Poisson point processes
  - Instantaneous copying is assumed at a meeting instant
- At a meeting instant where a have node meets a have-not node
  - A "copy" or "do-not-copy" decision has to be made

Content Popularity and Dissemination 2/2

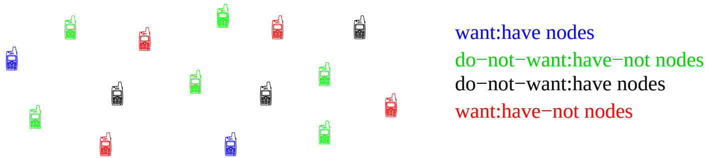

- Random meeting instants between pairs of nodes
  - Modeled as independent Poisson point processes
  - Instantaneous copying is assumed at a meeting instant
- At a meeting instant where a have node meets a have-not node
  - A "copy" or "do-not-copy" decision has to be made
- Delay-cost tradeoff
  - Copying to relays (do-not-wants) (i) promotes early delivery to destinations, and (ii) they could become destinations themselves
  - may need to pay an incentive to relays to help in content forwarding

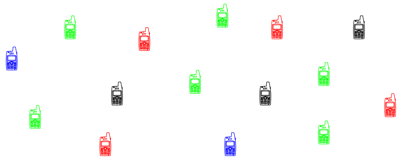

want:have nodes
do−not−want:have−not nodes
do−not−want:have nodes
want:have−not nodes

- Random meeting instants between pairs of nodes
  - Modeled as independent Poisson point processes
  - Instantaneous copying is assumed at a meeting instant
- At a meeting instant where a have node meets a have-not node
  - A "copy" or "do-not-copy" decision has to be made
- Delay-cost tradeoff
  - Copying to relays (do-not-wants) (i) promotes early delivery to destinations, and (ii) they could become destinations themselves
  - may need to pay an incentive to relays to help in content forwarding
- Objective: Quickly spread content to a large fraction of destinations nodes while minimizing the residual number of relays that have the content

