

Deterministic Oblivious Local Broadcast in the SINR Model

Tomasz Jurdziński, Michał Róžański


Institute of Computer Science
University of Wrocław, Poland

FCT 2017
Bordeaux, France

Local Broadcast

Definition


- A graph (network) is the input to the problem.
- Each node (station) has a message that has to be delivered to all the neighbors in the graph.



Local Broadcast

Definition


- A graph (network) is the input to the problem.
- Each node (station) has a message that has to be delivered to all the neighbors in the graph.



Local Broadcast

Definition


- A graph (network) is the input to the problem.
- Each node (station) has a message that has to be delivered to all the neighbors in the graph.



Local Broadcast

Definition


- A graph (network) is the input to the problem.
- Each node (station) has a message that has to be delivered to all the neighbors in the graph.



Local Broadcast

Definition


- A graph (network) is the input to the problem.
- Each node (station) has a message that has to be delivered to all the neighbors in the graph.



Local Broadcast

Definition


- A graph (network) is the input to the problem.
- Each node (station) has a message that has to be delivered to all the neighbors in the graph.



Local Broadcast

Definition


- A graph (network) is the input to the problem.
- Each node (station) has a message that has to be delivered to all the neighbors in the graph.



Local Broadcast

Definition

- A graph (network) is the input to the problem.
- Each node (station) has a message that has to be delivered to all the neighbors in the graph.



Local Broadcast

Definition

- A graph (network) is the input to the problem.
- Each node (station) has a message that has to be delivered to all the neighbors in the graph.

Environment

- Nodes of the network are located on a plane.

Local Broadcast

Definition

- A graph (network) is the input to the problem.
- Each node (station) has a message that has to be delivered to all the neighbors in the graph.

Environment

- Nodes of the network are located on a plane.
- Algorithm is executed in a distributed manner by all the nodes.

Local Broadcast

Definition

- A graph (network) is the input to the problem.
- Each node (station) has a message that has to be delivered to all the neighbors in the graph.

Environment

- Nodes of the network are located on a plane.
- Algorithm is executed in a distributed manner by all the nodes.
- All stations start the execution of the algorithm at the same time.

Local Broadcast

Definition

- A graph (network) is the input to the problem.
- Each node (station) has a message that has to be delivered to all the neighbors in the graph.

Environment

- Nodes of the network are located on a plane.
- Algorithm is executed in a distributed manner by all the nodes.
- All stations start the execution of the algorithm at the same time.
- Messages are exchanged in rounds, according to the SINR constraints.

Local Broadcast

Definition

- A graph (network) is the input to the problem.
- Each node (station) has a message that has to be delivered to all the neighbors in the graph.

Environment

- Nodes of the network are located on a plane.
- Algorithm is executed in a distributed manner by all the nodes.
- All stations start the execution of the algorithm at the same time.
- Messages are exchanged in rounds, according to the SINR constraints.
- Each station has a unique identifier from the set $\{1, \dots, N\}$.

Local Broadcast


Definition

- A graph (network) is the input to the problem.
- Each node (station) has a message that has to be delivered to all the neighbors in the graph.

Environment

- Nodes of the network are located on a plane.
- Algorithm is executed in a distributed manner by all the nodes.
- All stations start the execution of the algorithm at the same time.
- Messages are exchanged in rounds, according to the SINR constraints.
- Each station has a unique identifier from the set $\{1, \dots, N\}$.
- The algorithm is oblivious – all actions of a node are predetermined by its identifier.

The model: signal power



- sender: v
- receiver: u

Power of signal from v received by u

$$P(v, u) = \frac{P_v}{d(u, v)^\alpha}$$

- $\alpha > 2$ – path loss exponent (environment dependent)
- P_v – transmission power of station v ,
we assume that every station has the same power P

Signal to Interference and Noise Ratio (SINR)

Interference

$$I_T(u) = \sum_{w \in \mathcal{T}} P(w, u) \quad (1)$$

Signal to Interference and Noise Ratio (SINR)

Interference

$$I_{\mathcal{T}}(u) = \sum_{w \in \mathcal{T}} P(w, u) \quad (1)$$

SINR

$$SINR(v, u, \mathcal{T}) = \frac{P(v, u)}{\mathcal{N} + I_{\mathcal{T} \setminus \{v\}}(u)} \quad (2)$$

