Lecture 16

Min Cut and Karger's Algorithm

Announcements

- HW8 Due Wednesday
- Midterm 4 on Mon-Tue next week.
 - Cannot be dropped
 - Material covered this week is included

Last time

- Minimum Spanning Trees!
 - Prim's Algorithm
 - Kruskal's Algorithm

Today

- Minimum Cuts!
 - Karger's algorithm
 - Karger-Stein algorithm
 - Back to randomized algorithms!

*For today, all graphs are **undirected and unweighted**.

Recall: cuts in graphs

• A cut is a partition of the vertices into two nonempty parts.



Recall: cuts in graphs

• A cut is a partition of the vertices into two nonempty parts.

Part 2

This is not a cut



This is a cut



This is a cut

These edges cross the cut.

• They go from one part to the other.



A (global) minimum cut is a cut that has the fewest edges possible crossing it.



A (global) minimum cut is a cut that has the fewest edges possible crossing it.



Why might we care about global minimum cuts?



12 *For the rest of today edges aren't weighted; but the algorithm can be adapted to deal with edge weights.

- Finds **global minimum cuts** in undirected graphs
- Randomized algorithm
- Karger's algorithm **might be wrong**.
 - Compare to QuickSort, which just might be slow.
- Why would we want an algorithm that might be wrong?
 - With high probability it won't be wrong.
 - Maybe the stakes are low and the cost of a deterministic algorithm is high.

Different sorts of gambling

- QuickSort is a Las Vegas randomized algorithm
 - It is always correct.
 - It might be slow.

Yes, this is a technical term.



Different sorts of gambling

- Karger's Algorithm is a Monte Carlo randomized algorithm
 - It is always fast.
 - It might be wrong.



- Pick a random edge.
- Contract it.
- Repeat until you only have two vertices left.



Why is this a good idea? We'll see shortly.





































a,b,c,d

Now stop!

• There are only two nodes left.

The **minimum cut** is given by

the remaining super-nodes:

• {a,b,c,d} and {e,h,f,g}

e,h,g,f

The **minimum cut** is given by the remaining super-nodes:

• {a,b,c,d} and {e,h,f,g}



• Does it work?



How do we implement this?

- See Lecture 16 IPython Notebook for one way
 - This maintains a secondary "superGraph" which keeps track of superNodes and superEdges
 - There's a skipped slide with pseudocode
- Running time?
 - We contract n-2 edges
 - Each time we contract an edge we get rid of a vertex, and we get rid of n 2 vertices total.
 - Naively each contraction takes time O(n)
 - Maybe there are Ω(n) nodes in the superNodes that we are merging. (We can do better with fancy data structures).
 - So total running time O(n²).
 - We can do $O(m \cdot \alpha(n))$ with a union-find data structure, but $O(n^2)$ is good enough for today.

• Does it work?



Think-share! 1 minute think 1 minute share



- Does it work? No?
- Is it fast? Create a supernode! • O(n²) b g а Create a superedge! e С Create a d superedge! h Algorithm: Randomly contract edges until there are only 37 two supernodes left.
Why did that work?

- We got really lucky!
- This could have gone wrong in so many ways.



Karger's algorithm

Say we had chosen this edge



Karger's algorithm

Say we had chosen this edge

Now there is **no way** we could return a cut that separates b and e.



Even worse

If the algorithm EVER chooses either of these edges, it will be wrong.



How likely is that?



• For this particular graph, I did it 10,000 times:



That doesn't sound good

• Too see why it's good after all, we'll first do a case study of this graph. Then we'll generalize.

The plan:

- See that 20% chance of correctness is actually nontrivial.
- Use repetition to boost an algorithm that's correct 20% of the time to an algorithm that's correct 99% of the time.
- To see the first point, let's compare Karger's algorithm to the algorithm:

Choose a completely random cut and hope that it's a minimum cut. e

Uniformly random cut in 🌽

• Pick a random way to split the vertices into two parts:



Uniformly random cut in

- Pick a random way to split the vertices into two parts:
- The probability of choosing the minimum cut is*... number of min cuts in that graph number of ways to split 8 vertices in 2 parts $=\frac{2}{2^8-2} \approx 0.008$
- Aka, we get a minimum cut 0.8% of the time.

Karger is better than completely random!



What's going on?

• Which is more likely?

Thing 1: It's unlikely that Karger will hit the min cut since it's so small!



