Towards Efficient Computation of the Age of Information in Situation Awareness Gossip Networks

Yoni Nazarathy,
The University of Queensland,

joint work with,

Jori Selen, Eindhoven University of Technology,
Lachlan Andrew, Swinburne University of Technology,
Hai Vu, Swinburne University of Technology.

Talk Outline

1. Scope
2. A general age of information model
3. The simplest non-trivial example
4. A stylized ring network
   - Approximations
   - Efficient numeric evaluation
5. Outlook

Scope

Situation Awareness Gossip Communication Networks

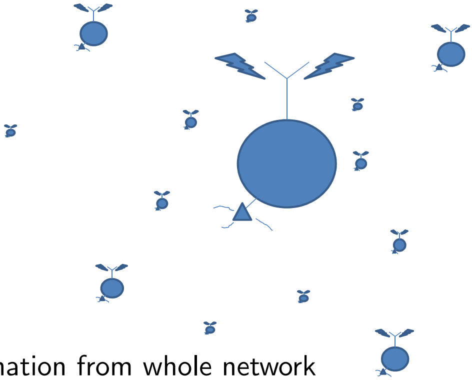

Nodes:

- Sense
- Transmit
- Receive
- Gossip
- Are users of information from whole network

Goal: Performance analysis and control with respect to

$$X(i, j) = \textbf{the age of } j\textbf{-info at node } i$$

for all node pairs i, j

A General Age of Information Model

A General Age of Information Model

$$X_{n+1}(i,j) = \begin{cases} \left(X_n(i,j) \wedge \bigwedge_{\{k:R_n(k,i)U_n(k,j)=1\}} X_n(k,j)\right) + 1 & i \neq j, \\ 0 & i = j. \end{cases}$$

- $i, j \in \{1, \ldots, M\}$
- Discrete time, $n = 0, 1, 2, \ldots$
- **State** $X_n(i,j)$ is age of the the j-info at node i
- **Control** $U_n(k,j)$ is indicator of j-info in message from k
- $T_n(k) = \mathbf{1}_{\{\sum_{j=1}^{M} U_n(k,j) \geq 1\}}$ is indicator of message transmit at k
- **Environment** $R_n(k,i)$ is indicator of reception at i of message from k. Can only equal 1 if $T_n(k) = 1$

Bernoulli Channels and Policies (Controls)

$$X_{n+1}(i,j) = \begin{cases} \Big(X_n(i,j) \wedge \bigwedge_{\{k: R_n(k,i) U_n(k,j)=1\}} X_n(k,j)\Big) + 1 & i \neq j, \\ 0 & i = j. \end{cases}$$

Bernoulli channels: $R_n(\cdot,\cdot)$ are "i.i.d" depending only on transmissions in current time: $\Big(T_n(1), \dots, T_n(M)\Big)$
Local Bernoulli policies: $U_n(\cdot,\cdot)$ are i.i.d not depending on anything yet with possible constraints such as $\sum_{j=1}^{M} U_n(i,j) \leq K$

This makes $\{X_n\}_{n=0}^{\infty}$ a Markov chain on state space $\mathbb{Z}_+^{M\times M} \setminus M$

The following performance measures are of interest:

$$\pi_\ell(i,j) = \lim_{n\to\infty} P\Big(X_n(i,j) = \ell\Big), \quad m(i,j) = \sum_{\ell=0}^{\infty} \ell \, \pi_\ell(i,j).$$

as well as joint distributions

The Simplest Non-trivial Example

The Simplest Nontrivial Example: $M = 2$

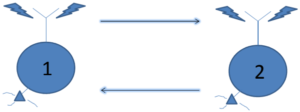

$q_1, q_2 \in (0,1)$: probs' of Tx.
$p_1, p_2 \in (0,1)$: probs' of Rx without interference
$p_{1*}, p_{2*} \in (0,1)$: probs' of Rx with interference (during Tx)

$\{X_n\}_{n=0}^{\infty}$ is a Markov chain on $\mathbb{Z}_+^2$ with four types of transitions:

