Mean-Field Analysis for the Evaluation of Gossip Protocols

Rena Bakhshi[1], Lucia Cloth[2],
Wan Fokkink[1], Boudewijn Haverkort[2,3]

[1]Department of Computer Science, Vrije Universiteit Amsterdam, The Netherlands
[2]Centre for Telematics & Information Technology, University of Twente, Enschede, Netherlands
[3]Embedded Systems Institute, Eindhoven, The Netherlands

TCS seminar

October 9, 2009

Why mean-field approximation?

Observations

- Very costly simulations/analysis of large-scale prob. systems
- Interactions between nodes
- Modelling of local behaviour of node
- Union of them results in a large (stochastic) model

Assuming "$N \to \infty$" instead of "large N"

- Small deterministic process [Le Boudec et al. 2007]
- Limit behaviour of complete system
- Simple matrix-vector multiplications
- Limit of the measure of interest

Gossip protocols naturally fit for mean-field analysis

- Operate in large-scale, decentralized network
- Symmetrical behaviour of nodes
- Data exchange in a random fashion

Why mean-field approximation?

Observations

- Very costly simulations/analysis of large-scale prob. systems
- Interactions between nodes
- Modelling of local behaviour of node
- Union of them results in a large (stochastic) model

Assuming "$N \to \infty$" instead of "large N"

- Small deterministic process [Le Boudec et al. 2007]
- Limit behaviour of complete system
- Simple matrix-vector multiplications
- Limit of the measure of interest

Gossip protocols naturally fit for mean-field analysis

- Operate in large-scale, decentralized network
- Symmetrical behaviour of nodes
- Data exchange in a random fashion

Why mean-field approximation?

Observations

- Very costly simulations/analysis of large-scale prob. systems
- Interactions between nodes
- Modelling of local behaviour of node
- Union of them results in a large (stochastic) model

Assuming "$N \to \infty$" instead of "large N"

- Small deterministic process [Le Boudec et al. 2007]
- Limit behaviour of complete system
- Simple matrix-vector multiplications
- Limit of the measure of interest

Gossip protocols naturally fit for mean-field analysis

- Operate in large-scale, decentralized network
- Symmetrical behaviour of nodes
- Data exchange in a random fashion

Mean-field approximation

N interacting objects, identically defined

Example

Gossiping nodes, each with state space $S = \{0, 1, \ldots K - 1\}$

Time is discrete $t \in \mathbb{N}$

Occupancy measure $M_i^N(t)$

- fraction of N objects in state i at time t

Example

30% of nodes informed at time 10

State transition probabilities

$P_{i,j}^N(\mathbf{m}) = \Pr\{X_n^N(t+1) = j \mid X_n^N(t) = i, M^N(t) = \mathbf{m}\}, i, j \in S, \ \mathbf{m} \in S_M^N.$

Mean-field approximation

N interacting objects, identically defined

Example

Gossiping nodes, each with state space $S = \{0, 1, \dots K-1\}$

Time is discrete $t \in \mathbb{N}$

Occupancy measure $M_i^N(t)$

- fraction of N objects in state i at time t

Example

30% of nodes informed at time 10

State transition probabilities

$P_{i,j}^N(\mathbf{m}) = \Pr\{X_n^N(t+1) = j \mid X_n^N(t) = i, M^N(t) = \mathbf{m}\}, i, j \in S,\ \mathbf{m} \in S_M^N.$

Mean-field approximation

N interacting objects, identically defined

Example

Gossiping nodes, each with state space $S = \{0, 1, \ldots K-1\}$

Time is discrete $t \in \mathbb{N}$

Occupancy measure $M_i^N(t)$

- fraction of N objects in state i at time t

Example

30% of nodes informed at time 10

State transition probabilities

$$P_{i,j}^N(\mathbf{m}) = \Pr\{X_n^N(t+1) = j \mid X_n^N(t) = i, M^N(t) = \mathbf{m}\}, i, j \in S,\ \mathbf{m} \in S_M^N.$$

