Information Freshness in Dynamic Gossip Networks

Arunabh Srivastava, **T. J. M.** Sennur Ulukus

Department of Electrical and Computer Engineering
University of Maryland

November 19, 2025

Overview

Motivation

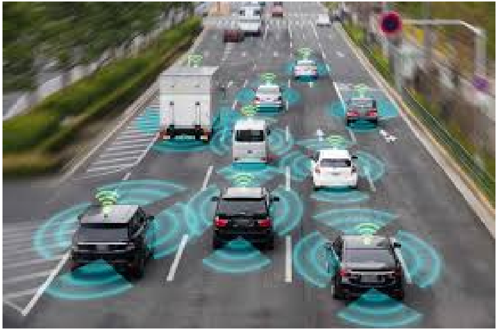

System Model

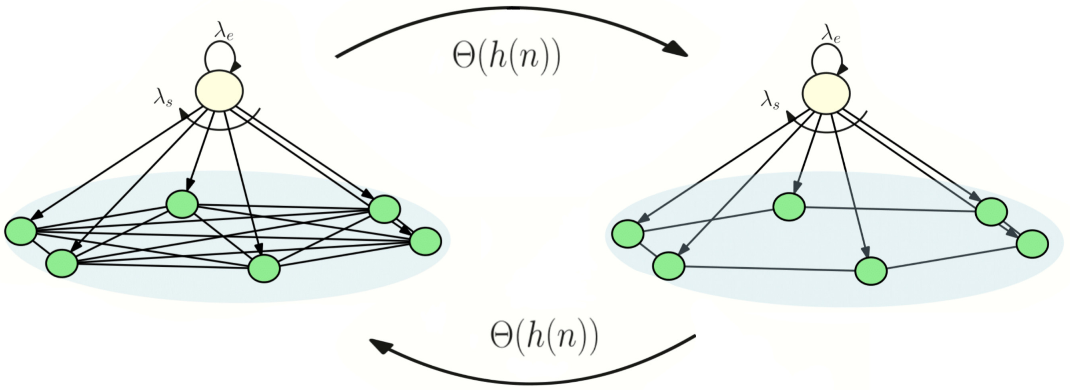

Model and Age of Information

- Packets are sent in the network and hold versions of some data.

Model and Age of Information

- Packets are sent in the network and hold versions of some data.
  - Think of software versions distributed to users over time

Model and Age of Information

- Packets are sent in the network and hold versions of some data.
  - Think of software versions distributed to users over time
- Nodes randomly send updates to neighbors via independent Poisson Processes

Model and Age of Information

- Packets are sent in the network and hold versions of some data.
  - Think of software versions distributed to users over time
- Nodes randomly send updates to neighbors via independent Poisson Processes
- Nodes uniformly communicate to neighbors - i.e. no preferred communication direction

Model and Age of Information

- Packets are sent in the network and hold versions of some data.
    - Think of software versions distributed to users over time
- Nodes randomly send updates to neighbors via independent Poisson Processes
- Nodes uniformly communicate to neighbors - i.e. no preferred communication direction
- Nodes always accept the newest version of the data.

Model and Age of Information

- Packets are sent in the network and hold versions of some data.
  - Think of software versions distributed to users over time
- Nodes randomly send updates to neighbors via independent Poisson Processes
- Nodes uniformly communicate to neighbors - i.e. no preferred communication direction
- Nodes always accept the newest version of the data.
- $X_u(t)$ is the number of versions node u is behind the source n_0 at time t.

Model and Age of Information

- Packets are sent in the network and hold versions of some data.
    - Think of software versions distributed to users over time
- Nodes randomly send updates to neighbors via independent Poisson Processes
- Nodes uniformly communicate to neighbors - i.e. no preferred communication direction
- Nodes always accept the newest version of the data.
- $X_u(t)$ is the number of versions node u is behind the source n_0 at time t.
- $X_S(t)$ is the newest version available in the subset of nodes S.