Event	State Change	Probability
no Rx	$(++1, \quad ++1)$	$\lambda(\emptyset) = \bar{q}_1\bar{q}_2 + q_1\bar{q}_2\bar{p}_2 + \bar{q}_1 q_2 \bar{p}_1 + q_1 q_2 \bar{p}_{2*}\bar{p}_{1*}$
Rx 1 only	$(=1, \quad ++1)$	$\lambda(1) = \bar{q}_1 q_2 p_1 + q_1 q_2 p_{1*}$
Rx 2 only	$(++1, \quad =1)$	$\lambda(2) = q_1 \bar{q}_2 p_2 + q_1 q_2 p_{2*}$
Rx both	$(=1, \quad =1)$	$\lambda(1,2) = q_1 q_2 p_{2*} p_{1*}$

Want: $\tilde{\pi}_{\ell_1,\ell_2} = \lim_{n\to\infty} P\Big(X_n(1,2) - 1 = \ell_1,\ X_n(2,1) - 1 = \ell_2\Big)$

The Simplest Nontrivial Example: $M = 2$ cont.

$$\lambda(\emptyset) = \bar{q}_1\bar{q}_2 + q_1\bar{q}_2\bar{p}_2 + \bar{q}_1 q_2 \bar{p}_1 + q_1 q_2 \bar{p}_{2*}\bar{p}_{1*}, \; \lambda(1,2) = q_1 q_2 p_{2*} p_{1*}$$
$$\lambda(1) = \bar{q}_1 q_2 p_1 + q_1 q_2 p_{1*}, \quad \lambda(2) = q_1 \bar{q}_2 p_2 + q_1 q_2 p_{2*}$$

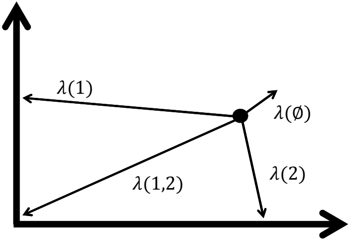

Proposition:

$$\tilde{\pi}_{\ell_1,\ell_2} = \begin{cases} \lambda(\emptyset)^{\ell_1}\lambda(1,2) & \text{for } \ell_1 = \ell_2, \\ \lambda(\emptyset)^{\ell_2}\lambda(2) c_2^{\ell_1-\ell_2-1}(1-c_2) & \text{for } \ell_1 > \ell_2, \\ \lambda(\emptyset)^{\ell_1}\lambda(1) c_1^{\ell_2-\ell_1-1}(1-c_1) & \text{for } \ell_1 < \ell_2, \end{cases}$$

with,
$$c_1 = \lambda(\emptyset) + \lambda(1), \quad c_2 = \lambda(\emptyset) + \lambda(2).$$

Derivation of Joint Distribution

Look first at marginal distributions,

$$\tilde{\pi}_{\ell_1,\cdot} = \sum_{\ell_2=0}^{\infty} \tilde{\pi}_{\ell_1,\ell_2}, \quad \tilde{\pi}_{\cdot,\ell_2} = \sum_{\ell_1=0}^{\infty} \tilde{\pi}_{\ell_1,\ell_2},$$

The associated Markov chains are well known:

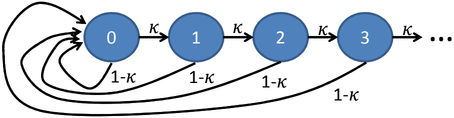

For the i'th marginal $(i = 1, 2)$, take $\kappa = c_i = \lambda(\emptyset) + \lambda(i)$

Balance equations are:

$$\hat{\pi}_\ell = \kappa \hat{\pi}_{\ell-1}, \quad \ell = 1, 2, \ldots .$$

as well as $\sum_{\ell=0}^{\infty} \hat{\pi}_\ell = 1$.

The solution is: $\hat{\pi}_\ell = \kappa^\ell (1 - \kappa), \quad \ell = 0, 1, 2, \ldots,$

Derivation of Joint Distribution cont.

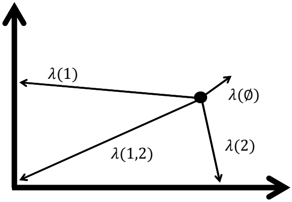

Look now at the balance equations:

$$\pi_{0,0} = \lambda(1,2) \sum_{\ell_1,\ell_2} \pi_{\ell_1,\ell_2},$$

$$\pi_{\ell_1,0} = \lambda(2) \sum_{\ell_2=0}^{\infty} \pi_{(\ell_1-1),\ell_2} \quad , \ell_1 \geq 1,$$

$$\pi_{0,\ell_2} = \lambda(1) \sum_{\ell_1=0}^{\infty} \pi_{\ell_1,(\ell_2-1)} \quad , \ell_2 \geq 1,$$

$$\pi_{\ell_1,\ell_2} = \lambda(\emptyset)\pi_{(\ell_1-1),(\ell_2-1)} \quad , \ell_1, \ell_2 \geq 1.$$