Mean-field approximation

N interacting objects, identically defined

Example

Gossiping nodes, each with state space $S = \{0, 1, \ldots K - 1\}$

Time is discrete $t \in \mathbb{N}$

Occupancy measure $M_i^N(t)$

- fraction of N objects in state i at time t

Example

30% of nodes informed at time 10

State transition probabilities

$$P_{i,j}^N(\mathbf{m}) = \Pr\{\underbrace{X_n^N(t+1) = j}_{\text{next state}} \mid X_n^N(t) = i, M^N(t) = \mathbf{m}\}, i, j \in S,\ \mathbf{m} \in S_M^N.$$

Mean-field approximation

N interacting objects, identically defined

Example

Gossiping nodes, each with state space $S = \{0, 1, \dots K-1\}$

Time is discrete $t \in \mathbb{N}$

Occupancy measure $M_i^N(t)$

- fraction of N objects in state i at time t

Example

30% of nodes informed at time 10

State transition probabilities

$$P_{i,j}^N(\mathbf{m}) = \Pr\{\underbrace{X_n^N(t+1) = j}_{\text{next state}} \mid \underbrace{X_n^N(t) = i}_{\text{current state}}, M^N(t) = \mathbf{m}\}, i, j \in S,\ \mathbf{m} \in S_M^N.$$

Mean-field approximation

N interacting objects, identically defined

Example

Gossiping nodes, each with state space $S = \{0, 1, \ldots K - 1\}$

Time is discrete $t \in \mathbb{N}$

Occupancy measure $M_i^N(t)$

- fraction of N objects in state i at time t

Example

30% of nodes informed at time 10

State transition probabilities

$$P_{i,j}^N(\mathbf{m}) = \Pr\{\underbrace{X_n^N(t+1) = j}_{\text{next state}} \mid \underbrace{X_n^N(t) = i}_{\text{current state}}, \underbrace{M^N(t) = \mathbf{m}}_{\text{occupancy measure}}\}, i, j \in S,\ \mathbf{m} \in S_M^N.$$

Mean-field approximation

- Presented by Le Boudec et al. [QEST'07]
- Approximation for the occupancy measure for large N

Theorem (Mean-field convergence)

Fix $M^N(0) = \mu(0)$ *for* $\forall N$, *define*

$$P(\mathbf{m}) = \lim_{N\to\infty} P^N(\mathbf{m}),$$

and the deterministic process

$$\mu(t+1) = \mu(t) \cdot P(\mu(t)).$$

Then

$$\lim_{N\to\infty} M^N(t) = \mu(t), \text{ with probability 1.}$$

Mean-field approximation

- Presented by Le Boudec et al. [QEST'07]
- Approximation for the occupancy measure for large N

Theorem (Mean-field convergence)

Fix $M^N(0) = \mu(0)$ *for* $\forall N$*, define*

$$P(\mathbf{m}) = \lim_{N\to\infty} P^N(\mathbf{m}),$$

and the deterministic process

$$\mu(t+1) = \mu(t) \cdot P(\mu(t)).$$

Then

$$\lim_{N\to\infty} M^N(t) = \mu(t), \text{ with probability 1.}$$

Mean-field approximation

- Presented by Le Boudec et al. [QEST'07]
- Approximation for the occupancy measure for large N

Theorem (Mean-field convergence)

Fix $M^N(0) = \mu(0)$ *for* $\forall N$*, define*

$$P(\mathbf{m}) = \lim_{N\to\infty} P^N(\mathbf{m}),$$

and the deterministic process

$$\mu(t+1) = \mu(t) \cdot P(\mu(t)).$$

Then

$$\lim_{N\to\infty} M^N(t) = \mu(t), \text{ with probability 1}.$$

i.e., $\mu(t)$ is the deterministic limit occupancy measure for $N \to \infty$.

