Improved Massively Parallel Computation Algorithms for MIS, Matching, and Vertex Cover

Themis Gouleakis

MIT

June 14, 2018

Joint work with: Mohsen Ghaffari(ETH), Christian Konrad(University of Bristol), Slobodan Mitrović(EPFL) and Ronitt Rubinfeld(MIT and Tel Aviv University)

CONGESTED CLIQUE MODEL [LOTKER PAVLOV PATT-SHAMIR PELEG 03]

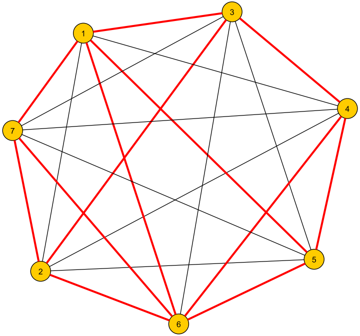

- *Complete* communication graph: $G' = K_{|V|}$.
- *Synchronous* messages of size $O(\log n)$ bits
- complexity=♯ rounds.

Massively Parallel Computation (MPC) model [Karloff Suri Vassilvitskii 10]

- Inspired by MapReduce
- Input size=N
- m machines
- Space $S \in \left[\frac{N}{m}, N\right]$ per machine
- unbounded internal computation
- Total communication per node bounded by S each round.
- complexity=$\sharp$ rounds.

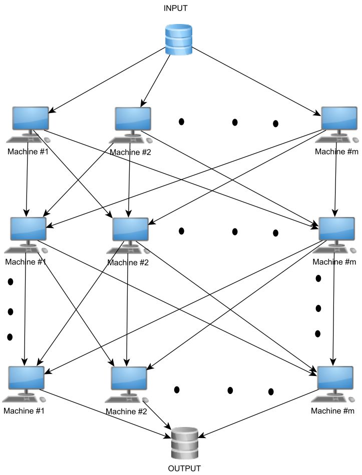

RESULTS

Theorem: There is an algorithm that with high probability computes an MIS in $O(\log\log\Delta)$ rounds of the MPC model, with $\tilde{O}(n)$-bits of memory per machine.

Theorem: There is an algorithm that with high probability computes a $(2+\varepsilon)$-approximate integral maximum matching and a $(2+\varepsilon)$-approximate integral minimum vertex cover in $O(\log\log n)$ rounds of the MPC model, with $\tilde{O}(n)$-bits of memory per machine.

Improved Massively Parallel Computation Algorithms for MIS, Matching, and Vertex Cover

Mohsen Ghaffari[1], Themis Gouleakis[2], Christian Konrad[3], Slobodan Mitrović[4] and Ronitt Rubinfeld[2,5]

ETH[1], MIT[2], University of Bristol[3], EPFL[4], Tel Aviv University[5]

Models

CONGESTED CLIQUE model
[Lotker Pavlov Patt-Shamir Peleg 03]

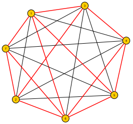

- *Complete* communication graph: $G' = K_{|V|}$.
- *Synchronous* messages of size $O(\log n)$ bits
- complexity=♯ rounds.

Massively Parallel Computation model
[Karloff Suri Vassilvitskii 10]

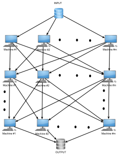

- Input size=N
- m machines
- Space $S \in [\frac{N}{m}, N]$ per machine
- Total communication per node per round: $\leq S$
- complexity=♯ rounds.

Greedy randomized Algorithm for MIS

Greedy Randomized Maximal Independent Set

- Choose a permutation $\pi : [n] \to [n]$ u.a.r.
- Repeat until the next rank is at least $n/\log^{10} n$ and the maximum degree is at most $\log^{10} n$:
  (A) Add smallest rank vertex v to the MIS.
  (B) Remove all the neighbors of v.
- Run $O(\log\log\Delta)$ rounds of the Sparsified MIS Algorithm of [Ghaffari 17] in the remaining graph. Remove from the graph the constructed MIS and its neighborhood.
- Find MIS of the residual graph (single machine)
- Output the union of the constructed MIS sets.

