Stateless Information Dissemination Algorithms

Volker Turau

27th Int. Colloquium on Structural Information and Communication Complexity

June 29th, 2020

Institute of Telematics
Hamburg University of Technology

TUHH

1

Introduction

Stateless Protocols

Definition (Stateless Protocols)

A *stateless protocol* is a communications protocol in which no session information is retained by participating nodes.

- Stateless protocols do not utilize local storage
- Big advantage in high volume applications, increasing performance by removing the load caused by retention of session information

This paper:
Stateless information dissemination algorithms for distributed systems

Deterministic Flooding

- Originator of information sends message with information to all neighbors
- Whenever a node receives message for the first time, it sends it to all its neighbors
- Flooding is a stateful algorithm
  - Each node keeps a record of which messages have already arrived
- Terminates in $\epsilon_G(v_0) + 1$ rounds ($\epsilon_G(v_0)$ is eccentricity of v_0)
- Requires storage proportional to number of disseminated messages per node
- Since termination of flooding cannot be detected by nodes, storage requirements grow over time

Amnesiac Flooding

- Variant of flooding by Hussak & Trehan [PODC19]
- Only for synchronous systems

  Every time a node receives a message, it forwards it to those neighbors from which it didn't receive message in current round

- Difference to classic flooding, a node may forward a message several times
- Amnesiac flooding is stateless

Amnesiac Flooding

Algorithm 1: Algorithm $\mathcal{A}_{AF}$ distributes a message in the graph G

input : A graph $G = (V, E)$, a subset S of V, and a message m

Round 1: Each node $v \in S$ sends message m to each neighbor;
Round $i > 1$: Each node v executes

```
    M := N(v);
    foreach receive(w, m) do
        M := M \ {w}
    if M ≠ N(v) then
        forall u ∈ M do send(u, m);
```

Amnesiac Flooding: Example

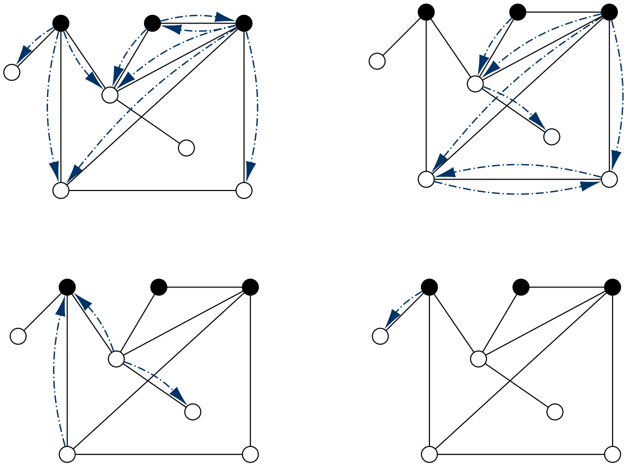

Amnesiac Flooding

- Amnesiac flooding terminates on any finite graph
- Bipartite graphs: $\epsilon_G(v_0)$ rounds
- Non-bipartite graphs: At most $\epsilon_G(v_0) + \mathit{Diam}(G) + 1$ rounds
- Bounds are sharp

- Big gap to classic flooding
- Does not work for asynchronous systems

Contributions

- Stateless information dissemination algorithms for synchronous and asynchronous systems
- Synchronous systems
  - New stateless flooding algorithm $\mathcal{A}_{SF}$ with same termination time as classic flooding
  - Also works for groups of initiators
- Asynchronous systems
  - There exists no deterministic stateless information dissemination algorithm that can only update a constant number of bits of message
  - There exists a stateless information dissemination algorithm that is allowed to update $O(\log n)$ bits of message

2

Definitions

Stateless Information Dissemination

Definition (Truly Stateless Dissemination Algorithm)

A synchronous information dissemination algorithm is called *truly stateless* if

- nodes decides only on basis of messages received in current round which messages to send in this round
- nodes are not allowed to change content of a received message before forwarding.

Stateless Information Dissemination

Definition ($f(n)$-Stateless Dissemination Algorithm)

Let f be a function from $\mathbb{N}$ to $\mathbb{N}$. An asynchronous information dissemination algorithm is called $f(n)$-*stateless* if

- nodes decide only on basis of each received message which messages to send as a reaction
- nodes are allowed to update up to $O(f(n))$ bits of a message before forwarding it.

3

Synchronous Systems

$\mathcal{A}_{SF}$: Stateless Flooding

- Task: Disseminate information stored at nodes of a set S
- First round: Each node of S sends message to all neighbors
- Second round: Nodes in S that do not receive a messages sent in round one again send message to all neighbors
- In each of following rounds – including round two – each node that receives a message forwards message to all neighbors from which it did not receive this message in this round

$\mathcal{A}_{SF}$: Stateless Flooding

Algorithm 2: Algorithm $\mathcal{A}_{SF}$ distributes a message in the graph G

input : A graph $G = (V, E)$, a subset S of V, and a message m.

