Online Degree-Bounded Steiner Network Design*

The problem of satisfying connectivity demands on a graph while respecting given constraints has been a pillar of the area of network design since the early seventies [45, 14, 13, 12, 41]. The problem of DEGREE-BOUNDED SPANNING TREE, introduced in Garey and Johnson's *Black Book* of NP-Completeness [37], was first investigated in the pioneering work of Fürer and Raghavachari [17] (Allerton'90). In the DEGREE-BOUNDED SPANNING TREE problem, the goal is to construct a spanning tree for a graph $G = (V, E)$ with n vertices whose maximal degree is the smallest among all spanning trees. Let b^* denote the maximal degree of an optimal spanning tree. Fürer and Raghavachari [17] give a parallel approximation algorithm which produces a spanning tree of degree at most $O(\log(n)b^*)$.

Agrawal, Klein, and Ravi ([1]) consider the following generalizations of the problem. In the DEGREE-BOUNDED STEINER TREE problem we are only required to connect a given subset $T \subseteq V$. In the even more general DEGREE-BOUNDED STEINER FOREST problem the demands consist of vertex pairs, and the goal is to output a subgraph in which for every demand there is a path connecting the pair. They design an algorithm that obtains a multiplicative approximation factor of $O(\log(n))$. Their main technique is to reduce the problem to minimizing congestion under integral concurrent flow restrictions and to then use the randomized rounding approach due to Raghavan and Thompson ([43]).

Shortly after the work of Agrawal *et al.*, Fürer and Raghavachari [18] significantly improved the result for DEGREE-BOUNDED STEINER FOREST by presenting an algorithm which produces a Steiner forest with maximum degree at most $b^* + 1$. They show that the same guarantee carries over to the *directed* variant of the problem as well. Their result is based on reducing the problem to that of computing a sequence of maximal matchings on certain auxiliary graphs. This result settles the approximability of the problem, as computing an optimal solution is NP-hard even in the spanning tree case.

In this paper, we initiate the study of degree-bounded network design problems in an *online* setting, where connectivity demands appear over time and must be immediately satisfied. We first design a deterministic algorithm for ONLINE DEGREE-BOUNDED STEINER FOREST with a logarithmic competitive ratio. Then we show that this competitive ratio is asymptotically best possible by proving a matching lower bound for randomized algorithms that already holds for the Steiner tree variant of the problem.

In the offline scenario, the results of Fürer, Raghavachari [17, 18] and Agrawal, Klein, Ravi [1] were the starting point of a very popular line of work on various degree-bounded network design problems [34, 20, 40, 31, 27, 15]. We refer the reader to the next sections for a brief summary. One particular variant that has been extensively studied is the *edge-weighted* DEGREE-BOUNDED SPANNING TREE. Initiated by Marathe *et al.* ([34]), in this version, we are given a weight function over the edges and a bound b on the maximum degree of a vertex. The goal

*A preliminary version of the paper appeared in ACM-SIAM Symposium on Discrete Algorithms (SODA'16).

is to find a minimum-weight spanning tree with maximum degree at most b. The groundbreaking results obtained by Goemans ([20]) and Singh and Lau ([44]) settle the problem by giving an algorithm that computes a minimum-weight spanning tree with degree at most $b+1$. A slightly worse result is obtained by Singh and Lau ([31]) for the Steiner tree variant. Unfortunately, in the online setting it is not possible to obtain a comparable result. We show that for any (randomized) algorithm $\mathcal{A}$ there exists a request sequence such that $\mathcal{A}$ outputs a sub-graph that either has weight $\Omega(n) \cdot \mathrm{OPT}_b$ or maximum degree $\Omega(n) \cdot b$.

In the online variant of DEGREE-BOUNDED STEINER FOREST , we are given the graph G in advance, however, demands arrive in an online fashion. At online step i, a new demand (s_i, t_i) arrives. Starting from an empty subgraph, at each step the online algorithm should augment its solution so that the endpoints of the new demand s_i and t_i are connected. The goal is to minimize the maximum degree of the solution subgraph. In the *non-uniform* variant of the problem, a degree bound $b_v \in \mathbb{R}^+$ is given for every vertex v. For a subgraph H and a vertex v, let $\deg_H(v)$ denote the degree of v in H. The *load* of a vertex is defined as the ratio $\deg_H(v)/b_v$. In the non-uniform variant of ONLINE DEGREE-BOUNDED STEINER FOREST, the goal is to find a subgraph satisfying the demands while minimizing the maximum load of a vertex.