Knowledge of the marginals gives us the solution

The General Idea for “Efficient Computation”

In the previous example we saw that the stationary distribution can be expressed in terms of the marginal distributions

The General Idea for "Efficient Computation"

In the previous example we saw that the stationary distribution can be expressed in terms of the marginal distributions

It turns out that this is also the case for more complicated examples

The General Idea for "Efficient Computation"

In the previous example we saw that the stationary distribution can be expressed in terms of the marginal distributions

It turns out that this is also the case for more complicated examples

We can exploit this relationship for efficient numeric computation of more complicated examples

A Ring Network

A Ring Network

Consider networks of M nodes organized in a ring:

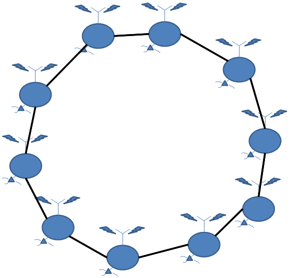

Assume:

- $R_n(i,j) = \mathbf{1}_{\{\text{i and j neighbors}\}}$
- For each i, $U_n(i,j)$ is uniformly distributed on $j = 1, \ldots, M$

This implies that the law of $\{X_n(1,i), X_n(2,i), \ldots, X_n(M,i)\}_{n=0}^{\infty}$ is the same for all i. What do sample paths "look like"?

A Ring Network

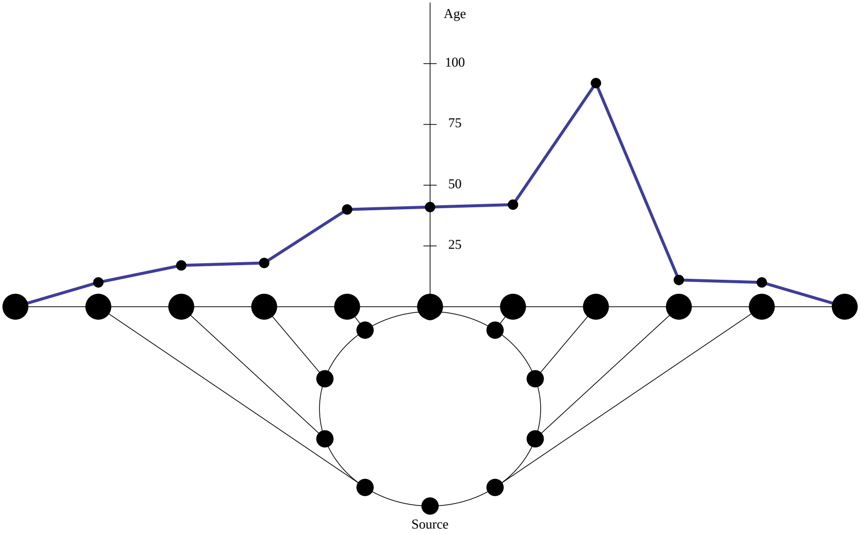

Play Movie

An Iteration Scheme for Solving the Ring

Initialization:

"Exact" 4-d distribution up to bounding box of face-size N, $O(N^6)$ in time $O(N^4)$ in memory

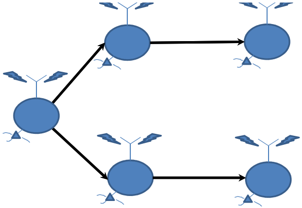

Iteration:

Move from distribution of red box to that of green box with truncation, $O(N^6)$ in time, $O(N^4)$ in memory

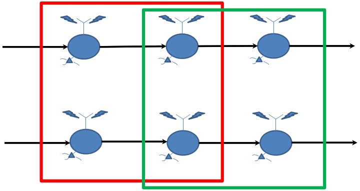

An Iteration Scheme cont.

- Iterate **clockwise** and **counter-clockwise** in parallel
- This yields for each node the joint distribution of the "clockwise age" and "counter-clockwise age"
- Use the joint distribution to find the **minimum**

For ring of length M, this yields the $M-1$ joint-clockwise-counter-clockwise distributions in $O(MN^6)$ time

An Iteration Scheme cont.