Mean-field approximation

- Presented by Le Boudec et al. [QEST'07]
- Approximation for the occupancy measure for large N

Theorem (Mean-field convergence)

Fix $M^N(0) = \mu(0)$ *for* $\forall N$, *define*

$$P(\mathbf{m}) = \lim_{N \to \infty} P^N(\mathbf{m}),$$

and the deterministic process

$$\mu(t+1) = \mu(t) \cdot P(\mu(t)).$$

Then

$$\lim_{N \to \infty} M^N(t) = \mu(t), \text{ with probability 1.}$$

i.e., for $\forall i$, $M_i^N(t)$ in state i at time t is known with probability 1.

Our methodology

1. Description of the protocol
2. Local states and transitions
3. Transition probabilities
4. Mean-field requirements
5. Mean-field limit

Example

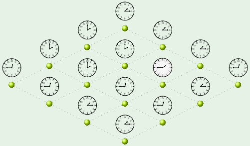

Figure: GTP network

Our methodology

1. Description of the protocol
2. Local states and transitions
3. Transition probabilities
4. Mean-field requirements
5. Mean-field limit

Example

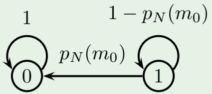

Figure: 1D ST diagram

Our methodology

1. Description of the protocol
2. Local states and transitions
3. Transition probabilities
4. Mean-field requirements
5. Mean-field limit

Example

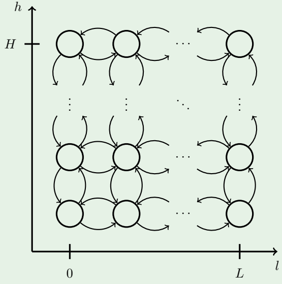

Figure: 2D ST diagram

Our methodology

1. Description of the protocol
2. Local states and transitions
3. Transition probabilities
4. Mean-field requirements
5. Mean-field limit

Example

$$\Pr\{X_n^N(t+1) = j \mid X_n^N(t) = i, M^N(t) = \mathbf{m}\}$$

Figure: Trans. probabilities

Our methodology

1. Description of the protocol
2. Local states and transitions
3. Transition probabilities
4. Mean-field requirements
5. Mean-field limit

Example

$$P^N(\mathbf{m}) = \begin{pmatrix} \cdot & \cdots & \cdot \\ \vdots & P^N_{i,j}(\mathbf{m}) & \vdots \\ \cdot & \cdots & \cdot \end{pmatrix}$$

Figure: Transition prob. matrix

Our methodology

1. Description of the protocol
2. Local states and transitions
3. Transition probabilities
4. Mean-field requirements
5. Mean-field limit

Example

$$M^N(t) = \Big(M_0^N(t), \ldots, M_{K-1}^N(t)\Big)$$

$$M_i^N(t) = \frac{1}{N} \sum_{n=1}^{N} 1_{\{X_n^N(t)=i\}}$$

Figure: Occupancy measure

Our methodology

1. Description of the protocol
2. Local states and transitions
3. Transition probabilities
4. Mean-field requirements
5. Mean-field limit

Example

$$p^N(m_0) = g \cdot \frac{m_0 \cdot N}{N - 1}$$

Figure: Prob. of transition $0 \rightarrow 1$

Our methodology

1. Description of the protocol
2. Local states and transitions
3. Transition probabilities
4. Mean-field requirements
5. Mean-field limit

Example

$$\lim_{N\to\infty} p^N(m_0) = \lim_{N\to\infty} g \cdot \frac{m_0 \cdot N}{N-1}$$

Figure: Limit of $P(0|1)$

Our methodology

1. Description of the protocol
2. Local states and transitions
3. Transition probabilities
4. Mean-field requirements
5. Mean-field limit

Example

$$p(m_0) = g \cdot m_0$$

Figure: Limit of $P(0|1)$

Our methodology

1. Description of the protocol
2. Local states and transitions
3. Transition probabilities
4. Mean-field requirements
5. Mean-field limit

Example

$$P(m_0) = \begin{pmatrix} 1 & 0 \\ p(m_0) & 1 - p(m_0) \end{pmatrix}$$

Figure: Limit of prob. matrix

Our methodology

1. Description of the protocol
2. Local states and transitions
3. Transition probabilities
4. Mean-field requirements
5. Mean-field limit

