Massively Parallel Algorithms for Small Subgraph Counting









Amartya Shankha Biswas MIT CSAIL Talya Eden Boston University/ MIT → Bar-Ilan U Quanquan C. Liu Northwestern Slobodan Mitrović UC Davis Ronitt Rubinfeld MIT CSAIL

To Appear in APPROX 2022

Massively Parallel Computation (MPC)

- Massively parallel systems
 - Distributed cluster of
 multiple machines



Massively Parallel Computation (MPC)

- Massively parallel systems
 - Distributed cluster of
 multiple machines
 - Communicate with each other via rounds of communication



Massively Parallel Computation (MPC)

- Massively parallel systems
 - Distributed cluster of multiple machines
 - Communicate with each other via rounds of communication
 - Limited space in each individual machine



Commercial Data Centers



Google Kubernetes Engine



Commercial Data Centers





Machine 1 Machine 2 Machine 3

Massively Parallel Computation (MPC) Model

• Theoretical standard for studying parallel frameworks such as MapReduce, Hadoop, Spark, Dryad, and Google Cloud Dataflow



Graph Algorithms in MPC Model

- Matching and MIS [BBDFHKU19, BHH19, GGKMR19, CLMMOS18, NO21]
- Connectivity [ASSWZ18, BDELM19, DDKPSS19]
- Graph sparsification [GU19, CDP20]
- Vertex cover [Assadi17, GGKMR18]
- MST and 2-edge connectivity [NO21]
- Well-connected components [ASW18, ASW19]
- Coloring [BDHKS19, CFGUZ19]

- M machines
- Synchronous rounds

- M machines
- Synchronous rounds



- M machines
- Synchronous rounds



- M machines
- Synchronous rounds



- M machines
- Synchronous rounds



- M machines
- Synchronous rounds



- *M* machines
- Synchronous rounds



- Strongly sublinear memory:
 - $S = n^{\delta}$ for some constant $\delta \in (0, 1)$

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

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

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

Also want: $O(\log \log n)$ or O(1) rounds

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

Also want: $O(\log \log n)$ or O(1) rounds

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

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

Also want: $O(\log \log n)$ or O(1) rounds

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

All are **sublinear in number of edges** *m* in graph

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

$\delta > 0$ is any constant

[SV11]: Suri and Vassilvitski, WWW '11 [CC11]: Chu and Cheng KDD '11 [CN85]: Chiba and Nishizeki SICOMP '85

FODSI Sublinear Algorithms Workshop 2022

	Exact	Setting	Arboricity α : number of
Previous Work	MPC Rounds	Space Pe	forests that edges can be
[SV11]	1	0(m	partitioned into
[CC11]	O(n)	0(
Folklore [CN85]	$O(\log n)$	0((Real-world graphs:
BELMR22	$O(\log \log n)$	O (1	nolv(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

FODSI Sublinear Algorithms Workshop 2022

	Exact	Arboricity Setting	α: number of forests that edg be partitioned into
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	$O(\log \log n)$	$oldsymbol{O}(n^{\delta})$	0(mα)

Better space per machine or better total space when $\alpha \leq m^{1/2-\varepsilon}$, but worse number of rounds

	Eve et (Arboricity	α: number of forests that end of the partitioned into	dges c
	Exact	Setting			
Previous Work	MPC Rounds	Space Pe	r Machine	Total Space	
[SV11]	1	0(m	$/ ho^2)$	O(ho m)	
[CC11]	O(n)	0(<i>n</i>)	<i>O</i> (<i>m</i>)	
Folklore [CN85]	$O(\log n)$	0(0	$\chi^2)$	$O(m\alpha)$	
BELMR22	$O(\log \log n)$	0(1	n^{δ})	Ο (m α)	

Better rounds and space per machine, but total space when $\alpha = \omega(1)$

	Exact	A Setting	rboricity α : number of forests that edge be partitioned into
Previous Work	MPC Rounds	Space Per M	lachine Total Space
[SV11]	1	0(m/µ	$O^2) O(\rho m)$
[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 $(m\alpha)$

Smaller number of rounds, but worse space per machine when $\alpha < n^{o(1)}$

	Exact	Setting	be partitioned into
Previous Work	MPC Rounds	Space Per M	lachine Total Space
[SV11]	1	0(m/µ	$O^2) O(\rho m)$
[CC11]	O(n)	O(n)) $O(m)$
Folklore [CN85]	$O(\log n)$	$O(\alpha^2$) $O(m\alpha)$
BELMR22	$O(\log \log n)$	$O(n^{\delta}$	$\boldsymbol{O}(\boldsymbol{m}\boldsymbol{\alpha})$

