

CS 761 : Randomized Algorithms , Spring 2018 , Waterloo

Lecture 4 : Approximation algorithms

We will see different ways of using randomization in designing approximation algorithms for graph problems.

The first example is graph sparsification, with a striking application in designing a near linear time exact algorithm for minimum cuts.

The second and third examples are by randomized rounding, one on linear program and one on semidefinite program, for the congestion minimization problem and the maximum cut problem.

Graph Sparsification

Given an edge weighted undirected graph $G=(V,E,w)$, for a subset of vertices $S \subseteq V$,

let $\delta_G(S)$ be the set of edges with one endpoint in S and another endpoint in $V-S$, and

let $w(\delta_G(S))$ be the total weight of the edges in $\delta_G(S)$.

We say H is a $(1 \pm \epsilon)$ -cut-approximator of G if $(1-\epsilon)w(\delta_G(S)) \leq w(\delta_H(S)) \leq (1+\epsilon)w(\delta_G(S))$

for all $S \subseteq V$. Note that H is on the same vertex set but may have different edge weights.

We are interested in finding a $(1 \pm \epsilon)$ -cut-approximator of G that is sparse (having few edges).

This is an interesting problem that leads to many beautiful results.

Today is our first time to see this problem, and we will just prove a simple first result, but as we will see this simple result already leads to a significant and surprising application.

Assumption : We consider a simple setting in which G is unweighted and has min-cut value $\Omega(\log n)$.

Algorithm: In this simple setting, the algorithm is very simple. Set a sampling probability p .

For every edge $e \in E(G)$, put it in H with weight $\frac{1}{p}$ with probability p .

The idea is to set the expectation right - so that we expect to have p fraction of edges, and hence every cut the expected weight is the same as in the original graph.

Of course, it is not enough to just have the expected values right, as we need to ensure that all cuts have approximately the same weights simultaneously, for which we will use Chernoff bound and the assumption that the min-cut value is at least $\Omega(\log n)$ (recall that $\Omega(\log n)$ is the regime that we can expect tight concentration). The following is the precise statement.

Theorem Set $p = 9 \ln n / (\epsilon^2 c)$ where c is the min-cut value of G .

Then H is a $(1 \pm \epsilon)$ -cut-approximator of G with $O(p \cdot |E(G)|)$ edges with prob $\geq 1 - \frac{4}{n}$.

Proof Consider a subset $S \subseteq V$. Say $\delta_G(S)$ has k edges. Note that $k \geq c$ by definition.

By linearity of expectation, $E[|\delta_H(S)|] = pk$ and thus $E[w(\delta_H(S))] = p \cdot k \cdot \frac{1}{p} = k = |\delta_G(S)|$.

So, the expected value is correct.

Next, we bound the probability that the actual value is "far" from the expected value.

Since each edge is an independent 0-1 random variable, by Chernoff, we have.

$$\Pr[| |\delta_H(S)| - pk | > \epsilon pk] \leq 2e^{-pk \frac{\epsilon^2}{3}} = 2e^{-\left(\frac{9 \ln n}{\epsilon^2 c}\right) \left(\frac{k \epsilon^2}{3}\right)} = 2e^{-\frac{3k \ln n}{c}}.$$

Since $k \geq c$ by definition, it follows that this probability is at most $2/n^3$.

So, the probability that one cut is violated is pretty small but there are exponentially many cuts, and a naive union bound would not work.

The important observation is that the probability that a large cut is violated is much smaller, and there are not many small cuts!

Claim: The number of cuts with at most αc edges for $\alpha \geq 1$ is at most $n^{2\alpha}$.

proof: It follows from the same analysis in the random contraction algorithm that a cut with αc edges survive with probability at least $1/n^{2\alpha}$. (We've seen the argument for $\alpha=1$) \square

Now, with the claim, we can bound

$$\begin{aligned} & \Pr(\text{some cut } S \text{ is violated}) \\ & \leq \sum_{S \subseteq V} \Pr(\text{cut } S \text{ is violated}) && // \text{union bound} \\ & = \int_{\alpha \geq 1} \sum_{S \subseteq V: |\delta_G(S)| = \alpha c} \Pr(\text{cut } S \text{ is violated}) && // \text{grouped by size} \\ & \leq \int_{\alpha \geq 1} n^{2\alpha} \cdot \Pr(\text{cut } S \text{ is violated} \mid |\delta_G(S)| = \alpha c) && // \text{by the claim} \\ & \int_{\alpha \geq 1} n^{2\alpha} \cdot 2e^{-3\alpha \ln n} && // \text{by Chernoff} \\ & \int_{\alpha \geq 1} 2n^{-\alpha} \leq 4/n. \end{aligned}$$

Therefore, with probability $\geq 1 - \frac{4}{n}$, all cuts are within $(1 \pm \epsilon)$ -factor as in the original graph.

