Sparsest Cut

1 Introduction

SPARSEST CUT is a fundamental problem in graph algorithms with connections to various cut related problems.

Problem 1 (NON-UNIFORM SPARSEST CUT) *The input is a graph $G = (V, E)$ with edge capacities $c : E \to \mathbb{R}_+$ and a set of vertex pairs $\{s_1, t_1\}, \dots, \{s_k, t_k\}$ along with demand values $D_1, \dots, D_k \in \mathbb{R}_+$. The goal is to find a cut $\delta(S)$ of G such that $\frac{c(\delta(S))}{\sum_{i:|S\cap\{s_i,t_i\}|=1} D_i}$ is minimized.*

In other words, NON-UNIFORM SPARSEST CUT finds the cut that minimizes its capacity divided by the sum of demands of the vertex pairs it separates. There are two important varients of NON-UNIFORM SPARSEST CUT. Note that we always consider unordered pair $\{s_i, t_i\}$, i.e., we do not distinguish $\{s_i, t_i\}$ and $\{t_i, s_i\}$.

UNIFORM SPARSEST CUT is the uniform version of NON-UNIFORM SPARSEST CUT. The demand is 1 for every possible vertex pair $\{s_i, t_i\}$. In this case, we can remove from the input the pairs and demands. The goal becomes to minimize $\frac{c(\delta(S))}{|S||V\setminus S|}$.

EXPANSION further simplifies the objective of UNIFORM SPARSEST CUT to $\min_{|S|\le n/2} \frac{c(\delta(S))}{|S|}$.

These problems are interesting since they are related to central concepts in graph theory and help to design algorithms for hard problems on graph. One connections is expander graphs. The importance of expander graphs is thoroughly surveyed in [HLW06]. The optimum of EXPANSION is also known as Cheeger constant or conductance of a graph. UNIFORM SPARSEST CUT provides a 2-approximation of EXPANSION, which is especially important in the context of expander graphs as it is a way to measure the edge expansion of a graph. NON-UNIFORM SPARSEST CUT is related to other cut problems such as Multicut and Balanced Separator. From a more mathematical perspective, the techniques developed for approximating SPARSEST CUT are deeply related to metric embedding, which is another fundamental problem in geometry. Besides theoretical interests, SPARSEST CUT is useful in practical scenarios such as in image segmentation and in some machine leaning algorithms.

1.1 related works

NON-UNIFORM SPARSEST CUT is APX-hard [CK09] and, assuming the Unique Game Conjecture, has no polynomial time constant factor aproximation algorithm [CKK+05]. UNIFORM SPARSEST CUT admits no PTAS [AMS07], assuming that NP-complete problems cannot be solved in randomized subexponential time. The currently best approximation algorithm for UNIFORM SPARSEST CUT has ratio $O(\sqrt{\log n})$ and running time $\tilde{O}(n^2)$ [AHK10]. For NON-UNIFORM SPARSEST CUT the best approximation is $O(\sqrt{\log n}\log\log n)$ [ALN05, ALN07]. There are also works concerning approximating SPARSEST CUT on special graph classes such as planar graphs [LS10], graphs with low treewidth [CKR10, GTW13, CKM+24].

The seminal work of [LR88, LR99] starts this line of research. They studied multicommodity flow problem and proved a $O(\log n)$ flow-cut gap for Uniform Sparsest Cut. They developed a $O(\log n)$ approximation algorithm for Uniform Sparsest Cut. The technique is called region growing. They also discovered a lowerbound of $\Omega(\log n)$. Note that the flow-cut gap describes the ratio of the max concurrent flow to the min sparsity of a cut. [GVY96] studied the flow-cut gap for min multicut and max multicommodity flow, which is also $\Theta(\log n)$. The result of Garg, Vazirani and Yannakakis [GVY96] provides an $O(\log n)$ approximation algorithm for Multicut, which implies a $O(\log^2 n)$ approximation for Non-Uniform Sparsest Cut. Although [LR99] showed an $\Omega(\log n)$ lowerbound for flow-cut gap, better approximation for Sparsest Cut is still possible through other methods.

