14 May 2021

Randomized Min Cut

<u>Announcement.</u> Final exam scheduling poll

Alt time slots on May 20–24.

Link to poll is in a pinned post on Ed.

Responses due Saturday (tomorrow) night, 23:59.

No response $\Longrightarrow$ May 25, 9:30 a.m.

Today is last day to drop courses.

Global Min Cut in Undirected Graphs.

Undirected $G = (V, E)$.

All edges have capacity 1.

<u>Goal.</u> Find a partition $V = A \cup B$ with A and B non-empty such that $c(A, B)$ is minimized.

Distinct from the minimum s-t cut problem where we require $s \in A$, $t \in B$.

<u>Example application</u>. You want to make sure that your fuel pipeline network would remain connected even if a cyberattacker shut down k pipelines....
compute global min cut, test if $c(A,B) > k$.

One way to solve: reduce to many instances of s-t min cut.

<u>Ex.</u> for all s in V:
    for all t in $V \setminus \{s\}$:
        compute s-t min cut
    output best cut discovered.

This takes $n \cdot (n-1) \cdot T(\text{min cut})$.
If m = # edges, n = # vertices,
Ford-Fulk computes one s-t min cut in $O(mn) \Longrightarrow$ running time $O(mn^3)$.

Faster: Choose any s in V,
    For all $t \in V \setminus \{s\}$:
        compute s-t min cut
    Output best cut discovered.

Improves running time to $O(mn^2)$.

Karger's Randomized Contraction Algorithm

A multigraph is a graph that may have 2 or more edges between the same pair of vertices. 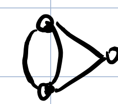

If G is a multigraph and e is an edge of G, $e = (u,v)$, then

$\mathrm{CONTRACT}(G,e)$ is an operation that outputs new graph G' with

$$V(G') = \big(V(G) \setminus \{u,v\}\big) \cup \{w\}$$

$$\begin{aligned} E(G') = &\big\{ (x,y) \in E(G) \;\big|\; x,y \notin \{u,v\} \big\} \\ &\cup \big\{ (x,w) \;\big|\; (x,u) \in E(G),\ x \notin \{u,v\} \big\} \\ &\cup \big\{ (x,w) \;\big|\; (x,v) \in E(G),\ x \notin \{u,v\} \big\} \end{aligned}$$

union of multisets, adding multiplicities

In words: merge u,v together into a new node w. The neighbors of w are all the neighbors of u,v combined. If G has edges between u & v delete them.

G 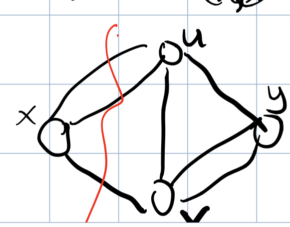

$\Longrightarrow$ G' 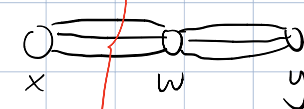

Karger's Algorithm

repeat s times: // s is a parameter that determines success probability

let H = fresh copy of G
while H has more than 2 vertices:
  sample random edge $e \in E(H)$
  replace H with CONTRACT(H, e)
// now H has 2 vertices.
let a, b be the 2 vertices of H.
let $A = \{$vertices of G contracted to $a\}$
let $B = \{$ — — — — — $b\}$
store (A, B) in our collection of potential min cuts.

output the cut with smallest capacity among potential min-cuts discovered.

Implementation: Contract(H, e) can be done in $O(\log n)$ time per operation using union-find to implicitly store $V(H)$.
(Find each endpoint of e, take union of those 2 sets.)

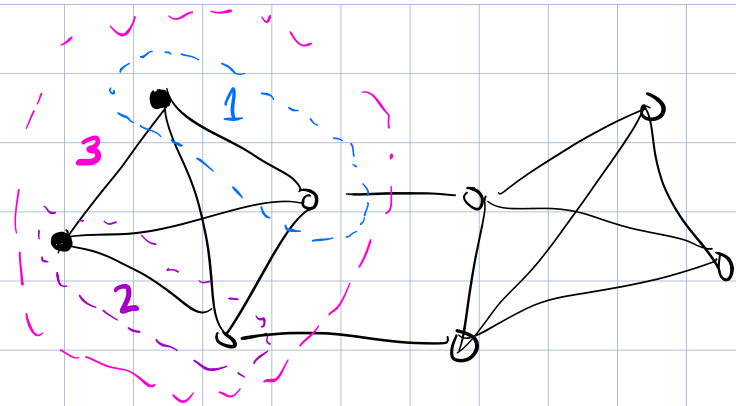

One iteration of main loop: $O(n \log n)$.
All s iterations: $O(n s \log n)$.

Success probability? The hope is that we make it through $n-1$ iterations of edge contraction without ever contracting an edge of the min cut, (A^*, B^*).

Let $k = c(A^*, B^*)$.
Consider an iteration where we have a contracted graph H with m_H edges, n_H vertices, and suppose that no edge from A^* to B^* is yet contracted.
$\Longrightarrow$ min cut capacity of $H = k$.

Degree of every vertex in $H \geq k$.

$$2 \cdot m_H = \sum \text{degrees of vertices in } H \geq n_H \cdot k$$

$$\frac{k}{m_H} \leq \frac{2}{n_H}.$$

Prob of randomly sampling an edge from A^* to B^*.

Prob (an edge from A^* to B^* is not sampled)

$$\geq \frac{n_H - 2}{n_H}.$$

What is the probability of getting from n vertices down to 2 without ever contracting an edge from A^* to B^*?

It's at least

$$\left(\frac{n-2}{n}\right) \cdot \left(\frac{\cancel{n-3}}{n-1}\right) \cdot \left(\frac{\cancel{n-4}}{\cancel{n-2}}\right) \cdot \left(\frac{\cancel{n-5}}{\cancel{n-3}}\right) \cdots \left(\frac{1}{\cancel{3}}\right)$$

$$= \frac{2}{n(n-1)}.$$

Each outer loop iteration succeeds

in finding (A^*, B^*) with prob at least $\frac{1}{\binom{n}{2}}$.

Repeat outer loop $s = \binom{n}{2} \cdot \ln(n)$ times.

$$\Pr(\text{none of the iterations succeed})$$

$$\leq \left(1 - \frac{1}{\binom{n}{2}}\right)^{\binom{n}{2} \cdot \ln(n)} \qquad 1 - x \leq e^{-x}$$

$$< e^{-\left(\cancel{\frac{1}{\binom{n}{2}}}\right) \cdot \cancel{\binom{n}{2}} \cdot \ln(n)}$$

$$= e^{-\ln(n)} = \frac{1}{n}.$$

Running time $O(s \cdot n \cdot \log n) = O(n^3 \log^2 n)$.

Karger & Stein found a beautiful divide-conquer way to improve the running time by factor n.