- Iterate **clockwise** and **counter-clockwise** in parallel
- This yields for each node the joint distribution of the "clockwise age" and "counter-clockwise age"
- Use the joint distribution to find the **minimum**

For ring of length M, this yields the $M-1$ joint-clockwise-counter-clockwise distributions in $O(MN^6)$ time

Honesty Box

For small to moderate M, simple Monte-Carlo simulation "out performs" our current iteration scheme :-(

For large M we have a simple asymptotic approximation.... :-)

Asymptotic Approximation

- The **clockwise** and **counter-clockwise** ages are distributed as negative binomial random variables (asymptotically Gaussian)
- If we assume that the **clockwise age** and **counter-clockwise age** are independent, then we can find the distribution (and mean) of the minimum easily

Asymptotic Approximation

- The **clockwise** and **counter-clockwise** ages are distributed as negative binomial random variables (asymptotically Gaussian)
- If we assume that the **clockwise age** and **counter-clockwise age** are independent, then we can find the distribution (and mean) of the minimum easily

Approximation based on independence assumption:
Let $F_k^{(M)}(\ell)$ be the age distribution at distance k from the source in a ring of size M. Assume $k \leq M/2$.

Then,

$$F_k^{(M)}(\ell) \approx P\Big(N_1 \wedge N_2 \leq \ell\Big),$$

with independent,

$$N_1 \sim \text{Gaussian}\Big(kM,\ k\frac{M^3}{2}\Big), \quad N_2 \sim \text{Gaussian}\Big((M-k)M,\ (M-k)\frac{M^3}{2}\Big).$$

In progress: Proof of the corresponding weak convergence theorem together with uniform integrality conditions.

