

Will look at many graph algorithms from parallel perspective

- Connectivity
- Shortest Paths
- Matchings
- Densest Subgraph
- Maximal Ind. Set

Main question(s) as usual: can we design work-efficient algorithms?


- Can we make the algorithms practical and theoretically-efficient?

Start with graph connectivity.

Input: $G(V, E)$ (undirected), $s, t \in V$

Let $n = |V|$, $m = |E|$

Output: Are s, t connected in G ?



Usual idea: compute components of each $v \in V$ $C[v]$

Check if $C[s] = C[t]$.

Sequentially, easy to solve in $O(n+m)$ time, eg using BFS/DFS.

→ Also practical in essentially linear time using union-find.

Parallel BFS: Understand some basic ideas in graph algs + impls using BFS.

BFS(G, r): ^{not nec. undirected}

- compute a BFS tree rooted at r i.e. compute an array of Parents, $P[v]$

s.t. $P[v] = \text{parent of } v \text{ in BFS tree } (v \neq r)$

$P[r] = r$

How do you solve BFS in parallel?

One way: just emulate sequential BFS and do it level-by-level.

• edge-map: primitive for traversal (useful beyond BFS).

edge-map(U, u, f) =

perform $i \in [0: |U|]$

$N[i] = \{v \in N^+(U[i]) \mid f(u, v)\}$

return Flatten(N)


- Take in a subset of vertices U , and an update function f .

- Return all vertices $v \in V$ s.t.

- $(u, v) \in E$

- $u \in U$

- $f(u, v) = \text{true}$.



Output = [3, 7, 3, 8]

• Usually ensure the output of edge-map is a set by ensuring that a vertex $v \in N(u)$ is atomically acquired by only one vertex in U .

• Note: This flatten-based impl is not very practical. Next Tuesday's lecture will discuss more practical ideas.

Work: $O(|U| + \sum_{u \in U} \text{deg}^+(u))$

Depth: $O(\log(n))$

Back to BFS: we will need the test-and-set primitive from HW1.

• Test-And-Set(x) (TS): take a reference to memory location x , check if $\text{value}(x) = 0$ and if so atomically set it to 1, returning 0; if $\text{value}(x) = 1$, it returns 1.

BFS output: parents array parent s.t. the tree formed by all $(u, \text{parent}[u])$ edges $\forall u \in V, u \neq r$ is a valid BFS tree

i.e. any non-tree edge (v, w) either lies in the same level, or between adjacent levels.

BFS(G, r) =

$F \leftarrow \{r\}$ // initially only root in the frontier

perform $u \in V$:

$X[u] = 0$

$P[u] = u$

$X_r \leftarrow 1$ // mark the root as visited

while ($|F| > 0$):

$F \leftarrow \text{edge-map}(G, F, \lambda(u, v).$

if ($\text{test-and-set}(X[u]) == 0$) then $P[u] = u$ and return true;

return false;

return P

Def: The shortest path between $u, v \in V$ in G is denoted $\delta_G(u, v)$

Def: the radius of a vertex $u \in V$ is $\max_{v \in V} \delta_G(u, v) = \text{radius}_G(u)$

Def: the diameter of a graph G ($\text{diam}(G)$) is $\max_{v \in V} \text{radius}_G(v)$

- The parallel BFS algorithm runs in $O(n+m)$ work and $O(\text{diam}(G) \log n)$ depth.

Unfortunately, in general no good bounds on $\text{diam}(G)$ in RW graphs.

Often small, so BFS-based approaches can be good, but many important graphs have high diameter.

- Currently we don't know how to ^{efficiently} solve BFS (or ^{many} other ^{related} SSSP problems) in $\text{polylog}(n)$ depth in general graphs.

What about graph probs. we used BFS for?

- For many probs, another powerful technique is applicable.

Low-Diameter Decomposition: (LDD)

• idea: ^{instead of searching from a single vertex,} break graph up into some number of connected clusters s.t.

(1) few edges are cut

(2) internal diameters of each cluster are bounded (ensures

that clusters can be efficiently explored in parallel; and

generally a useful alg. property that algorithms using LDD can exploit).


More formally:

A (β, d) -decomposition for $0 < \beta < 1$ is a partition of V into

clusters V_1, \dots, V_k s.t.

- Shortest path between any $u, v \in V_i$ using only vertices in V_i is at most d (strong diameter)
- The # edges $(u, v) \in E$ s.t. $u \in V_i, v \in V_j$ and $i \neq j$ is at most $\beta \cdot m$ (few inter-component edges).

pictorially:



$$\# \text{intercomp} \leq \beta \cdot m$$

Apps include

- connectivity, spanners, hop-sets, LSSTs (useful in sparse system solvers).
- First used in Dist. Comp; apps in metric embedding, linear system solves, parallel algo.


