# *SENDING AND FORGETTING: TERMINATION OF AMNESIAC FLOODING ON A GRAPH**

WALTER HUSSAK AND
AMITABH TREHAN
LOUGHBOROUGH UNIVERSITY

- *On the Termination of Flooding. STACS 2020*
- *On Termination of a Flooding Process (Brief Announcement). PODC 2019*

www.amitabhtrehan.net
www.huntforthetowel.wordpress.com

# THE AMNESIAC AMBITIOUS **WHATSAPPERS!**

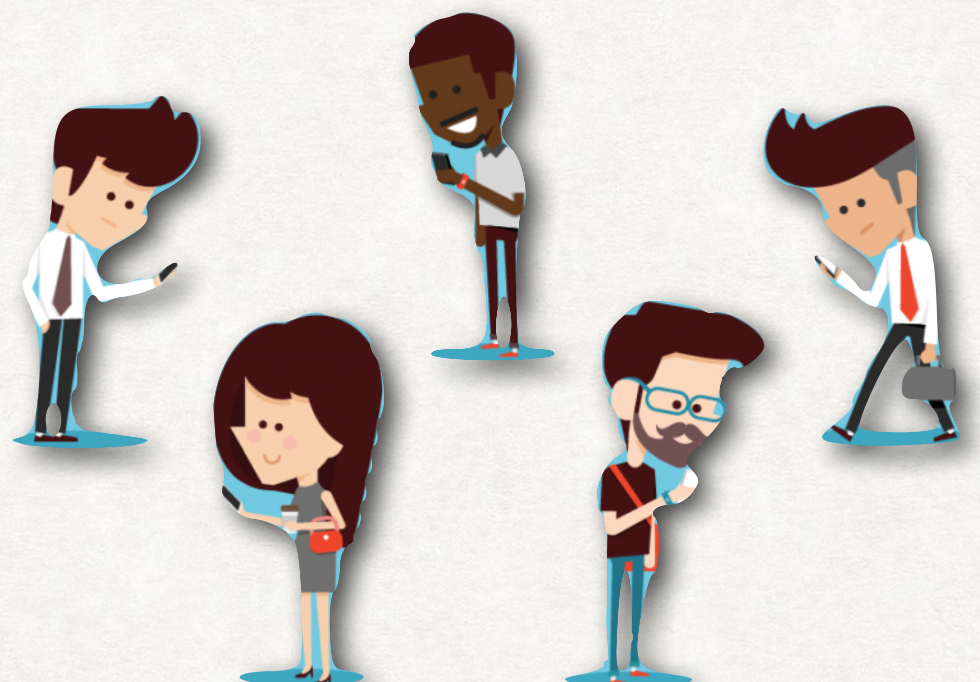

*Creatures that exists amongst us - maybe right now (in the audience!) -*

*Often with a deep interest in politics and unwittingly purveyors of fake news!*

# THE AMNESIAC AMBITIOUS **WHATSAPPERS!**

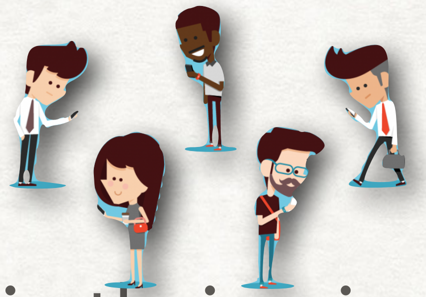

- Forward every message they receive!
- However, <u>very discerning persons</u> (at least in their mind!):
  - Do not forward messages back to the person received from!
- But <u>forgetful </u>- too many messages/too little time!
  - Forget if the person had sent the message some time before!
  - Will flood again if `asked'!

Q: Will that annoying WhatsApp message ever stop?

# FORMAL MODEL

- A graph **G(V,E)** is a formal model for a network
- Graph G: The network
- V: Vertices are the nodes
- E: Edges are the connections

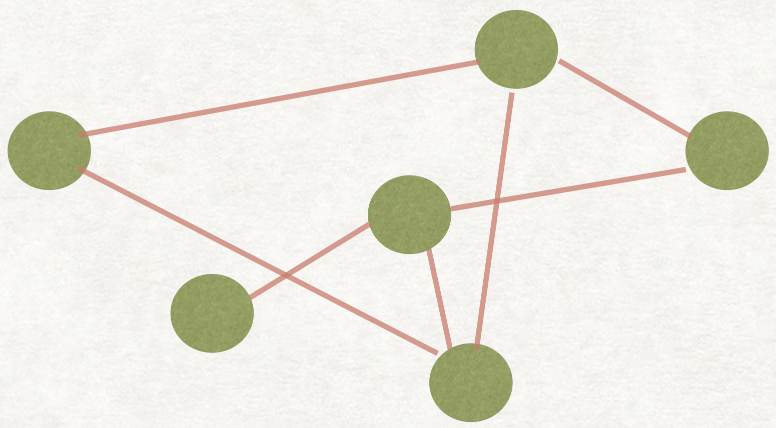

# FORMAL MODEL

- The graph **G(V,E)** is a formal model for a network.

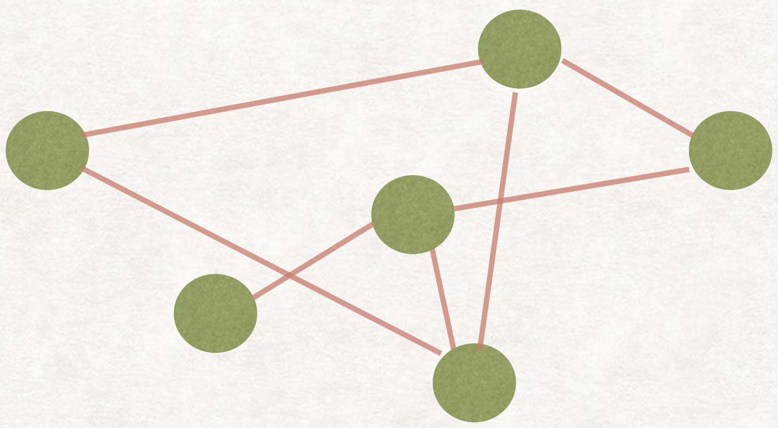

- Message passing: nodes only communicate by sending messages.

  A. Synchronous and Reliable: communication in synchronous rounds, messages delivered by end of the round sent in.

  B. Adaptive Round asynchronous: 'global' rounds but adaptive adversary decides the delay on each edge

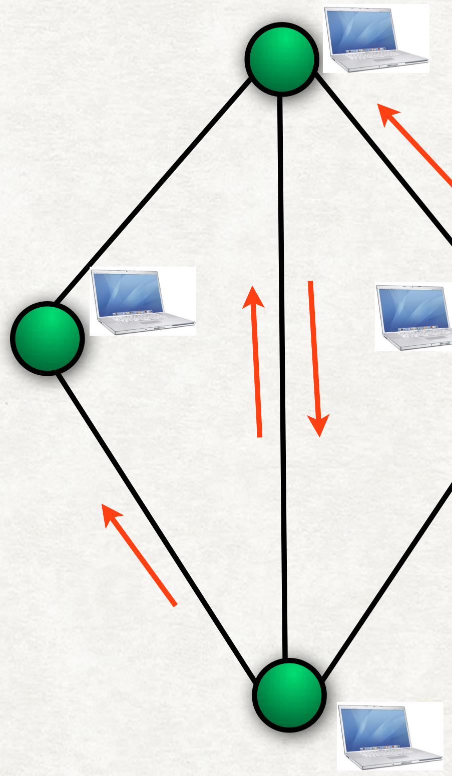

# GAME:THE AMNESIAC AMBITIOUS **WHATSAPPER!**

## WORD OF MOUTH ANNOUNCEMENT

- Problem: Inform everyone about class = Message M!
- Solution: Send/ Broadcast/Flood M from a source to every node in the network!

