

Brief Announcement: Improved Massively Parallel Triangle Counting in $O(1)$ Rounds

Quanquan C. Liu

Yale UniversityTM

quanquan.liu@yale.edu



C. Seshadhri

UC SANTA CRUZ


sesh@ucsc.edu

MPC Model Definition

- M machines
- Synchronous rounds




MPC Model Definition

- M machines
- Synchronous rounds



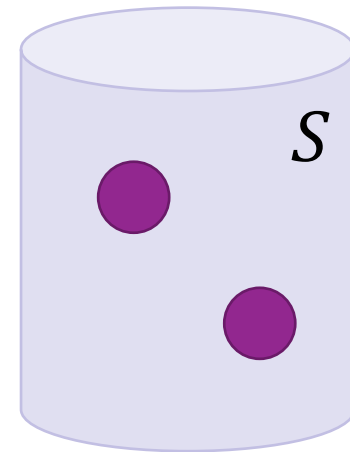


MPC Model Definition

- M machines
- Synchronous rounds





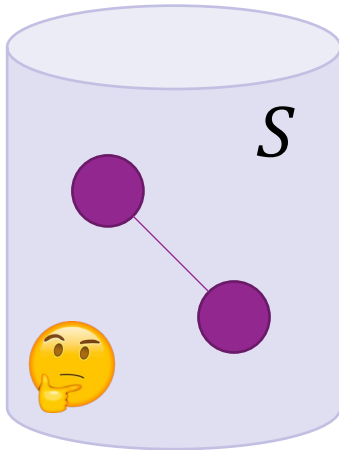
MPC Model Definition

- M machines
- Synchronous rounds




MPC Model Definition

- M machines
- Synchronous rounds



MPC Model Definition


- M machines
- Synchronous rounds



MPC Model Definition

- M machines
- Synchronous rounds

Total Space: $M \cdot S$




MPC Model Definition

- M machines
- Synchronous rounds

Complexity measures:

- Total Space
- Space Per Machine
- Rounds of communication

Total Space: $M \cdot S$



Space per Machine in MPC

- **Strongly sublinear memory:**
 - $S = n^\delta$ for some constant $\delta \in (0, 1)$
- **Near-linear memory:**
 - $S = \tilde{\Theta}(n)$ (ignoring $\text{poly}(\log(n))$ factors)
- **Strongly superlinear memory:**
 - $S = n^{1+\delta}$ for some constant $\delta > 0$

Also want: $O(1)$ rounds

Space per Machine in MPC

- **Strongly sublinear memory:**
 - $S = n^\delta$ for some constant $\delta \in (0, 1)$
- **Near-linear memory:**
 - $S = \tilde{\Theta}(n)$ (ignoring $\text{poly}(\log(n))$ factors)
- **Strongly superlinear memory:**
 - $S = n^{1+\delta}$ for some constant $\delta > 0$

Also want: $O(1)$ rounds

**Also want: $\tilde{O}(n + m)$
total space**

Space per Machine in MPC

- **Strongly sublinear memory:**

- $S = n^\delta$ for some constant $\delta \in (0, 1)$

Also want: $O(1)$ rounds

- **Near-linear memory:**

- $S = \tilde{\Theta}(n)$ (ignoring $\text{poly}(\log(n))$ factors)

Also want: $\tilde{O}(n + m)$
total space

- **Strongly superlinear memory:**

- $S = n^{1+\delta}$ for some constant $\delta > 0$

All space per machine are **sublinear in number of edges m** in graph

Space per Machine in MPC

- **Strongly sublinear memory:**

- $S = n^\delta$ for some constant $\delta \in (0, 1)$

Also want: $O(1)$ rounds

- **Near-linear memory:**

- $S = \tilde{\Theta}(n)$ (ignoring $\text{poly}(\log(n))$ factors)

Also want: $\tilde{O}(n + m)$

- **Strongly superlinear memory:**


- $S = n^{1+\delta}$ for some

total space


$$\text{Often } \delta = \frac{1}{2}$$

All space per machine are **sublinear in number of edges m** in graph

Triangle Counting



3 triangles



- Clustering Algorithms
- Identifying thematic structures of networks
 - Spam and fraud detection
- Link classification and recommendation
- Joining three relations in a database
 - Database query optimization