For Non-Uniform Sparsest Cut the $O(\log^2 n)$ approximation is further improved by [LLR95] and [AR98]. [AR98] applied metric embedding to Non-Uniform Sparsest Cut and obtained a $O(\log n)$ flow-cut gap as well as a $O(\log n)$ approximation algorithm for Non-Uniform Sparsest Cut. The connections between metric embedding and Non-Uniform Sparsest Cut is influential. Non-Uniform Sparsest Cut can be formulated as an integer program. [AR98], [AR95] and [LLR95] considered the metric relaxation of the IP. They observed that Non-Uniform Sparsest Cut is polynomial time solvable for trees and more generally for all ℓ_1 metrics. The $O(\log n)$ gap follows from the $O(\log n)$ distortion in the metric embedding theorem.

[ARV04] and [AHK10] further improved the approximation ratio for Uniform Sparsest Cut to $O(\sqrt{\log n})$ via semidefinite relaxation. This is currently the best approximation ratio for Uniform Sparsest Cut on general undirected graphs. For Non-Uniform Sparsest Cut, the approximation is improved to $O(\sqrt{\log n}\log\log n)$ [ALN05, ALN07]. Later [GS13] gives a $\frac{1+\delta}{\epsilon}$ approximation in time $2^{r/(\delta\epsilon)}\operatorname{poly}(n)$ provided that $\lambda_r \geq \mathrm{OPT}/(1-\delta)$.

There is also plenty of research on Sparsest Cut in some graph classes, for example [BBPP12]. One of the most popular class is graphs with constant treewidth. [CKM⁺24] gave a $O(k^2)$ approximation algorithm with complexity $2^{O(k)}\operatorname{poly}(n)$. [CAMV24] obtained a 2-approximation algorithm for sparsest cut in treewidth k graph with running time $2^{2^{O(k)}}\operatorname{poly}(n)$.

Sparsest Cut is easy on trees and the flow-cut gap is 1 for trees. One explaination[1] is that shortest path distance in trees is an ℓ_1 metric. There are works concerning planar graphs and more generally graphs with constant genus. [LR99] provided a $\Omega(\log n)$ lowerbound for flow-cut gap for Sparsest Cut. However, it is conjectured that the gap is $O(1)$, while currently the best upperbound is still $O(\sqrt{\log n})$ [Rao99]. For graphs with constant genus, [LS10] gives a $O(\sqrt{\log g})$ approximation for Sparsest Cut, where g is the genus of the input graph. For flow-cut gap in planar graphs the techniques are mainly related to metric embedding theory[2].

2 Approximations

Techniques for approximating Sparsest Cut.

[1] https://courses.grainger.illinois.edu/cs598csc/fa2024/Notes/lec-sparsest-cut.pdf
[2] https://home.ttic.edu/~harry/teaching/teaching.html

2.1 LP $\Theta(\log n)$ - Non-Uniform Sparsest Cut

$$
\begin{aligned}
\min \quad & \frac{\sum_e c_e x_e}{\sum_i D_i y_i} \\
\text{s.t.} \quad & \sum_{e \in p} x_e \geq y_i \qquad \forall p \in \mathcal{P}_{s_i,t_i}, \forall i \\
& x_e, y_i \in \{0, 1\}
\end{aligned} \tag{1}
$$

$$
\begin{aligned}
\min \quad & \sum_e c_e x_e \\
\text{s.t.} \quad & \sum_i D_i y_i = 1 \\
& \sum_{e \in p} x_e \geq y_i \quad \forall p \in \mathcal{P}_{s_i,t_i}, \forall i \\
& x_e, y_i > 0
\end{aligned} \tag{2}
$$

$$
\begin{aligned}
\max \quad & \lambda \\
\text{s.t.} \quad & \sum_{p \in \mathcal{P}_{s_i,t_i}} y_p \geq \lambda D_i \quad \forall i \\
& \sum_i \sum_{p \in \mathcal{P}_{s_i,t_i}, p \ni e} y_p \geq c_e \quad \forall e \\
& y_p \geq 0
\end{aligned} \tag{3}
$$

$$
\begin{aligned}
\min \quad & \sum_{uv \in E} c_{uv} d(u, v) \\
\text{s.t.} \quad & \sum_i D_i d(s_i, t_i) = 1 \\
& d \text{ is a metric on } V
\end{aligned} \tag{4}
$$

1. IP1 ≥ LP2. Given any feasible solution to IP1, we can scale all x_e and y_i simultaneously with factor $1/\sum_i D_i y_i$. The scaled solution is feasible for LP2 and gets the same objective value.