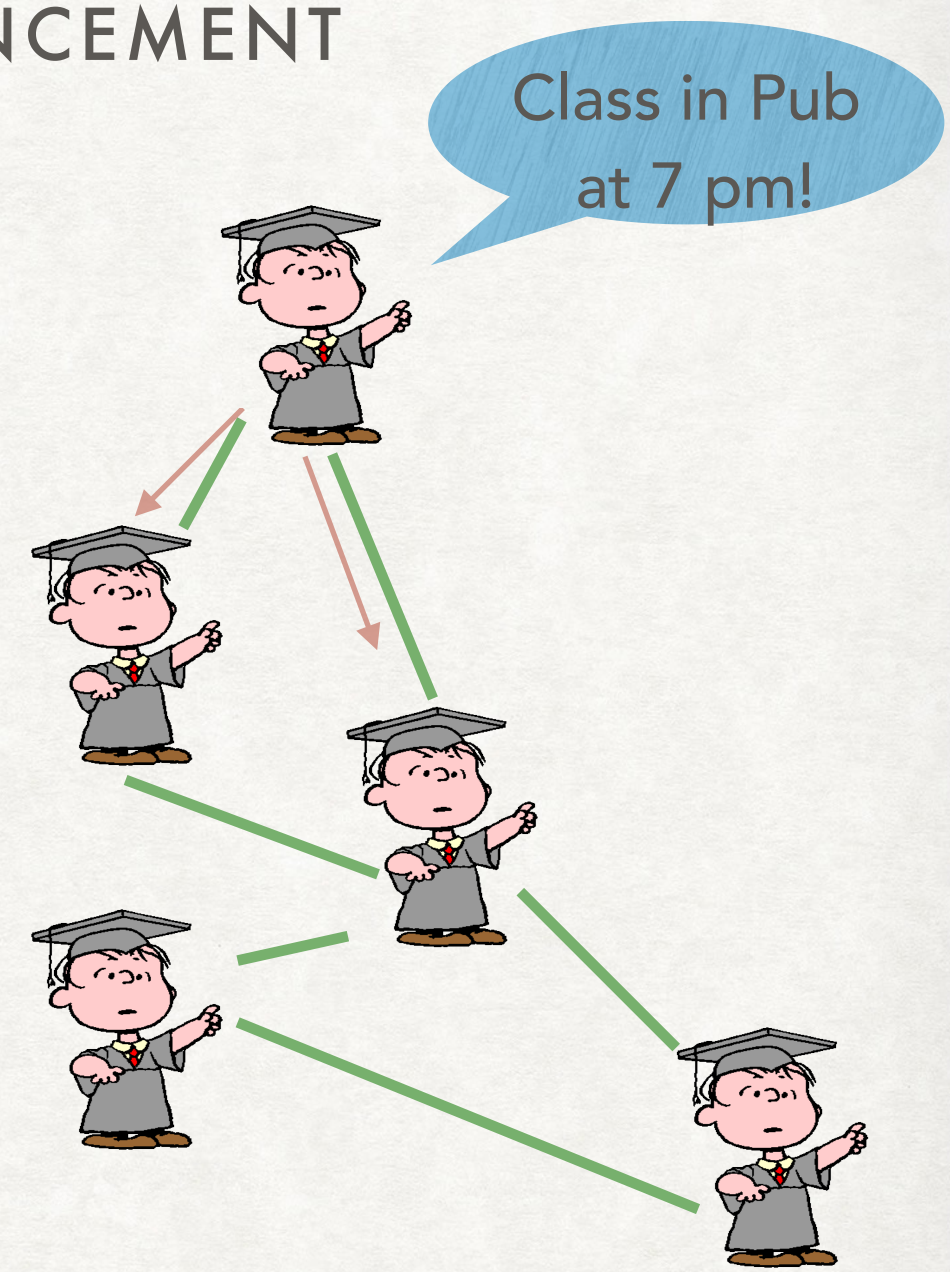

# AMNESIAC FLOODING (AF)!

- Flooding: `dumb' but most fundamental of distributed algorithms
- AF: Flooding with a slight twist!
- Ambitious WhatsApper :

  *If I have a message, I will forward!*

- Amnesiac WhatsApper:

  *I shalt not remember the past!*

- Polite WhatsApper:

  *If you have just sent me the message, I will not flood it back!*

  - Nodes only remember previous round (no explicit message flags)
  - Send messages to exactly all those who did not send to it in the previous round

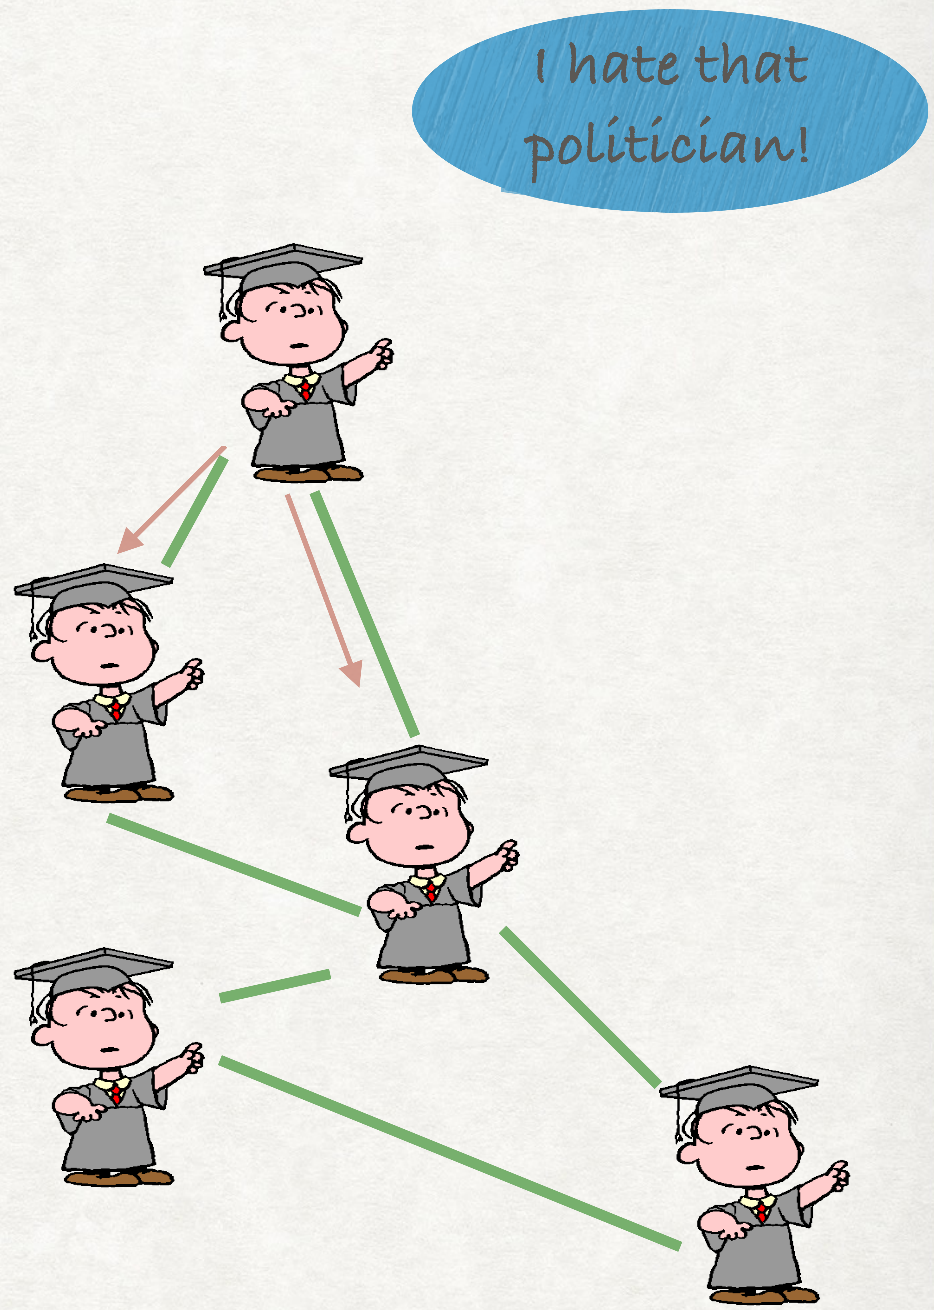

- More formally: ***Amnesiac Flooding (AF)***
- Start: A distinguished node $l$
- *Round 1:* $l$ sends message $M$ to all neighbours
- *Round i ( > 1):* If node $v$ receives $M$ from neighbours $R$ in round *i-1*, $v$ floods to $n(v)/R$ (i.e. neighbours besides $R$)

Q1: Does AF terminate?

Q2: If AF terminates, how long does it take?