Lucky the lackadaisical lemur

B: The algorithm never chooses any of the edges in **this big cut**.

A: The algorithm never chooses either of the edges in **the minimum cut**.



• Neither A nor B are very likely, **but A is more likely than**₄₇**B**.



What's going on?

Thing 2: By only contracting edges we are ignoring certain really-not-minimal cuts.



Lucky the lackadaisical lemur





This cut actually separates the graph into three pieces, so it's not minimal – either half of it is a smaller cut.



A: This cut can be returned by Karger's algorithm.

Why does that help?

- Okay, so it's better than completely random...
- We're still wrong about 80% of the time.
- The main idea: repeat!
 - If I'm wrong 80% of the time, then if I repeat it a few times I'll eventually get it right.



Thought experiment from pre-lecture exercise

- Suppose you have a magic button that produces one of 5 numbers, {a,b,c,d,e}, uniformly at random when you push it.
- You don't know what {a,b,c,d,e} are.
- Q: What is the minimum of a,b,c,d,e?



3 5 5 3 2 2

How many times do you have to push the button, in expectation, before you see the minimum value?

What is the probability that you have to push it more than 5 times? 10 times?



In this context

- Run Karger's! The cut size is 6!
- Run Karger's! The cut size is 3!
 - Run Karger's! The cut size is 3!
- Run Karger's! The cut size is 2! 🔶
 - !

Correct!



• Run Karger's! The cut size is 5!

If the success probability is about 20%, then if you run Karger's algorithm 5 times and take the best answer you get, that will likely be correct! (with probability about 0.67)

For this particular graph

- a b f g c d e h
- Repeat Karger's algorithm about 5 times, and we will get a min cut with decent probability.
 - In contrast, we'd have to choose a random cut about 1/0.008 = 125 times!

Hang on! This "20%" figure just came from running experiments on this particular graph. What about general graphs? Can we prove something?



Also, we should be a bit more precise about this "about 5 times" statement.

Plucky the pedantic penguin





1. What is the probability that Karger's algorithm returns a minimum cut in a general graph?

- 2. How many times should we run Karger's algorithm to "probably" succeed?
 - Say, with probability 0.99?
 - Or more generally, probability $1-\delta$?

Answer to Question 1

Claim:

The probability that Karger's algorithm returns a minimum cut on a graph with n vertices is

at least
$$\frac{1}{\binom{n}{2}}$$



In this case, $\frac{1}{\binom{8}{2}} = 0.036$, so we are guaranteed to win at least 3.6% of the time.

Questions



1. What is the probability that Karger's algorithm returns a minimum cut?

According to the claim, at least $\frac{1}{\binom{n}{2}}$

- 2. How many times should we run Karger's algorithm to "probably" succeed?
 - Say, with probability 0.99?
 - Or more generally, probability 1δ ?

Before we prove the Claim

2. How many times should we run Karger's algorithm to succeed with probability $1 - \delta$?



A computation

Punchline: If we repeat $\mathbf{T} = \binom{n}{2} \ln(1/\delta)$ times, we win with probability at least $1 - \delta$.

- Suppose :
 - the probability of successfully returning a minimum cut is $p \in [0, 1]$,
 - we want failure probability at most $\delta \in (0,1)$.
- Pr[don't return a min cut in T trials] = $(1 p)^T$
- The claim says $p = 1/\binom{n}{2}$. Let's choose $T = \binom{n}{2} \ln(1/\delta)^{-1}$
- Pr[don't return a min cut in T trials]
 - = $(1 p)^T$
 - $\leq (e^{-p})^T$
 - = e^{-pT}

•
$$= e^{-\ln\left(\frac{1}{\delta}\right)}$$

• = δ

 e^{-p} 1-p 1

 $1 - p \le e_{58}^{-p}$





1. What is the probability that Karger's algorithm returns a minimum cut?

According to the claim, at least $\frac{1}{\binom{n}{2}}$

- 2. How many times should we run Karger's algorithm to "probably" succeed?
 - Say, with probability 0.99?
 - Or more generally, probability 1δ ?

 $\binom{n}{2}\ln\left(\frac{1}{\delta}\right)$ times.

Theorem

Assuming the claim about $1/\binom{n}{2}$...