Strictly sublinear setting

$(1 + \varepsilon)$ -Approximate Setting					
Previous Work	MPC Rounds	Space Per Machine	Total Space	Triangles Lower Bound	
[PT12]	0(1)	$O\left(\frac{m\Delta_e}{T}\right)$	O(m)	$\Omega(d_{avg})$	
[SPK13]	0(1)	$O(n^{\delta})$	O(m)	$\Omega\left(\sum_{v\in V} \deg(v)^2\right)$	
BELMR22	0 (1)	$\widetilde{\boldsymbol{o}}(\boldsymbol{n})$	$\widetilde{\boldsymbol{o}}(\boldsymbol{m})$	$\Omega(\sqrt{d_{avg}})$	

[PT12]: Pagh and Tsourakakis, IPL '12 [SPK13]: Seshadhri, Pinar, Kolda, ICDM '13

$(1+oldsymbol{arepsilon}$ -Approximate Setting					
Previous Work	MPC Rounds	Space Per Machine	Total Space	Triangles Lower Bound	
[PT12]	0(1)	$O\left(\frac{m\Delta_e}{T}\right)$	<i>O</i> (<i>m</i>)	$\Omega(d_{avg})$	
[SPK13]	0(1)	$O(n^{\delta})$	<i>O(m)</i>	$\Omega\left(\sum_{\nu\in V} \deg(\nu)^2\right)$	
BELMR22	0 (1)	$\widetilde{\boldsymbol{o}}(\boldsymbol{n})$	$\widetilde{\boldsymbol{o}}(\boldsymbol{m})$	$\Omega(\sqrt{d_{avg}})$	

Better triangle lower bounds, but slightly worse total space

$(1+oldsymbol{arepsilon}$ -Approximate Setting					
Previous Work	MPC Rounds	Space Per Machine	Total Space	Triangles Lower Bound	
[PT12]	0(1)	$O\left(\frac{m\Delta_e}{T}\right)$	<i>O</i> (<i>m</i>)	$\Omega(d_{avg})$	
[SPK13]	0(1)	$O(n^{\delta})$	<i>O</i> (<i>m</i>)	$\Omega\left(\sum_{\nu\in V} \deg(\nu)^2\right)$	
BELMR22	0 (1)	$\widetilde{\boldsymbol{o}}(\boldsymbol{n})$	$\widetilde{\boldsymbol{0}}(\boldsymbol{m})$	$\Omega(\sqrt{d_{avg}})$	

Worse space per machine than SPK13

$(1 + \varepsilon)$ -Approximate Setting					
Previous Work	MPC Rounds	Space Per Machine	Total Space	Triangles Lower Bound	
[PT12]	0(1)	$O\left(\frac{m\Delta_e}{T}\right)$	<i>O</i> (<i>m</i>)	$\Omega(d_{avg})$	
[SPK13]	0(1)	$O(n^{\delta})$	<i>O</i> (<i>m</i>)	$\Omega\left(\sum_{\nu\in V} \deg(\nu)^2\right)$	
BELMR22	0 (1)	$\widetilde{\boldsymbol{o}}(\boldsymbol{n})$	$\widetilde{\boldsymbol{o}}(\boldsymbol{m})$	$\Omega(\sqrt{d_{avg}})$	

Better space per machine than PT12 when $n = o\left(\frac{m\Delta_e}{T}\right)$

- Strongly sublinear memory:
 - **Exact** triangle counting:
 - Bounded arboricity
 - $O(\log \log n)$ rounds
 - $O(m\alpha)$ total space

- Strongly sublinear memory:
 - **Exact** triangle counting:
 - Bounded arboricity
 - $O(\log \log n)$ rounds
 - $O(m\alpha)$ total space
- Near-linear memory:
 - Approximate triangle counting
 - $(1 + \varepsilon)$ -approximation when $T \ge \sqrt{m/n}$
 - O(1) rounds, $\tilde{O}(n+m)$ total space

- Strongly sublinear memory:
 - **Exact** triangle counting:
 - Bounded arboricity
 - $O(\log \log n)$ rounds
 - $O(m\alpha)$ total space
- Near-linear memory:
 - Approximate triangle counting
 - $(1 + \varepsilon)$ -approximation when $T \ge \sqrt{m/n}$
 - O(1) rounds, $\tilde{O}(n+m)$ total space