*Termination: M is not sent in a round (and subsequent rounds) by any node in the network

Termination time: The number of rounds AF takes i.e. last round with a transmission

# AMNESIAC FLOODING: TERMINATION

*Amnesiac Flooding*

- Start: A distinguished node $l$
- *Round 1:* $l$ sends message $M$ to all neighbours
- *Round i ( > 1):* If node $v$ receives $M$ from neighbours $R$ in round *i-1*, $v$ floods to $n(v)/R$ (i.e. neighbours besides $R$)

Q1: Does this process **terminate?**

- Let's try some examples

**1. Line Graph:**

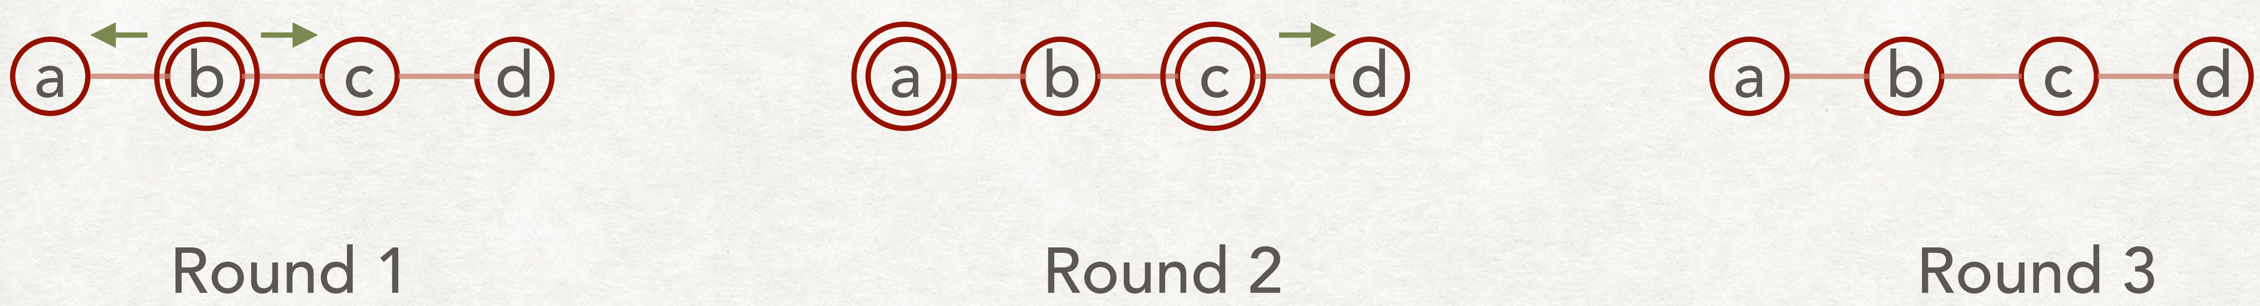

# AMNESIAC FLOODING: TERMINATION

***Amnesiac Flooding***

- Start: A distinguished node $l$
- *Round 1:* $l$ sends message $M$ to all neighbours
- *Round i ( > 1):* If node $v$ receives $M$ from neighbours $R$ in round $i$-$1$, $v$ floods to $n(v)/R$ (i.e. neighbours besides $R$)

Q1: Does this process **terminate?**

## <u>2. Triangle/Clique/Odd Cycle:</u>

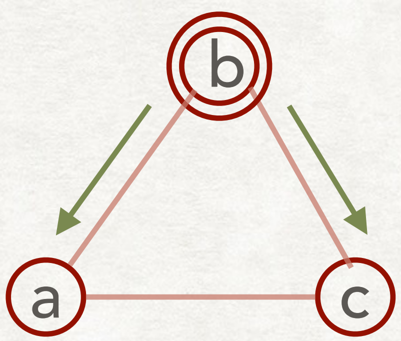


Round 1

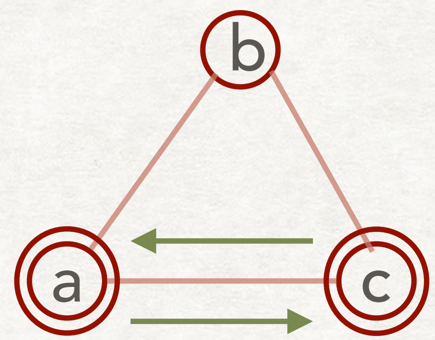


Round 2

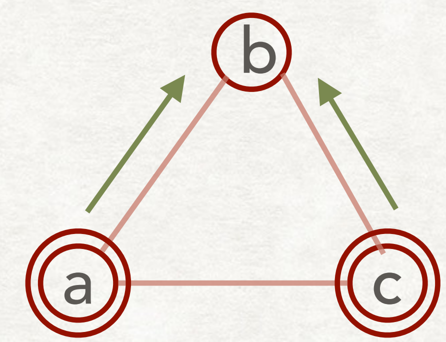


Round 3

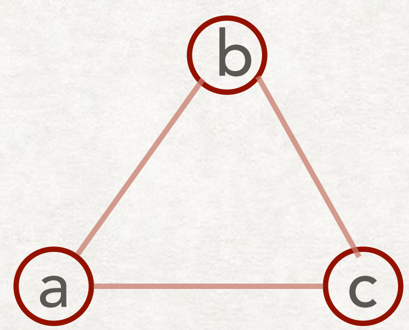


Round 4

# AMNESIAC FLOODING: TERMINATION

*Amnesiac Flooding*

- Start: A distinguished node *l*
- *Round 1: l* sends message *M* to all neighbours
- *Round i ( > 1):* If node *v* receives *M* from neighbours *R* in round *i-1*, *v* floods to *n(v)/R* (i.e. neighbours besides *R*)

Q1: Does this process terminate?

## 3. Even Cycle:

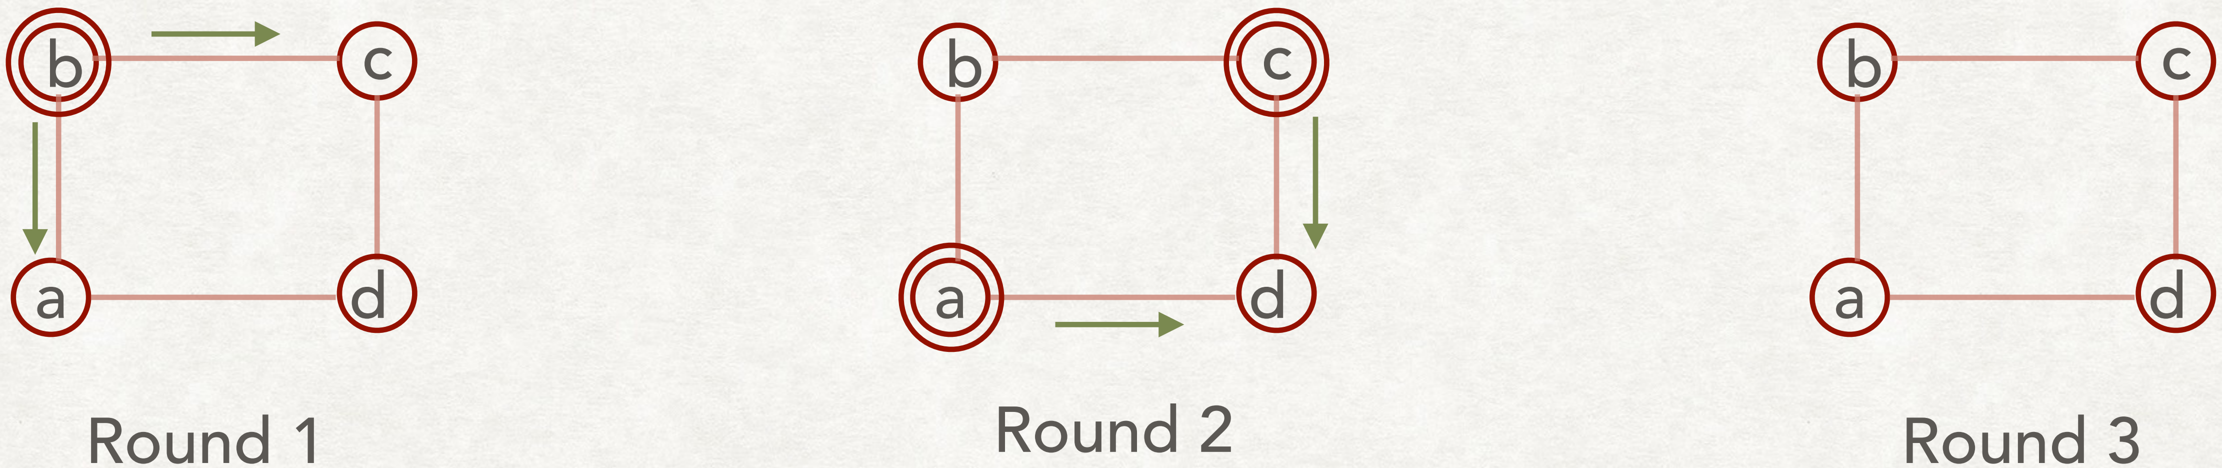

# AMNESIAC FLOODING: TERMINATION

> *Amnesiac Flooding*
> - Start: A distinguished node $l$
> - *Round 1:* $l$ sends message $M$ to all neighbours
> - *Round i* ( $> 1$): If node $v$ receives $M$ from neighbours $R$ in round $i$-$1$, $v$ floods to $n(v)/R$ (i.e. neighbours besides $R$)

Q1: Does this process terminate?