Example

$$\mu(0) = (0.01, 0.99)$$

$$P(m_0) = \begin{pmatrix} 1 & 0 \\ m_0 & 1 - m_0 \end{pmatrix}$$

$$\mu(1) = \mu(0) \cdot P(m_0)$$

Our methodology

1. Description of the protocol
2. Local states and transitions
3. Transition probabilities
4. Mean-field requirements
5. Mean-field limit

Example

$$\mu(0) = (0.01, 0.99)$$

$$P(m_0) = \begin{pmatrix} 1 & 0 \\ m_0 & 1 - m_0 \end{pmatrix}$$

$$\mu(2) = \mu(1) \cdot P(m_0)$$

Our methodology

1. Description of the protocol
2. Local states and transitions
3. Transition probabilities
4. Mean-field requirements
5. Mean-field limit

Example

$$\mu(0) = (0.01, 0.99)$$

$$P(m_0) = \begin{pmatrix} 1 & 0 \\ m_0 & 1 - m_0 \end{pmatrix}$$

$$\mu(3) = \mu(2) \cdot P(m_0) \ldots$$

Our methodology

1. Description of the protocol
2. Local states and transitions
3. Transition probabilities
4. Mean-field requirements
5. Mean-field limit

Example

$$\mu(0) = (0.01, 0.99)$$

$$P(m_0) = \begin{pmatrix} 1 & 0 \\ m_0 & 1 - m_0 \end{pmatrix}$$

$$\mu(t+1) = \mu(t) \cdot P(m_0)$$

Gossip Time Protocol

- Self-managing time synchronization protocol
- Synchronization of clocks in both time and frequency
- Nodes use peer-sampling service
- Presence of “time source”

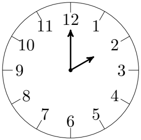

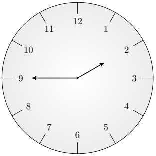

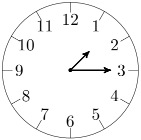

Gossip Time Protocol

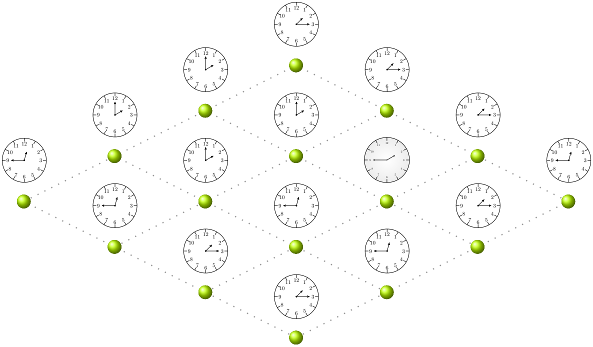

- Network of nodes, each equipped with a local clock
- Nodes periodically exchange time info in random fashion
- Node with the worse-quality time adopts the higher-quality time of its peer

Gossip Time Protocol

A initiates a gossip with random peer *B*

- *A* is *active*
- *B* is *passive*

In *Basic GTP*

- time sample based on hop count metric
- sample with higher hop count is rejected
- update is enforced if last-update timer expires

In *Gradual GTP*

- node may adjust gossip frequency

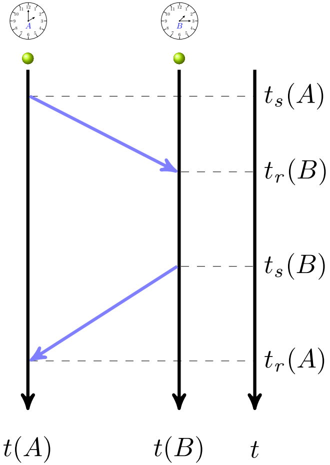

GTP Analysis

Node's behaviour := state + occupancy measure

States

Represented by a triple (g, l, h)

- gossip delay g
- last update counter l
- hop count h

Occupancy measure

$\mathbf{m}_{(g,l,h)}$ is fraction of nodes in state (g, l, h)

