



# Breaking the Linear-Memory Barrier in Massively Parallel Computing

MIS on Trees with Strongly Sublinear Memory

Sebastian Brandt, **Manuela Fischer**, Jara Uitto  
ETH Zurich

Model:

**Massively Parallel Computing (MPC)**

Model:

# Massively Parallel Computing (MPC)

parallel computing framework

Model:

# Massively Parallel Computing (MPC)

parallel computing framework  
inspired by MapReduce


Model:

# Massively Parallel Computing (MPC)

parallel computing framework  
inspired by MapReduce


*Karloff, Suri, Vassilvitskii [SODA'10]*

# Massively Parallel Computing (MPC) Model



**$M$  machines**  
 **$S$  memory per machine**

# Massively Parallel Computing (MPC) Model




$M$  machines

$S$  memory per machine

Synchronous Rounds

# Massively Parallel Computing (MPC) Model



$M$  machines


$S$  memory per machine

**Synchronous Rounds**

**1. Local Computation**  
at every machine



# Massively Parallel Computing (MPC) Model



$M$  machines

$S$  memory per machine

**Synchronous Rounds**


**1. Local Computation**

at every machine

**2. Global Communication**

between machines

# Massively Parallel Computing (MPC) Model



$M$  machines


$S$  memory per machine

**Synchronous Rounds**

**1. Local Computation**  
at every machine

**2. Global Communication**  
between machines

# Massively Parallel Computing (MPC) Model



$M$  machines

$S$  memory per machine

**Synchronous Rounds**


**1. Local Computation**

at every machine

**2. Global Communication**

between machines

# Massively Parallel Computing (MPC) Model



$M$  machines

$S$  memory per machine

**Synchronous Rounds**

**1. Local Computation**

at every machine


**2. Global Communication**

between machines

**Complexity:**

number of rounds

# Massively Parallel Computing (MPC) Model



$M$  machines

$S$  memory per machine

**Synchronous Rounds**


**1. Local Computation**  
at every machine

**2. Global Communication**  
between machines

**Complexity:**

number of rounds

# Massively Parallel Computing (MPC) Model



$M$  machines

$S$  memory per machine

**Synchronous Rounds**

**1. Local Computation**

at every machine


**2. Global Communication**

between machines

**Complexity:**

number of rounds

# Massively Parallel Computing (MPC) Model



$M$  machines

$S$  memory per machine

$$M \cdot S = \tilde{O}(m + n)$$

**Synchronous Rounds**

**1. Local Computation**

at every machine

**2. Global Communication**

between machines

**Complexity:**

number of rounds

# Local Memory in MPC

$M$  machines

$S$  memory per machine

$$M \cdot S = \tilde{O}(m + n)$$



# Local Memory in MPC

$M$  machines

$S$  memory per machine

$$M \cdot S = \tilde{O}(m + n)$$

$S$



# Local Memory in MPC

$M$  machines

$S$  memory per machine

$$M \cdot S = \tilde{O}(m + n)$$

$S$

$$\tilde{O}(n^\delta)$$

$$\tilde{\Theta}(n)$$

$$\tilde{\Omega}(n^{1+\delta})$$

# Local Memory in MPC

$M$  machines

$S$  memory per machine

$$M \cdot S = \tilde{O}(m + n)$$

$S$

$$\tilde{O}(n^\delta)$$

$$\tilde{\Theta}(n)$$

$$\tilde{\Omega}(n^{1+\delta})$$

**Superlinear Memory:**

$$S = \tilde{O}(n^{1+\delta}), 0 < \delta \leq 1$$


**Machines see all nodes.**

# Local Memory in MPC

$M$  machines

$S$  memory per machine

$$M \cdot S = \tilde{O}(m + n)$$



**Superlinear Memory:**

$$S = \tilde{O}(n^{1+\delta}), 0 < \delta \leq 1$$

**Machines see all nodes.**


often trivial

# Local Memory in MPC

$M$  machines

$S$  memory per machine

$$M \cdot S = \tilde{O}(m + n)$$



**Superlinear Memory:**

$$S = \tilde{O}(n^{1+\delta}), 0 < \delta \leq 1$$

**Machines see all nodes.**

often trivial

for many problems,  
admits  $O(1)$ -round algorithms  
based on very simple  
sampling approach


*Lattanzi et al. [SPAA'11]*

# Local Memory in MPC

$M$  machines

$S$  memory per machine

$$M \cdot S = \tilde{O}(m + n)$$



**Linear Memory:**

$$S = \tilde{O}(n)$$

**Machines see all nodes.**

**Superlinear Memory:**

$$S = \tilde{O}(n^{1+\delta}), 0 < \delta \leq 1$$

**Machines see all nodes.**

often trivial

for many problems,  
admits  $O(1)$ -round algorithms  
based on very simple  
sampling approach

*Lattanzi et al. [SPAA'11]*

# Local Memory in MPC

$M$  machines

$S$  memory per machine

$$M \cdot S = \tilde{O}(m + n)$$

$S$

$$\tilde{O}(n^\delta)$$

$$\tilde{\Theta}(n)$$

$$\tilde{\Omega}(n^{1+\delta})$$

**Linear Memory:**

$$S = \tilde{O}(n)$$

**Machines see all nodes.**

usual assumption

**Superlinear Memory:**

$$S = \tilde{O}(n^{1+\delta}), 0 < \delta \leq 1$$

**Machines see all nodes.**

often trivial

for many problems,  
admits  $O(1)$ -round algorithms  
based on very simple  
sampling approach