A Possible Application Framework

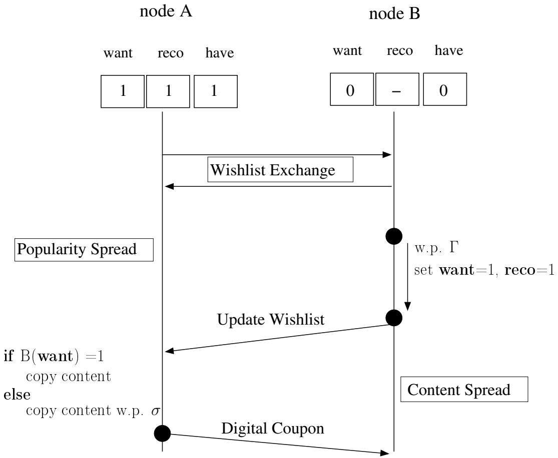

A Possible Application Framework

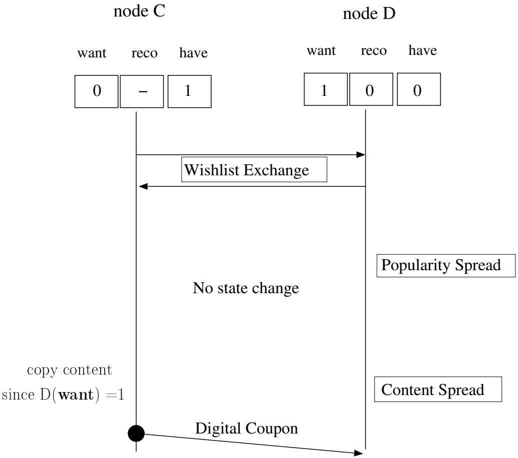

SIR-SI States and Evolution

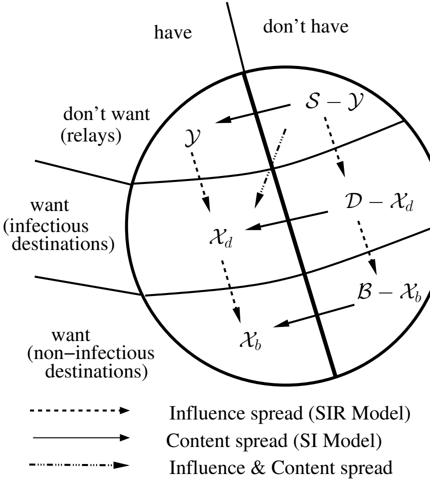

Snapshot of partition of the population over the states

- $\mathcal{X}(t)$: destinations that have the content
- $\mathcal{X}_d(t) = \mathcal{X}(t) \cap \mathcal{D}(t)$
  - destinations (**want** $=$**1**)
  - infectious (**reco** $=$**1**)
  - have the content (**have** $=$**1**)
- $\mathcal{X}_b(t) = \mathcal{X}(t) \cap \mathcal{B}(t)$
  - destinations (**want** $=$**1**)
  - non-infectious (**reco** $=$**0**)
  - have the content (**have** $=$**1**)

SIR-SI Model: Parameters for Population of Size N

- 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)
  - We first take σ to be fixed over time
  - A destination (want node) is always copied to

Outline

Fluid limits

- Kurtz[5] provided a technique to approximate pure jump Markov process by differential equations

[5] Thomas G. Kurtz, "Solutions of Ordinary Differential Equations as Limits of Pure Jump Markov Processes", J. Appl. Prob. 7, 49-58 (1970)

Fluid limits

- Kurtz[5] provided a technique to approximate pure jump Markov process by differential equations
- The process involves writing down the expected drift rates of the Markov process
- Under some regularity conditions and appropriate scaling (with system size N), one can show the convergence

[5] Thomas G. Kurtz, "Solutions of Ordinary Differential Equations as Limits of Pure Jump Markov Processes", J. Appl. Prob. 7, 49-58 (1970)

An Example: Discrete Time M/M/1 Queue

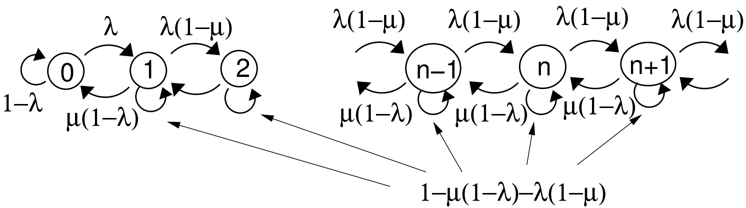

- Let $X(k), k \geq 0$, be the queue length process

An Example: Discrete Time M/M/1 Queue

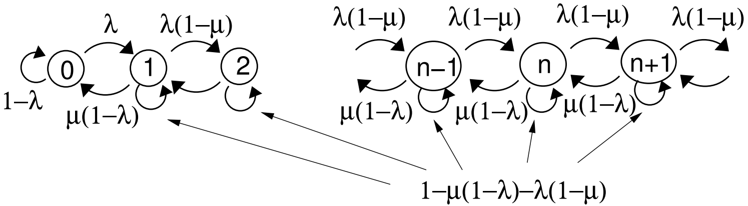

- Let $X(k), k \geq 0$, be the queue length process
- Define $\bar{X}^{(N)}(t) = \frac{1}{N} X(\lfloor Nt \rfloor)$

An Example: Discrete Time M/M/1 Queue

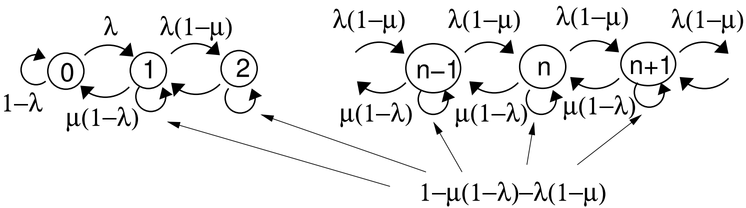

- Let $X(k), k \geq 0$, be the queue length process
- Define $\bar{X}^{(N)}(t) = \frac{1}{N}X(\lfloor Nt \rfloor)$
- Consider the ODE

$$\dot{x}(t) = (\lambda - \mu)I_{\{x(t)>0\}}$$

with initial condition such that, for each N, $x(0) = \bar{X}^{(N)}(0)$

How Does the ODE Approximate the M/M/1 Process?