AoI Example

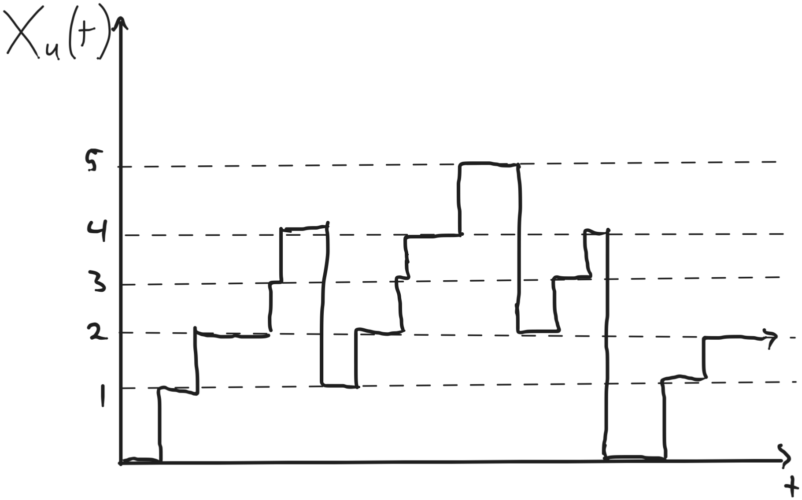

Network Dynamics

- The network randomly switches between two topologies G_1, G_2

Network Dynamics

- The network randomly switches between two topologies G_1, G_2
- Switching is mediated by a continuous time Markov chain (CTMC)

Network Dynamics

- The network randomly switches between two topologies G_1, G_2
- Switching is mediated by a continuous time Markov chain (CTMC)
- Time spent in each topology is $\Theta(h(n))$, and we will consider different regimes of $h(n)$

Network Dynamics

- The network randomly switches between two topologies G_1, G_2
- Switching is mediated by a continuous time Markov chain (CTMC)
- Time spent in each topology is $\Theta(h(n))$, and we will consider different regimes of $h(n)$
- Each vertex gossips uniformly to its neighbors in each of G_1, G_2

Main Question

Given two networks G_1 and G_2, how does the average version age $\frac{1}{n}\sum_u \mathbb{E}X_u(t)$ vary with the holding time $h(n)$?

Previous Combinatorial Work

- Yates (2021) showed the version age $v_G(S) = \lim_{t\to\infty} \mathbb{E}X_S(t)$ exists and can be computed combinatorially:

$$v_G(S) = \frac{\lambda_e + \sum_{i \notin S} \lambda_i(S) v_G(S \cup \{i\})}{\lambda_0(S) + \sum_{i \notin S} \lambda_i(S)}$$

Previous Combinatorial Work

- Yates (2021) showed the version age $v_G(S) = \lim_{t\to\infty} \mathbb{E}X_S(t)$ exists and can be computed combinatorially:

$$v_G(S) = \frac{\lambda_e + \sum_{i\notin S} \lambda_i(S) v_G(S \cup \{i\})}{\lambda_0(S) + \sum_{i\notin S} \lambda_i(S)}$$

- Our group has studied other models of AoI in dynamic gossip networks which are a subset of this work

Previous Combinatorial Work

- Yates (2021) showed the version age $v_G(S) = \lim_{t\to\infty} \mathbb{E} X_S(t)$ exists and can be computed combinatorially:

$$v_G(S) = \frac{\lambda_e + \sum_{i \notin S} \lambda_i(S) v_G(S \cup \{i\})}{\lambda_0(S) + \sum_{i \notin S} \lambda_i(S)}$$

- Our group has studied other models of AoI in dynamic gossip networks which are a subset of this work
- **Issue:** When the network can vary between topologies G_1 and G_2, the SHS approach no longer works due to no large time-limit convergence/ stability.

First-Passage Percolation and Hopcount

- Let $G = (V, E, w)$ be a network with random edge weights

First-Passage Percolation and Hopcount

- Let $G = (V, E, w)$ be a network with random edge weights
- For any two vertices $u, v \in V$, define the *first passage time* $T_G(u, v)$ as the weight of the shortest path between u and v (found, for instance, using Dijkstra's algorithm)

A General Distribution Characterization for Static Networks

- The AoI of u is essentially the time it takes for a packet to travel from the source to u

A General Distribution Characterization for Static Networks

- The AoI of u is essentially the time it takes for a packet to travel from the source to u
- This travel time is captured by the language of first-passage percolation, which is the (random) distance between two points in a graph with random edge weights.