Finally, again by Chernoff bound, it is easy to show that the number of edges in H is $O(p|E(G)|)$.

This completes the proof. \square

Remark: Just using Chernoff bound and the union bound can already solve many interesting problems.

Applications: One natural application is to design fast approximation algorithms for cut problems.

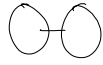
Take your favorite cut problem, say minimum s-t cut.

The running time of the algorithms usually depends on $|E|$, which could be $\Omega(N^2)$.

To speedup, we can first "sparsify" G by constructing a $(1 \pm \epsilon)$ -cut approximator H with fewer edges.

Then, a minimum s-t cut in $S \subseteq V(H)$ is a $(1 \pm 3\epsilon)$ -approximation of the minimum s-t cut in G .

This gives a tradeoff between the approximation guarantee and the running time.

Improvements: Without the minimum cut assumption, then it is easy to see that a uniform sampling algorithm won't work, e.g. in , it is very likely that the cut edge is not picked.

Benczur and Karger designed a very clever non-uniform sampling algorithm, where the sampling probability for each edge is inversely proportional to the "connectivity" of the two endpoints.

They defined a notion called "strong connectivity" and proved that sampling inversely proportional to it will result in a $(1 \pm \epsilon)$ -cut-approximator with $O(n \log n)$ edges.

Furthermore, they showed that there is an almost linear-time algorithm to estimate the strong connectivity which is good enough for the purpose, leading to an $\tilde{O}(n^2)$ -algorithm for approx min s-t cut.

The definition of "strong connectivity" is a bit unnatural, and they conjectured that it can be replaced by the more natural edge-connectivity (i.e. the max u-v flow value for edge uv).

This conjecture is proven by Fung, Hariharan, Harvey and Panigrahi in 2011.

Spectral sparsification: Spielman and Tang defined a notion called "spectral sparsifier" that is stronger than that of a "cut sparsifier". Spielman and Srivastava proved that a spectral sparsifier with $O(n \log n)$ edges always exists. The algorithm is also by random sampling, by assigning each edge a probability proportional to the effective resistance of the edge.

I plan to prove this result in the last part of the course, which has implications in solving linear equations. We won't discuss Benczur-Karger in more details.

Minimum Cut in Near Linear Time (Sketch)

Minimum Cut in Near Linear Time (Sketch)

It is an amazing result that uses graph sparsification to design faster exact algorithms.

We consider the problem of finding a global minimum cut in an undirected graph, i.e.

to remove a minimum number of edges to disconnect the input graph.

Let k be the optimal value of this problem. So, the input graph is k -edge-connected.

Karger cleverly used a classical result on spanning tree packing.

Theorem (Tutte) If a graph is k -edge-connected, then it has $\lfloor k/2 \rfloor$ edge disjoint spanning trees.

Suppose we have edge disjoint spanning trees $T_1, T_2, \dots, T_{\lfloor k/2 \rfloor}$. Then, at least one tree T that

crosses a minimum cut $S \subseteq V$ at most two times.

Karger observed that having such a T would help find a minimum cut efficiently, as one

just needs to consider the cuts formed by removing at most two edges from T .

He showed that the minimum cut over these cuts can be computed in $\tilde{O}(m)$ time using

dynamic programming (the case that T only crosses a minimum cut once is easier).

Okay, but how to find such a tree T ?

There is a known algorithm to find c edge disjoint spanning trees in $\tilde{O}(m \cdot c)$ time.

But this is not fast enough for this purpose.

The idea here, of course, is graph sparsification.

Just using the above graph sparsification result, we can sparsify the graph so that its

min-cut is $O(\log n)$ while approximately preserving all the cuts.

Now, we can compute edge-disjoint spanning trees $T_1, \dots, T_{O(\log n)}$ in the sparsifier,

and one of these trees would cross a minimum cut at most two times.