*Lattanzi et al. [SPAA'11]*

# Local Memory in MPC

$M$  machines

$S$  memory per machine

$$M \cdot S = \tilde{O}(m + n)$$



**Linear Memory:**

$$S = \tilde{O}(n)$$

**Machines see all nodes.**

usual assumption

often unrealistic

**Superlinear Memory:**

$$S = \tilde{O}(n^{1+\delta}), 0 < \delta \leq 1$$

**Machines see all nodes.**

often trivial

for many problems,  
admits  $O(1)$ -round algorithms  
based on very simple  
sampling approach

*Lattanzi et al. [SPAA'11]*




# Local Memory in MPC

$M$  machines

$S$  memory per machine

$$M \cdot S = \tilde{O}(m + n)$$



**Linear Memory:**

$$S = \tilde{O}(n)$$

**Machines see all nodes.**

usual assumption

often unrealistic

- $\tilde{O}(n)$  prohibitively large

**Superlinear Memory:**

$$S = \tilde{O}(n^{1+\delta}), 0 < \delta \leq 1$$

**Machines see all nodes.**

often trivial

for many problems,  
admits  $O(1)$ -round algorithms  
based on very simple  
sampling approach


*Lattanzi et al. [SPAA'11]*

# Local Memory in MPC

$M$  machines

$S$  memory per machine

$$M \cdot S = \tilde{O}(m + n)$$



**Linear Memory:**

$$S = \tilde{O}(n)$$

**Machines see all nodes.**

usual assumption

often unrealistic

- $\tilde{O}(n)$  prohibitively large
- sparse graphs trivial

**Superlinear Memory:**

$$S = \tilde{O}(n^{1+\delta}), 0 < \delta \leq 1$$

**Machines see all nodes.**

often trivial

for many problems,  
admits  $O(1)$ -round algorithms  
based on very simple  
sampling approach


*Lattanzi et al. [SPAA'11]*

# Local Memory in MPC

$M$  machines

$S$  memory per machine

$$M \cdot S = \tilde{O}(m + n)$$



## Strongly Sublinear Memory:

$$S = \tilde{O}(n^\delta), 0 \leq \delta < 1$$

**No machine sees all nodes.**

## Linear Memory:

$$S = \tilde{O}(n)$$

**Machines see all nodes.**

## Superlinear Memory:

$$S = \tilde{O}(n^{1+\delta}), 0 < \delta \leq 1$$

**Machines see all nodes.**

usual assumption

often unrealistic

- $\tilde{O}(n)$  prohibitively large
- sparse graphs trivial

often trivial

for many problems,  
admits  $O(1)$ -round algorithms  
based on very simple  
sampling approach


*Lattanzi et al. [SPAA'11]*

# Local Memory in MPC

$M$  machines

$S$  memory per machine

$$M \cdot S = \tilde{O}(m + n)$$



## Strongly Sublinear Memory:

$$S = \tilde{O}(n^\delta), 0 \leq \delta < 1$$

**No machine sees all nodes.**

for most problems,  
only direct simulation of  
LOCAL/PRAM algorithms  
known

## Linear Memory:

$$S = \tilde{O}(n)$$

**Machines see all nodes.**

usual assumption

often unrealistic

- $\tilde{O}(n)$  prohibitively large
- sparse graphs trivial

## Superlinear Memory:

$$S = \tilde{O}(n^{1+\delta}), 0 < \delta \leq 1$$

**Machines see all nodes.**

often trivial

for many problems,  
admits  $O(1)$ -round algorithms  
based on very simple  
sampling approach


*Lattanzi et al. [SPAA'11]*

# Local Memory in MPC

$M$  machines

$S$  memory per machine

$$M \cdot S = \tilde{O}(m + n)$$



## Strongly Sublinear Memory:

$$S = \tilde{O}(n^\delta), 0 \leq \delta < 1$$

**No machine sees all nodes.**

for most problems,  
only direct simulation of  
LOCAL/PRAM algorithms  
known

## Linear Memory:

$$S = \tilde{O}(n)$$

**Machines see all nodes.**

usual assumption

often unrealistic

- $\tilde{O}(n)$  prohibitively large
- sparse graphs trivial

## Superlinear Memory:

$$S = \tilde{O}(n^{1+\delta}), 0 < \delta \leq 1$$

**Machines see all nodes.**

often trivial

for many problems,  
admits  $O(1)$ -round algorithms  
based on very simple  
sampling approach


*Lattanzi et al. [SPAA'11]*

# Local Memory in MPC

$M$  machines

$S$  memory per machine

$$M \cdot S = \tilde{O}(m + n)$$



## Strongly Sublinear Memory:

$$S = \tilde{O}(n^\delta), 0 \leq \delta < 1$$

No machine sees all nodes.

for most problems,  
only direct simulation of  
LOCAL/PRAM algorithms  
known

**Algorithms have been stuck at this linear-memory barrier!**

## Linear Memory:

$$S = \tilde{O}(n)$$

Machines see all nodes.

usual assumption

often unrealistic

- $\tilde{O}(n)$  prohibitively large
- sparse graphs trivial

## Superlinear Memory:

$$S = \tilde{O}(n^{1+\delta}), 0 < \delta \leq 1$$

Machines see all nodes.

often trivial

for many problems,  
admits  $O(1)$ -round algorithms  
based on very simple  
sampling approach