Results in Our Paper

- Strongly sublinear memory:
 - Extensions to clique counting

- Strongly sublinear memory:
 - **Exact** triangle counting:
 - Bounded arboricity
 - $O(\log \log n)$ rounds
 - $O(m\alpha)$ total space
- Near-linear memory:
 - Approximate triangle counting
 - $(1 + \varepsilon)$ -approximation when $T \ge \sqrt{m/n}$
 - O(1) rounds, $\tilde{O}(n+m)$ total space

Results in Our Paper

- Strongly sublinear memory:
 - Extensions to clique counting
- Linear memory:
 - Extensions to clique counting

- Strongly sublinear memory:
 - **Exact** triangle counting:
 - Bounded arboricity
 - $O(\log \log n)$ rounds
 - $O(m\alpha)$ total space
- Near-linear memory:
 - Approximate triangle counting
 - $(1 + \varepsilon)$ -approximation when $T \ge \sqrt{m/n}$
 - O(1) rounds, $\tilde{O}(n+m)$ total space

Results in Our Paper

- Strongly sublinear memory:
 - Extensions to clique counting
- Linear memory:
 - Extensions to clique counting
- Counting all small subgraphs of size at most 5 using Bera, Pashanasangi and Seshadhri [ITCS 2020]
Results in This Presentation

- Strongly sublinear memory:
 - **Exact** triangle counting:
 - Bounded arboricity
 - $O(\log \log n)$ rounds
 - $O(m\alpha)$ total space
- Near-linear memory:
 - Approximate triangle counting
 - $(1 + \varepsilon)$ -approximation when $T \ge \sqrt{m/n}$
 - O(1) rounds, $\tilde{O}(n+m)$ total space

Results in Our Paper

- Strongly sublinear memory:
 - Extensions to clique counting
- Linear memory:
 - Extensions to clique counting
- Counting all small subgraphs of size at most 5 using Bera, Pashanasangi and Seshadhri [ITCS 2020]
- Simulations on real-world graphs:
 - Improvements in number of rounds
 - Improvements in approximation

Results in This Presentation

- Strongly sublinear memory:
 - **Exact** triangle counting:
 - Bounded arboricity
 - $O(\log \log n)$ rounds
 - $O(m\alpha)$ total space

• Near-linear memory:

- Approximate triangle counting
- $(1 + \varepsilon)$ -approximation when $T \ge \sqrt{m/n}$
- O(1) rounds, $\tilde{O}(n + m)$ total space

Results in Our Paper

- Strongly sublinear memory:
 - Extensions to clique counting

• Linear memory:

- Extensions to clique counting
- Counting all small subgraphs of size at most 5
- Simulations on real-world graphs:
 - Improvements in number of rounds
 - Improvements in approximation

Exact Triangle Counting Bounded Arboricity

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

Exact Triangle Counting Bounded Arboricity

There exists a MPC algorithm that outputs the exact count of triangles in a graph with arboricity α in $O(\log \log n)$ rounds, $O(n^{\delta})$ space per machine for any constant $\delta > 0$ and $O(m\alpha)$ total space.

Massively Parallel Algorithms for Small Subgraph Counting

Amartya Shankha Biswas, Talya Eden, Quanquan C. Liu, Slobodan Mitrovic, Ronitt Rubinfeld [arxiv.org/2002.08299]

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

Exact Triangle Counting Bounded Arboricity

There exists a MPC algorithm that outputs the exact count of triangles in a graph with arboricity α in $O(\log \log n)$ rounds, $O(n^{\delta})$ space per machine for any constant $\delta > 0$ and $O(m\alpha)$ total space.

Massively Parallel Algorithms for Small Subgraph Counting

Amartya Shankha Biswas, Talya Eden, Quanquan C. Liu, Slobodan Mitrovic, Ronitt Rubinfeld

[arxiv.org/2002.08299]

Standard Triangle Counting: $O(\log n)$ rounds $\Omega(\alpha^2)$ space per machine $O(m\alpha)$ total space



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



- Successively remove vertices with degree less than 2α and count number of triangles adjacent to the removed vertices
 - Maintain total count



- Successively remove vertices with degree less than 2α and count number of triangles adjacent to the removed vertices
 - Maintain total count





• Successively remove vertices with degree less than 2α and count number of triangles adjacent to the removed vertices



- Successively remove vertices with degree less than 2α and count number of triangles adjacent to the removed vertices
 - Maintain total count



- Successively remove vertices with degree less than 2α and count number of triangles adjacent to the removed vertices
 - Maintain total count

Maximum number of edges in the graph: $m \leq n\alpha$

Number of vertices remaining: $\frac{n\alpha}{2\alpha} = \frac{n}{2}$

Number of rounds needed: $O(\log n)$



• Successively remove vertices with degree less than 2α and count number of triangles adjacent to the removed vertices