Our algorithm is a simple and intuitive greedy algorithm. Upon the arrival of a new demand (s_i, t_i), the greedy algorithm (GA) satisfies the demand by choosing an (s_i, t_i)-path P_i such that after augmenting the solution with P_i, the maximum load of a vertex in P_i is minimum. A main result of our paper is to prove that the maximum load of a vertex in the output of GA is within a logarithmic factor of OPT, the maximum load of a vertex in an optimal offline solution which knows all the demands in advance.

Theorem 0.1 *The algorithm GA produces an output with maximum load at most $O(\log n) \cdot \mathrm{OPT}$.*

The crux of our analysis is establishing several structural properties of the solution subgraph. First we group the demands according to the maximum load of the bottleneck vertex at the time of arrival of the demand. We then show that for every threshold $r > 0$, vertices with load at least r at the end of the run of GA, form a cut set that well separates the group of demands with load at least r at their bottleneck vertex. Since the threshold value can be chosen arbitrarily, this leads to a series of cuts that form a chain. The greedy nature of the algorithm indicates that each cut highly disconnects the demands. Intuitively, a cut that highly disconnects the graph (or the demands) implies a lower bound on the number of branches of every feasible solution.

We use a natural dual-fitting argument to show that for every cut set, the ratio between the number of demands in the corresponding group, over the total degree bound of the cut, is a lower bound for OPT. Hence, the problem comes down to showing that one of the cuts in the series has ratio at least $1/O(\log n)$ fraction of the maximum load h of the output of GA. To this end, we first partition the range of $r \in (0, h]$ into $O(\log n)$ layers based on the total degree bound of the corresponding cut. We then show that the required cut can be found in an interval with maximum range of r.

We complement our first theorem by giving an example for a special case of ONLINE DEGREE-BOUNDED STEINER TREE in which no online (randomized) algorithm can achieve a (multiplicative) competitive ratio $o(\log n)$. This also implies that obtaining (non-trivial) additive competitiveness is not possible in the online setting.

Theorem 0.2 *Any (randomized) online algorithm for the degree bounded online Steiner tree problem has (multiplicative) competitive ratio $\Omega(\log n)$. This already holds when $b_v = 1$ for every node.*

The previously known techniques. As discussed before, the majority of techniques used for solving the offline variants of degree-bounded problems involve rounding an optimal fractional solution of a relaxed linear program. Since one may need to buy a long path to connect the endpoints of a demand, independent rounding of a fractional solution is hardly efficient. Instead, dependent and iterative rounding methods are usually used for attacking degree-bounded problems. In the online paradigm, one can maintain a competitive fractional solution for these problems, however, it is inherently difficult to apply the aforementioned rounding techniques in an online setting: the underlying online fractional solution changes in between the rounding steps, thus breaking the chain of dependencies.

In contrast to the works on the offline paradigm, in this paper we propose a simple combinatorial algorithm with a dual-fitting analysis. We use the structural properties of the output of our algorithm to show the existence of a chain of cuts that well separates the demand endpoints. When restricted to the case of uniform bounds, these cuts imply an upper bound on the *toughness* of the graph. The toughness of a graph is defined as $\min_{X \subseteq V} \frac{|X|}{|\mathrm{CC}(G \setminus X)|}$; where for a graph H, $\mathrm{CC}(H)$ denotes the collection of connected components of H. It can be shown that the reciprocal of the toughness gives a lower bound for OPT. Therefore we use a combinatorial argument to show that the minimum of this ratio over the cuts in our chain of cuts is within $O(\log(n))$ approximation of the reciprocal of the maximum load of a vertex in our solution.

We would like to emphasize that although the concept of toughness is well-studied in the literature, this line of research is mainly focused on relating toughness conditions to the existence of cycle structures, see for example a comprehensive survey by Bauer et al. [9]. The relation between the graph toughness and degree-bounded problems have been previously observed by Win [45] and Agrawal et al. [1]. However as mentioned in the introduction, Agrawal et al. use a completely different argument for analyzing the problem when reduced to a congestion minimization problem. We hope that the structural properties introduced in this paper together with the dual interpretation of our analysis, paves the way for solving the classical problems in the family of degree-bounded problems.