*Lattanzi et al. [SPAA'11]*

# Local Memory in MPC

$M$  machines

$S$  memory per machine

$$M \cdot S = \tilde{O}(m + n)$$



## Strongly Sublinear Memory:

$$S = \tilde{O}(n^\delta), 0 \leq \delta < 1$$

No machine sees all nodes.

for most problems,  
only direct simulation of  
LOCAL/PRAM algorithms  
known

**Algorithms have been stuck at this linear-memory barrier!**  
**Fundamentally?**

## Linear Memory:

$$S = \tilde{O}(n)$$

Machines see all nodes.

usual assumption

often unrealistic

- $\tilde{O}(n)$  prohibitively large
- sparse graphs trivial

## Superlinear Memory:

$$S = \tilde{O}(n^{1+\delta}), 0 < \delta \leq 1$$

Machines see all nodes.

often trivial

for many problems,  
admits  $O(1)$ -round algorithms  
based on very simple  
sampling approach  
*Lattanzi et al. [SPAA'11]*

Breaking the Linear-Memory Barrier:



Breaking the Linear-Memory Barrier:

# **Efficient MPC Graph Algorithms with Strongly Sublinear Memory**

Breaking the Linear-Memory Barrier:

## Efficient MPC Graph Algorithms with Strongly Sublinear Memory

$S = O(n^\delta)$  local memory

$M = O(m/n^\delta)$  machines

poly  $\log \log n$  rounds

Breaking the Linear-Memory Barrier:

## Efficient MPC Graph Algorithms with Strongly Sublinear Memory

$S = O(n^\delta)$  local memory

$M = O(m/n^\delta)$  machines

poly  $\log \log n$  rounds

*Ghaffari, Kuhn, Uitto* [FOCS'19]

**Conditional Lower Bound**

$\Omega(\log \log n)$  rounds

Breaking the Linear-Memory Barrier:

## Efficient MPC Graph Algorithms with Strongly Sublinear Memory

$S = O(n^\delta)$  local memory

$M = O(m/n^\delta)$  machines

poly  $\log \log n$  rounds

Breaking the Linear-Memory Barrier:

## Efficient MPC Graph Algorithms with Strongly Sublinear Memory

$S = O(n^\delta)$  local memory

$M = O(m/n^\delta)$  machines

poly log log  $n$  rounds

### Imposed Locality:

machines see only subset of nodes,  
regardless of sparsity of graph

Breaking the Linear-Memory Barrier:

## Efficient MPC Graph Algorithms with Strongly Sublinear Memory


$S = O(n^\delta)$  local memory

$M = O(m/n^\delta)$  machines

poly  $\log \log n$  rounds

### Imposed Locality:

machines see only subset of nodes,  
regardless of sparsity of graph



Breaking the Linear-Memory Barrier:

## Efficient MPC Graph Algorithms with Strongly Sublinear Memory


$S = O(n^\delta)$  local memory

$M = O(m/n^\delta)$  machines

poly  $\log \log n$  rounds

### Imposed Locality:

machines see only subset of nodes,  
regardless of sparsity of graph



Breaking the Linear-Memory Barrier:

## Efficient MPC Graph Algorithms with Strongly Sublinear Memory


$S = O(n^\delta)$  local memory

$M = O(m/n^\delta)$  machines

poly  $\log \log n$  rounds

### Imposed Locality:

machines see only subset of nodes,  
regardless of sparsity of graph





Breaking the Linear-Memory Barrier:

## Efficient MPC Graph Algorithms with Strongly Sublinear Memory

$S = O(n^\delta)$  local memory


$M = O(m/n^\delta)$  machines

poly  $\log \log n$  rounds

### Imposed Locality:

machines see only subset of nodes,  
regardless of sparsity of graph

### Our Approach to Cope with Locality:



Breaking the Linear-Memory Barrier:

## Efficient MPC Graph Algorithms with Strongly Sublinear Memory

$S = O(n^\delta)$  local memory

$M = O(m/n^\delta)$  machines


poly  $\log \log n$  rounds

### Imposed Locality:

machines see only subset of nodes,  
regardless of sparsity of graph

### Our Approach to Cope with Locality:

enhance **LOCAL algorithms** with **global communication**



Breaking the Linear-Memory Barrier:

## Efficient MPC Graph Algorithms with Strongly Sublinear Memory

$S = O(n^\delta)$  local memory

$M = O(m/n^\delta)$  machines


poly  $\log \log n$  rounds

### Imposed Locality:

machines see only subset of nodes,  
regardless of sparsity of graph

### Our Approach to Cope with Locality:

enhance **LOCAL algorithms** with **global communication**



Breaking the Linear-Memory Barrier:

## Efficient MPC Graph Algorithms with Strongly Sublinear Memory

$S = O(n^\delta)$  local memory

$M = O(m/n^\delta)$  machines


poly  $\log \log n$  rounds

### Imposed Locality:

machines see only subset of nodes,  
regardless of sparsity of graph

### Our Approach to Cope with Locality:

enhance **LOCAL algorithms** with **global communication**



Breaking the Linear-Memory Barrier:

## Efficient MPC Graph Algorithms with Strongly Sublinear Memory

$S = O(n^\delta)$  local memory

$M = O(m/n^\delta)$  machines

poly log log  $n$  rounds


### Imposed Locality:

machines see only subset of nodes,  
regardless of sparsity of graph

### Our Approach to Cope with Locality:

enhance **LOCAL algorithms** with **global communication**

- exponentially faster than LOCAL algorithms due to shortcuts



Breaking the Linear-Memory Barrier:

## Efficient MPC Graph Algorithms with Strongly Sublinear Memory

$S = O(n^\delta)$  local memory

$M = O(m/n^\delta)$  machines

poly log log  $n$  rounds


### Imposed Locality:

machines see only subset of nodes,  
regardless of sparsity of graph

### Our Approach to Cope with Locality:

enhance **LOCAL algorithms** with **global communication**

- exponentially faster than LOCAL algorithms due to shortcuts



best we can hope for  
*GKU* [FOCS'19]

Breaking the Linear-Memory Barrier:

## Efficient MPC Graph Algorithms with Strongly Sublinear Memory

$S = O(n^\delta)$  local memory

$M = O(m/n^\delta)$  machines

poly log log  $n$  rounds


### Imposed Locality:

machines see only subset of nodes,  
regardless of sparsity of graph

### Our Approach to Cope with Locality:

enhance **LOCAL algorithms** with **global communication**

- exponentially faster than LOCAL algorithms due to shortcuts



Breaking the Linear-Memory Barrier:

## Efficient MPC Graph Algorithms with Strongly Sublinear Memory

$S = O(n^\delta)$  local memory

$M = O(m/n^\delta)$  machines

poly  $\log \log n$  rounds


### Imposed Locality:

machines see only subset of nodes,  
regardless of sparsity of graph

### Our Approach to Cope with Locality:

enhance **LOCAL algorithms** with **global communication**

- exponentially faster than LOCAL algorithms due to shortcuts
- polynomially less memory than most MPC algorithms






Problem:

# **Maximal Independent Set (MIS)**


# Maximal Independent Set (MIS)

# Maximal Independent Set (MIS)






# Maximal Independent Set (MIS)



**Independent Set:**  
set of non-adjacent nodes

# Maximal Independent Set (MIS)



**Independent Set:**  
set of non-adjacent nodes


**Maximal:**  
no node can be added  
without violating independence

# MIS: State of the Art

$M$  machines

$S$  memory per machine

$$M \cdot S = \tilde{O}(m + n)$$



**Strongly Sublinear Memory:**

$$S = \tilde{O}(n^\delta), 0 \leq \delta < 1$$

**No machine sees all nodes.**

**Linear Memory:**

$$S = \tilde{O}(n)$$

**Machines see all nodes.**

**Superlinear Memory:**

$$S = \tilde{O}(n^{1+\delta}), 0 < \delta \leq 1$$


**Machines see all nodes.**

# MIS: State of the Art

$M$  machines

$S$  memory per machine

$$M \cdot S = \tilde{O}(m + n)$$



**Strongly Sublinear Memory:**

$$S = \tilde{O}(n^\delta), 0 \leq \delta < 1$$

**No machine sees all nodes.**

**Linear Memory:**

$$S = \tilde{O}(n)$$

**Machines see all nodes.**

**Superlinear Memory:**

$$S = \tilde{O}(n^{1+\delta}), 0 < \delta \leq 1$$

**Machines see all nodes.**

*Lattanzi et al. [SPAA'11]*

$$O(1)$$




# MIS: State of the Art

$M$  machines

$S$  memory per machine

$$M \cdot S = \tilde{O}(m + n)$$



**Strongly Sublinear Memory:**

$$S = \tilde{O}(n^\delta), 0 \leq \delta < 1$$

**No machine sees all nodes.**

**Linear Memory:**

$$S = \tilde{O}(n)$$

**Machines see all nodes.**

**Superlinear Memory:**

$$S = \tilde{O}(n^{1+\delta}), 0 < \delta \leq 1$$

**Machines see all nodes.**

*Ghaffari et al. [PODC'18]*

$$O(\log \log n)$$

*Lattanzi et al. [SPAA'11]*


$$O(1)$$

# MIS: State of the Art

$M$  machines

$S$  memory per machine

$$M \cdot S = \tilde{O}(m + n)$$



**Strongly Sublinear Memory:**

$$S = \tilde{O}(n^\delta), 0 \leq \delta < 1$$

**No machine sees all nodes.**

*Luby's Algorithm*  
 $O(\log n)$

**Linear Memory:**

$$S = \tilde{O}(n)$$

**Machines see all nodes.**

*Ghaffari et al. [PODC'18]*  
 $O(\log \log n)$

**Superlinear Memory:**

$$S = \tilde{O}(n^{1+\delta}), 0 < \delta \leq 1$$

**Machines see all nodes.**


*Lattanzi et al. [SPAA'11]*  
 $O(1)$

# MIS: State of the Art

$M$  machines

$S$  memory per machine

$$M \cdot S = \tilde{O}(m + n)$$



**Strongly Sublinear Memory:**

$$S = \tilde{O}(n^\delta), 0 \leq \delta < 1$$

**No machine sees all nodes.**

*Luby's Algorithm*

$$O(\log n)$$

*Ghaffari and Uitto [SODA'19]*

$$\tilde{O}(\sqrt{\log n})$$

**Linear Memory:**

$$S = \tilde{O}(n)$$

**Machines see all nodes.**

*Ghaffari et al. [PODC'18]*

$$O(\log \log n)$$

**Superlinear Memory:**

$$S = \tilde{O}(n^{1+\delta}), 0 < \delta \leq 1$$

**Machines see all nodes.**

*Lattanzi et al. [SPAA'11]*


$$O(1)$$

# MIS: State of the Art **on Trees**

$M$  machines

$S$  memory per machine

$$M \cdot S = \tilde{O}(m + n)$$



**Strongly Sublinear Memory:**

$$S = \tilde{O}(n^\delta), 0 \leq \delta < 1$$

**No machine sees all nodes.**

*Luby's Algorithm*

$$O(\log n)$$

*Ghaffari and Uitto [SODA'19]*

$$\tilde{O}(\sqrt{\log n})$$

**Linear Memory:**

$$S = \tilde{O}(n)$$

**Machines see all nodes.**

*Ghaffari et al. [PODC'18]*

$$O(\log \log n)$$

**Superlinear Memory:**

$$S = \tilde{O}(n^{1+\delta}), 0 < \delta \leq 1$$

**Machines see all nodes.**

*Lattanzi et al. [SPAA'11]*


$$O(1)$$

# MIS: State of the Art **on Trees**

$M$  machines

$S$  memory per machine

$$M \cdot S = \tilde{O}(m + n)$$



**Strongly Sublinear Memory:**

$$S = \tilde{O}(n^\delta), 0 \leq \delta < 1$$

**No machine sees all nodes.**

**Linear Memory:**

$$S = \tilde{O}(n)$$

**Machines see all nodes.**

**Superlinear Memory:**

$$S = \tilde{O}(n^{1+\delta}), 0 < \delta \leq 1$$

**Machines see all nodes.**

*Luby's Algorithm*

$$O(\log n)$$

*Trivial solution*

$$O(1)$$

*Trivial solution*

$$O(1)$$

*Ghaffari and Uitto [SODA'19]*


$$\tilde{O}(\sqrt{\log n})$$

# MIS: State of the Art **on Trees**

$M$  machines

$S$  memory per machine

$$M \cdot S = \tilde{O}(m + n)$$



**Strongly Sublinear Memory:**

$$S = \tilde{O}(n^\delta), 0 \leq \delta < 1$$

**No machine sees all nodes.**

**Linear Memory:**

$$S = \tilde{O}(n)$$

**Machines see all nodes.**

**Superlinear Memory:**

$$S = \tilde{O}(n^{1+\delta}), 0 < \delta \leq 1$$

**Machines see all nodes.**

*Our Result*

$$O(\log^3 \log n)$$

*Trivial solution*

$$O(1)$$

*Trivial solution*

$$O(1)$$

*Ghaffari and Uitto [SODA'19]*

$$\tilde{O}(\sqrt{\log n})$$

# Our Result

**$O(\log^3 \log n)$** -round MPC algorithm

with  **$S = \tilde{O}(n^\delta)$**  memory that w.h.p. computes MIS on trees.

# Our Result

$\tilde{O}(\sqrt{\log n})$  rounds

$\mathbf{S} = \tilde{O}(n^\delta)$  memory

*Ghaffari and Uitto [SODA'19]*

$O(\log^3 \log n)$ -round MPC algorithm

with  $\mathbf{S} = \tilde{O}(n^\delta)$  memory that w.h.p. computes MIS on trees.



# Our Result

$\tilde{O}(\sqrt{\log n})$  rounds

$\mathbf{S} = \tilde{O}(n^\delta)$  memory

*Ghaffari and Uitto [SODA'19]*

$O(\log \log n)$  rounds

$\mathbf{S} = \tilde{O}(n)$  memory

*Ghaffari et al. [PODC'18]*

$O(\log^3 \log n)$ -round MPC algorithm

with  $\mathbf{S} = \tilde{O}(n^\delta)$  memory that w.h.p. computes MIS on trees.

# Our Result

$\tilde{O}(\sqrt{\log n})$  rounds

$\mathbf{S} = \tilde{O}(n^\delta)$  memory

*Ghaffari and Uitto* [SODA'19]

$O(\log \log n)$  rounds

$\mathbf{S} = \tilde{O}(n)$  memory

*Ghaffari et al.* [PODC'18]

$O(\log^3 \log n)$ -round MPC algorithm

with  $\mathbf{S} = \tilde{O}(n^\delta)$  memory that w.h.p. computes MIS on trees.


Conditional  $\Omega(\log \log n)$ -round lower bound for  $\mathbf{S} = \tilde{O}(n^\delta)$

*Ghaffari, Kuhn, and Uitto* [FOCS'19]

# Algorithm


# Algorithm Outline

# Algorithm Outline



# Algorithm Outline


## 1) Shattering



# Algorithm Outline

## 1) Shattering

main LOCAL technique  
*Beck* [RSA'91]




# Algorithm Outline

## 1) Shattering

break graph into small components

main LOCAL technique

*Beck* [RSA'91]






# Algorithm Outline

## 1) Shattering

break graph into small components

main LOCAL technique

*Beck* [RSA'91]




# Algorithm Outline

## 1) Shattering

break graph into small components

main LOCAL technique

*Beck* [RSA'91]



# Algorithm Outline

## 1) Shattering


break graph into small components

main LOCAL technique

*Beck* [RSA'91]