Maximum number of edges in the graph: $m \leq n\alpha$

Number of vertices remaining: $\frac{n\alpha}{2\alpha} = \frac{n}{2}$

Number of rounds needed: $O(\log n)$



• Successively remove vertices with degree less than 2α and count number of triangles adjacent to the removed vertices

Total space used: $O(m\alpha)$

Maximum number of edges in the graph: $m \le n\alpha$

Number of vertices remaining: $\frac{n\alpha}{2\alpha} = \frac{n}{2}$

Number of rounds needed: $O(\log n)$



• Successively remove vertices with degree less than 2α and count number of triangles adjacent to the removed vertices







$\deg(v) \le 4\alpha$



$\deg(v) \le 4\alpha$



$\deg(v) \le 6\alpha$



$\deg(v) \le 6\alpha$

5 Triangles

i = 1



$\deg(v) \le 10\alpha$

5 Triangles

i = 2

$O(\log \log n)$ i=2

 $\deg(v) \le 10\alpha$

5 Triangles

• Number of vertices left after first round: X

- Number of vertices left after first round: X
- Total number of edges left after first round:

$$m \ge \frac{1}{2} \cdot X \cdot 4\alpha = 2\alpha X$$

- Number of vertices left after first round: X
- Total number of edges left after first round:

$$m \ge \frac{1}{2} \cdot X \cdot 4\alpha = 2\alpha X$$

 $m_1 \leq X\alpha$

- Number of vertices left after first round: X
- Total number of edges left after first round:

$$m \ge \frac{1}{2} \cdot X \cdot 4\alpha = 2\alpha X$$

$$m_1 \le X\alpha$$
$$m_1 \le \frac{m}{2}$$

- Number of vertices left after i-th round: X
- Total number of edges left after first round:

$$m \ge \frac{1}{2} \cdot X \cdot 4\alpha = 2\alpha X$$

$$m_{i-1} \ge \frac{1}{2} \cdot X \cdot 2^{\left(\frac{3}{2}\right)^{i-1}} \cdot 2\alpha$$

$$m_i \le X\alpha$$

$$m_1 \le \frac{m}{2}$$

$$m_i \le \frac{m_{i-1}}{2^{\left(\frac{3}{2}\right)^{i-1}}} < \frac{m}{2^{\left(\frac{3}{2}\right)^i}}$$

• Number of vertices left after i-th round: X



• Number of vertices left after i-th round: X



• Number of vertices left after i-th round: X



 Last Challenge: Cannot count on one machine because that is too much space

- Last Challenge: Cannot count on one machine because that is too much space
 - Solution: Reduce to a problem where we merge several lists, sort, and find duplicates

- Last Challenge: Cannot count on one machine because that is too much space
 - Solution: Reduce to a problem where we merge several lists, sort, and find duplicates
 - Every removed node sends its adjacency list to its neighbors

- Last Challenge: Cannot count on one machine because that is too much space
 - Solution: Reduce to a problem where we merge several lists, sort, and find duplicates
 - Every removed node sends its adjacency list to its neighbors
 - Each neighbor which receives adjacency lists merges received lists with its own adjacency list












- MPC sorting algorithm of [GSZ11] to sort lists in O(1) rounds
- Find duplicates using new MPC primitive











Exact Triangle Counting

- Challenge: Cannot count on one machine because that is too much space
 - Need to have an MPC specific counting procedure
 - Removed nodes send list of neighbors to all neighbors
 - MPC sorting algorithm of [GSZ11] to sort lists
 - Find duplicates using new MPC primitive

There exists a MPC algorithm that outputs the exact count of triangles in a graph with arboricity α in $O(\log \log n)$ rounds, $O(n^{\delta})$ space per machine for any constant $\delta > 0$ and $O(m\alpha)$ total space.

Exact Triangle Counting

- Challenge: Cannot count on one machine because that is too much space
 - Need to have an MPC specific counting procedure
 - Removed nodes send list of neighbors to all neighbors
 - MPC sorting algorithm of [GSZ11] to sort lists
 - Find duplicates using new MPC primitive

Somewhat resembles **round compression** technique although simpler on bounded arboricity graphs and deterministic: do not need to do sampling

Results in This Presentation

- Strongly sublinear memory:
 - Exact triangle counting:
 - Bounded arboricity
 - O(log log n) rounds
 O(mα) total space
- Near-linear memory:
 - Approximate triangle counting
 - $(1 + \varepsilon)$ -approximation when $T \ge \sqrt{m/n}$
 - O(1) rounds, $\tilde{O}(n+m)$ total space