So, doing dynamic programming on each tree would work, and total complexity is $\tilde{O}(m)$.

This is just a sketch of the main ideas.

Congestion minimization

This is probably the first example of doing randomized rounding on linear programming solution.

Input : A directed or undirected graph $G=(V,E)$, K source-sink pairs (s_i, t_i)

Task : For each pair (s_i, t_i) , find a path P_i connecting s_i to t_i , so that each edge e is used by at most C paths.

Objective : Minimize the congestion C .

Integer linear program formulation

This problem is NP-complete and we will design an approximation algorithm for it.

The general approach is to formulate the problem as an integer linear program and then "relax" it.

Let \mathcal{P}_i be the set of all s_i - t_i paths in the graph. (There could be exponentially many.)

For each path $P \in \mathcal{P}_i$, we create a variable f_P^i , with the intention of setting it to be one if we choose this path to connect s_i and t_i , otherwise we set it to zero.

The congestion minimization problem can be formulated as an integer linear program.

minimize C

subject to $\sum_{P \in \mathcal{P}_i} f_P^i = 1 \quad \forall 1 \leq i \leq K$ // exactly one path in \mathcal{P}_i is chosen to connect s_i and t_i

$\sum_{i=1}^K \sum_{P \in \mathcal{P}_i: e \in P} f_P^i \leq C \quad \forall e \in E$ // each edge e is used in at most C times

$f_P^i \in \{0,1\} \quad \forall i \quad \forall P \in \mathcal{P}_i$

Since this is an exact formulation of the NP-complete problem, we do not know how to solve it in polytime.

Linear programming relaxation

So, we relax the discrete constraint $f_P^i \in \{0,1\}$ by $0 \leq f_P^i \leq 1$, and then the linear program can be solved in polynomial time. (Even though there are exponentially many variables, as there is an equivalent linear program with only polynomial number of variables. Not our focus; details omitted.)

However, in the relaxation, we allow "fractional solutions" and these may not correspond to the solution that we wanted (i.e. s_i - t_i path).

Rounding linear programming solution

Nevertheless, the optimum of this LP serves as a lower bound on the optimum of the original problem.

Suppose we could find an integral solution which has congestion at most $\alpha \cdot \text{OPT}_{LP}$, then it is a k -approximate solution to the original problem.

$$\begin{array}{c}
 \xleftarrow{\text{a factor of } \alpha} \\
 \text{---|---|---} \\
 \text{OPT}_{LP} \quad \text{OPT}_{INT} \quad \text{SOL}_{INT}
 \end{array}
 \Rightarrow \frac{\text{SOL}_{INT}}{\text{OPT}_{LP}} \leq \alpha \Rightarrow \frac{\text{SOL}_{INT}}{\text{OPT}_{INT}} \leq \alpha.$$

Randomized rounding: The idea is to interpret f_p^i as the probability of choosing path P to connect s_i and t_i . Then, the algorithm is for each i , pick one path P according to the probability f_p^i .

Theorem The maximum congestion is $O(C \ln n / \ln \ln n)$ with probability $\geq 1 - \frac{1}{n}$.

proof Let $X_e^i = \begin{cases} 1 & \text{if the path connecting } s_i \text{ and } t_i \text{ uses the edge } e. \\ 0 & \text{otherwise} \end{cases}$.

Let $X_e = \sum_{i=1}^k X_e^i$, the congestion on e .

Note that $E[X_e^i] = \sum_{P \in \mathcal{P}_i: e \in P} f_p^i$.

And thus $E[X_e] = \sum_{i=1}^k E[X_e^i] = \sum_{i=1}^k \sum_{P \in \mathcal{P}_i: e \in P} f_p^i \leq C$ by the LP constraint.

Since each X_e^i is an independent 0-1 variable, we can use Chernoff bound to obtain

$$\Pr(X > (1+\delta)C) < \left(\frac{e^\delta}{(1+\delta)^{1+\delta}}\right)^C \leq \frac{e^\delta}{(1+\delta)^{1+\delta}} \text{ as } C \geq 1.$$

If we set $\delta = \Theta(\ln n / \ln \ln n)$, then $(1+\delta) \ln(1+\delta) = \Omega(\ln n)$.

We can choose the hidden constant to be large enough so that $(1+\delta)^{1+\delta} \geq n^4$.