Means of the Ring Network, $M = 10$

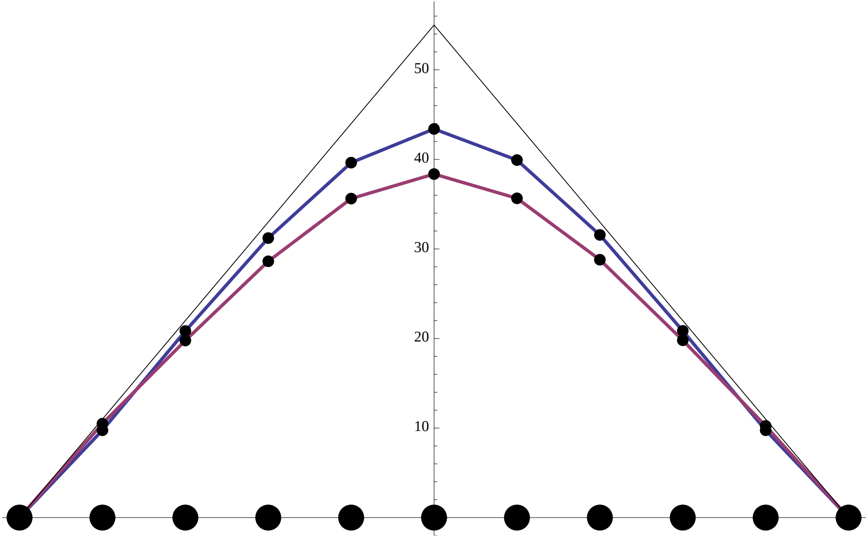

Means of the Ring Network, $M = 20$

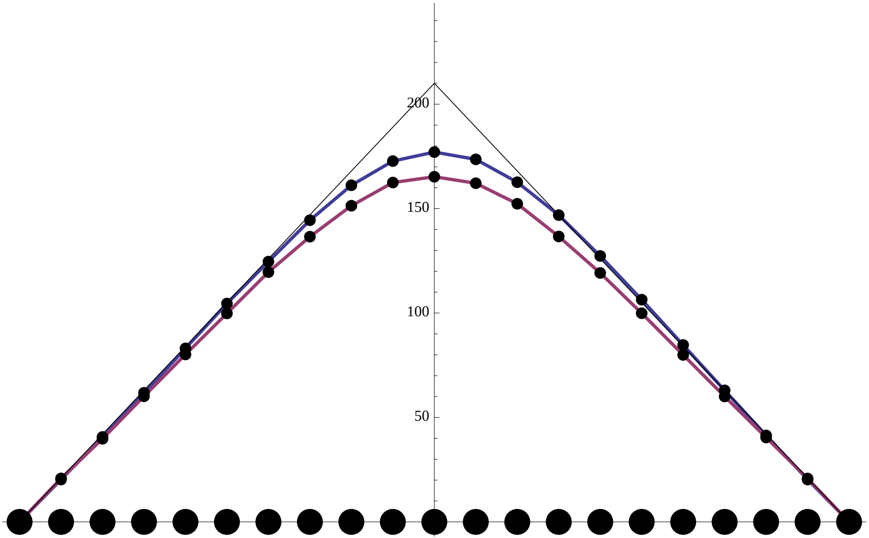

Means of the Ring Network, $M = 30$

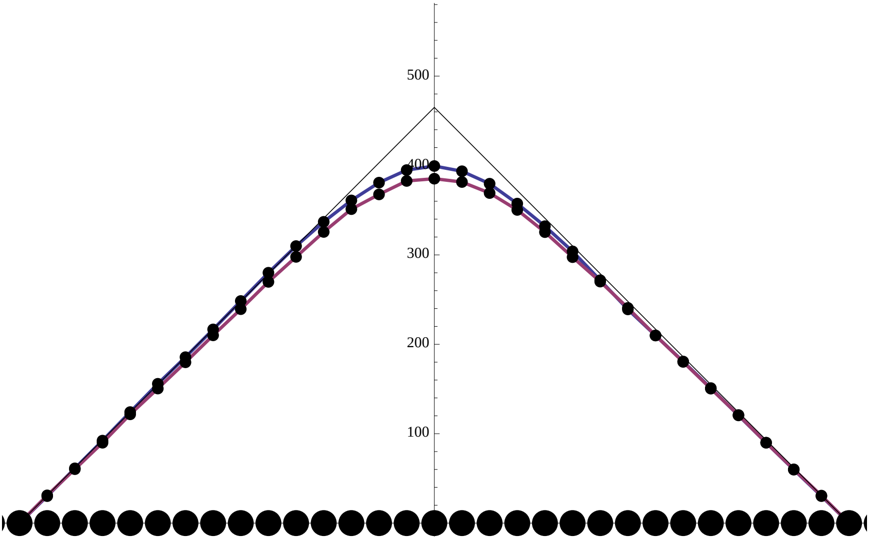

Means of the Ring Network, $M = 100$

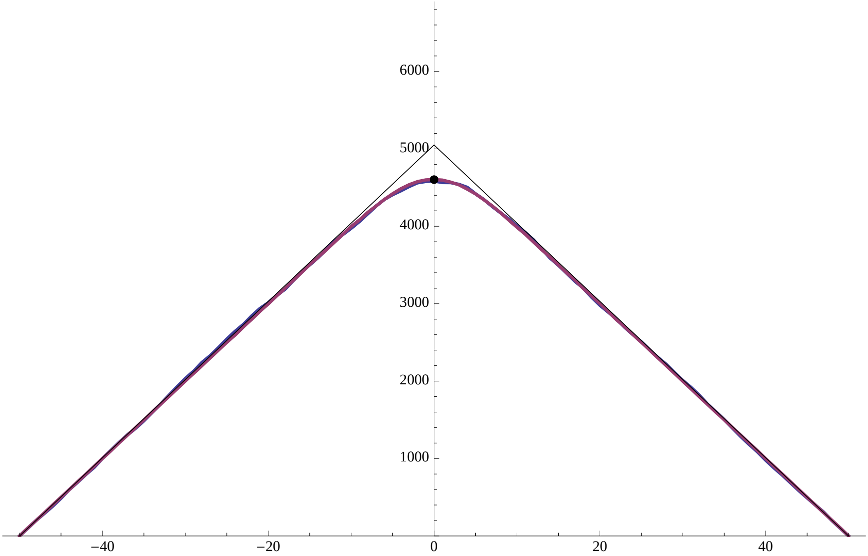

Outlook

Summary

- We have introduced a model for age of information in gossip networks
- In general performance analysis of this model yields intractable Markov chains
- To overcome this we are pursuing two paths:
  - Asymptotic approximations
  - Efficient approximation (or exact) algorithms
- A general goal is to be able to compute the age of information over regular grid networks

Extras

Tree Networks

Tree Networks