A General Distribution Characterization for Static Networks

- The AoI of u is essentially the time it takes for a packet to travel from the source to u
- This travel time is captured by the language of first-passage percolation, which is the (random) distance between two points in a graph with random edge weights.

Theorem (M. and Michelen 2024)

There exists a static 'dual graph' $\bar{G}$ with random edge weights such that for any node u, for large enough t the version age of u at time t has distribution given by

$$X_u(t) = \mathrm{Pois}(\lambda_e T(\bar{u}, \bar{n}_0))$$

Where $T(\bar{a}, \bar{b})$ is the minimum random distance between $\bar{a}$ and $\bar{b}$ in $\bar{G}$

A General Distribution Characterization for Static Networks

- The AoI of u is essentially the time it takes for a packet to travel from the source to u
- This travel time is captured by the language of first-passage percolation, which is the (random) distance between two points in a graph with random edge weights.

Theorem (M. and Michelen 2024)

There exists a static 'dual graph' $\bar{G}$ with random edge weights such that for any node u, for large enough t the version age of u at time t has distribution given by

$$X_u(t) = \mathrm{Pois}(\lambda_e T(\bar{u}, \bar{n}_0))$$

Where $T(\bar{a}, \bar{b})$ is the minimum random distance between $\bar{a}$ and $\bar{b}$ in $\bar{G}$

- We will use this characterization of AoI in place of the SHS approach

Fast Switching

Theorem

Let G_1 and G_2 be two networks on n vertices. Suppose G_i has long term average version age $\lim_{t\to\infty}\sum_u \mathbb{E}X_u^{G_i}(t) = \Theta(f_i(n))$, $i \in \{1, 2\}$, with $f_1(n) = o(f_2(n))$. If $h(n) = O(f_1(n))$, then the long term average version age for the time varying system is $\Theta(f_1(n))$.

Fast Switching

Theorem

Let G_1 and G_2 be two networks on n vertices. Suppose G_i has long term average version age $\lim_{t\to\infty}\sum_u \mathbb{E}X_u^{G_i}(t) = \Theta(f_i(n))$, $i \in \{1, 2\}$, with $f_1(n) = o(f_2(n))$. If $h(n) = O(f_1(n))$, then the long term average version age for the time varying system is $\Theta(f_1(n))$.

- In words, if the holding time is at most the time it takes for a source generated packet to reach u in G_1, then the overall version age is controlled by the age in G_1, which we denote by f_1.

Proof Idea

1. Lower bound follows by setting $G_2 = G_1$, and noting this is not time-varying and has age f_1.

Proof Idea

1. Lower bound follows by setting $G_2 = G_1$, and noting this is not time-varying and has age f_1.
2. Upper bound follows by setting G_2 to the empty graph.

Proof Idea

1. Lower bound follows by setting $G_2 = G_1$, and noting this is not time-varying and has age f_1.
2. Upper bound follows by setting G_2 to the empty graph.
   - It is intuitively clear that the version age scales as f_1 (age increases linearly for $O(f(1))$ time)

Proof Idea

1. Lower bound follows by setting $G_2 = G_1$, and noting this is not time-varying and has age f_1.
2. Upper bound follows by setting G_2 to the empty graph.
   - It is intuitively clear that the version age scales as f_1 (age increases linearly for $O(f(1))$ time)
   - Proof is a bit technical, since we need to bound the bad events that the age is too large or too small for too much time

Proof Idea

1. Lower bound follows by setting $G_2 = G_1$, and noting this is not time-varying and has age f_1.
2. Upper bound follows by setting G_2 to the empty graph.
   - It is intuitively clear that the version age scales as f_1 (age increases linearly for $O(f(1))$ time)
   - Proof is a bit technical, since we need to bound the bad events that the age is too large or too small for too much time
   - Uses some basic facts about CTMCs and the connection between AoI and percolation