Triangle Counting in MPC Model

Exact Setting			
Previous Work	MPC Rounds	Space Per Machine	Total Space
[SV11]	1	$O(m/\rho^2)$	$O(\rho m)$
[CC11]	$O(n)$	$O(n)$	$O(m)$
Folklore [CN85]	$O(\log n)$	$O(\alpha^2)$	$O(m\alpha)$
[BELMR22]	$O(\log \log n)$	$O(n^\delta)$	$O(m\alpha)$

$\delta > 0$ is **any constant**

[SV11]: Suri and Vassilvitski, WWW '11

[CC11]: Chu and Cheng KDD '11

[CN85]: Chiba and Nishizeki SICOMP '85

[BELMR22]: Biswas, Eden, L, Mitrović, Rubinfeld APPROX '22

Triangle Counting in MPC Model

Exact Setting			
Previous Work	MPC Rounds	Space Per Machine	Total Space
[SV11]	1	$O(m/\rho^2)$	$O(\rho m)$
[CC11]	$O(n)$	$O(n)$	$O(m)$
Folklore [CN85]	$O(\log n)$	$O(\alpha^2)$	$O(m\alpha)$
[BELMR22]	$O(\log \log n)$	$O(n^\delta)$	$O(m\alpha)$

$\delta > 0$ is **any constant**

[SV11]: Suri and Vassilvitski, WWW '11

[CC11]: Chu and Cheng KDD '11

[CN85]: Chiba and Nishizeki SICOMP '85

[BELMR22]: Biswas, Eden, L, Mitrović, Rubinfeld APPROX '22

Triangle Counting in MPC Model

Exact Setting		
Previous Work	MPC Rounds	Space Per Processor
[SV11]	1	$O(m, n)$
[CC11]	$O(n)$	$O(m)$
Folklore [CN85]	$O(\log n)$	$O(m)$
[BELMR22]	$O(\log \log n)$	$O(m)$

Arboricity α : number of forests that edges can be partitioned into

Real-world graphs:
arboricity generally
poly($\log n$)

$\delta > 0$ is **any constant**

[SV11]: Suri and Vassilvitski, WWW '11


[CC11]: Chu and Cheng KDD '11

[CN85]: Chiba and Nishizeki SICOMP '85

[BELMR22]: Biswas, Eden, L, Mitrović, Rubinfeld APPROX '22


Arboricity of the Graph

- **Arboricity** of the graph
 - Minimum number of forests to decompose the graph




Chiba-Nishizeki Sum of Minimum Degree Endpoints [CN85]

- Given an input graph $G = (V, E)$ with arboricity α , it holds sum of minimum degrees endpoints of every edge is at most $2m\alpha$



Chiba-Nishizeki Sum of Minimum Degree Endpoints [CN85]


- Given an input graph $G = (V, E)$ with arboricity α , it holds sum of minimum degrees endpoints of every edge is at most $2m\alpha$



Chiba-Nishizeki Sum of Minimum Degree Endpoints [CN85]

- Given an input graph $G = (V, E)$ with arboricity α , it holds sum of minimum degrees endpoints of every edge is at most $2m\alpha$


Sum of minimum degrees of endpoints = **29**



Chiba-Nishizeki Sum of Minimum Degree Endpoints [CN85]


- Given an input graph $G = (V, E)$ with arboricity α , it holds sum of minimum degrees endpoints of every edge is at most $2m\alpha$