## 3. More complex topologies:

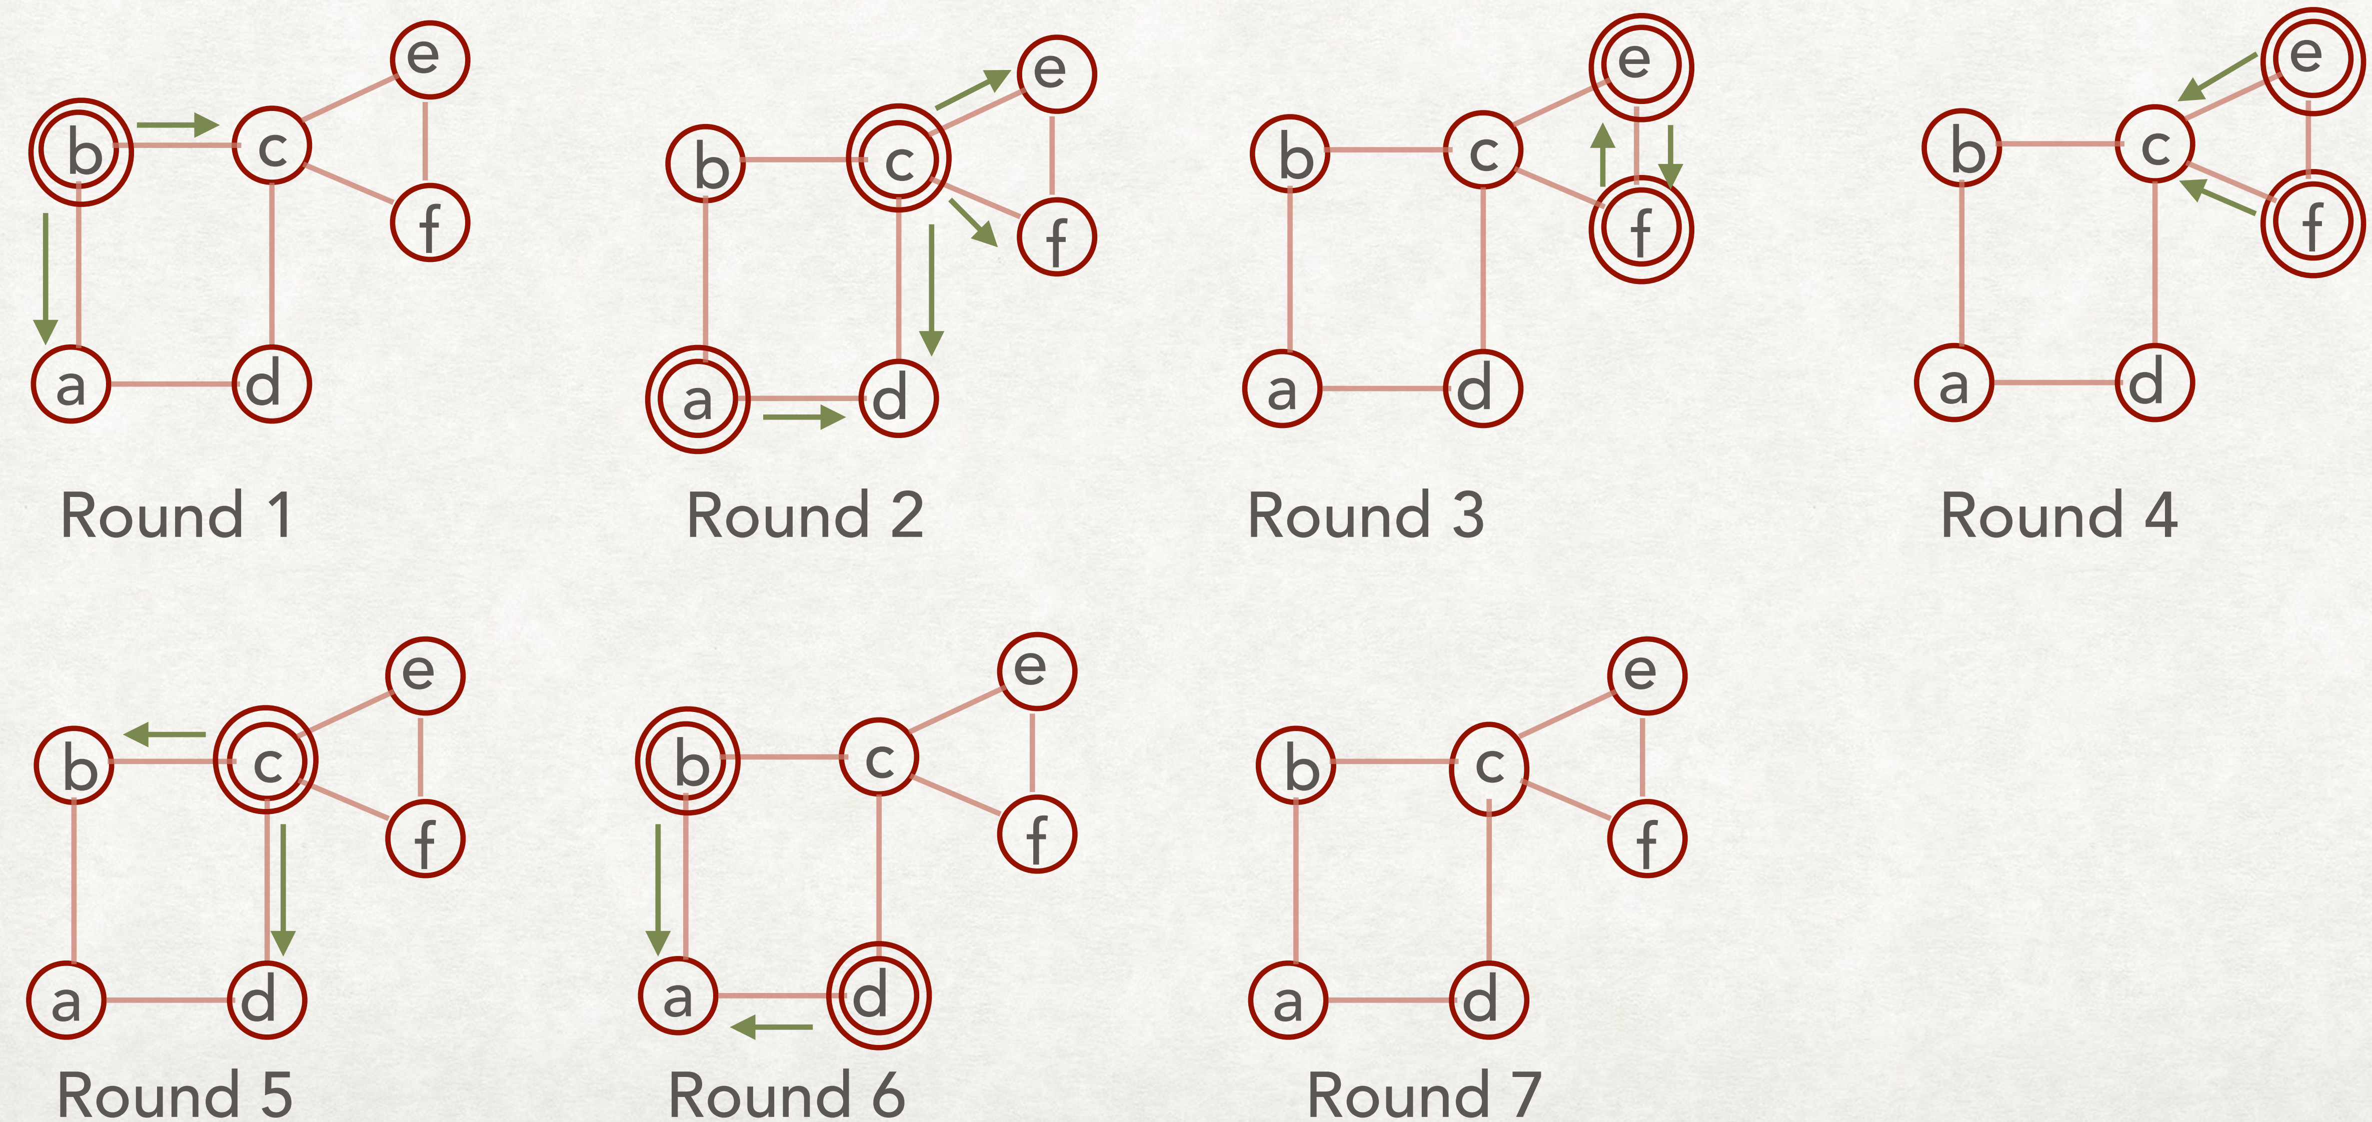

# AMNESIAC FLOODING: TERMINATION

***Amnesiac Flooding***

- Start: A distinguished node $l$
- *Round 1:* $l$ sends message $M$ to all neighbours
- *Round i* ( $> 1$): If node $v$ receives $M$ from neighbours $R$ in round *i-1*, $v$ floods to $n(v)/R$ (i.e. neighbours besides $R$)

Q1: Does this process terminate?