Round 1: Each node $v \in S$ sends m to each neighbor;

Round 2: Each node $v \in S$ that does not receive m in round 1 sends m to each neighbor in G;

Round $i > 1$: Each node v executes

```
M := N(v);
foreach receive(w, m) do
    M := M \ {w}
if M ≠ N(v) then
    forall u ∈ M do send(u, m);
```

$\mathcal{A}_{SF}$: Examples

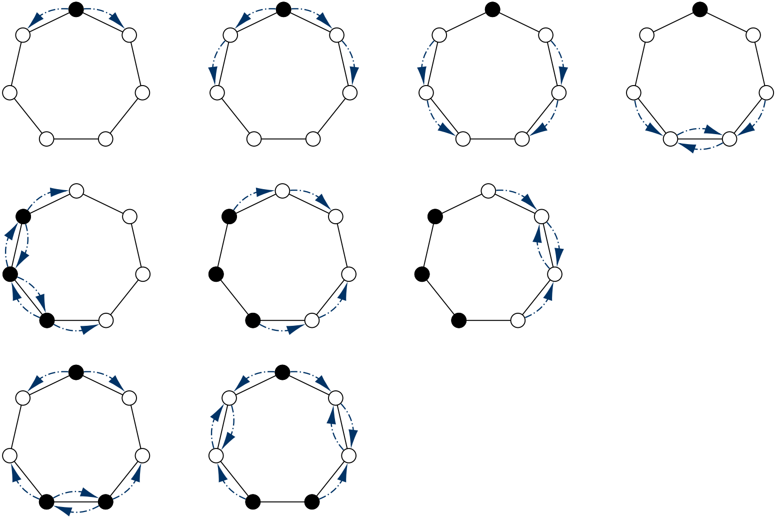

Results

Theorem

Let $G = (V, E)$ be a connected graph and $S \subseteq V$. Algorithm $\mathcal{A}_{SF}$ is truly stateless, distributes a message stored at the nodes of S to all nodes, and terminates after $d_G(S, V) + 1$ rounds.

Results

Definition

Denote by $SF_G(S)$ the number of rounds algorithm $\mathcal{A}_{SF}$ requires to terminate for graph G when started by all nodes in S. For $k \leq n$ define

$$SF_k(G) = \min\{SF_G(S) \mid S \subseteq V \text{ with } |S| = k\}.$$

Theorem

Let G be a connected graph with $n > 2$ and $k < n$. Then $SF_k(G) = r_k(G) + 1$. In particular $SF_1(G) = Rad(G) + 1$.

Results

Theorem

The time complexity of $\mathcal{A}_{SF}$ is optimal unless G is bipartite with $V = V_1 \cup V_2$ such that V_1 or V_2 contains a k-center. In this case $\mathcal{A}_{AF}$ requires one round less.