- Plots of $\bar{X}^{(N)}(t)$ and $x(t)$ vs. t

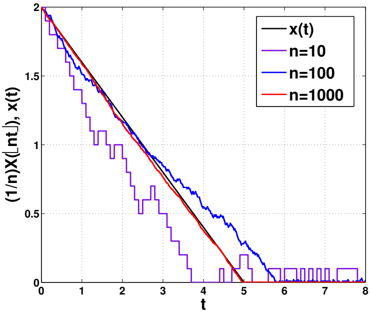

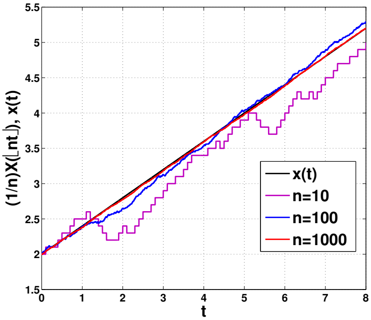

- Left panel: $\lambda = 0.3, \mu = 0.7$; positive recurrent DTMC
- Right panel: $\lambda = 0.7, \mu = 0.3$; transient DTMC

SIR-SI Markov Chain: Transitions at Meeting Epochs

- We model the evolution of SIR-SI as a continuous time Markov chain
- Let $k = 0, 1, 2, \cdots$ index the meeting/recovery epochs at times t_k
- System state: $Z(k) = (B(k), D(k), X_b(k), X_d(k), Y(k))$

- We model the evolution of SIR-SI as a continuous time Markov chain
- 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$(1 - \Gamma)\sigma$

SIR-SI Markov Chain: Expected Drift Rates

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

SIR-SI Markov Chain: O.D.E. Limit

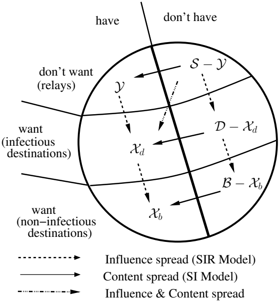

$$
\begin{aligned}
\dot{b} &= \beta d \\
\dot{d} &= -\beta d + \lambda \Gamma d s \\
\dot{x}_b &= \beta x_d + \lambda (b - x_b)(x + y) \\
\dot{x}_d &= \lambda (d - x_d)(x + y) + \lambda \Gamma d y + \Gamma \lambda x_d (s - y) - \beta x_d \\
\dot{y} &= -\Gamma \lambda d y + \lambda \sigma (s - y)(x_b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

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

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

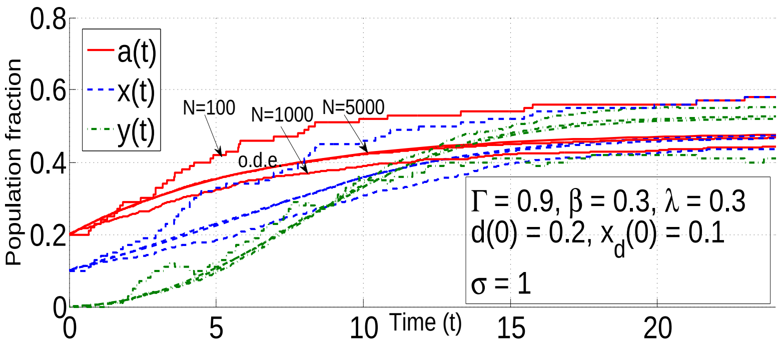

Simulation of SIR-SI CTMC for various N; shown with the fluid limit; $\sigma = 1$