Gossip delay function

$$G : \{0, \ldots, H, \infty\} \mapsto \{0, \ldots, G_{\max}\}$$

GTP Analysis

Node's behaviour $:=$ state + occupancy measure

States

Represented by a triple (g, l, h)

- gossip delay g
- last update counter l
- hop count h

Occupancy measure

$\mathbf{m}_{(g,l,h)}$ is fraction of nodes in state (g, l, h)

Gossip delay function

$$G : \{0, \ldots, H, \infty\} \mapsto \{0, \ldots, G_{\text{max}}\}$$

GTP Analysis

Node's behaviour := state + occupancy measure

States

Represented by a triple (g, l, h)

- gossip delay g
- last update counter l
- hop count h

Occupancy measure

$\mathbf{m}_{(g,l,h)}$ is fraction of nodes in state (g, l, h)

Gossip delay function

$$G : \{0, \ldots, H, \infty\} \mapsto \{0, \ldots, G_{\text{max}}\}$$

GTP Analysis

States

1. Time sources $\Leftrightarrow$ states $(g \geq 0, l, h = 0)$
2. Active nodes $\Leftrightarrow$ states $(g = 0, l, h > 0)$
3. Passive nodes $\Leftrightarrow$ states $(g > 0, l, h > 0)$

For active and passive nodes:

Transition probabilities

- Successful update $\Leftrightarrow (g, \cdot, h) \to (G(h' + 1), L, h' + 1)$
- Clock did not update $\Leftrightarrow (g, l, h) \to (G(h), l - 1, h)$

GTP Analysis

States

1. Time sources $\Leftrightarrow$ states $(g \geq 0, l, h = 0)$
2. Active nodes $\Leftrightarrow$ states $(g = 0, l, h > 0)$
3. Passive nodes $\Leftrightarrow$ states $(g > 0, l, h > 0)$

For active and passive nodes:

Transition probabilities

- Successful update $\Leftrightarrow (g, \cdot, h) \rightarrow (G(h' + 1), L, h' + 1)$
- Clock did not update $\Leftrightarrow (g, l, h) \rightarrow (G(h), l - 1, h)$

GTP Analysis

A *time source*

- has state $(g, L, 0)$
- is independent of the enviroment

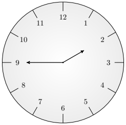

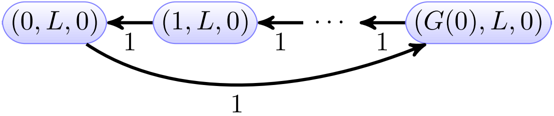

Transition probabilities

$$P^N_{(g,l,0|g>0),(g-1,L,0)}(\mathbf{m}) = 1$$
$$P^N_{(0,l,0),(G(0),L,0)}(\mathbf{m}) = 1$$

GTP Analysis

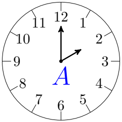

An *active node*

- has state $(0, l, h)$
- initiates gossip with random peer

Update $A \leftarrow (G(h_B + 1), L, h_B + 1)$ is

- enforced, if $l_A = 0 \wedge h_B \neq \infty$
- optional, if $l_A \neq 0 \wedge h_A > h_B$

Transition probabilities

Probability of successful update:

$$P^N_{s(l=0),s'}(\mathbf{m}) = \mathbf{m}_{(g',l',h' \mid g'>0)} \cdot \frac{N}{N-1} \cdot \mathrm{noc}^N(\mathbf{m}), \quad \forall h' < \infty,$$

$$P^N_{s(l>0),s'}(\mathbf{m}) = \mathbf{m}_{(g',l',h' \mid g'>0)} \cdot \frac{N}{N-1} \cdot \mathrm{noc}^N(\mathbf{m}), \quad \forall h' < h.$$

$$s(\cdot) = (0, \cdot, h \mid h > 0), s' = (G(h'+1), L, h'+1)$$

GTP Analysis

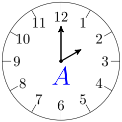

An *active node*

- has state $(0, l, h)$
- initiates gossip with random peer

Update $A \leftarrow (G(h_B + 1), L, h_B + 1)$ is

- enforced, if $l_A = 0 \wedge h_B \neq \infty$
- optional, if $l_A \neq 0 \wedge h_A > h_B$