- Suppose G has n vertices.
- Consider the following algorithm:
 - bestCut = None
 - for $t = 1, ..., {n \choose 2} ln\left(\frac{1}{\delta}\right)$:
 - candidateCut ← Karger(G)
 - if candidateCut is smaller than bestCut:
 - bestCut ← candidateCut
 - return bestCut

How many repetitions would you need if instead of Karger we just chose a uniformly random cut?

60

• Then Pr[this doesn't return a min cut] $\leq \delta$.

What's the running time?

• $\binom{n}{2} \ln \left(\frac{1}{\delta}\right)$ repetitions, and O(n²) per repetition. • So, $O\left(n^2 \cdot \binom{n}{2} \ln \left(\frac{1}{\delta}\right)\right) = O(n^4)$ Treating δ as constant.

> Again we can do better with a union-find data structure. Write pseudocode for—or better yet, implement—a fast version of Karger's algorithm! How fast can you make the asymptotic running time?



Ollie the over-achieving ostrich

Theorem Assuming the claim about $1/\binom{n}{2}$...

Suppose G has n vertices. Then [repeating Karger's algorithm a bunch of times] finds a min cut in G with probability at least 0.99 in time O(n⁴).

Now let's prove the claim...

Break

Claim

The probability that Karger's algorithm returns a minimum cut in a graph with n vertices is

at least
$$\frac{1}{\binom{n}{2}}$$

. . .

- Suppose the edges that we choose are e₁, e₂, ..., e_{n-2}
- **PR**[return S*] = **PR**[none of the e_i cross S*]
 - = $PR[e_1 \text{ doesn't cross } S^*]$ $\times PR[e_2 \text{ doesn't cross } S^* | e_1 \text{ doesn't cross } S^*]$
 - × **PR**[e_{n-2} doesn't cross S* | e_1 ,..., e_{n-3} don't cross S*]



Focus in on: **PR**[e_j doesn't cross S* | e₁,...,e_{j-1} don't cross S*]

- Suppose: After j-1 iterations, we haven't messed up yet!
- What's the probability of messing up now?



Focus in on: **PR**[e_j doesn't cross S* | e₁,...,e_{j-1} don't cross S*]

- Suppose: After j-1 iterations, we haven't messed up yet!
- What's the probability of messing up now?
- Say there are k edges that cross S*
- Every supernode has at least k (original) edges coming out.
 - Otherwise we'd have a smaller cut.
- Thus, there are at least (n-j+1)k/2 edges total.
 - b/c there are n j + 1 supernodes left, each with at least k edges.

So the probability that we choose one of the k edges crossing S* at step j is at most:

$$\frac{k}{\left(\frac{(n-j+1)k}{2}\right)} = \frac{2}{n-j+1}$$



Focus in on: **PR**[e_j doesn't cross S* | e₁,...,e_{j-1} don't cross S*]

 So the probability that we choose one of the k edges crossing S* at step j is at most:

$$\frac{k}{\left(\frac{(n-j+1)k}{2}\right)} = \frac{2}{n-j+1}$$

• The probability we **don't** choose one of the k edges is at least:

. . .

- Suppose the edges that we choose are e₁, e₂, ..., e_{n-2}
- **PR**[return S*] = **PR**[none of the e_i cross S*]
 - = $PR[e_1 \text{ doesn't cross } S^*]$ $\times PR[e_2 \text{ doesn't cross } S^* | e_1 \text{ doesn't cross } S^*]$
 - × **PR**[e_{n-2} doesn't cross S* | e_1 ,..., e_{n-3} don't cross S*]



- Suppose the edges that we choose are e_1 , e_2 , ..., e_{n-2}
- **PR**[return S*] = **PR**[none of the e_i cross S*]

$$= \left(\frac{n-2}{n}\right) \left(\frac{n-3}{n-1}\right) \left(\frac{n-4}{n-2}\right) \left(\frac{n-5}{n-3}\right) \left(\frac{n-6}{n-4}\right) \cdots \left(\frac{4}{6}\right) \left(\frac{3}{5}\right) \left(\frac{2}{4}\right) \left(\frac{1}{3}\right)$$



- Suppose the edges that we choose are e₁, e₂, ..., e_{n-2}
- PR[return S*] = PR[none of the e_i cross S*]



Theorem Assuming the claim about $1/\binom{n}{2}$...

Suppose G has n vertices. Then [repeating Karger's algorithm a bunch of times] finds a min cut in G with probability at least 0.99 in time O(n⁴).

That proves this Theorem!

What have we learned?