Sum of minimum degrees of endpoints = **29**




$$29 \leq 2 \cdot 12 \cdot 2 = 48$$

MPC Triangle Counting Algorithm



Send adjacency list of
smaller degree endpoint
of each edge to
machines


MPC Triangle Counting Algorithm



Each machine generates
wedge queries

Send adjacency list of
smaller degree endpoint
of each edge to
machines

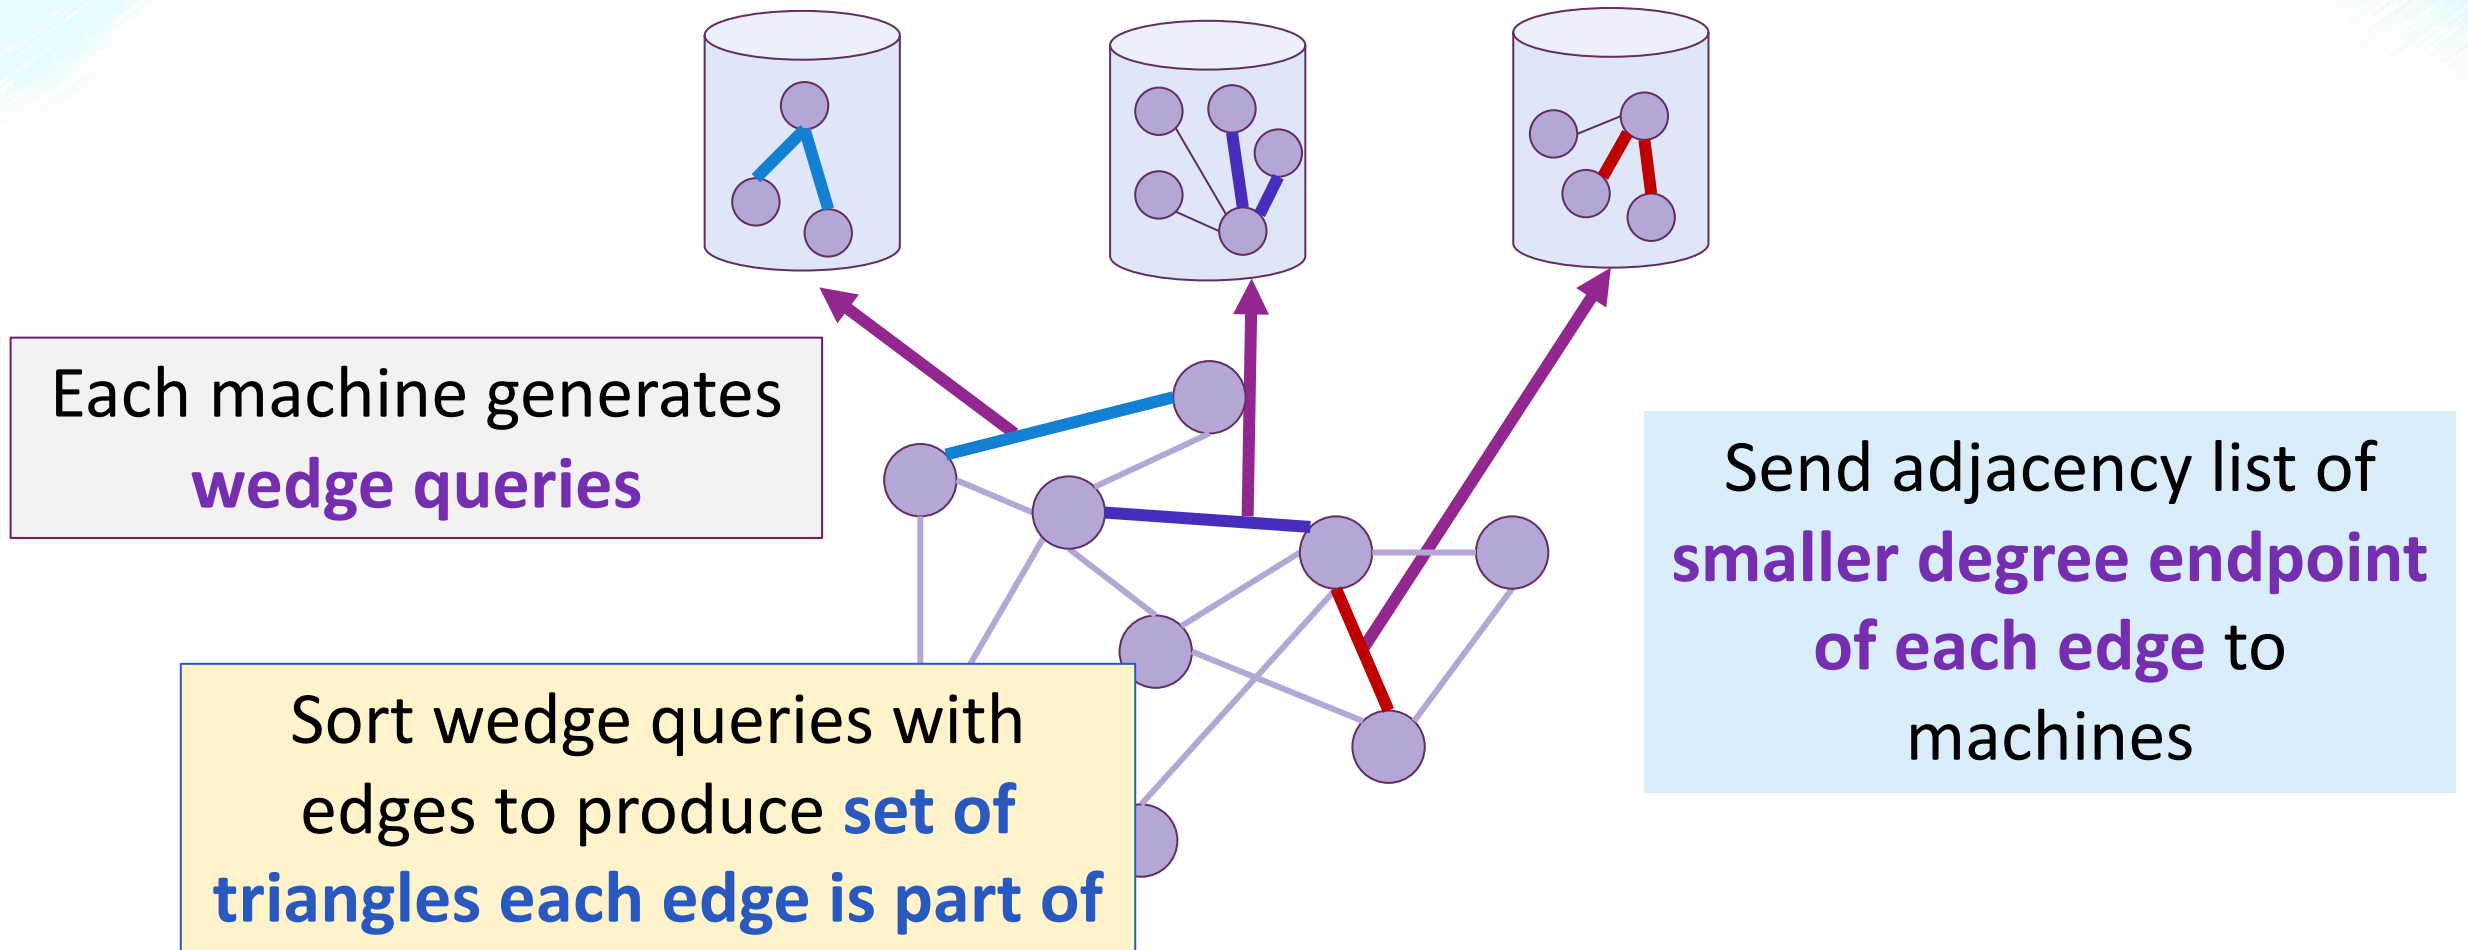
MPC Triangle Counting Algorithm



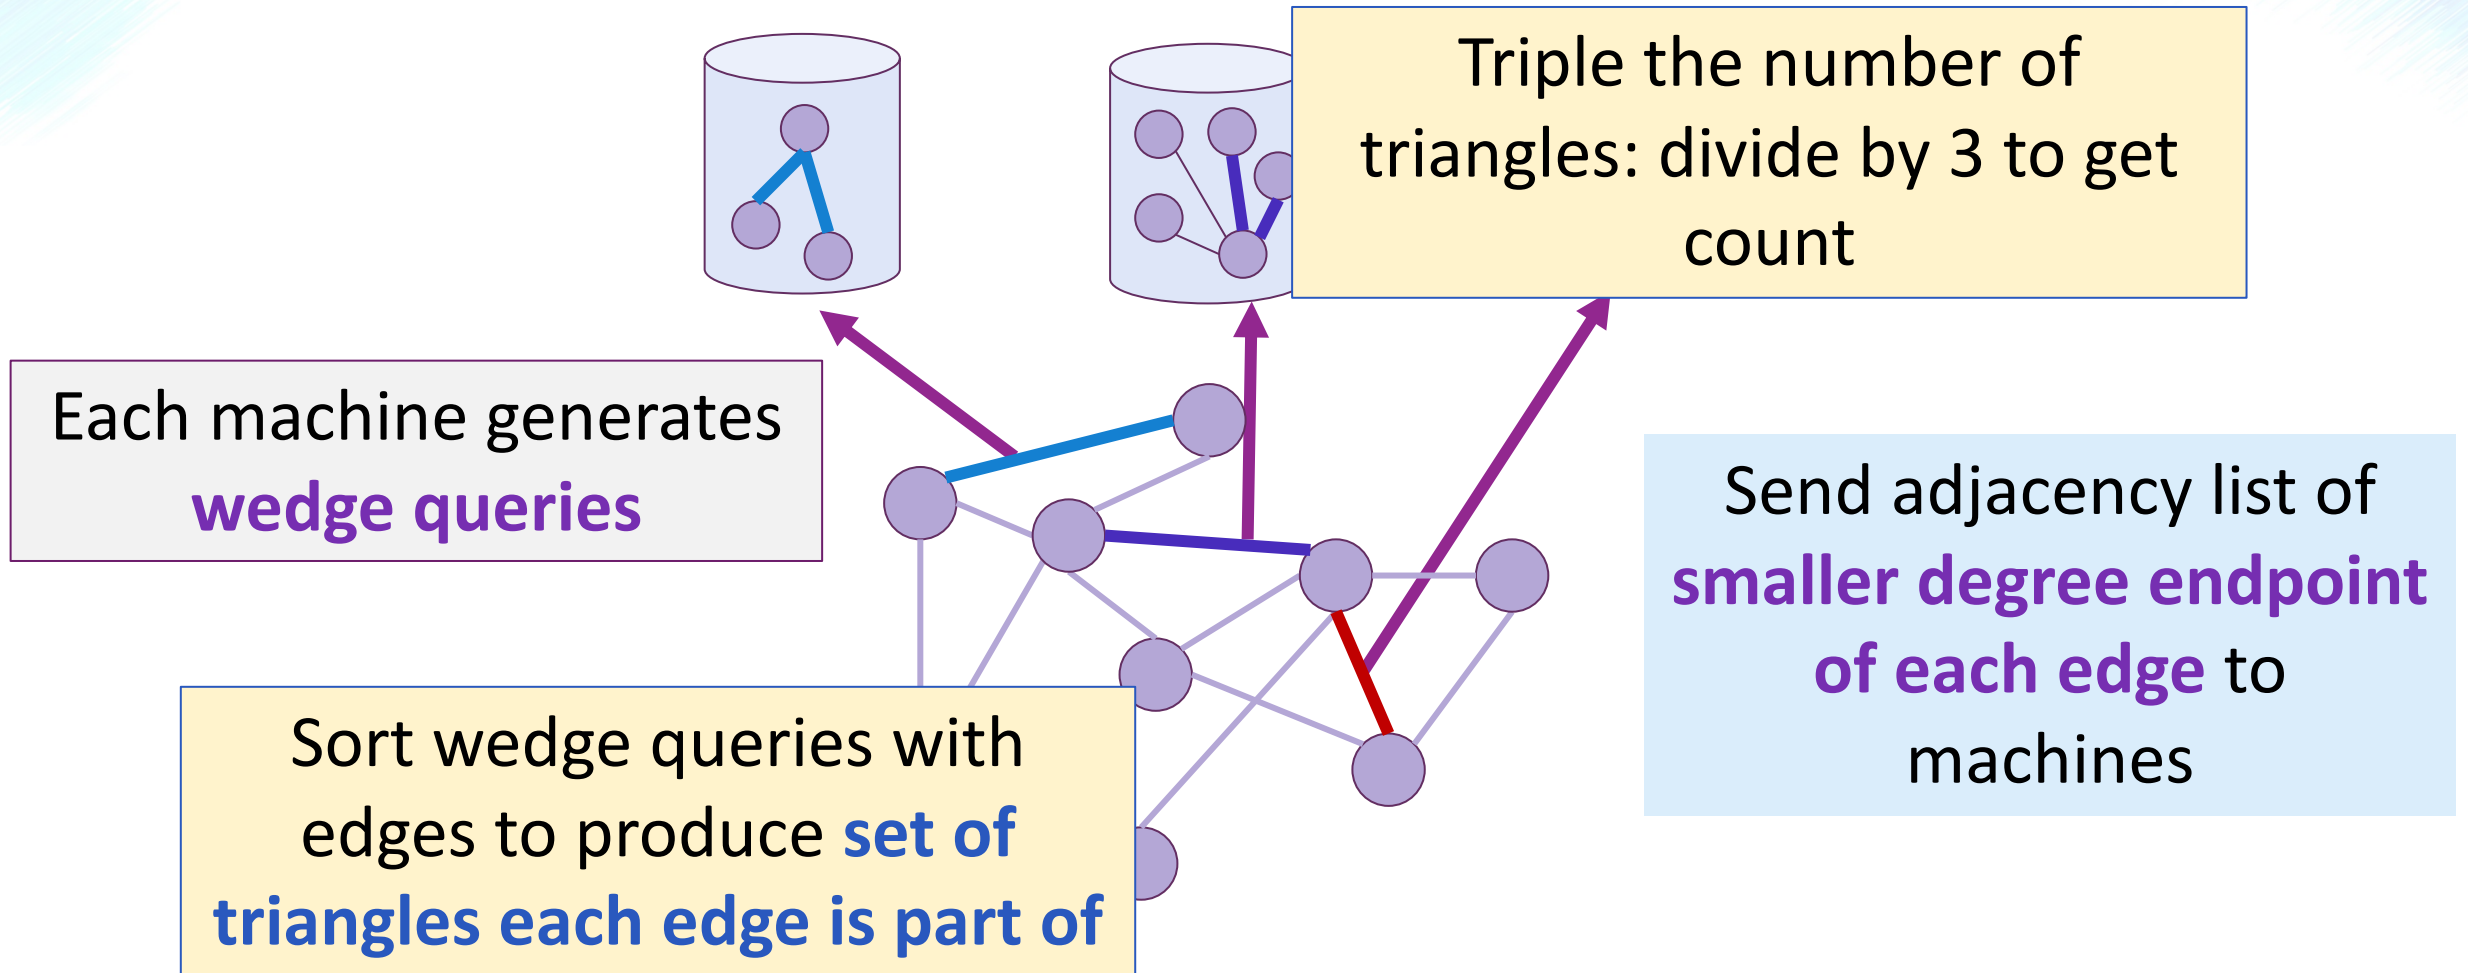
Each machine generates
wedge queries

Send adjacency list of
smaller degree endpoint
of each edge to
machines



MPC Triangle Counting Algorithm



MPC Triangle Counting Algorithm



MPC Triangle Counting Algorithm



Triple the number of triangles: divide by 3 to get count

Each machine
wedge