## 2) Post-Shattering

solve problem on remaining components



# Algorithm Outline

## 1) Shattering

break graph into small components


main LOCAL technique

*Beck [RSA'91]*

## 2) Post-Shattering

solve problem on remaining components

### i) Gathering of Components



# Algorithm Outline

## 1) Shattering

break graph into small components



main LOCAL technique

*Beck [RSA'91]*

## 2) Post-Shattering

solve problem on remaining components

### i) Gathering of Components



# Algorithm Outline

## 1) Shattering

break graph into small components



main LOCAL technique

*Beck [RSA'91]*

## 2) Post-Shattering

solve problem on remaining components

### i) Gathering of Components



# Algorithm Outline

## 1) Shattering

break graph into small components

main LOCAL technique



*Beck [RSA'91]*

## 2) Post-Shattering

solve problem on remaining components

i) **Gathering of Components**

ii) **Local Computation**



# Algorithm Outline

## 1) Shattering

break graph into small components

main LOCAL technique


*Beck [RSA'91]*

## 2) Post-Shattering

solve problem on remaining components

i) **Gathering of Components**

ii) **Local Computation**





# Algorithm Outline

## 1) Shattering


break graph into small components

## 2) Post-Shattering

solve problem on remaining components

i) Gathering of Components

ii) Local Computation



# Algorithm Outline

## 1) Shattering

break graph into small components


ii) **LOCAL Shattering** *Ghaffari* [SODA'16]

## 2) Post-Shattering

solve problem on remaining components

i) **Gathering of Components**

ii) **Local Computation**



# Algorithm Outline

## 1) Shattering

break graph into small components

i) Degree Reduction


ii) LOCAL Shattering *Ghaffari* [SODA'16]