2. LP2 = LP3. by duality.

3. LP4 = LP2. It is easy to see LP4 ≥ LP2 since any feasible metric to LP4 induces a feasible solution to LP2. In fact, the optimal solution to LP2 also induces a feasible metric. Consider a solution x_e, y_i to LP2. Let d_x be the shortest path metric on V using edge length x_e. It suffices to show that $y_i = d_x(s_i, t_i)$. This can be seen from a reformulation of LP2. The constraint $\sum_i D_i y_i = 1$ can be removed and the objective becomes $\sum_e c_e x_e / \sum_i D_i y_i$. This reformulation does not change the optimal solution. Now suppose in the optimal solution to LP2 there is some y_i which is strictly smaller than $d_x(s_i, t_i)$. Then the denominator $\sum_i D_i y_i$ in the objective of our reformulation can be larger, contradicting to the optimality of solution x_e, y_i.

Theorem 2.1 (Japanese Theorem) *D is a demand matrix. D is routable in G iff $\forall l : E \to \mathbb{R}_+$, $\sum_e c_e l(e) \geq \sum_{uv} D(u, v) d_l(u, v)$, where $d_l(s, t)$ is the short path distance induced by $l(e)$.*

Note that D is routable iff the optimum of the LPs is at least 1. Then the theorem follows directly from LP4.

$\Theta(\log n)$ flow-cut gap The flow-cut gap is defined as OPT(IP1)/ OPT(LP2) and the $\Theta(\log n)$ bound is proven in [LR99].

Suppose that G satisfies the cut condition, that is, $c(\delta(S))$ is at least the demand separated by $\delta(S)$ for all $S \subset V$. This implies OPT(IP1) ≥ 1 and in this case the largest integrality gap is 1/ OPT(LP2). For 1 and 2-commodity flow problem the gap is 1 [FF56, Hu63]. However, for $k \geq 3$ the gap becomes larger[3]. It is mentioned in [LR99] that [Sch90] proved if the demand graph does not contain either three disjoint edges or a triangle and a disjoint edge, then the gap is 1.

[3] https://en.wikipedia.org/wiki/Approximate_max-flow_min-cut_theorem

For the $\Omega(\log n)$ lowerbound consider an Uniform Sparsest Cut instance on some 3-regular graph G with unit capacity. In [LR99] they further required that for any $S \subset V$ and small constant c, $|\delta(S)| \geq c \min(|S|, |\bar{S}|)$. Then the value of the sparsest cut is at least $\frac{c}{n-1}$. Observe that for any fixed vertex v, there are at most $n/2$ vertices within distance $\log n - 3$ of v. Thus at least half of the $\binom{n}{2}$ demand pairs are connected with shortest path of length at least $\log n - 2$. To sustain a flow f we need at least $\frac{1}{2}\binom{n}{2}(\log n - 2)f \leq 3n/2$. Any feasible flow satisfies $f \leq \frac{3n}{\binom{n}{2}(\log n - 2)}$ and the gap is therefore $\Omega(\log n)$.

For the upperbound it suffices to show there exists a cut of ratio $O(f \log n)$. [LR99] gave an algorithmic proof based on LP4. This can also be proven using metric embedding results. We can solve LP4 in polynomial time and get a metric on V. Then there is an embedding of V into $\mathbb{R}^d$ with ℓ_1 metric such that the distortion is $O(\log n)$. Since ℓ_1 metric is in the cut cone, our metric on $\mathbb{R}^d$ is a conic combination of cut metrics, which implies[4] that there is a cut in the conic combination with value at most $O(\log n)\,\mathrm{OPT}(\mathrm{LP}4)$. To find such a cut it suffices to compute a conic combination of cut metrics which is exactly our ℓ_1 metric in $\mathbb{R}^d$. One way to do this is test $(n-1)d$ cuts by observing the followings,

1. Every coordinate of $\mathbb{R}^d$ corresponds to a line metric;
2. ℓ_1 metric in $\mathbb{R}^d$ is the sum of those line metrics;
3. Every line metric on n points can be represented as some conic combination of $n - 1$ cut metrics.