- If we randomly contract edges:
 - It's unlikely that we'll end up with a min cut.
 - But it's not **TOO** unlikely
 - By repeating, we likely will find a min cut.

```
Here I chose \delta = 0.01 just for concreteness.
```

- Repeating this process:
 - Finds a global min cut in time O(n⁴), with probability 0.99.
 - We can run a bit faster if we use a **union-find** data structure.

More generally

- If we have a Monte-Carlo algorithm with a small success probability,
- and we can check how good a solution is,
- Then we can **boost** the success probability by repeating it a bunch and taking the best solution.


Can we do better?

- Repeating O(n²) times is pretty expensive.
 - O(n⁴) total runtime to get success probability 0.99.
- The Karger-Stein Algorithm will do better!
 - The trick is that we'll do the repetitions in a clever way.
 - O(n²log²(n)) runtime for the same success probability.
 - Warning! This is a tricky algorithm! We'll sketch the approach here: the important part is the high-level idea, not the details of the computations.

To see how we might save on repetitions, let's run through Karger's algorithm again.



Probability that we didn't mess up: 12/14

There are 14 edges, 12 of which are good to contract.







Probability that we didn't mess up: 11/13

Now there are only 13 edges, since the edge between a and b disappeared.







Probability that we didn't mess up: 10/12

Now there are only 12 edges, since the edge between e and h disappeared.







Probability that we didn't mess up: 9/11











Probability that we didn't mess up: 3/5







Now stop!

• There are only two nodes left.



Probability of not messing up

- At the beginning, it's pretty likely we'll be fine.
- The probability that we mess up gets worse and worse over time.



Moral: Repeating the stuff from the beginning of the algorithm is wasteful!



In words

- Run Karger's algorithm on G for a bit. • Until there are $\frac{n}{\sqrt{2}}$ supernodes left.
- Then split into two independent copies, G₁ and G₂
- Run Karger's algorithm on each of those for a bit.
 - Until there are $\frac{\left(\frac{n}{\sqrt{2}}\right)}{\sqrt{2}} = \frac{n}{2}$ supernodes left in each.
- Then split each of those into two independent copies...

In pseudocode

- KargerStein(G = (V,E)):
 - n ← |V|
 - if n < 4:
 - find a min-cut by brute force \\ time O(1)
 - Run Karger's algorithm on G with independent repetitions until $\left|\frac{n}{\sqrt{2}}\right|$ nodes remain.
 - G₁, G₂ ← copies of what's left of G
 - S₁ = KargerStein(G₁)
 - S₂ = KargerStein(G₂)
 - **return** whichever of S₁, S₂ is the smaller cut.



Recursion tree

- depth is $\log_{\sqrt{2}}(n) = \frac{\log(n)}{\log(\sqrt{2})} = 2\log(n)$
- number of leaves is $2^{2\log(n)} = n^2$



Two questions

- Does this work?
- Is it fast?



- The amount of work per level is the amount of work needed to reduce the number of nodes by a factor of $\sqrt{2}$.
- That's at most O(n²).
 - since that's the time it takes to run Karger's algorithm once, cutting down the number of supernodes to two.
- Our recurrence relation is... $T(n) = 2T(n/\sqrt{2}) + O(n^{2})$

The Master Theorem says... $T(n) = O(n^2 \log(n))$

Jedi Master Yoda

Two questions

• Does this work?



- Is it fast?
 - Yes, O(n²log(n)).

Why $n/\sqrt{2}$?

...

Suppose we contract n – t edges, until there are t supernodes remaining.

Suppose the first n-t edges that we choose are

 $e_{1}, e_{2}, ..., e_{n-t}$ • PR[none of $e_{1}, e_{2}, ..., e_{n-t} \operatorname{cross} S^{*}]$ $= PR[e_{1} \operatorname{doesn't} \operatorname{cross} S^{*}]$ $\times PR[e_{2} \operatorname{doesn't} \operatorname{cross} S^{*} | e_{1} \operatorname{doesn't} \operatorname{cross} S^{*}]$

× **PR**[e_{n-t} doesn't cross S* | e₁,...,e_{n-t-1} don't cross S*]

Why
$$n/\sqrt{2}$$
 ?

Suppose we contract n – t edges, until there are t supernodes remaining.