Sequentially, easy to compute using "ball growing"

- pick an uncovered vertex v
- repeatedly grow by BFS'ing from v until # edges incident to current frontier is $\leq \beta$ fraction of the # internal edges. Remove this ball from G once this condition is met.

• Not hard to see that b.c. of the stopping condition, the # cut edges is guaranteed to be $\leq \beta m$. (few inter-comp edges)

• What about diameter?

- Each new BFS level is added b.c. stopping criteria not yet applicable; i.e. more than β -fraction of internal edges on the boundary




$$\# \text{boundary} > \beta m' \Rightarrow m'' \geq m' + \beta m' = (1 + \beta)m'$$

\Rightarrow in each BFS level, the num. of internal nodes grows by a $(1 + \beta)$ factor.

$$\text{This can happen at most } \log_{1+\beta} m = O\left(\frac{\log m}{\beta}\right)$$

\Rightarrow Seq alg gives a $(\beta, O(\log n / \beta))$ decomposition

using $O(n + m)$ time.



Since the balls are inspected one after the other, the algorithm (in the worst case) is fully sequential.

Bad example for seq?

Good example for seq?

Next: we will see a fully parallel, work-efficient alg. by Miller, Peng, Xu [MPX]

MPX gives a $(\beta, O(\log n / \beta))$ -decomp in

$O(m+n)$ expected work and $O(\log^3 n)$ depth whp.

Main idea is to grow multiple balls (carefully) in parallel, from different vertices. We will mimic the sequential process, but carefully handle conflicts.

• Challenge (1): which balls do we grow in parallel?

Idea: use "exponentially-shifted" start times to start growing from v .


• Challenge (2): how to prove few-cut-edges property? strong diameter?

Idea: use properties of exp. dist.

Exp dist quick refresher:

- X non-negative C.R.V. is exponential(β) if

$$\Pr[X=x] = \beta e^{-\beta x}$$




- CDF of X is (integrate)

$$\Pr[X \leq x] = 1 - e^{-\beta x}$$

- Memorylessness property:

$$\Pr[X > a+b \mid X > a] = \Pr[X > b]$$

i.e. if we look at the tail of the distribution past a point a and rescale it so that the total prob = 1, then dists are the same.



How shall we decide when to start growing a ball at v ?

→ sample $\delta_v \sim \text{Exp}(\beta)$. If we use δ_v as start times, then almost all vertices start in the first few timesteps

Instead, use "shifted" start times: $T_v = \delta_{\max} - \delta_v$

→ Exp dist is subtracted, so start times have a backwards exponential dist (increases over time).

In particular,

- very few vertices start in first unit of time

- exponentially growing number start on each subsequent timestep.

Which ball should v be assigned to?

$$C_u \leftarrow \arg \min_{v \in V} (T_v + d(u, v))$$

shifted distance.

- Ignoring whether this generates a (β, d) -decomp, can we impl efficiently? Turns out to be relatively easy, using simultaneous BFS (aka multi-BFS).

parfor $v \in V$: $S_v \leftarrow \text{Exp}(\beta)$

$$\delta_{\max} \leftarrow \max_{v \in V} S_v$$

parfor $v \in V$:

$$T_v \leftarrow \delta_{\max} - S_v$$

$C_v \leftarrow \text{unvisited}$ // initially all vertices unvisited

$\gamma_v \leftarrow \infty$ // how the balls compete to see which one wins a vertex.
(based on earliest T_v)

$l \leftarrow 0$ // # visited / started balls

$r \leftarrow 1$

while ($l < |V|$):

$$F \leftarrow F \cup \{v \in V \mid (T_v < r) \wedge (C_v == \text{unvisited})\}$$

vertices "wake up" and start a ball if $\lfloor T_v \rfloor$ steps elapsed and not yet visited by another BFS / ball.

$$l \leftarrow l + |F|$$

edge-map($G, F, \lambda(u, v)$).

if ($C_v == \text{unvisited}$)

$$c \leftarrow C_u$$

WriteMin($\gamma_v, T_c - \lfloor T_c \rfloor$)

fractional part of T_c

$F \leftarrow \text{edge-map}(G, F, \lambda(u, v))$

$$r \leftarrow r + 1$$

$$c \leftarrow C_u$$

if ($\gamma_v == T_c - \lfloor T_c \rfloor$)

$$C_v \leftarrow c$$

return true

return false

return C

Let's first analyze work/span:

- # rounds (r) $\leq \lceil \delta_{\max} \rceil$ since all vertices will be either removed or added by this point.
- Sorting centers by start times easy to do in parallel in $O(n)$ work and $O(\log n + \delta_{\max})$ depth (eg using parallel counting sort

(see earlier lecture:

n ints in $[0:m]$ in
 $O(n)$ work, $O(\log n + m)$ span.
 \uparrow
 δ_{\max}

- Each vertex appears in the frontier once, so each edge visited at most twice

\Rightarrow $\text{total work is } O(n+m)$

- Each round has $O(\log n)$ depth. $\text{Therefore overall depth is } O(\delta_{\max} \log n)$

Let's bound δ_{\max} .