## 3. More complex topologies:

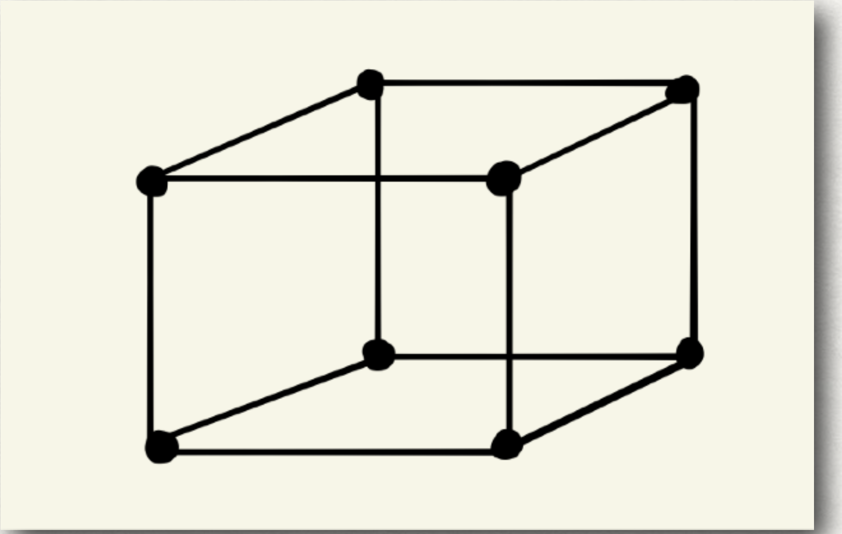

Hypercube

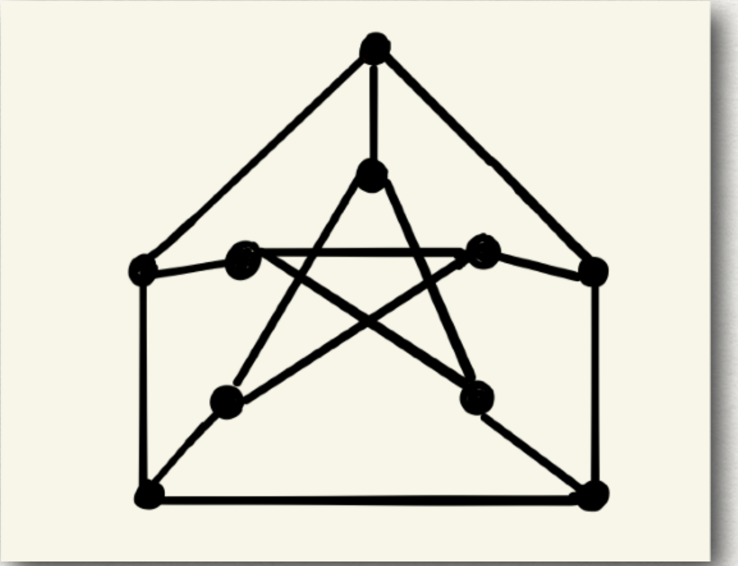

Petersen Graph

# AMNESIAC FLOODING: TERMINATION

*Amnesiac Flooding*

- Start: A distinguished node $l$
- *Round 1:* $l$ sends message $M$ to all neighbours
- *Round i* ( $>1$ ): If node $v$ receives $M$ from neighbours $R$ in round *i-1*, $v$ floods to $n(v)/R$ (i.e. neighbours besides $R$)

Q1: Does this process terminate?

## 3. More complex topologies:

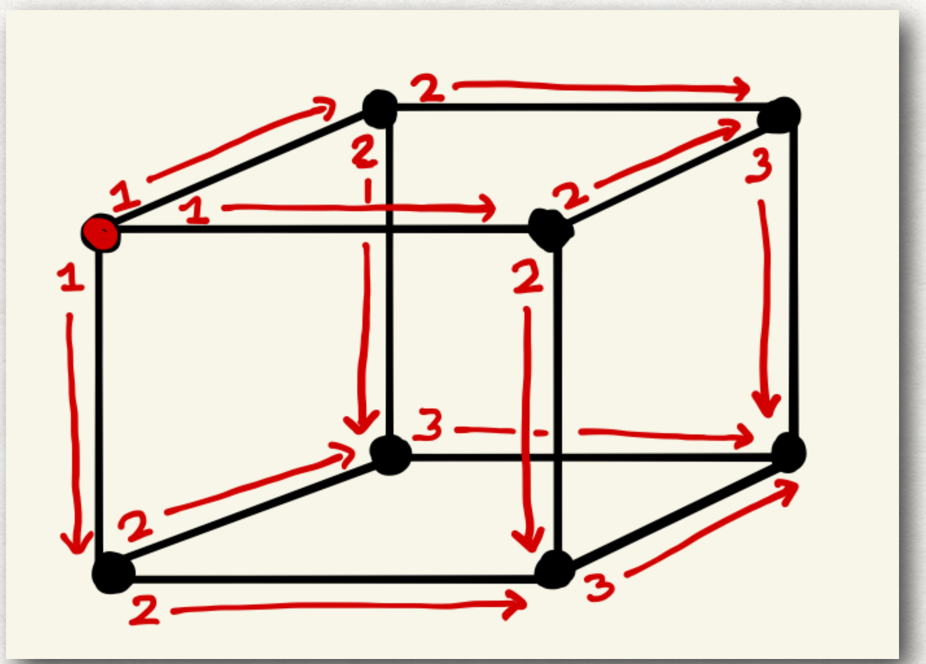

Hypercube
T = 3 rounds
Diameter = 3

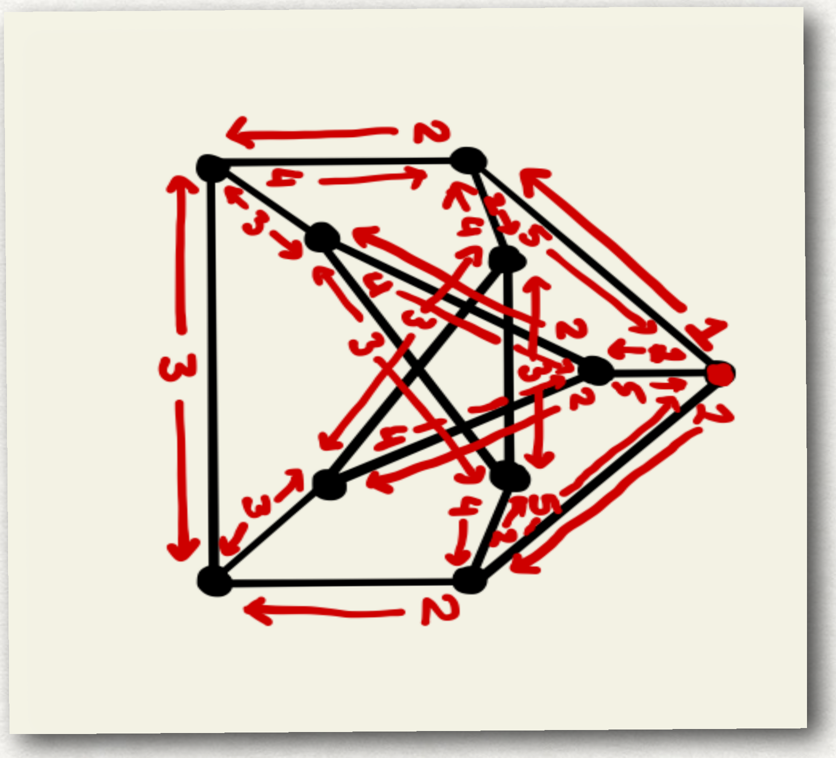

Petersen Graph
T = 5 rounds
Diameter = 2

Why such different termination times?

# AF TERMINATION

DOES AF TERMINATE?

ON EVERY GRAPH?

IS IT QUICK?

# AF TERMINATION

DOES AF TERMINATE?

YES

ON EVERY GRAPH?