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

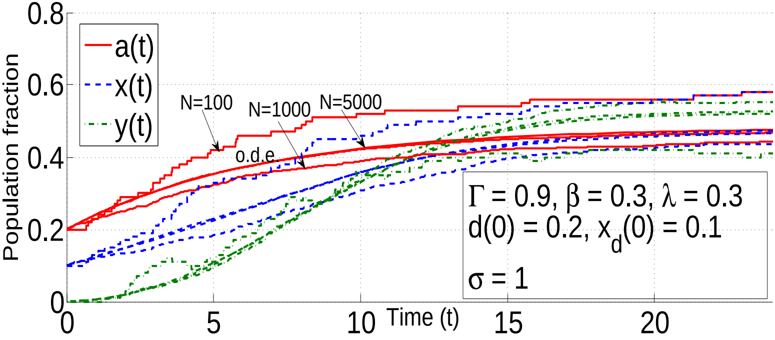

Simulation of SIR-SI CTMC for various N; shown with the fluid limit; $\sigma = 1$

- 95% of the "wants" have the content by $t = 15$
- But, by then, almost all the "do-not-wants" have the content too
  - This could be a waste, due to the incentive cost

Convergence of Scaled Markov Chain to O.D.E. Limit

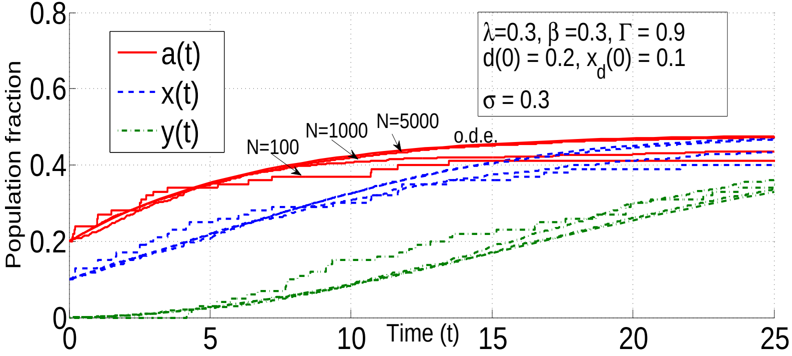

Simulation of SIR-SI CTMC for various N; shown with fluid limit; $\sigma = 0.3$

Convergence of Scaled Markov Chain to O.D.E. Limit

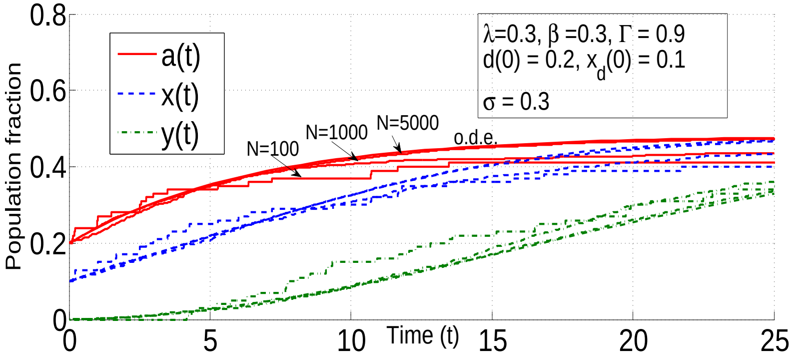

Simulation of SIR-SI CTMC for various N; shown with fluid limit; $\sigma = 0.3$

- 95% of the "wants" have the content by $t = 19$
- At this time, just 40% of the eventual "do-not-wants" have the content

- Let the control σ be a function of the current state $Z(k)$

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

SIR-SI Model: Dynamic Control of Copying

- Let the control σ be a function of the current state $Z(k)$
- For each N we have a continuous time Markov Decision Process
  - Optimal control appears difficult to obtain
- Replace the constant copying probability in the ODE limit by $\sigma(t)$
  - which yields a controlled ODE
- Obtain the optimal (deterministic) control for the controlled ODE

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

SIR-SI Model: Dynamic Control of Copying

- Let the control σ be a function of the current state $Z(k)$
- For each N we have a continuous time Markov Decision Process
  - Optimal control appears difficult to obtain
- Replace the constant copying probability in the ODE limit by $\sigma(t)$
  - which yields a controlled ODE
- Obtain the optimal (deterministic) control for the controlled ODE
- What value does this control have for the original controlled Markov chain?