Hardness under more general constraints. We further investigate the following extensions of the online degree bounded Steiner tree problem. First, we consider the edge-weighted variant of the degree-bounded Steiner tree problem. Second, we consider the group Steiner tree version in which each demand consists of a subset of vertices, and the goal is to find a tree that covers at least one vertex of each demand group. The following theorems show that one cannot obtain a non-trivial competitive ratio for these versions in their general form.[1]

Theorem 0.3 *Consider the edge weighted variant of* ONLINE DEGREE-BOUNDED STEINER TREE. *For any (randomized) online algorithm $\mathcal{A}$, there exists an instance and a request sequence such that either* $\mathrm{E}[\mathrm{maxdegree}(\mathcal{A})] \geq \Omega(n) \cdot b$ *or* $\mathrm{E}[\mathrm{weight}(\mathcal{A})] \geq \Omega(n) \cdot \mathrm{OPT}_b$, *where* OPT_b *denotes the minimum weight of a Steiner tree with maximum degree b.*

Theorem 0.4 *There is no deterministic algorithm with competitive ratio $o(n)$ for the* DEGREE-BOUNDED GROUP STEINER TREE *problem.*

[1] Our lower bound results imply that one needs to restrict the input in order to achieve competitiveness. In particular for the edge-weighted variant, our proof does not rule out the existence of a competitive algorithm when the edge weights are polynomially bounded.

The classical family of degree-bounded network design problems have various applications in broadcasting information, package distribution, decentralized communication networks, etc. (see e.g. [19, 23]). Marathe *et al.* ([34]), first considered the general *edge-weighted* variant of the problem. They give a bi-criteria $(O(\log n), O(\log n) \cdot b)$-approximation algorithm, i.e., the degree of every node in the output tree is $O(\log n) \cdot b$ while its total weight is $O(\log n)$ times the optimal weight. A long line of work (see e.g. [28] and [29]) was focused on this problem until a groundbreaking breakthrough was obtained by Goemans ([20]); his algorithm computes a minimum-weight spanning tree with degree at most $b + 2$. Later on, Singh and Lau ([44]) improved the degree approximation factor by designing an algorithm that outputs a tree with optimal cost while the maximum degree is at most $b + 1$.

In the *degree-bounded survivable network design* problem, a number d_i is associated with each demand (s_i, t_i). The solution subgraph should contain at least d_i edge-disjoint paths between s_i and t_i. Indeed this problem has been shown to admit bi-criteria approximation algorithms with constant approximation factors (e.g. [31]). We refer the reader to a recent survey in [30]. This problem has been recently considered in the node-weighted variant too (see e.g. [40, 15]). The degree-bounded variant of several other problems such as k-MST and k-arborescence has also been considered in the offline setting for which we refer the reader to [27, 8] and references therein.

Online network design problems have attracted substantial attention in the last decades. The online edge-weighted Steiner tree problem, in which the goal is to find a minimum-weight subgraph connecting the demand nodes, was first considered by Imase and Waxman ([24]). They showed that a natural greedy algorithm has a competitive ratio of $O(\log n)$, which is optimal up to constants. This result was generalized to the online edge-weighted Steiner forest problem by Awerbuch *et al.* ([6]) and Berman and Coulston ([10]). Later on, Naor, Panigrahi, and Singh ([39])) and Hajiaghayi, Liaghat, and Panigrahi ([22]), designed poly-logarithmic competitive algorithms for the more general *node-weighted* variant of Steiner connectivity problems. This line of work has been further investigated in the prize-collecting version of the problem, in which one can ignore a demand by paying its given penalty. Qian and Williamson ([42]) and Hajiaghayi, Liaghat, and Panigrahi ([21]) develop algorithms with a poly-logarithmic competitive algorithms for these variants.

The online b-matching problem is another related problem in which vertices have degree bounds but the objective is to maximize the size of the solution subgraph. In the worst case model, the celebrated result of Karp *et al.* ([26]) gives a $(1 - 1/e)$-competitive algorithm. Different variants of this problem have been extensively studied in the past decade, e.g., for the random arrival model see [16, 25, 32, 36], for the full information model see [33, 38], and for the prophet-inequality model see [4, 2, 3]. We also refer the reader to the comprehensive survey by Mehta [35].