MPC simulation

1. Choose a permutation $\pi : [n] \to [n]$ u.a.r
2. **Repeat** for each chunk V_k:
   a. Machine ♯1 receives all $G[V_k]$ edges.
   b. Broadcasts $MIS(G[V_k])$
   c. Update local memories

Until: $\Delta < \log^{10} n$

- Run $O(\log\log\Delta)$ rounds of sparsified MIS algorithm [Ghaffari 17]

CONGESTED CLIQUE vs MPC

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

- *CONGESTED CLIQUE* $\Rightarrow$ *MPC* [BDH18]
  (1 MPC machine per graph node)
  - **Assign** each MPC machine to a vertex.
  - **For** each edge $\{u, v\}$, **send** message to machines u, v.
  - **Simulate** *CONGESTED CLIQUE* algorithm.
    ($< n$ MPC machines: assign multiple vertices to each one)
- *MPC* $\Rightarrow$ *CONGESTED CLIQUE*
  - Each node needs to send/receive up to $O(n)$ bits.
  - Split messages into $O(\log n)$ bit chunks.
  - Exploit unused edges of the communication graph using routing scheme in [Lenzen13]
    - We can simulate one MPC round using $O(1)$ congested clique rounds.

Theorem (inf.) [GGKMR18], [BDH18]

CONGESTED CLIQUE $\approx$ *MPC* model with $S = O(n)$

Analysis

Lemma: Let G_r be the remaining graph after r vertices are simulated. Then $\Delta_r = O(\frac{n \log n}{r})$ w.h.p.

- $O(n)$ edges sent to the central machine each round w.h.p.
  - 1st round: For $\alpha < 1/2$, we get:

$$Pr[\{i, j\} \in G[V_1]] = \left(\frac{1}{\Delta^{\alpha}}\right)^2 \leq \frac{1}{n} \Rightarrow O(n) \text{ edges in } G[V_1]$$

  - k-th round:

$$\mathbb{E}[\sharp\text{edges}] = \Delta_r \cdot \left(\frac{1}{\Delta^{\alpha^i}}\right)^2 = O(n)$$

- After $i = O(\log\log\Delta)$-rounds the max degree $\Delta_i \leq \log^{10} n$
  - **Key idea:** High degree vertices are much more likely to be removed!

Results

Theorem: There is an algorithm that with high probability computes an MIS in $O(\log\log\Delta)$ rounds of the MPC model, with $\tilde{O}(n)$-bits of memory per machine.

Theorem: There is an algorithm that with high probability computes a $(2+\varepsilon)$-approximate integral maximum matching and a $(2+\varepsilon)$-approximate integral minimum vertex cover in $O(\log\log n)$ rounds of the MPC model, with $\tilde{O}(n)$-bits of memory per machine.

References

[1] Soheil Behnezhad, Mahsa Derakhshan, and MohammadTaghi Hajiaghayi.
Brief announcement: Semi-mapreduce meets congested clique.
arXiv preprint arXiv:1802.10297, 2018.

[2] Mohsen Ghaffari.
Distributed mis via all-to-all communication.
In *Proceedings of the ACM Symposium on Principles of Distributed Computing*, pages 141–149. ACM, 2017.

[3] Howard Karloff, Siddharth Suri, and Sergei Vassilvitskii.
A model of computation for mapreduce.
pages 938–948, 2010.

[4] Christoph Lenzen.
Optimal deterministic routing and sorting on the congested clique.
In *Proceedings of the 2013 ACM symposium on Principles of distributed computing*, pages 42–50. ACM, 2013.

[5] Zvi Lotker, Elan Pavlov, Boaz Patt-Shamir, and David Peleg.
MST construction in O($\log\log n$) communication rounds.
pages 94–100. ACM, 2003.