\mathcal{T} – set of transmitters

Model: signal reception


SINR, signal reception

$$\text{SINR}(v, u, \mathcal{T}) = \frac{P(v, u)}{\mathcal{N} + I_{\mathcal{T} \setminus \{v\}}(u)} \geq \beta \quad (3)$$

Parameters

- Path Loss Exponent $\alpha > 2$
- Background Noise $\mathcal{N} > 0$
- Threshold $\beta \geq 1$

SINR example



Node u receives a message from v iff

$$\text{SINR}(v, u, \{w_1, w_2\}) = \frac{P(v, u)}{\mathcal{N} + P(w_1, u) + P(w_2, u)} \geq \beta.$$

SINR Diagrams




Figure: Reception zones

- Blue dots - transmitting stations
- Yellow area - region of audibility
- Green area - too noisy to decode anything

The model: network topology

Communication graph

$$G = (V, E)$$

V - the set of stations

r - transmission range

$$(u, v) \in E \Leftrightarrow d(u, v) \leq (1 - \varepsilon)r$$

ε - sensitivity parameter

The model: network topology

Communication graph

$$G = (V, E)$$

V - the set of stations

r - transmission range

$$(u, v) \in E \Leftrightarrow d(u, v) \leq (1 - \varepsilon)r$$

ε - sensitivity parameter

Without loss of generality we assume that $r = 1$.

The model: network topology

Communication graph

$$G = (V, E)$$

V - the set of stations

r - transmission range

$$(u, v) \in E \Leftrightarrow d(u, v) \leq (1 - \varepsilon)r$$

ε - sensitivity parameter

Without loss of generality we assume that $r = 1$.

Complexity parameters

- Δ - density of the network = maximal number of nodes per unit disc
- n - number of stations

Local broadcast

Previous results

- Randomized algorithm
 $O(\Delta \log N)$ [Goussevskaia, Moscibroda, Wattenhofer '08]
- Randomized algorithm with the feedback mechanism
 $O(\Delta + \log N \log \log N)$ [Barenboim, Peleg '15]
- Randomized algorithm without the knowledge of Δ
 $O(\Delta \log N + \log^2 N)$ [Yu, Hua, Wang, Lao '12]

Our results (deterministic)

- Oblivious algorithm – $O(\Delta^{2+2/(\alpha-2)} \log N)$
- Semi-adaptive algorithm with the feedback mechanism – $O(\Delta \log N)$

Oblivious algorithms

Oblivious algorithm

An oblivious algorithm \mathcal{A}_Δ solving a local broadcast in m rounds in networks of density Δ can be represented by a sequence (R_1, \dots, R_m) , where $R_i \subseteq [N]$ is the set of nodes transmitting in round i .

First approach

Observation 1

Let v be any node. There exists a constant I_{max} such that, if maximal interference in $B(v, 1 - \varepsilon)$ is less than I_{max} then, if v transmits, all neighbors of v hear its message.

First approach

Observation 1

Let v be any node. There exists a constant I_{max} such that, if maximal interference in $B(v, 1 - \epsilon)$ is less than I_{max} then, if v transmits, all neighbors of v hear its message.

In other words, if interference in any point of $B(v, 1 - \epsilon)$ is at most I_{max} , then for any neighbor u of v we have:

$$SINR(v, u) = \frac{P(v, u)}{\mathcal{N} + I(v)} \geq \beta,$$

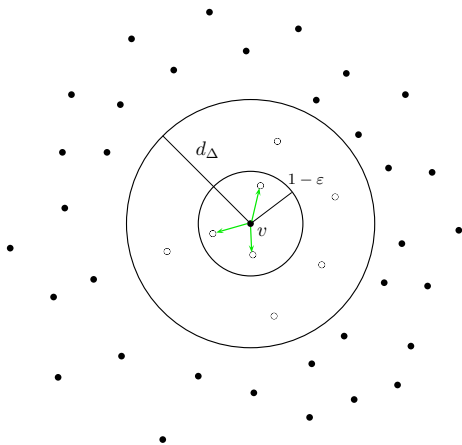
where

$$I(v) = \sum_{\substack{w\text{-transmitting} \\ w \neq v}} P(w, u) \leq I_{max}.$$

First approach

Observation 2

For all networks of density Δ there exists $d_\Delta \in \mathbb{R}$ such that if no node transmits in $B(v, d_\Delta)$ (apart from v), then all neighbors of v hear its message.



First approach

Observation 2

For all networks of density Δ there exists $d_\Delta \in \mathbb{R}$ such that if no node transmits in $B(v, d_\Delta)$ (apart from v), then all neighbors of v hear its message.