The gap can be improved to $\log k$ through a stronger metric embedding theorem (k is the number of demand pairs).

Remark I believe the later method is more general and works for Non-Uniform Sparsest Cut, while the former method is limited to Uniform Sparsest Cut. However, the proof in [LR99] may have connections with the proof of Bourgain's thm? Why does their method fail to work on Non-Uniform Sparsest Cut?

2.2 SDP $O(\sqrt{\log n})$ - Uniform Sparsest Cut

This $O(\sqrt{\log n})$ approximation via SDP is developed in [ARV04]. This is also described in [WS11, section 15.4].

$$
\begin{aligned}
\min \quad & \frac{\sum_{ij \in E} c_{ij}(x_i - x_j)^2}{\sum_{ij \in V \times V}(x_i - x_j)^2} \\
\text{s.t.} \quad & (x_i - x_j)^2 + (x_j - x_k)^2 \geq (x_i - x_k)^2 \quad \forall i, j, k \in V \\
& x_i \in \{+1, -1\} \quad \forall i \in V
\end{aligned}
$$

This SDP models Uniform Sparsest Cut since every assignment of x corresponds to a cut and the objective is the sparsity of the cut (up to a constant factor, but we don't care since we cannot achieve a constant factor approximation anyway). Consider a relaxation which is similar to LP2.

[4]for details see Thm11 in https://courses.grainger.illinois.edu/cs598csc/fa2024/Notes/lec-sparsest-cut.pdf

$$\begin{aligned}
\min \quad & \sum_{ij\in E} c_{ij}\|v_i - v_j\|^2 \\
s.t. \quad & \sum_{ij\in V\times V} \|v_i - v_j\|^2 = 1 \\
& \|v_i - v_j\|^2 + \|v_j - v_k\|^2 \geq \|v_i - v_k\|^2 \quad \forall i,j,k \in V \\
& v_i \in \mathbb{R}^n \qquad \forall i \in V
\end{aligned}$$

To get a $O(\sqrt{\log n})$ (randomized) approximation algorithm we need to first solve the SDP and then round the solution to get a cut $\delta(S)$ with $c(\delta(S)) = |S|\,\mathrm{OPT}(SDP)O(n\sqrt{\log n})$. If there are two sets $S, T \subset V$ both of size $\Omega(n)$ that are well-separated, in the sense that for any $s \in S$ and $t \in T$, $\|v_s - v_t\|^2 = \Omega(1/\sqrt{\log n})$, then the SDP gap follows from

$$\frac{c(\delta(S))}{|S||V-S|} \leq \frac{\sum_{ij\in E} c_{ij}\|v_i - v_j\|^2}{\sum_{i\in S, j\in T}\|v_i - v_j\|^2} \leq \frac{\sum_{ij\in E} c_{ij}\|v_i - v_j\|^2}{n^2} O(\sqrt{\log n}) \leq O(\sqrt{\log n})\,\mathrm{OPT}(SDP).$$

This is the framework of the proof in [ARV04]. I think the intuition behind this SDP relaxation is almost the same as LP4. ℓ_1 metrics are good since they are in the cut cone. However, if we further require that the metric in LP4 is an ℓ_1 metric in $\mathbb{R}^d$, then resulting LP is NP-hard, since the integrality gap becomes 1. [LR99] showed that the $\Theta(\log n)$ gap is tight for LP4, but add extra constraints to LP4 (while keeping it to be a relaxation of Sparsest Cut and to be polynomially solvable) may provides better gap. The SDP relaxation is in fact trying to enforce the metric to be ℓ_2^2 in $\mathbb{R}^n$.

Remark $O(\sqrt{\log n})$ is likely to be the optimal bound for the above SDP. To get better gap one can stay with SDP and add more additional constraints (like Sherali-Adams, Lovász-Schrijver and Lasserre relaxations); or think distance as variables in an LP and force feasible solution to be certain kind of metrics. [AGS13] is following the former method and considers Lasserre relaxations. For the later method, getting a cut from the optimal metric is the same as embedding it to ℓ_1. Thus it still relies on progress in metric embedding theory. Note that both methods need to satisfy

1. the further constrained programs is polynomially solvable,
2. it remains a relaxation of Sparsest Cut,
3. the gap is better.