Transition probabilities

Probability of successful update:

$$P^N_{s(l=0),s'}(\mathbf{m}) = \mathbf{m}_{(g',l',h'|g'>0)} \cdot \frac{N}{N-1} \cdot \mathrm{noc}^N(\mathbf{m}), \quad \forall h' < \infty,$$

$$P^N_{s(l>0),s'}(\mathbf{m}) = \mathbf{m}_{(g',l',h'|g'>0)} \cdot \frac{N}{N-1} \cdot \mathrm{noc}^N(\mathbf{m}), \quad \forall h' < h.$$

$$s(\cdot) = (0, \cdot, h \mid h > 0), s' = (G(h'+1), L, h'+1)$$

GTP Analysis

Limits

The “no-collision” probability $\text{noc}^N(\mathbf{m})$ converges for $N \to \infty$:

$$\text{noc}(\mathbf{m}) = \lim_{N\to\infty} \text{noc}^N(\mathbf{m}) = \lim_{N\to\infty} \left(\frac{N-3}{N-1}\right)^{\mathbf{m}_{(0,l,h)}\cdot N-1} = e^{-2\cdot\mathbf{m}_{(0,l,h)}}.$$

Moreover,

$$\lim_{N\to\infty} \frac{N}{N-1} = 1$$

Example

$$P^N_{s(0),s'}(\mathbf{m}) = \mathbf{m}_{(g',l',h'|g'>0)} \cdot \frac{N}{N-1} \cdot \text{noc}^N(\mathbf{m})$$

GTP Analysis

Limits

The "no-collision" probability $\text{noc}^N(\mathbf{m})$ converges for $N \to \infty$:

$$\text{noc}(\mathbf{m}) = \lim_{N\to\infty} \text{noc}^N(\mathbf{m}) = \lim_{N\to\infty} \left(\frac{N-3}{N-1}\right)^{\mathbf{m}_{(0,l,h)}\cdot N-1} = e^{-2\cdot\mathbf{m}_{(0,l,h)}}.$$

Moreover,

$$\lim_{N\to\infty} \frac{N}{N-1} = 1$$

Example

$$P^N_{s(0),s'}(\mathbf{m}) = \mathbf{m}_{(g',l',h'|g'>0)} \cdot \frac{N}{N-1} \cdot \text{noc}^N(\mathbf{m})$$

GTP Analysis

Limits

The “no-collision” probability $\text{noc}^N(\mathbf{m})$ converges for $N \to \infty$:

$$\text{noc}(\mathbf{m}) = \lim_{N\to\infty} \text{noc}^N(\mathbf{m}) = \lim_{N\to\infty} \left(\frac{N-3}{N-1}\right)^{\mathbf{m}_{(0,l,h)}\cdot N-1} = e^{-2\cdot\mathbf{m}_{(0,l,h)}}.$$

Moreover,

$$\lim_{N\to\infty} \frac{N}{N-1} = 1$$

Example

$$\lim_{N\to\infty} P^N_{s(0),s'}(\mathbf{m}) = \mathbf{m}_{(g',l',h'|g'>0)} \cdot \quad 1 \quad \cdot e^{-2\cdot\mathbf{m}_{(0,l,h)}}$$

Mean-field vs. Emulation

Emulation results taken from [Iwanicki]

Setup

- $N = 1500$, 1 time source
- $G_{max} = 25$ sec, $L = 25$ sec, $H = 15$

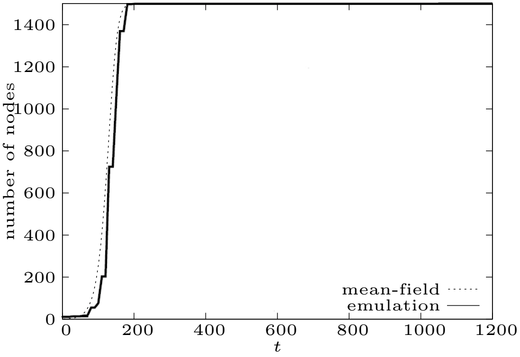

Figure: # nodes discovering the time source over time

Mean-field vs. Emulation

Emulation results taken from [Iwanicki]