- There are $O(i\Delta)$ nodes within distance $[i, i + 1]$ from v .

First approach

Observation 2

For all networks of density Δ there exists $d_\Delta \in \mathbb{R}$ such that if no node transmits in $B(v, d_\Delta)$ (apart from v), then all neighbors of v hear its message.

- There are $O(i\Delta)$ nodes within distance $[i, i + 1]$ from v .
- The interference from them, assuming that all are transmitting, is at most $O(i\Delta \cdot \frac{P}{i^\alpha})$.

First approach

Observation 2

For all networks of density Δ there exists $d_\Delta \in \mathbb{R}$ such that if no node transmits in $B(v, d_\Delta)$ (apart from v), then all neighbors of v hear its message.

- There are $O(i\Delta)$ nodes within distance $[i, i + 1]$ from v .
- The interference from them, assuming that all are transmitting, is at most $O(i\Delta \cdot \frac{P}{i^\alpha})$.
- The total interference is $\sum_{i \geq d_\Delta} O(i\Delta \cdot \frac{P}{i^\alpha}) = O(\Delta d_\Delta^{2-\alpha})$.

First approach

Observation 2

For all networks of density Δ there exists $d_\Delta \in \mathbb{R}$ such that if no node transmits in $B(v, d_\Delta)$ (apart from v), then all neighbors of v hear its message.

- There are $O(i\Delta)$ nodes within distance $[i, i + 1]$ from v .
- The interference from them, assuming that all are transmitting, is at most $O(i\Delta \cdot \frac{P}{i^\alpha})$.
- The total interference is $\sum_{i \geq d_\Delta} O(i\Delta \cdot \frac{P}{i^\alpha}) = O(\Delta d_\Delta^{2-\alpha})$.
- Which is less than I_{max} for $d_\Delta = \Theta(\Delta^{\frac{1}{\alpha-2}})$

Observation 3

If every node v transmits in a round, when all other nodes in $B(v, d_\Delta)$ are silent, then local broadcast is accomplished.

First approach

Strongly selective family

(N, k) -ssf is a sequence of sets (S_1, \dots, S_m) , where $S_i \subseteq \{1, \dots, N\}$ such that for any set $X \subseteq [N]$ of size k , for each element $x \in X$ there is $i \leq m$ such that $X \cap S_i = \{x\}$.

There exists a (N, k) -ssf of size $O(k^2 \log N)$.

First approach


Strongly selective family

(N, k) -ssf is a sequence of sets (S_1, \dots, S_m) , where $S_i \subseteq \{1, \dots, N\}$ such that for any set $X \subseteq [N]$ of size k , for each element $x \in X$ there is $i \leq m$ such that $X \cap S_i = \{x\}$.


There exists a (N, k) -ssf of size $O(k^2 \log N)$.

- $B(v, d_\Delta)$ contains at most $\kappa = O(\Delta \cdot d_\Delta^2)$ nodes.
- An oblivious algorithm made from (N, κ) -ssf accomplishes local broadcast in $\kappa^2 \log N = O(\Delta^{2+4/(\alpha-2)} \log N)$ rounds.

Second approach



(a) Naive approach – all nodes in the gray area different from v stay silent.



(b) Balanced approach – All nodes in A but v stay silent, number of nodes allowed to transmit in B_i increases with i .

Balanced strongly selective family

Balanced strongly selective family

(N, k, \vec{b}) -bssf is a sequence of sets (S_1, \dots, S_m) , where $S_i \subseteq \{1, \dots, N\}$ such that for any set $X \subseteq [N]$ of size k , and $B_1, \dots, B_l \subseteq [N]$ of fixed size ($|B_j| = b_j$), for each element $x \in X$ there is $i \leq m$ such that:

- a) $S_i \cap X = \{x\}$,
- b) $|S_i \cap B_j| \leq |B_j|/k$ for all j .

Balanced strongly selective family

Balanced strongly selective family

(N, k, \vec{b}) -bssf is a sequence of sets (S_1, \dots, S_m) , where $S_i \subseteq \{1, \dots, N\}$ such that for any set $X \subseteq [N]$ of size k , and $B_1, \dots, B_l \subseteq [N]$ of fixed size ($|B_j| = b_j$), for each element $x \in X$ there is $i \leq m$ such that:

- a) $S_i \cap X = \{x\}$,
- b) $|S_i \cap B_j| \leq |B_j|/k$ for all j .