## 2) Post-Shattering

solve problem on remaining components

i) Gathering of Components

ii) Local Computation



Polynomial Degree Reduction:  
**Subsample-and-Conquer**

# Polynomial Degree Reduction: **Subsample-and-Conquer**


**Subsample**

**Conquer**

# Polynomial Degree Reduction: **Subsample-and-Conquer**

**Subsample**

**Conquer**




# Polynomial Degree Reduction: **Subsample-and-Conquer**

## **Subsample**

subsample **nodes** independently

## **Conquer**




# Polynomial Degree Reduction: **Subsample-and-Conquer**

## **Subsample**

subsample **nodes** independently

## **Conquer**






# Polynomial Degree Reduction: **Subsample-and-Conquer**

## **Subsample**

subsample **nodes** independently

## **Conquer**




# Polynomial Degree Reduction: **Subsample-and-Conquer**

## **Subsample**

subsample **nodes** independently

## **Conquer**

compute **random MIS** in **subsampled graph**




# Polynomial Degree Reduction: **Subsample-and-Conquer**

## **Subsample**

subsample **nodes** independently

## **Conquer**

compute **random MIS** in **subsampled graph**




# Polynomial Degree Reduction: **Subsample-and-Conquer**

## **Subsample**

subsample **nodes** independently

## **Conquer**

compute **random MIS** in **subsampled graph**




# Polynomial Degree Reduction: **Subsample-and-Conquer**

## **Subsample**

subsample **nodes** independently

## **Conquer**

compute **random MIS** in **subsampled graph**



# Polynomial Degree Reduction: **Subsample-and-Conquer**


## **Subsample**

subsample **nodes** independently

## **Conquer**

compute **random MIS** in **subsampled graph**

- gather connected components



# Polynomial Degree Reduction: **Subsample-and-Conquer**


## **Subsample**

subsample **nodes** independently

## **Conquer**

compute **random MIS** in **subsampled graph**

- gather connected components



# Polynomial Degree Reduction: **Subsample-and-Conquer**



## Subsample

subsample **nodes** independently

## Conquer

compute **random MIS** in **subsampled graph**

- gather connected components
- locally compute random 2-coloring  





# Polynomial Degree Reduction: **Subsample-and-Conquer**



## Subsample

subsample **nodes** independently

## Conquer

compute **random MIS** in **subsampled graph**

- gather connected components
- locally compute random 2-coloring



# Polynomial Degree Reduction: **Subsample-and-Conquer**

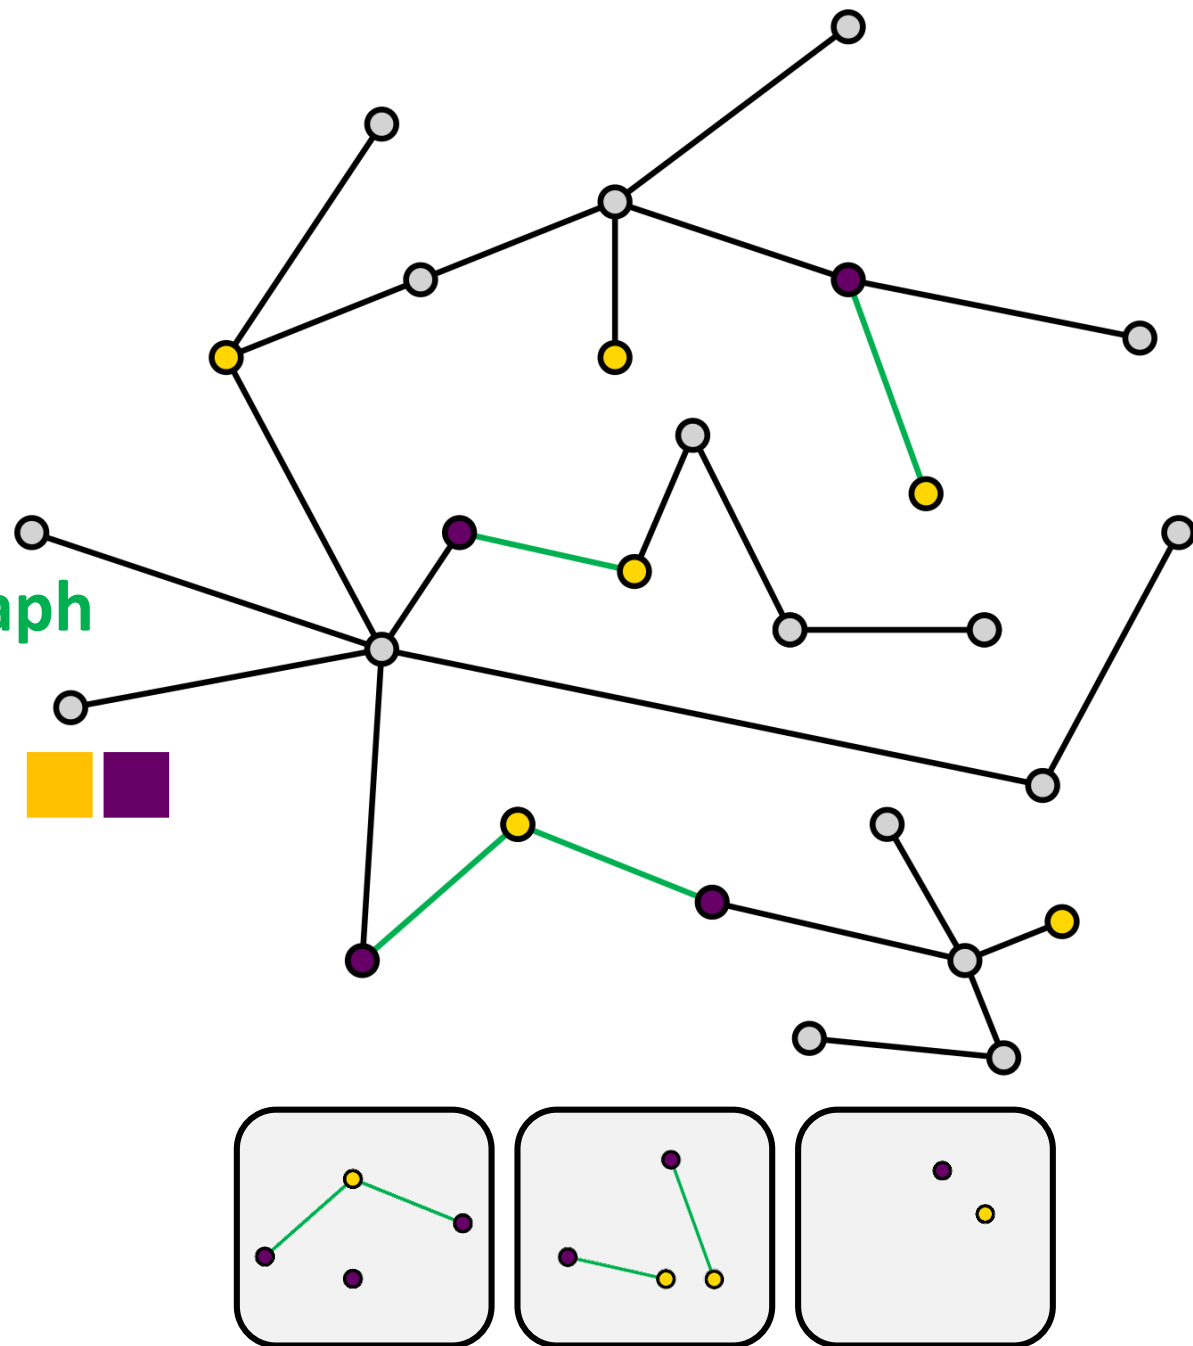
## Subsample

subsample **nodes** independently

## Conquer

compute **random MIS** in **subsampled graph**

- gather connected components
- locally compute random 2-coloring



# Polynomial Degree Reduction: **Subsample-and-Conquer**


## **Subsample**

subsample **nodes** independently

## **Conquer**

compute **random MIS** in **subsampled graph**

- gather connected components
- locally compute random 2-coloring






# Polynomial Degree Reduction: **Subsample-and-Conquer**


## Subsample

subsample **nodes** independently

## Conquer

compute **random MIS** in **subsampled graph**

- gather connected components
- locally compute random 2-coloring  
- add a color class to **MIS** 



# Polynomial Degree Reduction: **Subsample-and-Conquer**


## Subsample

subsample **nodes** independently

## Conquer

compute **random MIS** in **subsampled graph**

- gather connected components
- locally compute random 2-coloring
- add a color class to **MIS**






# Polynomial Degree Reduction: **Subsample-and-Conquer**

## **Subsample**

subsample **nodes** independently

## **Conquer**

compute **random MIS** in **subsampled graph**

- gather connected components
- locally compute random 2-coloring  
- add a color class to **MIS** 




# Polynomial Degree Reduction: **Subsample-and-Conquer**

## **Subsample**

subsample **nodes** independently

## **Conquer**

compute **random MIS** in **subsampled graph**

- gather connected components
- locally compute random 2-coloring  
- add a color class to **MIS** 

Non-sampled **High-Degree Node**

# Polynomial Degree Reduction: **Subsample-and-Conquer**

## **Subsample**


subsample **nodes** independently

## **Conquer**

compute **random MIS** in **subsampled graph**

- gather connected components
- locally compute random 2-coloring
- add a color class to **MIS**

