

Massively Parallel Communication (MPC)

Inspired by frameworks such as MapReduce, Hadoop, Dryad, Spark

Clean algorithmic framework for designing algos. that can be parallelized

[KSV '10] MapReduce Class MRC

[BKS '13] MPC

Input data size: N
machines: M
memory per machine: S

$S \geq N$: single machine suffices

typically $S = N^c$ $c < 1$

- Computation proceeds in synchronous rounds
- In each round, each machine operates on local data then sends/receives messages to/from other machines
- Communication sent/received $\leq S$ words per round
- ignore local computation
- Focus on # rounds

Other parameters: Replication factor: $\frac{M \cdot S}{N}$

Total work / Work Efficiency

* ignore asynchronous communication, fault tolerance

Graph Problems:

$$n = |V|, \quad m = |E|$$

$$\text{input size } N = m$$

$$\# \text{ machines } M = \Theta\left(\frac{N}{S}\right)$$

Three memory regimes:

① Strongly superlinear: $S = n^{1+\epsilon}$ $\epsilon > 0$

② Near linear $S = \Theta(n)$

③ Strongly sublinear $S = n^\alpha \quad \alpha \in (0, 1)$

Initial data distribⁿ: split across machines arbitrarily

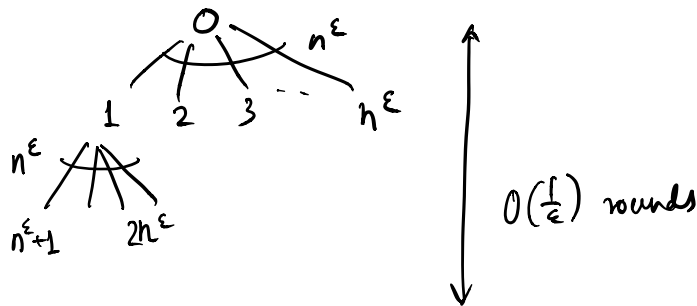
- can "load balance" using hashing in $O(1)$ rounds

S limits communication between machines in single round

How do we send/receive messages to all machines?

Say $S = n^{1+\epsilon}$, one machine needs to send n words to all M machines

$$S = n \cdot n^\epsilon$$



broadcast:

Converge Cast: All machines can send n word message to root with combination rule: union/intersection/sum

Computation output:

Output stored (in distributed fashion) on M machines

Sorting: machine holding item x knows rank/posⁿ of x

Matching: machine holding vertex v knows if v is an end-pt of matched edge (and if so, other end pt.)

Connectivity: machine holding vertex v knows id. of its connected component

If S large enough, entire output on one machine

Dense graph problems: $N = m = n^{1+c}$

Superlinear memory: $S = n^{1+\epsilon}$

[LM SV'11]: Filtering technique

MST, Matching

Minimum Spanning Tree

Idea: Partition edges into subsets of size $\leq S = n^{1+\epsilon}$
and send each subgraph to its own machine

- Each machine computes MST of local subgraph
- throw out edge if it's a heaviest edge on some cycle in local subgraph
- Any such edge is not in global MST either

Algorithm MST(V, E)

If $|E| < S$ then
compute $T^* = \text{MST}(E)$
return T^*

$l = 2|E|/S$

Partition E into $E_1 \dots E_l$ where $|E_i| < S$

using hash fn $h: E \rightarrow \{1, 2, \dots, l\}$

In parallel: compute T_i : min. spanning tree on $G(V, E_i)$

return $\text{MST}(V, \cup T_i)$

Each iteration reduces input size by n^ϵ

Lemma: Algo MST(V, E) terminates after $\lceil \frac{c}{\epsilon} \rceil$ iterations
& returns MST

Pf: Correctness (follows from earlier discussion)

random partitioning \Rightarrow wh.p. each machine gets S edges
 $\mathbb{E}[|E_i|] = \frac{S}{2}$. wh.p. $|E_i| < S$
(Chernoff)

$$|UT_i| \leq l \cdot (n-1) \leq \frac{2|E|}{S} \cdot n = O\left(\frac{|E|}{n^\epsilon}\right)$$

Terminates in $\lceil \frac{S}{\epsilon} \rceil$ rounds.

Matching:

Maximum Matching

Maximal Matching: $|\text{Maximal Matching}| \geq \frac{1}{2} |\text{Maximum Matching}|$

Idea: put subset $E' \subseteq E$ on single machine

Find maximal matching

Remove matched vertices from graph

Continue till remaining edges fit on single machine

E' : randomly selected: $|E'| \leq S$

edges decreases by $\sim n^\epsilon$ in each round

Algorithm

Machine 0 is free

Edges distributed amongst machines $1 \dots M$

$G_r(V, E_r)$: graph at round $r \in \{0, \dots, R\}$

G_0 : input graph

In each round r :

1. $m = |E_r|$

2. For $i \in \{1, \dots, M\}$ machine i marks each local edge indptly w. prob. $p = n^{1-\epsilon}/2m$

3. For $i \in \{1, \dots, M\}$ machine i sends marked edges to machine 0
4. Machine 0 computes maximal matching M_r on marked edges, announces matched vertices to machines 1 to M
5. For $i \in \{1, \dots, M\}$, machine i discards any local edge that has a matched vertex as an end pt.

Thm: Algorithm terminates in $O\left(\frac{c}{\epsilon}\right)$ rounds

Lemma: w.h.p. #marked edges fit onto single machine

$$\text{Edges sampled w. prob } p = \frac{n^{1+\epsilon}}{2m}$$

$$E[\text{\#marked edges}] = m \cdot p = \frac{n^{1+\epsilon}}{2}$$

Apply Chernoff.

Lemma: w.h.p. #remaining edges $\leq \frac{4m}{n^\epsilon}$

Pf: I : unmatched vertices at end of round

No marked edge between any pair of vertices in I

If \exists marked edge $\{u, v\}$, at least one of u, v will be matched. Contradiction

Consider arbitrary subset J of vertices

with $\geq \frac{4m}{n^\epsilon}$ induced edges

$$P[\text{All induced edges in } J \text{ unmarked}] \leq (1-p)^{\frac{4m}{n^\epsilon}} \quad \left[\text{Recall } p = \frac{n^{1+\epsilon}}{2m} \right]$$

$$\leq e^{-p \cdot \frac{4m}{n^\epsilon}} = e^{-2n}$$

Union bound over all such sets J (at most 2^n such)

$$Pr[\exists \text{ set with } \geq \frac{4m}{n^\epsilon} \text{ induced edges, s.t. all unmarked}] \leq 2^n \cdot e^{-2n}$$

\Rightarrow At most $\frac{4m}{n^\epsilon}$ remaining edges at end of round w.h.p.

Lemma: Algo terminates in $R = O(\frac{c}{\epsilon})$ rounds

$O(\log n)$ rounds with near linear memory

[Gheffan et al '18]: $O(1)$ approx in $O(\log \log n)$ rounds
w. near linear memory

[Gheffan, Utto '19]: $O(1)$ approx in $\tilde{O}(\sqrt{\log \Delta})$
with n^α memory $\alpha \in (0, 1)$

$O(1)$ approx can be turned into $(1+\epsilon)$ approx
using ideas of [McGregor '05]

$\sim (\frac{1}{\epsilon})^{1/\epsilon}$ overhead.