- There exists (N, k, \vec{b}) -bssf of size $O(k(k + s) \log N)$, where $s = b_1 + \dots + b_l$.

Balanced strongly selective family

Balanced strongly selective family

(N, k, \vec{b}) -bssf is a sequence of sets (S_1, \dots, S_m) , where $S_i \subseteq \{1, \dots, N\}$ such that for any set $X \subseteq [N]$ of size k , and $B_1, \dots, B_l \subseteq [N]$ of fixed size ($|B_j| = b_j$), for each element $x \in X$ there is $i \leq m$ such that:

- a) $S_i \cap X = \{x\}$,
- b) $|S_i \cap B_j| \leq |B_j|/k$ for all j .

- There exists (N, k, \vec{b}) -bssf of size $O(k(k + s) \log N)$, where $s = b_1 + \dots + b_l$.
- Application of bssf to oblivious local broadcast yields an algorithm of complexity $O(\Delta^{2+2/(\alpha-2)} \log N)$, which improves by a factor of $\Delta^{2/(\alpha-2)}$ over the naive approach.

Semi-adaptive algorithm

Semi-adaptive algorithms

A semi-adaptive algorithm is an oblivious algorithm, in which a node can quit the execution at some point.

Semi-adaptive algorithm

Semi-adaptive algorithms

A semi-adaptive algorithm is an oblivious algorithm, in which a node can quit the execution at some point.

Feedback mechanism

In a round, when a node transmits a message to all its neighbors, it receives a feedback informing about successful transmission.

Fractional balanced selector

Fractional balanced selector

(N, Δ) -fbs is a sequence of sets (S_1, \dots, S_m) , where $S_i \subseteq \{1, \dots, N\}$ such that for any set $X \subseteq [N]$ of size k , and $B_1, \dots, B_l \subseteq [N]$ of fixed size ($|B_i| = b_i$), for **at least half** elements $x \in X$ there is $i \leq m$ such that:

- a) $S_i \cap A = \{x\}$,
- b) $|S_i \cap B_j| \leq \delta^j |B_j| / \Delta$ for all j .

Fractional balanced selector

Fractional balanced selector

(N, Δ) -fbs is a sequence of sets (S_1, \dots, S_m) , where $S_i \subseteq \{1, \dots, N\}$ such that for any set $X \subseteq [N]$ of size k , and $B_1, \dots, B_l \subseteq [N]$ of fixed size ($|B_i| = b_i$), for **at least half** elements $x \in X$ there is $i \leq m$ such that:

- a) $S_i \cap A = \{x\}$,
- b) $|S_i \cap B_j| \leq \delta^j |B_j| / \Delta$ for all j .

- There exists (N, Δ) -fbs of size $O(\Delta \log N)$.

Fractional balanced selector

Fractional balanced selector

(N, Δ) -fbs is a sequence of sets (S_1, \dots, S_m) , where $S_i \subseteq \{1, \dots, N\}$ such that for any set $X \subseteq [N]$ of size k , and $B_1, \dots, B_l \subseteq [N]$ of fixed size ($|B_i| = b_i$), for **at least half** elements $x \in X$ there is $i \leq m$ such that:

- a) $S_i \cap A = \{x\}$,
- b) $|S_i \cap B_j| \leq \delta^j |B_j| / \Delta$ for all j .

- There exists (N, Δ) -fbs of size $O(\Delta \log N)$.
- After execution of (N, Δ) -fbs oblivious algorithm there are at most $\Delta/2$ nodes in any unit disc that did not send message to all its neighbors.

Fractional balanced selector

Semi-adaptive algorithm

Let \mathcal{A} be an algorithm formed by connecting $(N, \Delta/2^i)$ -fractional balanced selectors, for $i = 0, \dots, \log \Delta$, and the naive algorithm for constant density at the end.

If node receives a feedback at some point it quits the execution of the algorithm.

Fractional balanced selector

Semi-adaptive algorithm

Let \mathcal{A} be an algorithm formed by connecting $(N, \Delta/2^i)$ -fractional balanced selectors, for $i = 0, \dots, \log \Delta$, and the naive algorithm for constant density at the end.

If node receives a feedback at some point it quits the execution of the algorithm.

The complexity of the above algorithm is
 $O(\Delta \log N + \Delta/2 \log N + \dots + \log N) = O(\Delta \log N)$.

Thank you!