YES

# AF TERMINATION

*Theorem 1. Given a finite graph G, Amnesiac Flooding (AF) from a single source will terminate in a finite number of rounds*

**Proof.**

Proof is by contradiction

High level idea:

- Define ***round-sets*** as set of nodes receiving M in a particular round
- Consider sequences of round-sets E of even duration:

    *Condition of non-termination:*

    There must be at least one E having the same node repeated!

    - We show this is not possible!

E.g.

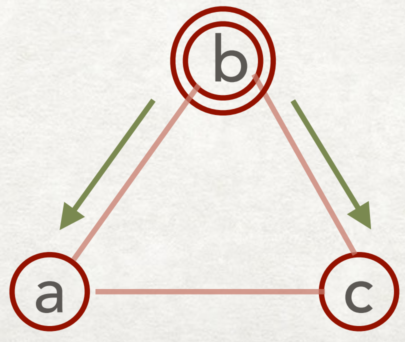


$R_0 = \{b\}$

$R1 = \{a,c\}$ $R2 = \{a,c\}$

$R3 = \{b\}$ $R4 = \{\}$

# AF TERMINATION

**Theorem 1.** *Given a finite graph G, Amnesiac Flooding (AF) from a single source will terminate in a finite number of rounds*

**Detailed Proof.**

Proof is by contradiction

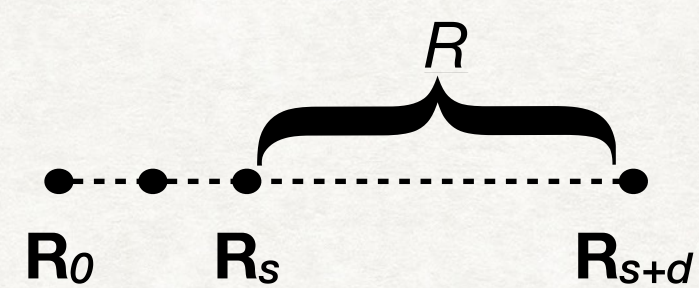


Definition. **Round-sets:**

$R_0, R_1, \ldots$

$R_0$ : Singleton with node $l$

$R_i$ : Set of nodes receiving a message at round i (>0)

Define **R** to be the set of finite sequences of the form

$\underline{R} = R_s, \ldots, R_{s+d}$ where $s \geq 0$, $d \geq 0$ and $R_s \cap R_{s+d} \neq \varnothing$

Consider $\mathbf{R}^e$ to be subset of R where d is even.

Claim 1. If AF does not terminate, $\mathbf{R}^e$ will be non-empty

# AF TERMINATION

**Theorem 1.** *Given a finite graph G, Amnesiac Flooding (AF) from a single source will terminate in a finite number of rounds*

Claim 1. If AF does not terminate, $\mathbf{R^e}$ will be non-empty

Proof. G is finite, therefore, if AF does not terminate, a node *x* must occur in infinitely many round-sets. Consider the first 3 such round-sets (e.g. 2,5, and 7);
Surely at least two of these are evenly spaced.

Thus, **$\mathbf{R^e}$ is non-empty.**

For the proof of Theorem 1, assume that $\mathbf{R^e}$ is non-empty and derive a contradiction.

# AF TERMINATION

**Theorem 1.** *Given a finite graph G, Amnesiac Flooding (AF) from a single source will terminate in a finite number of rounds*

**Detailed Proof.**

Claim 1. If AF does not terminate, $\mathbf{R^e}$ will be non-empty

Assume $\mathbf{R^e}$ is non-empty

Consider the *first smallest* sequence in the set $\mathbf{R^e}$!
i.e.

$$\underline{R^*} = R_{ms}, \ldots, R_{ms+md}$$

Where *md* is the shortest duration of any sequence
and ms is the earliest starting point of such sequences

Consider a node *x* common to $R_{ms}$ and $R_{ms+md}$ (exists by assumption)
and a node y which sent M to node x in round $R_{ms+md}$ :

**Case 1. y also sent M to x in round ms**

**OR**

**Case 2. x sent M to y in round ms+1**

# AF TERMINATION

**Theorem 1.** *Given a finite graph G, Amnesiac Flooding (AF) from a single source will terminate in a finite number of rounds*

Claim 1. If AF does not terminate, $\mathbf{R^e}$ will be non-empty

The *first smallest* sequence in the set $\mathbf{R^e}$!

$$\underline{R^*} = R_{ms}, \ldots, R_{ms+md}$$

node *x* in $R_{ms}$ and $R_{ms+md}$
node y: sent M to x in $R_{ms+md}$

**Case 1. y also sent M to x in round ms**


Thus, round ms-1 is either 0 and y is origin node $l$
Or y received M in round ms-1 (> 0):

Either way, there is a sequence $\underline{R^{*'}} = R_{ms-1}, \ldots, R_{ms+md-1}$ of even min-duration md but earlier start point ms-1 with $R_{ms-1} \cap R_{ms+md-1} \neq \emptyset$

Contradiction.

# AF TERMINATION

**Theorem 1.** *Given a finite graph G, Amnesiac Flooding (AF) from a single source will terminate in a finite number of rounds*

**Case 2. x sent M to y in round ms+1**

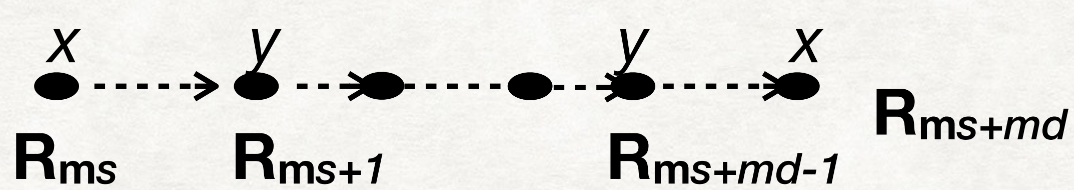


By definition, smallest value of md is 2 i.e.

Maybe

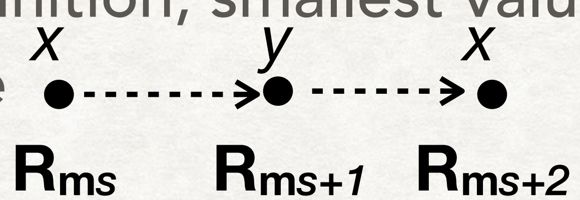


but this is not possible due to politeness!!

Thus, there is a sequence $\underline{R}^{*''} = R_{ms+1}, \ldots, R_{ms+md-1}$

Of even duration md - 2 with y repeating i.e. $R_{ms+1} \cap R_{ms+md-1} \neq \emptyset$

Contradiction.

Hence proved.

Claim 1. If AF does not terminate, $R^e$ will be non-empty

The *first smallest* sequence in the set $R^e$!

$\underline{R^*} = R_{ms}, \ldots, R_{ms+md}$

node *x* in $R_{ms}$ and $R_{ms+md}$

node y: sent M to x in $R_{ms+md}$

# AF TERMINATION

DOES (SYNCHRONOUS) AF TERMINATE?

YES

ON EVERY GRAPH?

YES

IS IT QUICK?

YES

# TERMINATION TIMES

*Revisiting the proof:*

**Theorem 1.** *Given a finite graph G, Amnesiac Flooding (AF) from a single source will terminate in a finite number of rounds*

Claim 1. If AF does not terminate, $\mathbf{R}^e$ will be non-empty

Proof. G is finite, therefore, if AF does not terminate, a node *x* must occur in infinitely many round-sets.

*Consider the first 3 such round-sets (e.g. 2,5, and 7);*

Surely at least two of these are evenly spaced.

*L1: In AF from a single source, a node can be visited at most twice!*

T2: AF from a single source terminates in at most 2n rounds where n is the number of nodes

However, we seek better bounds ....

# AF TERMINATION

HOW LONG DOES IT TAKE?