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

SIR-SI Model: Dynamic Control of Copying

- Let the control σ be a function of the current state $Z(k)$
- For each N we have a continuous time Markov Decision Process
  - Optimal control appears difficult to obtain
- Replace the constant copying probability in the ODE limit by $\sigma(t)$
  - which yields a controlled ODE
- Obtain the optimal (deterministic) control for the controlled ODE
- What value does this control have for the original controlled Markov chain?
- In some situations it can be shown that this control is asymptotically optimal
  - as the size of the finite size problem increases
  - *without* needing to solve the finite size problem[6]

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

Dynamic Control: Comparison of SIR-SI and the O.D.E.

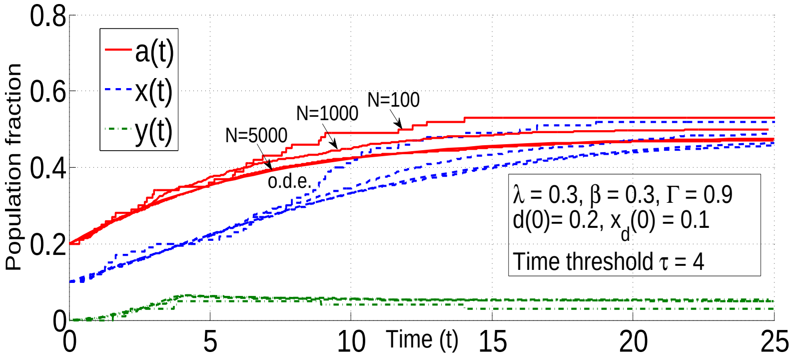

Simulation of SIR-SI CTMC for various N; shown with the fluid limit

Dynamic Control: Comparison of SIR-SI and the O.D.E.

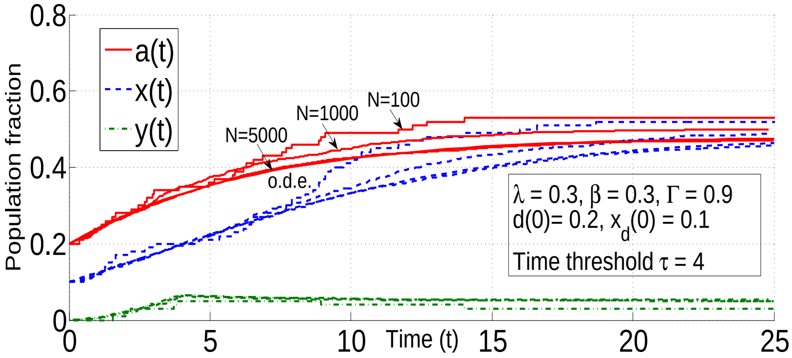

Simulation of SIR-SI CTMC for various N; shown with the fluid limit

- The control $\sigma(t) = 1, 0 \leq t \leq 4$, and $\sigma(t) = 0$ thereafter
- 95% of the "wants" have the content by $t = 22$
- Just 15% of the eventual "do-not-wants" have the content at this time
- Is there an *optimal control*?

Outline

Optimal Control of Co-evolution: An Objective

Fix α, and, for a control functon $\sigma(t)$, let

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

where

- $a(\infty) = b(\infty) + d(\infty)$ is the terminal fraction of destinations (SIR model)
- $x_\sigma(t)$ is the fraction of destinations that have the content at time t

Optimal Control of Co-evolution: An Objective

Fix α, and, for a control functon $\sigma(t)$, let

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

where

- $a(\infty) = b(\infty) + d(\infty)$ is the terminal fraction of destinations (SIR model)
- $x_\sigma(t)$ is the fraction of destinations that have the content at time t

The cost we are interested in optimizing is of the form

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

- $y_\sigma(T_\sigma)$ denotes the number of relays that have the content at time T_σ

Optimality of a Time Threshold Control

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

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

Theorem

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

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

Optimality of a Time Threshold Control

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

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

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 will require
  - Time stamping of the content when it is initially distributed
  - Zero skew between node clocks; or a distributed time-synch algorithm

Optimality of a Time Threshold Control: Sketch of Proof

- Define Kamke dominance: extension of Kamke condition for o.d.e.s[7]
- 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, $\forall t$,

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

[7] H. Smith, "Monotone dynamical systems: An introduction to the theory of competitive and cooperative systems." American Mathematical Soc., 1995.