Many of the aforementioned problems can be characterized as an online packing or covering linear program. Initiated by work of Alon *et al.* [5] on online set cover, Buchbinder and Naor developed a strong framework for solving packing/covering LPs fractionally online. For the applications of their general framework in solving numerous online problems, we refer the reader to the survey in [11]. Azar *et al.* [7] generalize this method for the fractional *mixed* packing and covering LPs. In particular, they show an application of their method for integrally solving a generalization of capacitated set cover. Their result is a bi-criteria competitive algorithm that violates the capacities by at most an $O(\log^2 n)$ factor while the cost of the ouput is within $O(\log^2 n)$ factor of optimum. We note that although the fractional variant of our problem is a special case of mixed packing/covering LPs, we do not know of any online rounding method for Steiner connectivity problems.

References

[1] A. Agrawal, P. N. Klein, and R. Ravi. How tough is the minimum-degree steiner tree?: A new approximate min-max equality. *Technical Report CS-91-49, Brown University*, 1991.

[2] S. Alaei, M. Hajiaghayi, and V. Liaghat. Online prophet-inequality matching with applications to ad allocation. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 18–35, 2012.

[3] S. Alaei, M. Hajiaghayi, and V. Liaghat. The online stochastic generalized assignment problem. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 11–25. 2013.

[4] S. Alaei, M. T. Hajiaghayi, V. Liaghat, D. Pei, and B. Saha. Adcell: Ad allocation in cellular networks. In *Algorithms–ESA 2011*, pages 311–322. 2011.

[5] N. Alon, B. Awerbuch, Y. Azar, N. Buchbinder, and J. Naor. The online set cover problem. *SIAM Journal on Computing*, 39(2):361–370, 2009.

[6] B. Awerbuch, Y. Azar, and Y. Bartal. On-line generalized steiner problem. In *Proceedings of the seventh annual ACM-SIAM symposium on Discrete algorithms*, pages 68–74, 1996.

[7] Y. Azar, U. Bhaskar, L. Fleischer, and D. Panigrahi. Online mixed packing and covering. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 85–100. SIAM, 2013.

[8] N. Bansal, R. Khandekar, and V. Nagarajan. Additive guarantees for degree-bounded directed network design. *SIAM J. Comput.*, 39(4):1413–1431, 2009.

[9] D. Bauer, H. Broersma, and E. Schmeichel. Toughness in graphs–a survey. *Graphs and Combinatorics*, 22(1):1–35, 2006.

[10] P. Berman and C. Coulston. On-line algorithms for steiner tree problems. In *Proceedings of the twenty-ninth annual ACM symposium on Theory of computing*, pages 344–353, 1997.

[11] N. Buchbinder and J. Naor. The design of competitive online algorithms via a primal: dual approach. *Foundations and Trends® in Theoretical Computer Science*, 3(2–3):93–263, 2009.

[12] P. M. Camerini and G. Galbiati. The bounded path tree problem. *SIAM Journal on Algebraic Discrete Methods*, 3(4):474–484, 1982.

[13] P. M. Camerini, G. Galbiati, and F. Maffioli. Complexity of spanning tree problems: Part i. *European Journal of Operational Research*, 5(5):346–352, 1980.

[14] V. Chvátal. Tough graphs and hamiltonian circuits. *Discrete Mathematics*, 5(3):215–228, 1973.

[15] A. Ene and A. Vakilian. Improved approximation algorithms for degree-bounded network design problems with node connectivity requirements. *STOC*, 2014.

[16] J. Feldman, A. Mehta, V. Mirrokni, and S. Muthukrishnan. Online stochastic matching: Beating 1-1/e. In *Foundations of Computer Science, 2009. FOCS'09. 50th Annual IEEE Symposium on*, pages 117–126, 2009.

[17] M. Fürer and B. Raghavachari. An NC approximation algorithm for the minimum degree spanning tree problem. In *Allerton Conf. on Communication, Control and Computing*, pages 274–281, 1990.

[18] M. Fürer and B. Raghavachari. Approximating the minimum-degree steiner tree to within one of optimal. *Journal of Algorithms*, 17(3):409–423, 1994.

[19] H. Garcia-Molina and B. Kogan. An implementation of reliable broadcast using an unreliable multicast facility. In *Reliable Distributed Systems, 1988. Proceedings., Seventh Symposium on*, pages 101–111, 1988.

[20] M. X. Goemans. Minimum bounded degree spanning trees. In *Foundations of Computer Science, 2006. FOCS'06. 47th Annual IEEE Symposium on*, pages 273–282, 2006.

[21] M. Hajiaghayi, V. Liaghat, and D. Panigrahi. Near-optimal online algorithms for prize-collecting steiner problems. In *Automata, Languages, and Programming*, pages 576–587. 2014.