This implies that each edge has congestion $\Omega(C \ln n / \ln \ln n)$ with probability $\leq \frac{1}{n^3}$.

By the union bound, some edge has congestion $\Omega(C \ln n / \ln \ln n)$ is $\leq \frac{1}{n}$, as there are at most n^2 edges in the graph. \square

Open question: This is still the best approximation algorithm for the congestion minimization problem.

Can you beat it?

Maximum cut

This is the first example of doing (any) rounding of semidefinite programming solutions.

Input: An undirected graph $G=(V,E)$.

Output: A subset $S \subseteq V$ that maximizes $|\delta(S)|$, the number of edges that crosses S .

Quadratic program formulation:
$$\max \sum_{ij \in E} \frac{1}{2} (1 - x_i x_j)$$

$$x_i^2 = 1 \quad \forall i \in V.$$

The constraint $x_i^2 = 1$ forces $x_i \in \{+1, -1\}$ and the objective is counting number of edges with different signs. So, this is an exact formulation for the maximum cut problem, which is NP-complete.

Semidefinite programming relaxation

$$\max \sum_{ij \in E} \frac{1}{2} (1 - \langle y_i, y_j \rangle)$$

$$\|y_i\|_2^2 = 1 \quad \forall i \in V$$

$$y_i \in \mathbb{R}^{|V|}$$


It turns out that this is a semidefinite program and can be solved in polynomial time (again, not our focus, details omitted).

It is not difficult to see that it is a relaxation, as any ± 1 solution can be embedded as $(\pm 1, 0, 0, \dots, 0)$ in a vector solution (but not vice versa), and so $\text{OPT}_{\text{SDP}} \geq \text{OPT}_{\text{INT}}$.

To get some intuition about this SDP, we can think of the solution geometrically.

An edge contributes a large value (close to 1) to the objective has $\langle y_i, y_j \rangle$ close to -1 , meaning that the vectors y_i, y_j are almost opposite to each other.

An optimal SDP solution tries to pull the two vectors of an edge as far as possible.



(picture from Gärtner-Matoušek)


Intuitively, we should use the vector solution as a guideline: we should cut an edge with a probability proportional to $1 - \langle y_i, y_j \rangle = 1 - \cos \theta_{ij}$, where θ_{ij} is the angle between y_i and y_j . If θ_{ij} is larger (closer to π), we should cut it with higher probability.

Randomized rounding: Goemans and Williamson used a random hyperplane to cut the vectors.


Let p be a random unit vector in \mathbb{R}^n .

The algorithm is simple: just set $S = \{i \mid \langle p, y_i \rangle \geq 0\}$. See the above picture for an illustration

To analyze its performance, we see that an edge ij is cut with probability θ_{ij} / π .



(picture from Gärtner-Matoušek)



(picture from Gärtner-Matoušek)

Therefore, $E[\# \text{ edges cut}] = \sum_{ij \in E} \Pr(\text{edge } ij \text{ is cut}) = \sum_{ij \in E} \theta_{ij} / \pi$.

Let's compare it to the performance of the SDP solution.

$$\frac{\text{sol}}{\text{OPT}_{\text{INT}}} \geq \frac{\text{sol}}{\text{OPT}_{\text{SDP}}} = \frac{\sum_{ij \in E} \theta_{ij} / \pi}{\sum_{ij \in E} \frac{1}{2}(1 - \cos \theta_{ij})} \geq \min_{\theta} \frac{\theta / \pi}{\frac{1}{2}(1 - \cos \theta)} \geq 0.8785672,$$

where the last inequality is obtained by plotting the curve of the function.

This completes the proof.

The algorithm performs very well in practice, producing almost optimal solution in most instances.

Remark: Very recent work has provided strong evidence that this 0.878-approximation algorithm is the optimal polynomial time approximation algorithm for the maximum cut problem.

References:

- Karger, Random sampling in cut, flow and network design problems, 1994.
- Benczur and Karger, Approximating s-t minimum cuts in $\tilde{O}(n^2)$ time, 1996.
- Fung, Hariharan, Harvey, Panigrahi, A general framework for graph sparsification, 2011.
- Karger. Minimum cuts in near linear time, 2000.
- Raghavan and Tompson. Randomized rounding: A technique for provably good algorithms and algorithmic proof, 1987.
- Goemans and Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming, 1995.
- Gärtner and Matoušek. Approximation algorithms and semidefinite programming, 2012.