BIPARTITE GRAPHS: <= D STEPS

NON-BIPARTITE GRAPHS: <= 2D + 1 STEPS

(where D is the diameter of the graph)

Note: A hypergraph is bipartite (termination takes D)
but Petersen graph is not (termination takes 2D+1)

# AF TERMINATION (MORE PRECISELY)

## HOW LONG DOES IT TAKE?

### BIPARTITE GRAPHS:

*A graph is Bipartite if and only if termination time t of AF from origin a is* $t = e(a) \leq D$ *rounds*

(where e(a) is the eccentricity of the origin node a and D the diameter of the graph)

# (SIMPLER CASE) TERMINATION IN BIPARTITE GRAPHS

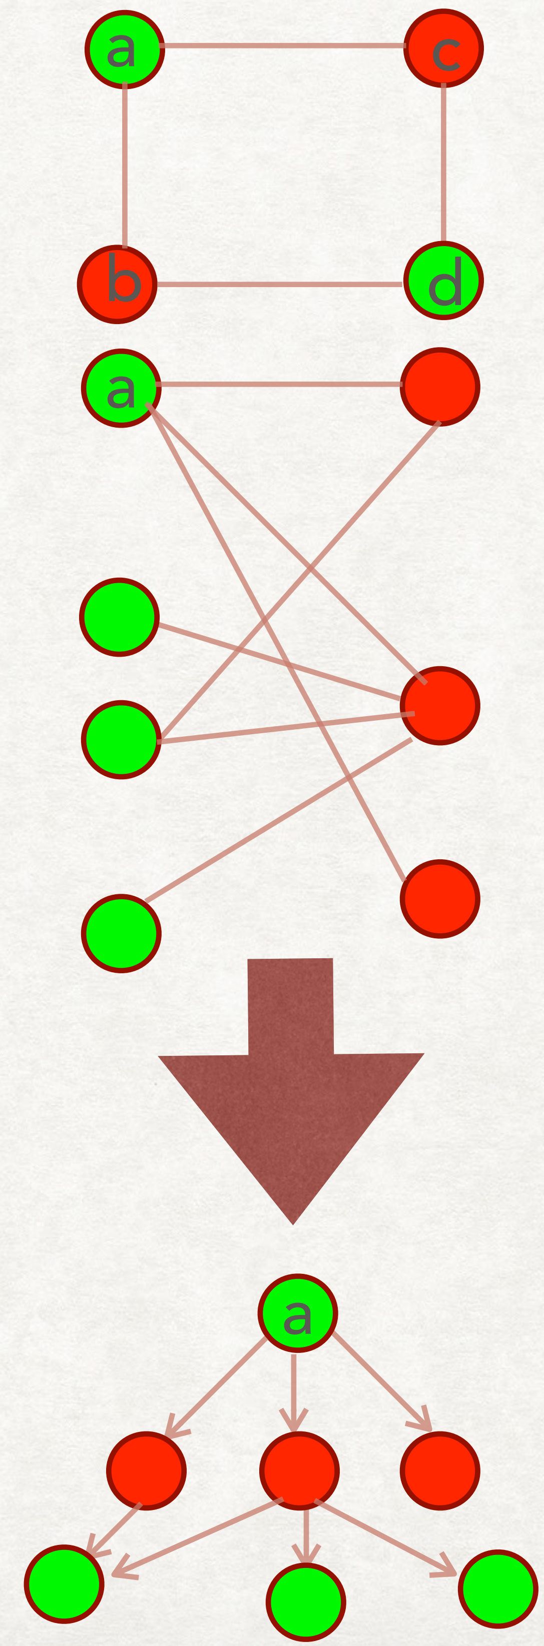


Bipartite Graph: A set of graph vertices decomposed into two disjoint sets (say, red and green) such that no two graph vertices within the same set are adjacent.

**Theorem.** In a connected bipartite graph, AF terminates in ***#rounds = e(a)***, where e(a) is the eccentricity of the vertex a, where a is the origin node.

**Proof Sketch.**

Consider the BFS traversal from the source!

There are <u>no cross edges</u>, therefore, nodes are explored at the earliest and by AF, there are no cycles.

# TERMINATION IN BIPARTITE GRAPHS

**Theorem.** In a connected bipartite graph, AF terminates in rounds = e(a), where e(a) is the eccentricity of the vertex a in graph B, where a is the origin node.

Since Diameter D is the largest possible e(a)

**Corollary.** In a connected bipartite graph, AF terminates in D rounds.

# AF TERMINATION (MORE PRECISELY)

## HOW LONG DOES IT TAKE?

### NON-BIPARTITE GRAPHS:

*In a non-bipartite graph, AF from an origin node a terminates in t rounds where*

$$e(a) < t \leq e(a) + D + 1 \leq 2D + 1$$

(where e(a) is the eccentricity of the origin node a and D the diameter of the graph)

# BIPARTITE VS. NON-BIPARTITE

- ***EC(Equidistantly-connected)* nodes** (from origin a): A node g is an ec node iff there exists another node g' such that distance(a,g) = distance(a,g') and g and g' are neighbours.

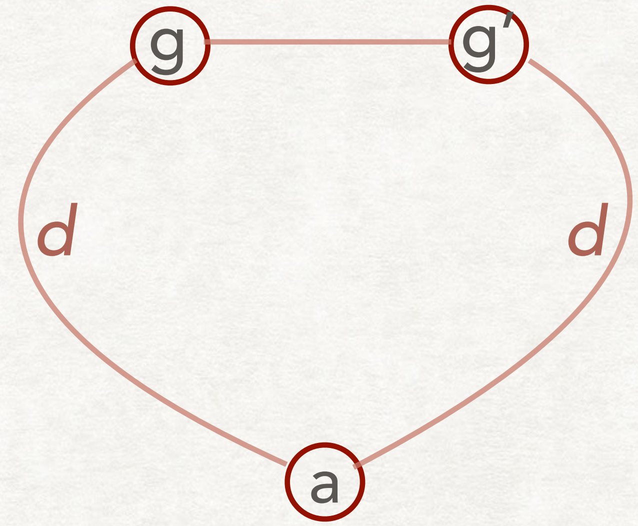


*G is bipartite iff it has no ec node*

Proof. Equidistant nodes belong to the same *partite set.*
A graph is bipartite iff no edge connects two such nodes!

Let us build on this to derive termination times for non-bipartite graphs:

# NON-BIPARTITE TERMINATION

## A VERY HIGH LEVEL VIEW....

Recall L1: A node can be visited at most twice.
Let $g^1$ be the first time node g is visited,
$g^2$ be when g is visited the second time.

Useful Technical Lemma:

*L2. For two nodes h and g in G; if $h^2 \in R_j$ and if g is a neighbour of h, then $g^2 \in R_{j-1}$ or $g^2 \in R_j$ or $g^2 \in R_{j+1}$*

# NON-BIPARTITE TERMINATION

## A VERY HIGH LEVEL VIEW….

In a non-bipartite graph, AF from an origin node a terminates in t rounds where $e(a) < t \leq e(a) + D + 1 \leq 2D + 1$

- Proof.

If G is non-bipartite, it has an ec node g.

$g^2 \in R_k$ where $k = d(a, g) + 1$ (L3: not shown here)

Let h be an arbitrary node in G other than g. There is a path:
$h_0 = g \rightarrow h_1 \rightarrow \ldots \rightarrow h_l = h$, where $l \leq d$;

Repeatedly using L2: $h_1^2 \in R_{j1}$where $k - 1 \leq j_1 \leq k + 1$; $h_2^2 \in R_{j2}$ where $j_1 - 1 \leq j_2 \leq j_1 + 1$, we get $h_l^2 \in R_{jl}$ where $j_{l-1} - 1 \leq j_l \leq j_{l-1} + 1$

i.e. $h_l^2 \in R_{jl}$ where $k - l \leq j_l \leq k + l$

Put $t = j_l$. From above and L3, it follows that $t \leq e + d + 1$

As G is non-bipartite, $t > e$. Hence proved.

*L2. For two nodes h and g in G; if $h^2 \in R_j$ and if g is a neighbour of h, then $g^2 \in R_{j-1}$ or $g^2 \in R_j$ or $g^2 \in R_{j+1}$*

# ASYNCHRONOUS FLOODING

## WHAT IF MESSAGES COULD BE DELAYED?

- Non-termination could be forced if an adversary could control delays!