Theorem: Exact triangle count in $O\left(\frac{1}{\delta}\right)$ MPC rounds, $O(n^\delta)$ space per machine, and $O(m\alpha)$ total space.

Priority list of
smaller degree endpoint
of each edge to
machines

Sort wedge queries with
edges to produce **set of
triangles** each edge is part of

Matching Lower Bound

Theorem: Exact triangle count requires $\Omega\left(\frac{1}{\delta}\right)$ MPC rounds, when given $O(n^\delta)$ space per machine, and $O(m\alpha)$ total space.

Matching Lower Bound

Theorem: Exact triangle count requires $\Omega\left(\frac{1}{\delta}\right)$ MPC rounds, when given $O(n^\delta)$ space per machine, and $O(m\alpha)$ total space.

Each machine contains disjoint count



m/n^δ machines

Matching Lower Bound

Theorem: Exact triangle count requires $\Omega\left(\frac{1}{\delta}\right)$ MPC rounds, when given $O(n^\delta)$ space per machine, and $O(m\alpha)$ total space.


**Aggregate
count into one
machine—tree
with m/n^δ
leaves**



m/n^δ machines

Matching Lower Bound


Aggregate
count into one
machine—tree
with m/n^δ
leaves



m/n^δ machines

Matching Lower Bound

Each machine can receive at most n^δ counts




Aggregate count into one machine—tree with m/n^δ leaves

m/n^δ machines

Matching Lower Bound

$$\log_{n^\delta} \left(\frac{m}{n^\delta} \right) \leq \log_{n^\delta} (n^2) = \Omega \left(\frac{1}{\delta} \right) \text{ depth for } \delta \in (0, 1)$$

Aggregate
count into one
machine—tree
with m/n^δ
leaves



m/n^δ machines

Open Questions and Future Directions

- Small subgraph counting for a **broader class of small subgraphs**

Open Questions and Future Directions

- Small subgraph counting for a **broader class of small subgraphs**
- **Decrease the total space usage** for exact triangle counting in $O(1)$ rounds to $O(m + n)$ (even \sqrt{n} number of rounds, linear space per machine not known)

Open Questions and Future Directions

- Small subgraph counting for a **broader class of small subgraphs**
- **Decrease the total space usage** for exact triangle counting in $O(1)$ rounds to $O(m + n)$ (even \sqrt{n} number of rounds, linear space per machine not known)
- Approximate triangle counting in $O(1)$ rounds and strictly sublinear space in **sparse graphs** where $m = \tilde{O}(n)$