Non-subsampled **High-Degree Node**





# Polynomial Degree Reduction: **Subsample-and-Conquer**

## **Subsample**

subsample **nodes** independently


## **Conquer**

compute **random MIS** in **subsampled graph**

- gather connected components
- locally compute random 2-coloring ■ ■
- add a color class to **MIS** ■

## Non-sampled **High-Degree Node**

- w.h.p. has many **subsampled neighbors**



# Polynomial Degree Reduction: **Subsample-and-Conquer**

## **Subsample**

subsample **nodes** independently


## **Conquer**

compute **random MIS** in **subsampled graph**

- gather connected components
- locally compute random 2-coloring ■ ■
- add a color class to **MIS** ■

## Non-sampled **High-Degree Node**

- w.h.p. has many **subsampled neighbors**



**independence due to restriction to trees!**

# Polynomial Degree Reduction: **Subsample-and-Conquer**

## Subsample

subsample **nodes** independently


## Conquer

compute **random MIS** in **subsampled graph**

- gather connected components
- locally compute random 2-coloring
- add a color class to **MIS**

## Non-sampled **High-Degree Node**

- w.h.p. has many **subsampled neighbors**



**independence due to restriction to trees!**



# Polynomial Degree Reduction: **Subsample-and-Conquer**

## **Subsample**

subsample **nodes** independently


## **Conquer**

compute **random MIS** in **subsampled graph**

- gather connected components
- locally compute random 2-coloring
- add a color class to **MIS**

## Non-subsampled **High-Degree Node**

- w.h.p. has many **subsampled neighbors**
- thus w.h.p. has at least one **MIS neighbor**



# Polynomial Degree Reduction: **Subsample-and-Conquer**

## **Subsample**

subsample **nodes** independently


## **Conquer**

compute **random MIS** in **subsampled graph**

- gather connected components
- locally compute random 2-coloring
- add a color class to **MIS**

## Non-sampled **High-Degree Node**

- w.h.p. has many **subsampled neighbors**
- thus w.h.p. has at least one **MIS neighbor**
- hence will be **removed** from the graph



# Algorithm Outline

## 1) Shattering

break graph into small components

i) **Degree Reduction** *Iterated Subsample-and-Conquer*

ii) **LOCAL Shattering** *Ghaffari [SODA'16]*

## 2) Post-Shattering

solve problem on remaining components

i) **Gathering of Components** *Distributed Union-Find*

ii) **Local Computation**

# Conclusion and Open Questions



MODEL:

**Sublinear-Memory MPC**




$S = \tilde{O}(n^\delta)$  local memory

poly  $\log \log n$  rounds

MODEL:

**Sublinear-Memory MPC**




$S = \tilde{O}(n^\delta)$  local memory  
poly  $\log \log n$  rounds



MODEL:

**Sublinear-Memory MPC**

$S = \tilde{O}(n^\delta)$  local memory  
poly  $\log \log n$  rounds




APPROACH:

LOCAL algorithms &  
global communication

MODEL:

**Sublinear-Memory MPC**

$S = \tilde{O}(n^\delta)$  local memory  
poly log log  $n$  rounds



APPROACH:

LOCAL algorithms &  
global communication




TECHNIQUE:

Shattering

MODEL:

**Sublinear-Memory MPC**

$S = \tilde{O}(n^\delta)$  local memory  
poly log log  $n$  rounds



APPROACH:

LOCAL algorithms &  
global communication

TECHNIQUE:

Shattering


PROBLEM:

MIS  
on trees

MODEL:

**Sublinear-Memory MPC**

$S = \tilde{O}(n^\delta)$  local memory  
poly log log  $n$  rounds



APPROACH:

LOCAL algorithms &  
global communication

TECHNIQUE:

Shattering

PROBLEM:


MIS  
on trees

other graph problems?  
more general graph families?

MODEL:

**Sublinear-Memory MPC**

$S = \tilde{O}(n^\delta)$  local memory  
poly log log  $n$  rounds



APPROACH:

LOCAL algorithms &  
global communication

TECHNIQUE:

Shattering

PROBLEM:

MIS  
on trees


other graph problems?  
more general graph families?

MIS & Matching for locally sparse graphs  
in follow-up work [PODC'19]

MODEL:

**Sublinear-Memory MPC**

$S = \tilde{O}(n^\delta)$  local memory  
poly log log  $n$  rounds



APPROACH:

LOCAL algorithms &  
global communication

TECHNIQUE:

Shattering

other LOCAL techniques?

PROBLEM:

MIS  
on trees

other graph problems?  
more general graph families?




MIS & Matching for locally sparse graphs  
in follow-up work [PODC'19]



MODEL:

**Sublinear-Memory MPC**

$S = \tilde{O}(n^\delta)$  local memory  
poly log log  $n$  rounds



APPROACH:

LOCAL algorithms &  
global communication

other approaches?

TECHNIQUE:

Shattering

other LOCAL techniques?

PROBLEM:

MIS  
on trees

other graph problems?  
more general graph families?




MIS & Matching for locally sparse graphs  
in follow-up work [PODC'19]

Thank you!

MODEL:

**Sublinear-Memory MPC**

$S = \tilde{O}(n^\delta)$  local memory  
poly log log  $n$  rounds



APPROACH:

LOCAL algorithms &  
global communication

other approaches?

TECHNIQUE:

Shattering

other LOCAL techniques?

PROBLEM:

MIS  
on trees

other graph problems?  
more general graph families?

MIS & Matching for locally sparse graphs  
in follow-up work [PODC'19]