[22] M. T. Hajiaghayi, V. Liaghat, and D. Panigrahi. Online node-weighted steiner forest and extensions via disk paintings. In *Foundations of Computer Science (FOCS), 2013 IEEE 54th Annual Symposium on*, pages 558–567, 2013.

[23] Y. Huang and H. Garcia-Molina. Publish/subscribe tree construction in wireless ad-hoc networks. In *Mobile Data Management*, pages 122–140, 2003.

[24] M. Imase and B. M. Waxman. Dynamic Steiner tree problem. *SIAM Journal on Discrete Mathematics*, 4(3):369–384, 1991.

[25] C. Karande, A. Mehta, and P. Tripathi. Online bipartite matching with unknown distributions. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pages 587–596, 2011.

[26] R. M. Karp, U. V. Vazirani, and V. V. Vazirani. An optimal algorithm for on-line bipartite matching. In *Proceedings of the twenty-second annual ACM symposium on Theory of computing*, pages 352–358, 1990.

[27] R. Khandekar, G. Kortsarz, and Z. Nutov. On some network design problems with degree constraints. *Journal of Computer and System Sciences*, 79(5):725–736, 2013.

[28] J. Könemann and R. Ravi. A matter of degree: Improved approximation algorithms for degree-bounded minimum spanning trees. In *Proceedings of the thirty-second annual ACM symposium on Theory of computing*, pages 537–546, 2000.

[29] J. Könemann and R. Ravi. Primal-dual meets local search: approximating msts with nonuniform degree bounds. *SIAM Journal on Computing*, 34(3):763–773, 2005.

[30] L. C. Lau, R. Ravi, and M. Singh. *Iterative methods in combinatorial optimization*, volume 46. Cambridge University Press, 2011.

[31] L. C. Lau and M. Singh. Additive approximation for bounded degree survivable network design. *SIAM Journal on Computing*, 42(6):2217–2242, 2013.

[32] M. Mahdian and Q. Yan. Online bipartite matching with random arrivals: an approach based on strongly factor-revealing LPs. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pages 597–606, 2011.

[33] V. H. Manshadi, S. Oveis Gharan, and A. Saberi. Online stochastic matching: Online actions based on offline statistics. *Mathematics of Operations Research*, 37(4):559–573, 2012.

[34] M. V. Marathe, R. Ravi, R. Sundaram, S. Ravi, D. J. Rosenkrantz, and H. B. Hunt III. Bicriteria network design problems. *Journal of Algorithms*, 28(1):142–171, 1998.

[35] A. Mehta. Online matching and ad allocation. *Theoretical Computer Science*, 8(4):265–368, 2012.

[36] A. Mehta, A. Saberi, U. Vazirani, and V. Vazirani. Adwords and generalized online matching. *Journal of the ACM (JACM)*, 54(5):22, 2007.

[37] R. G. Michael and S. J. David. Computers and intractability: a guide to the theory of np-completeness. *WH Freeman & Co., San Francisco*, 1979.

[38] V. S. Mirrokni, S. Oveis Gharan, and M. Zadimoghaddam. Simultaneous approximations for adversarial and stochastic online budgeted allocation. In *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1690–1701, 2012.

[39] J. Naor, D. Panigrahi, and M. Singh. Online node-weighted steiner tree and related problems. In *Foundations of Computer Science (FOCS), 2011 IEEE 52nd Annual Symposium on*, pages 210–219, 2011.

[40] Z. Nutov. Degree-constrained node-connectivity. In *LATIN 2012: Theoretical Informatics*, pages 582–593. 2012.

[41] C. H. Papadimitriou and M. Yannakakis. The complexity of restricted spanning tree problems. *Journal of the ACM*, 29(2):285–309, 1982.

[42] J. Qian and D. P. Williamson. An o (logn)-competitive algorithm for online constrained forest problems. In *Automata, Languages and Programming*, pages 37–48. 2011.

[43] P. Raghavan and C. D. Thompson. Provably good routing in graphs: regular arrays. In *Proceedings of the seventeenth annual ACM symposium on Theory of computing*, pages 79–87, 1985.

[44] M. Singh and L. C. Lau. Approximating minimum bounded degree spanning trees to within one of optimal. In *Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*, pages 661–670, 2007.

[45] S. Win. On a connection between the existence ofk-trees and the toughness of a graph. *Graphs and Combinatorics*, 5(1):201–205, 1989.