Setup

- $N = 1500$
- $G_{\text{max}} = 25$ sec, $L = 25$ sec, $H = 15$

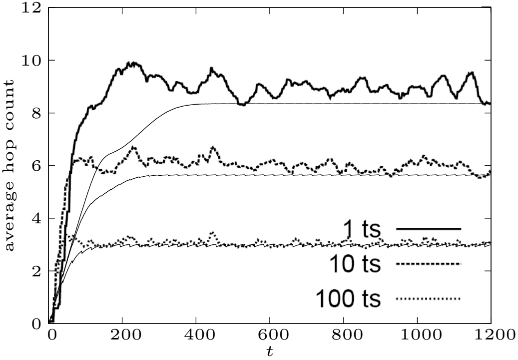

Figure: Average hop count for different # time sources

Mean-field vs. Emulation

Emulation results taken from [Iwanicki]

Setup

- $N = 1500$, 1 time source
- $G_{\max} = 25$ sec, $L = 25$ sec, $H = 15$

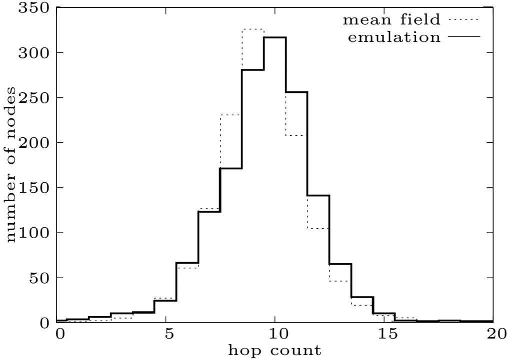

Figure: Hop count distribution after stabilization

Further Measurements

- $G_{\text{max}} = 25$ sec, $L = 25$ sec, $H = 15$

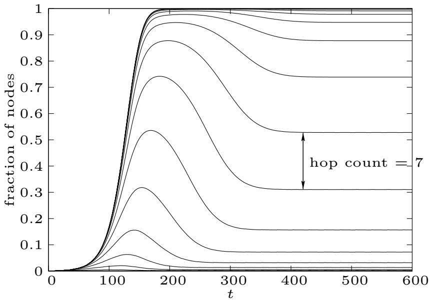

Figure: Hop count distribution over the first 10 min

Static gossip delay

Gossip delay function

$$G : \{0, \ldots, H, \infty\} \mapsto \{0, \ldots, G_{max}\}$$

- $G_{max} = 25$ sec, $L = 25$ sec, $H = 15$

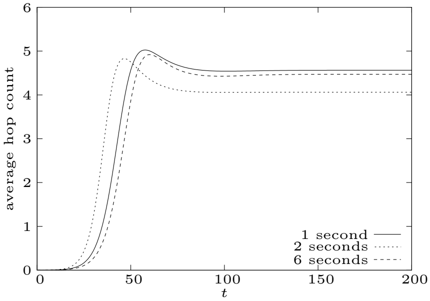

Figure: Average hop count for different gossip delays

Dynamic gossip delay

Adaptation of gossip delay $G(h)$

- nodes with “bad quality time” to gossip more
- $G_{\min}$ = minimal gossip delay

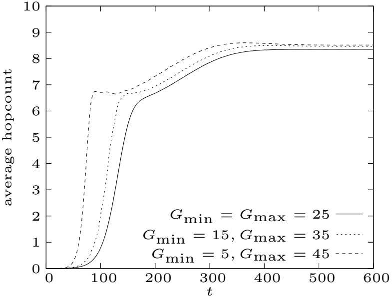

Figure: Average hop count for different $G_{\max}, G_{\min}$ in first 10 min