- What is the probability any vertex picks a shift $> \frac{c \log n}{\beta}$?

$$P\left[\delta_v > \frac{c \log n}{\beta}\right] = 1 - \underbrace{P\left[\delta_v \leq \frac{c \log n}{\beta}\right]}_{\text{Exp CDF}} = 1 - (1 - e^{-c \log n}) = \frac{1}{n^c}.$$

Now, take the union-bound over all vertices: Prob. that any vertex has $\delta_v > \frac{c \log n}{\beta}$

$$\text{i.e. } \Pr\left[\delta_{\max} > \frac{c \log n}{\beta}\right] \leq \frac{1}{n^{c-1}}$$


$$\Rightarrow \delta_{\max} \leq \frac{c \log n}{\beta} \text{ whp.}$$

Therefore overall depth is $O\left(\frac{\log^2 n}{\beta}\right)$ whp and $O(\log^2 n)$ whp for constant β .

This also shows the strong diameter result cluster radius bounded by $2\delta_{\max} = O\left(\frac{\log n}{\beta}\right)$ whp.

Next: only B_m edges cut in expectation.


Actually show something a bit stronger: the first k balls that arrive at an edge (its "midpoint") or at a vertex are unlikely to arrive within a small window of time.



Lemma: For any vertex or midpoint of an edge, The probability that the smallest and k -th smallest values of

$$\left\{ \underbrace{T_v}_{\text{starttime}} + \underbrace{d(v,u)}_{\text{dist}} : v \in V \right\}$$

differ by less than α is $\leq (\alpha\beta)^{k-1}$



$$T_v = \delta_{\max} - \delta_v$$

$$D_v(u) = T_v + d(v,u)$$


$$\hat{D}_v(u) = \delta_{\max} - D_v(u) = \delta_v - d(v,u)$$

$$\arg \min_{v \in V} \left(\delta_{\max} - \delta_v + \overbrace{d(v,u)}^{D_v(u)} \right)$$

$$\arg \max_{v \in V} \left(\delta_v - d(v,u) \right)$$

Consider the case of first and second balls. Analyze this like a "horse race" from all vertices to u .


- length the horse runs is δ_v . After δ_v units of time, it gets tired and stops.
- horse is penalized by $d(v, u)$



$$i = \operatorname{argmax}_{v \in V} \hat{D}(v)$$

Goal is to run as far from 0 as possible.

For an edge (u, v) $\hat{D}_i - \hat{D}_j < 1$ is exactly the event that (u, v) is cut.



We can analyze this more generally for the top k winners:

Consider running the ball-growing process "backwards" in time. i.e. we gradually shrink balls, with points that start later being removed first.

- Since we subtract exp dist to get start time, if we look at time going backwards, the distributions are proper exp. distr.

Consider removing the the k -th ball, and consider the time t at which it uncovers u . $k-1$ balls still remain, covering u .

By memorylessness property of exp dist, they each are still distributed as $\text{Exp}(\beta)$, starting at t . Therefore the prob. any one of them is removed within a units of time is the CDF:

$$1 - e^{-\beta a} < \beta a$$

$$1 - e^{-x} < x \quad \text{since } e^{-x} = 1 - x + \dots$$

The prob. that all $k-1$ other balls removed within a is therefore $< (\beta a)^{k-1}$.

For $k=2$, this is exactly the probability an edge is cut
 $a=1$

In forward time, if first ball arrives at time t , next ball at time $> t+1$ (at MP) then this edge is intra-cluster.

$$\Rightarrow \text{prob edge is inter-cluster} \leq (\beta a)^{2-1} = \beta$$

By lin. of expectations, # inter-cluster edges $\leq \beta m$.

KC: $(1, 2)$ -ND
 (r, s) -ND

$(1, 2)$ and above is P -complete
(for any constant $c \geq 3$)

ATLAS
vector retrieval models

$c' = (r, s)$ P -complete for which c ?

$$1 < r < s$$

$c' = (2, 3)$ -ND

map
k-core
directly

"u"

"v"