Results in Our Paper

- Strongly sublinear memory:
 - Extensions to clique counting

• Linear memory:

- Extensions to clique counting
- Counting all small subgraphs of size at most 5
- Simulations on real-world graphs:
 - Improvements in number of rounds
 - Improvements in approximation

Advantages and Disadvantages of Approximate Counting

- Main Advantage:
 - Small runtime, fast and requires little space
- Main Disadvantage:
 - Requires lower bound on the number of triangles

There exists a MPC algorithm that outputs a $(1 + \epsilon)$ -approximation for the number of triangles if the number of triangles $T \ge \sqrt{d_{avg}}$ and uses $\tilde{O}(m)$ total space and $\tilde{\Theta}(n)$ space per machine, O(1) MPC rounds.

Massively Parallel Algorithms for Small Subgraph Counting

Amartya Shankha Biswas, Talya Eden, Quanquan C. Liu, Slobodan Mitrovic, Ronitt Rubinfeld [arxiv.org/2002.08299]

There exists a MPC algorithm that outputs a $(1 + \epsilon)$ -approximation for the number of triangles if the number of triangles $T \ge \sqrt{d_{avg}}$ and uses $\tilde{O}(m)$ total space and $\tilde{\Theta}(n)$ space per machine, O(1) MPC rounds.

Massively Parallel Algorithms for Small Subgraph Counting

Amartya Shankha Biswas, Talya Eden, Quanquan C. Liu, Slobodan Mitrovic, Ronitt Rubinfeld [arxiv.org/2002.08299]

Previous: $T \ge d_{avg}$ [Pagh and Tsourakakis '12]

There exists a MPC algorithm that outputs a $(1 + \epsilon)$ -approximation for the number of triangles if the number of triangles $T \ge \sqrt{d_{avg}}$ and uses $\tilde{O}(m)$ total space and $\tilde{\Theta}(n)$ space per machine, O(1) MPC rounds.

Massively Parallel Algorithms for Small Subgraph Counting

Amartya Shankha Biswas, Talya Eden, Quanquan C. Liu, Slobodan Mitrovic, Ronitt Rubinfeld [arxiv.org/2002.08299]

Previous: $T \ge d_{avg}$ [Pagh and Tsourakakis '12]

[Seshadhri, Pinar, Kolda '13] can get better near-linear space per machine













 Challenge 1: Induced subgraphs do not exceed the space per machine

- Challenge 1: Induced subgraphs do not exceed the space per machine
- Challenge 2: How to compute the induced subgraph in each machine when one vertex *can appear on multiple machines*?

- Challenge 1: Induced subgraphs do not exceed the space per machine
- Challenge 2: How to compute the induced subgraph in each machine when one vertex *can appear on multiple machines*?
- Challenge 3: The number of triangles across the machines concentrates

Challenge 1: Induced subgraphs do not exceed the space per machine

Careful setting of p

• Challenge 2: How to compute the induced subgraph in each machine when one vertex *can appear on multiple machines*?

k-wise independent hash function for small k

Challenge 3: The number of triangles across the machines concentrates
 Constant probability of success and median trick

Challenge 1: Induced subgraphs do not exceed the space per machine

Careful setting of p

• Challenge 2: How to compute the induced subgraph in each machine when one vertex *can appear on multiple machines*?

k-wise independent hash function for small k

Challenge 3: The number of triangles across the machines concentrates
 Constant probability of success and median trick

Small subgraph counting for a broader class of small subgraphs

- Small subgraph counting for a broader class of small subgraphs
 - Recent works of Bressan '19 and Bera, Pashanasangi, and Seshadhri '21 use DAG tree decomposition

- Small subgraph counting for a broader class of small subgraphs
 - Recent works of Bressan '19 and Bera, Pashanasangi, and Seshadhri '21 use DAG tree decomposition
 - Can we implement in MPC?

- Small subgraph counting for a broader class of small subgraphs
 - Recent works of Bressan '19 and Bera, Pashanasangi, and Seshadhri '21 use DAG tree decomposition
 - Can we implement in MPC?
- Counting in the adaptive MPC model (AMPC)

- Small subgraph counting for a broader class of small subgraphs
 - Recent works of Bressan '19 and Bera, Pashanasangi, and Seshadhri '21 use DAG tree decomposition
 - Can we implement in MPC?
- Counting in the adaptive MPC model (AMPC)
- Approximate triangle counting in O(1) rounds and strictly sublinear space in sparse graphs where $m = \tilde{O}(n)$