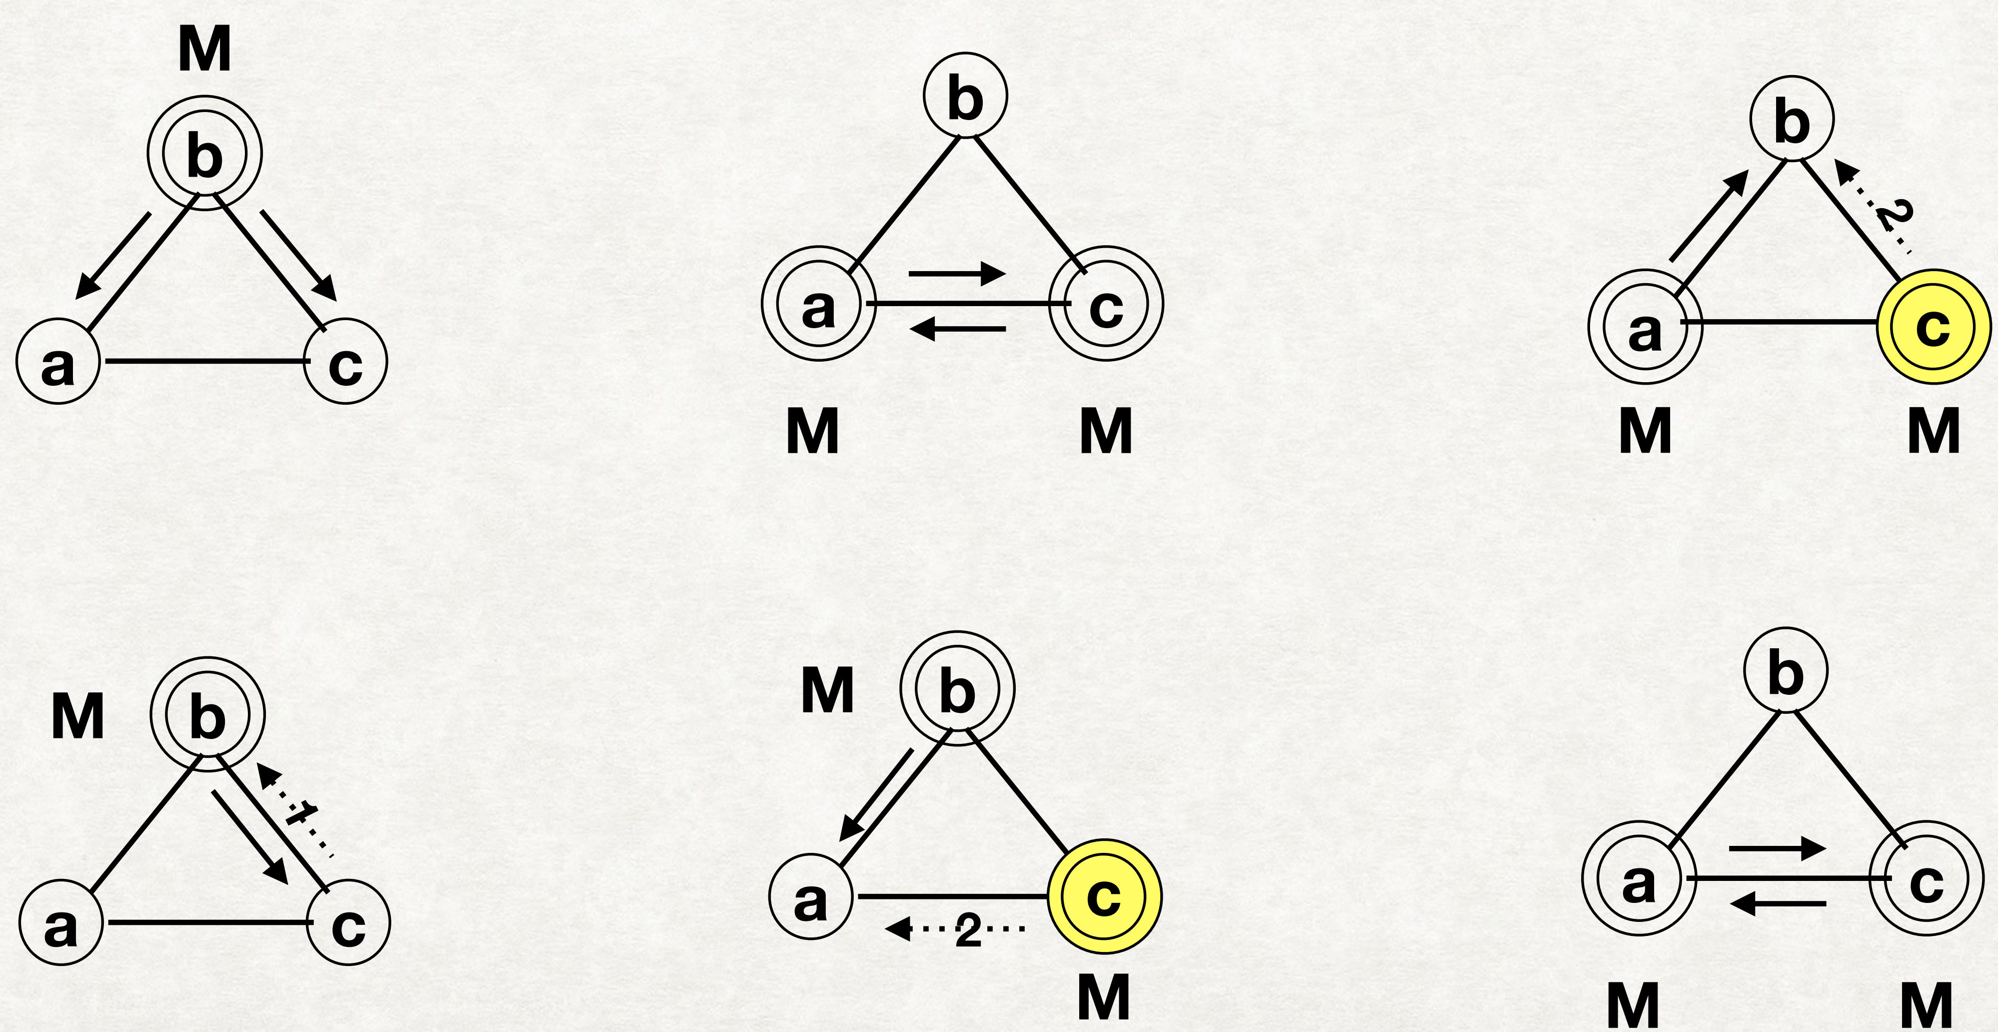

# RECENT FOLLOW UP WORK

- Multi-source termination: AF started simultaneously from multiple sources will terminate (Our journal submission).
- Alternate analysis by an auxiliary graph reduction from non-bipartite to bipartite graphs (Turau*)
- K-Amnesiac flooding problem: Given k starters, what is the placement of these starters that gives the smallest termination time (Turau*)

** Turau, Volker, Analysis of Amnesiac Flooding, Arxiv (https://arxiv.org/abs/2002.10752)*

# DISTRIBUTED AF VIS TEMPORAL GRAPH THEORY

- Temporal Graph Theory seems to be related to the Distributed *Dynamic graph model* with *T-interval connectivity* (Kuhn, Lynch, Oshman, ACM STOC, 2010).

- *Dynamic graph model:* Fixed set of nodes with changing edges with communication proceeding in synchronous rounds

- *T-interval connectivity:* For every block of T consecutive rounds, there exists a connected spanning subgraph that remains stable.

# DYNAMIC GRAPH WITH T-INTERVAL CONNECTIVITY

- *T-interval connectivity:* For every block of T consecutive rounds, there exists a connected spanning subgraph that remains stable.
  - 1-interval connectivity: Graph connected but can completely change every round
  - Infinite-interval connectivity: A permanent unchanging connected subgraph unknown to the algorithm
  - Surprisingly, nodes can still count network size and compute functions efficiently even with low stability!
- How would AF (or its variations) seem in the temporal graph/dynamic graph model?

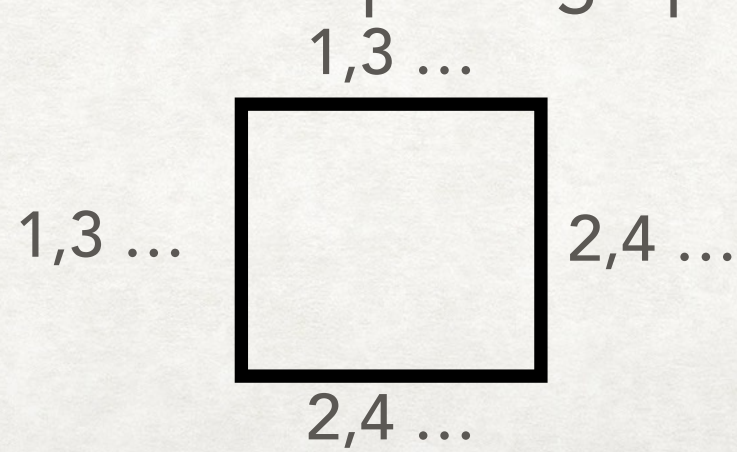

# CONCLUSIONS

- Synchronous (Amnesiac) flooding is a very lightweight communication method that achieves broadcast in almost optimal time.
- Termination times sharply differ in bipartite/non-bipartite graphs as a function of diameter suggesting possible topology tests!
- In asychronous networks, if an adaptive adversary is allowed to delay messages on edges, it can induce non-termination.
- Many open directions: More asynchronous settings, other graph parameters, randomised delays and links to random/coalescing walks/processes

# THANK YOU