Issue with generalizing to arbitrary h and the 'typical set' workaround

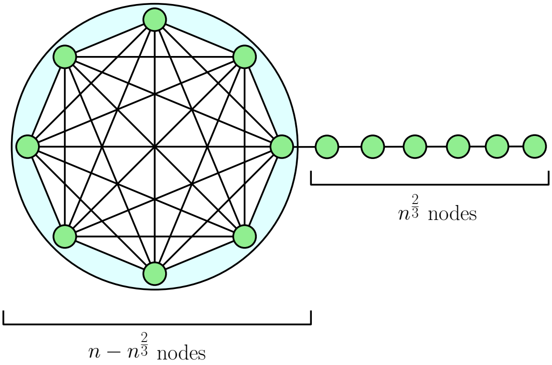

Figure: The average version age of this network is $\Theta(n^{1/3})$. Nodes in the complete graph have individual age $\Theta(\log n)$ age, while those in the path have $\Theta(n^{2/3})$ age.

- Given the issue above, define the 'Typical Set' of a network as those vertices with version age asymptotically less than the average.

Slow Switching

Theorem

Let G_1, G_2 be two networks on n vertices and f_1, f_2 be the long term average version age of the networks defined before. If $h(n) = \Omega(n \log n)$, then the long term average version age of the typical set scales as $\Theta(f_2(n))$.

Slow Switching

Theorem

Let G_1, G_2 be two networks on n vertices and f_1, f_2 be the long term average version age of the networks defined before. If $h(n) = \Omega(n \log n)$, then the long term average version age of the typical set scales as $\Theta(f_2(n))$.

- In words, if the holding time is at least $n \log n$, then the overall version age in the typical set is controlled by the age in G_2, which we denote by f_2.

Slow Switching

Theorem

Let G_1, G_2 be two networks on n vertices and f_1, f_2 be the long term average version age of the networks defined before. If $h(n) = \Omega(n \log n)$, then the long term average version age of the typical set scales as $\Theta(f_2(n))$.

- In words, if the holding time is at least $n \log n$, then the overall version age in the typical set is controlled by the age in G_2, which we denote by f_2.
- The previous example shows we can't hope to bound the average age for the entire network, as a small subset of vertices can disproportionately impact the overall age.

Proof Idea

1. Upper bound by replacing G_1 with G_2, and noting this network has age f_2.

Proof Idea

1. Upper bound by replacing G_1 with G_2, and noting this network has age f_2.
2. Lower bound by assuming every node has 0 age when in network G_1.

Proof Idea

1. Upper bound by replacing G_1 with G_2, and noting this network has age f_2.
2. Lower bound by assuming every node has 0 age when in network G_1.
   - The age process stabilizes in the typical set once every node has received an update from the source

Proof Idea

1. Upper bound by replacing G_1 with G_2, and noting this network has age f_2.
2. Lower bound by assuming every node has 0 age when in network G_1.
   - The age process stabilizes in the typical set once every node has received an update from the source
   - By a coupon collector type argument this takes at most $O(n \log n)$ time.

Proof Idea

1. Upper bound by replacing G_1 with G_2, and noting this network has age f_2.
2. Lower bound by assuming every node has 0 age when in network G_1.
    - The age process stabilizes in the typical set once every node has received an update from the source
    - By a coupon collector type argument this takes at most $O(n \log n)$ time.
    - Apply some technical lemmas to conclude

Proof Idea

1. Upper bound by replacing G_1 with G_2, and noting this network has age f_2.
2. Lower bound by assuming every node has 0 age when in network G_1.
   - The age process stabilizes in the typical set once every node has received an update from the source
   - By a coupon collector type argument this takes at most $O(n \log n)$ time.
   - Apply some technical lemmas to conclude
3. Note: This argument doesn't work if we replace the 'typical set' assumption with the whole network: a small subset of nodes can disproportionately increase the average age of the network.

Comments and Open Questions

- Our theorems can generalize to k networks. Is there a natural characterization in this case?

Comments and Open Questions

- Our theorems can generalize to k networks. Is there a natural characterization in this case?
- Can we obtain AoI characterization for other, more realistic dynamic network models?

Comments and Open Questions

- Our theorems can generalize to k networks. Is there a natural characterization in this case?
- Can we obtain AoI characterization for other, more realistic dynamic network models?
- Can we generalize the FPP framework to dynamic networks? This seems nontrivial.

Conclusion

Thanks!