Sketch of Proof

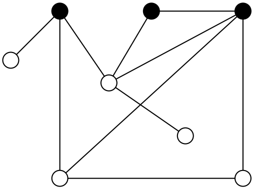

Sketch of Proof

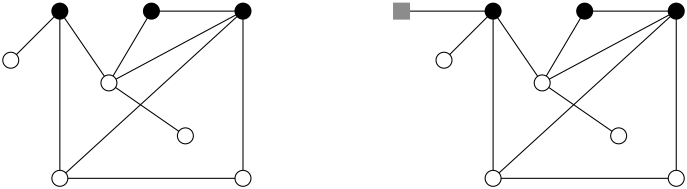

Sketch of Proof

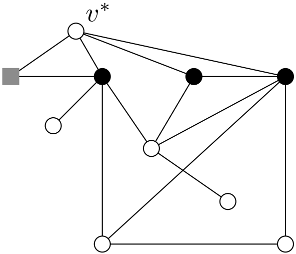

Sketch of Proof

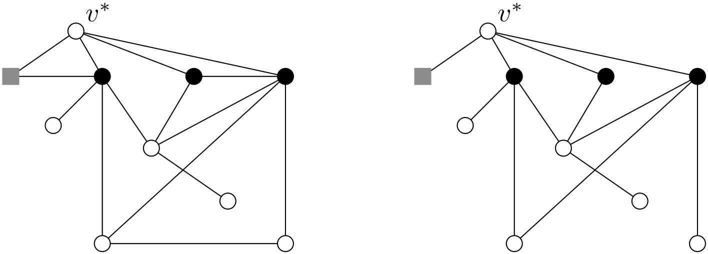

Sketch of Proof

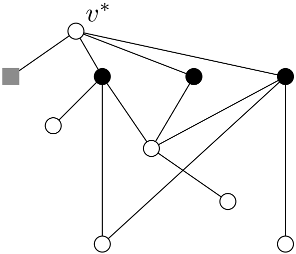

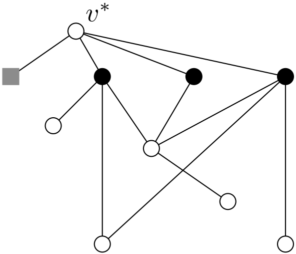

Sketch of Proof

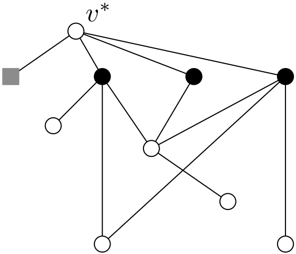

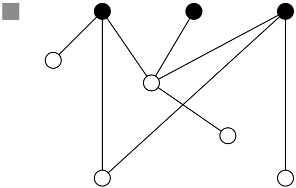

Auxiliary graph $\hat{G}$

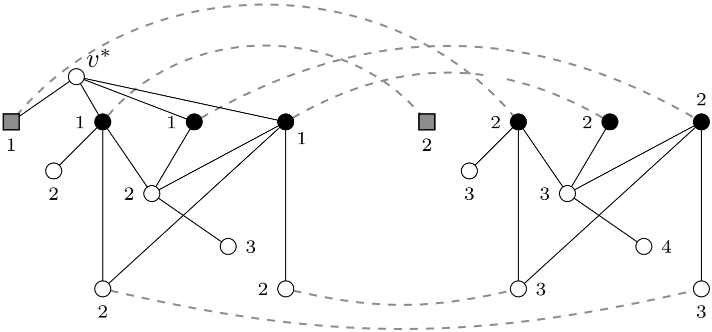

- Auxiliary graph $\hat{G}$ is bipartite
- Execution of $\mathcal{A}_{SF}$ on G with initiators S is equivalent to execution of $\mathcal{A}_{AF}$ on auxiliary graph $\hat{G}$ with initiator v^*
- $\mathcal{A}_{SF}$ terminates after $\epsilon_{\hat{G}}(v^*) = d_G(S, V) + 1$ rounds

4

Asynchronous Stateless Information Dissemination

1-stateless Information Dissemination

Theorem

There is no deterministic 1-stateless information dissemination algorithm for asynchronous systems.

Sketch of Proof.

Let $\mathcal{A}$ be a 1-stateless information dissemination algorithm that can update up to d bits in each message

- Let G be a graph that has a node v_0 with $\epsilon_G(v_0) > 2^d$
- Consider execution of $\mathcal{A}$ with initiator v_0
- There exists a message flow $\mathcal{S}$: $v_0 \xrightarrow{m_0} v_1 \xrightarrow{m_1} v_2 \xrightarrow{m_2} \ldots$ with nodes $v_0, v_1, \ldots$ and $v_i \in N(v_{i+1})$ such that v_i sends a message to v_{i+1} as a reaction of receiving a message from v_{i-1}

1-stateless Information Dissemination

Proof ctd.

- Length of $\mathcal{S}$ is greater than 2^d
- Thus, there are two nodes v_s and v_t in this flow with $s < t$ which receive identical messages
- Hence, as a reaction they also send identical messages
- Thus, $\mathcal{S}$ is infinite. This yields that $\mathcal{A}$ does not terminate
- Contradiction

□

log n-stateless Information Dissemination

Theorem

There exists a $\log n$*-stateless information dissemination algorithm for asynchronous systems terminating in* n^{c+1} *rounds provided each node has a unique identifier in the range* $0, \ldots, n^c$ *with* $c \geq 1$.

Sketch of Proof.

- Each message consists of two values each of size $O(\log n)$
- Originator v_0 sends pair $(v_0.id, v_0.id)$ to all neighbors
- If v receives a message (a, b):
  - If $v.id > a$ then v sends $(v.id, v.id)$ to all neighbors except the one from which message came
  - If $v.id < a$ and $b \neq 0$ then v sends $(a, b - 1)$ to all neighbors except the one from which message came

log n-stateless Information Dissemination

Proof ctd.

- Assume information does not reach all nodes
- Among uninformed nodes choose v such that $d(v, S)$ is minimal
- Let $w \in N(v)$ such that $d(w, S) < d(v, S)$
- Then w is informed with a message $(a, 0)$ and $a > w.id$ otherwise v would be informed
- Consider a shortest path from the node u with $u.id = a$ to w
- Since u sent messages (a, a) the second component was a-times decreased by nodes with an id less than a
- Thus, there must be $a + 1$ nodes with an id less than a
- Contradiction □

5

Conclusion

Conclusion & Outlook

- Optimal truly stateless information dissemination algorithm with k initiators for synchronous systems
- Algorithm terminates in $r_k(G) + 1$ rounds
- Unless $P = NP$ there is no approximation algorithm for the SF-problem with an approximation ratio less than $3/2$
- Open problems:
  - Design a $3/2$-approximation or disprove its existence
  - Number of messages of proposed $\log n$-stateless information dissemination algorithm grows exponentially with n: Design more efficient algorithm
  - (Dis)Prove: There exits a deterministic $f(n)$-stateless asynchronous information dissemination algorithm with $f \in o(\log n)$

Stateless Information Dissemination Algorithms

Volker Turau

27th Int. Colloquium on Struc … lexity

Volker Turau

Professor

Phone +49 / (0)40 428 78 3530

e-Mail turau@tuhh.de

http://www.ti5.tu-harburg.de/staff/turau