The Lasserre relaxation of SDP automatically satisfies 1 and 2. But I believe there may be some very strange kind of metric that embeds into ℓ_1 well?

Another possible approach for Non-Uniform Sparsest Cut would be making the number of demand vertices small and then applying a metric embedding (contraction) to ℓ_1 with better distortion on those vertices.

2.3 SDP $O(\sqrt{\log n}\log\log n)$ - Non-Uniform Sparsest Cut

Arora, Lee and Naor [ALN05,ALN07] proved that there is an embedding from ℓ_2^2 to ℓ_1 with distortion $O(\sqrt{\log n}\log\log n)$. This implies an approximation for Non-Uniform Sparsest Cut with the same ratio.

Recently the $O(\sqrt{\log n}\log\log n)$ gap has been improved to $\Theta(\sqrt{\log n})$[5].

3 What problem can I work on?

3.1 Nealy uniform Sparsest Cut

What is the best approximation ratio for Uniform Sparsest Cut instances where almost all demands are uniform? More formally, consider a Non-Uniform Sparsest Cut instance where only k vertices are associated with demand pairs with $D_i \neq 1$, we want to show that we can approximate nearly uniform Sparsest Cut in polynomial time to ratio $O(\sqrt{\log n} f(k))$, where $f(k) = O(\log\log n)$ when $k \to n$. Let those k non uniform vertices be outliers. [ARV04] shows that for non-outlier verteices the optimal solution to SDP (a metric) can be embedded into ℓ_1 with distortion $\sqrt{\log n}$. [CS23] is a recent result on getting approximate (k, c)-outlier embeddings.

This is not a interesting problem since if arbitary demand is allowed on the non-uniform vertices the approximation can be also arbitary. If only constant demand is allowed, then for any constant k we can obtain the same approximation as Uniform Sparsest Cut by ignoring the non-uniform part. This problem does not have much to do with outlier embeddings.

References

[AGS13] Sanjeev Arora, Rong Ge, and Ali Kemal Sinop. Towards a Better Approximation for Sparsest Cut? In *2013 IEEE 54th Annual Symposium on Foundations of Computer Science*, pages 270–279, Berkeley, CA, USA, October 2013. IEEE.

[AHK10] Sanjeev Arora, Elad Hazan, and Satyen Kale. $O(\sqrt{\log n})$ Approximation to SPARSEST CUT in $\tilde{O}(n^2)$ Time. *SIAM Journal on Computing*, 39(5):1748–1771, January 2010.

[ALN05] Sanjeev Arora, James R. Lee, and Assaf Naor. Euclidean distortion and the sparsest cut. In *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, STOC '05, pages 553–562, New York, NY, USA, May 2005. Association for Computing Machinery.

[ALN07] Sanjeev Arora, James R. Lee, and Assaf Naor. Fréchet Embeddings of Negative Type Metrics. *Discrete & Computational Geometry*, 38(4):726–739, December 2007.

[AMS07] Christoph Ambuhl, Monaldo Mastrolilli, and Ola Svensson. Inapproximability results for sparsest cut, optimal linear arrangement, and precedence constrained scheduling. In *48th Annual IEEE Symposium on Foundations of Computer Science (FOCS'07)*, pages 329–337, 2007.

[AR95] Yonatan Aumann and Yuval Rabani. Improved bounds for all optical routing. In *Proceedings of the Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '95, page 567–576, USA, 1995. Society for Industrial and Applied Mathematics.

[AR98] Yonatan Aumann and Yuval Rabani. An $O(\log k)$ approximate min-cut max-flow theorem and approximation algorithm. *SIAM Journal on Computing*, 27(1):291–301, 1998.

[ARV04] Sanjeev Arora, Satish Rao, and Umesh Vazirani. Expander flows, geometric embeddings and graph partitioning. In *Proceedings of the thirty-sixth annual ACM symposium on Theory of computing*, STOC '04, pages 222–231, New York, NY, USA, June 2004. Association for Computing Machinery.

[5] STOC '25 https://web.math.princeton.edu/~naor/homepagefiles/local-growth-STOC.pdf

[BBPP12] Paul Bonsma, Hajo Broersma, Viresh Patel, and Artem Pyatkin. The complexity of finding uniform sparsest cuts in various graph classes. *Journal of Discrete Algorithms*, 14:136–149, July 2012.