Suppose the first n-t edges that we choose are

$$PR[\text{ none of } e_1, e_2, ..., e_{n-t} \operatorname{cross} S^*]$$

$$= \left(\frac{n-2}{n}\right) \left(\frac{n-3}{n-1}\right) \left(\frac{n-4}{n-2}\right) \left(\frac{n-5}{n-3}\right) \left(\frac{n-6}{n-4}\right) \cdots \left(\frac{t+1}{t+3}\right) \left(\frac{t}{t+2}\right) \left(\frac{t-1}{t+1}\right)$$

$$= \frac{t \cdot (t-1)}{n \cdot (n-1)} \quad \text{Choose } t = n/\sqrt{2}$$

$$= \frac{\frac{n}{\sqrt{2}} \cdot \left(\frac{n}{\sqrt{2}} - 1\right)}{n \cdot (n-1)} \approx \frac{1}{2} \quad \text{when n is large}$$

e₁, **e**₂, ..., **e**_n,





The problem we need to analyze

- Let T be binary tree of depth 2log(n)
- Each node of T succeeds or fails independently with probability 1/2
- What is the probability that there's a path from the root to any leaf that's entirely successful?
- It turns out that this is $\Omega\left(\frac{1}{\log n}\right)$.
 - See skipped slides for proof, or try to do it yourself!
 - (Proof not covered on exam, but it's good practice with recurrence relations!)



Analysis

- Say the tree has height d.
- Let p_d be the probability that there's a path from the root to a leaf that **doesn't fail**.



Slide skipped in class

 $2^{d/2}$

nodes

 $2^{(d-1)/2}$

Contract a

bunch of

edges

It's a recurrence relation!

•
$$p_d = p_{d-1} - \frac{1}{2} \cdot p_{d-1}^2$$

•
$$p_0 = 1$$

- We are real good at those.
- In this case, the answer is:

• Claim: for all d,
$$p_d \ge \frac{1}{d+1}$$

Recurrence relation

•
$$p_d = p_{d-1} - \frac{1}{2} \cdot p_{d-1}^2$$

• $p_0 = 1$

• Claim: for all d,
$$p_d \ge \frac{1}{d+1}$$

- **Proof**: induction on d.
 - Base case: $1 \ge 1$. YEP.
 - Inductive step: say d > 0.

• Suppose that
$$p_{d-1} \ge \frac{1}{d}$$

•
$$p_d = p_{d-1} - \frac{1}{2} \cdot p_{d-1}^2$$

•
$$\geq \frac{1}{d} - \frac{1}{2} \cdot \frac{1}{d^2}$$
•
$$\geq \frac{1}{d} - \frac{1}{d(d+1)}$$

•
$$=\frac{1}{d+1}$$



What does that mean for Karger-Stein?

Claim: for all d, $p_d \ge \frac{1}{d+1}$

- For d = 2log(n)
 - that is, d = the height of the tree:

$$p_{2\log(n)} \ge \frac{1}{2\log(n) + 1}$$

• aka,

Pr[Karger-Stein is successful] =
$$\Omega\left(\frac{1}{\log(n)}\right)$$
Altogether now



- Karger-Stein succeeds with probability $\Omega\left(\frac{1}{\log n}\right)$.
- We can amplify the success probability by repetition:
 - Run Karger-Stein $O\left(\log(n) \cdot \log\left(\frac{1}{\delta}\right)\right)$ times to achieve success probability 1δ .
- Each iteration takes time $O(n^2 \log(n))$
 - That's what we proved before.
- Choosing $\delta = 0.01$ as before, the total runtime is $O(n^2 \log(n) \cdot \log(n)) = O(n^2 \log^2(n))$

What have we learned?

- Just repeating Karger's algorithm isn't the best use of repetition.
 - We're probably going to be correct near the beginning.
- Instead, Karger-Stein repeats when it counts.
 - If we wait until there are $\frac{n}{\sqrt{2}}$ nodes left, the probability that we fail is close to $\frac{1}{2}$.
- This lets us (probably) find a global minimum cut in an undirected graph in time O(n² log²(n)).
 - Notice that we can't do better than n² in a dense graph (we need to look at all the edges), so this is pretty good.

Recap

- Some algorithms:
 - Karger's algorithm for global min-cut
 - Improvement: Karger-Stein
- Some concepts:
 - Monte Carlo algorithms:
 - Might be wrong, are always fast.
 - We can boost their success probability with repetition.
 - Sometimes we can do this repetition very cleverly.

Next time

• More min-cuts...and max flows!

Before next time

• Pre-lecture exercise: routing on rickety bridges!