Optimality of a Time Threshold Control: Sketch of Proof

- Define Kamke dominance: extension of Kamke condition for o.d.e.s[7]
- 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 $\{\sigma(\cdot) : y_\sigma(T_\sigma) = \rho\}$

[7] H. Smith, "Monotone dynamical systems: An introduction to the theory of competitive and cooperative systems." American Mathematical Soc., 1995.

Optimality of a Time Threshold Control: Sketch of Proof

- Define Kamke dominance: extension of Kamke condition for o.d.e.s[7]
- 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

[7] H. Smith, "Monotone dynamical systems: An introduction to the theory of competitive and cooperative systems." American Mathematical Soc., 1995.

Optimality of a Time Threshold Control: Sketch of Proof

- Define Kamke dominance: extension of Kamke condition for o.d.e.s[7]
- 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

[7] H. Smith, "Monotone dynamical systems: An introduction to the theory of competitive and cooperative systems." American Mathematical Soc., 1995.

Obtaining the Optimal Control for the Running Example

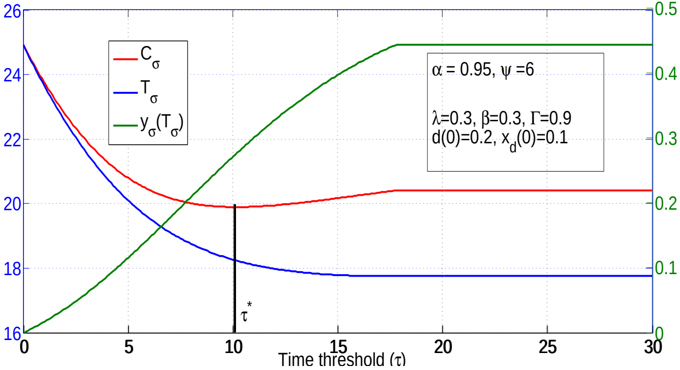

- $\alpha = 0.95$, $C_\sigma = 6y_\sigma(T_\sigma) + T_\sigma$
- The optimal control is to copy until $\tau = 10.1$ and then stop copying

Examples of the Optimal Control and Optimal Cost

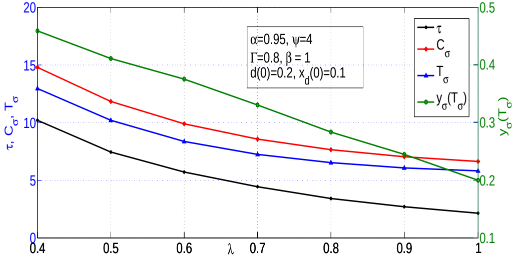

Optimal policy and costs vs. meeting rate parameter λ

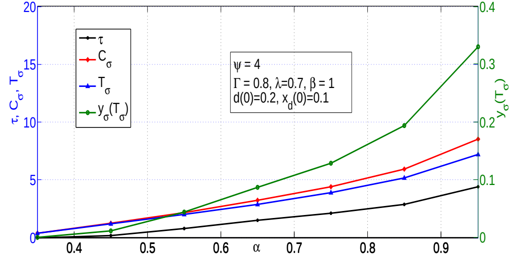

Optimal policy and costs vs. target spread fraction α

Conclusion

- A possible application framework for coupon delivery
- Delay-cost optimal forwarding
  - Joint evolution of content delivery and popularity
  - Modeled as CTMC and optimal policies were obtained in the fluid regime
- 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

Thank you