[CAMV24] Vincent Cohen-Addad, Tobias Mömke, and Victor Verdugo. A 2-approximation for the bounded treewidth sparsest cut problem in FPTtime. *Mathematical Programming*, 206(1):479–495, July 2024.

[CK09] Julia Chuzhoy and Sanjeev Khanna. Polynomial flow-cut gaps and hardness of directed cut problems. *J. ACM*, 56(2), April 2009.

[CKK+05] S. Chawla, R. Krauthgamer, R. Kumar, Y. Rabani, and D. Sivakumar. On the hardness of approximating MULTICUT and SPARSEST-CUT. In *20th Annual IEEE Conference on Computational Complexity (CCC'05)*, pages 144–153, June 2005. ISSN: 1093-0159.

[CKM+24] Parinya Chalermsook, Matthias Kaul, Matthias Mnich, Joachim Spoerhase, Sumedha Uniyal, and Daniel Vaz. Approximating sparsest cut in low-treewidth graphs via combinatorial diameter. *ACM Transactions on Algorithms*, 20(1):1–20, January 2024.

[CKR10] Eden Chlamtac, Robert Krauthgamer, and Prasad Raghavendra. Approximating Sparsest Cut in Graphs of Bounded Treewidth. In Maria Serna, Ronen Shaltiel, Klaus Jansen, and José Rolim, editors, *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 124–137, Berlin, Heidelberg, 2010. Springer.

[CS23] Shuchi Chawla and Kristin Sheridan. Composition of nested embeddings with an application to outlier removal, November 2023. arXiv:2306.11604 [cs].

[FF56] L. R. Ford and D. R. Fulkerson. Maximal flow through a network. *Canadian Journal of Mathematics*, 8:399–404, 1956.

[GS13] Venkatesan Guruswami and Ali Kemal Sinop. Approximating non-uniform sparsest cut via generalized spectra. In *Proceedings of the twenty-fourth annual ACM-SIAM symposium on Discrete algorithms*, SODA '13, pages 295–305, USA, January 2013. Society for Industrial and Applied Mathematics.

[GTW13] Anupam Gupta, Kunal Talwar, and David Witmer. Sparsest cut on bounded treewidth graphs: Algorithms and hardness results, 2013.

[GVY96] Naveen Garg, Vijay V. Vazirani, and Mihalis Yannakakis. Approximate Max-Flow Min-(Multi)Cut Theorems and Their Applications. *SIAM Journal on Computing*, 25(2):235–251, April 1996. Publisher: Society for Industrial and Applied Mathematics.

[HLW06] Shlomo Hoory, Nathan Linial, and Avi Wigderson. Expander graphs and their applications. *Bulletin of the American Mathematical Society*, 43(04):439–562, August 2006.

[Hu63] T. C. Hu. Multi-commodity network flows. *Operations Research*, 11(3):344–360, June 1963.

[LLR95] Nathan Linial, Eran London, and Yuri Rabinovich. The geometry of graphs and some of its algorithmic applications. *Combinatorica*, 15(2):215–245, June 1995.

[LR88] T. Leighton and S. Rao. An approximate max-flow min-cut theorem for uniform multicommodity flow problems with applications to approximation algorithms. In *[Proceedings 1988] 29th Annual Symposium on Foundations of Computer Science*, pages 422–431, 1988.

[LR99] Tom Leighton and Satish Rao. Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms. *J. ACM*, 46(6):787–832, November 1999.

[LS10] James R. Lee and Anastasios Sidiropoulos. Genus and the geometry of the cut graph: [extended abstract]. In *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 193–201. Society for Industrial and Applied Mathematics, January 2010.

[Rao99] Satish Rao. Small distortion and volume preserving embeddings for planar and Euclidean metrics. In *Proceedings of the fifteenth annual symposium on Computational geometry*, pages 300–306, Miami Beach Florida USA, June 1999. ACM.

[Sch90] Alexander Schrijver. Homotopic Routing Methods. *Paths, Flows, and VLSI-Layout*, pages 329 –371, 1990.

[WS11] David P. Williamson and David B. Shmoys. *The Design of Approximation Algorithms*. Cambridge University Press, 2011.