Information Processing Letters 174 (2022) 106209

Contents lists available at ScienceDirect

Information Processing Letters

www.elsevier.com/locate/ipl

Two-level hub Steiner trees ☆

Takuro Fukunaga [a], R. Ravi [b], Oleksandr Rudenko [c], Ziye Tang [b,*]

[a] Faculty of Science and Engineering, Chuo University, Japan
[b] Tepper School of Business, Carnegie Mellon University, USA
[c] Department of Mathematical Sciences, Carnegie Mellon University, USA

ARTICLE INFO

Article history:
Received 4 January 2021
Received in revised form 25 September 2021
Accepted 26 September 2021
Available online 8 October 2021
Communicated by Łukasz Kowalik

Keywords:
Approximation algorithms
Hierarchical network design
Spanning tree
Steiner tree
Set cover

ABSTRACT

We study a fundamental class of two-layer network design problems. A hub layer is configured by establishing hubs at selected nodes at considerable cost so that the routes between hubs can be operated cheaply. The remaining edges in the network are operated at regular cost. The resulting problem is to determine the set of nodes to open hubs and the set of edges to establish in order to find a network of minimum total cost.

We consider the case where the network is required to form a Steiner tree spanning a given set of terminal vertices. When edge costs are non-metric, we show logarithmic approximation hardness even for the special case of spanning trees. On the other hand, we show a polynomial-time reduction for Steiner trees to its corresponding node-weighted version thus proving a logarithmic approximation factor. When edge costs are metric, we show the problem is only a constant factor harder to approximate than its original version (with no hub installation) using a similar reduction.

1. Introduction

In contemporary supply chain management, logistics service providers typically operate at minimal margins, so that a high degree of consolidation and the resulting economies of scale are mandatory to run a profitable business. This leads to a typical layered design of logistics networks in which a subset of locations are opened as hubs. We operate high-volume lower-cost routes in the hub-level network and connect clients using the lower-level network that augments the backbone connectivity provided by the hub-level network. We term this the *hub network design* (HND) problem. We can define concrete formulations of the HND problem by specifying connectivity requirements that feasible solutions have to satisfy. One such problem is the *hub Steiner tree* problem that requires a subset of vertices specified as terminals to be connected.

Definition 1 *(Hub Steiner tree (HStT)).* Given an undirected graph $G = (V, E)$, a terminal set $R \subseteq V$, non-negative edge cost c_e for $e \in E$, non-negative hub opening cost f_v for $v \in V$ and a constant $\lambda \in [0, 1]$ reflecting the cost differential between two levels, a *hub Steiner tree* (HStT) is a tree $T = (V_T, E_T)$ spanning the terminal set R (i.e., $R \subseteq V_T$) along with a set of hubs $H \subseteq V_T$. Let T_H denote the set of edges in T induced by H (i.e., with both ends in H). We call T_H the set of *hub-level* edges and $T \setminus T_H$ the set of *lower-level* edges. We define $c(S) = \sum_{e \in S} c_e$ for $S \subseteq E$ and $f(U) = \sum_{v \in U} f_v$ for $U \subseteq V$. The cost of the hub Steiner tree (T, H) is $\lambda c(T_H) + c(T \setminus T_H) + f(H)$. The goal in the HStT problem is to find an HStT of minimum cost.

☆ This material is based upon research supported in part by the U.S. Office of Naval Research under award number N00014-21-1-2243, the Air Force Office of Scientific Research under award number FA9550-20-1-0080 and JST PRESTO under grant number JPMJPR1759.

* Corresponding author.
E-mail addresses: fukunaga@ise.chuo-u.ac.jp (T. Fukunaga), ravi@andrew.cmu.edu (R. Ravi), 4rudenko@gmail.com (O. Rudenko), ziyet@andrew.cmu.edu (Z. Tang).

https://doi.org/10.1016/j.ipl.2021.106209

When G is a complete graph and the edge costs c satisfy the triangle inequality (i.e., $c_{xy} + c_{yz} \geq c_{xy}$ for all $x, y, z \in V$), then we say that we have an instance of a *metric* HStT problem. Otherwise, we have an instance of the non-metric HstT.

Definition 2 *(Hub spanning tree (HST)).* The *hub spanning tree* (HST) problem is a special case of the HStT problem where the terminal set $R = V$.

We note that several two-level network design problems have been studied actively; e.g., access network design [1–3], buy-at-bulk network design [4–9], hub location problem [10,11], and location routing problem [12,13]. These problems and the HStT problem are closely related, but they differ in several aspects. For example, the existence of hubs is not considered or more complex connectivity is required in them. The HStT problem is simple, but captures an essential part of two-level network design that is as yet unstudied. Our contributions in this paper are summarized as follows.

1. We show NP-hardness and logarithmic approximation hardness of the HST problem by reductions from the set cover problem in Section 2. Since the HST problem is a special case of the HStT problem, the same results hold for the latter.
2. For the non-metric HStT problem in Section 3, we show a polynomial-time reduction to the node-weighted Steiner tree problem. This implies an $2 \log |R|$-approximation algorithm for this case.
3. For the metric HStT problem in Section 4, we show a polynomial-time reduction to the original version of the Steiner tree problem (with no hub installation). This implies a constant-factor approximation algorithm.

2. Hardness of hub spanning tree problem

In this section, we prove hardness results for the HST problem by reducing from the set cover problem. Similar reductions have been used in several related network design problems (e.g., [14–16]).

Definition 3 *(Set cover).* Let $S_1, \ldots, S_n$ be arbitrary subsets on a ground set of elements $x_1, \ldots, x_t$. The set cover problem is to find a minimum cardinality set of subsets whose union is the set of all elements.

The following hardness of approximation result for the set cover problem is due to Dinur and Steurer [17].

Theorem 2.1 *([17]). For every $\delta > 0$, it is NP-hard to approximate the set cover problem to within $(1 - \delta) \ln n$, where n is the size of the instance.*

Our hardness results presented in this section repeatedly use the following reduction from the set cover problem.

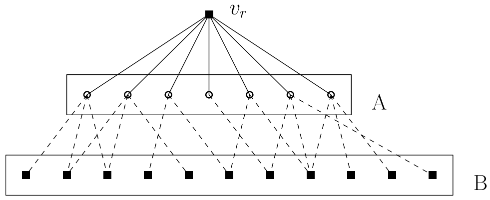

Fig. 1. An illustration of the reduction. The cost of a solid (dashed, resp.) edge is zero (β, resp).

Reduction 1. We construct an undirected weighted graph $G = (V, E)$ as follows: create a node v_{S_i} for each set S_i, a vertex v_j for each element x_j, and a new vertex v_r as the root. Let $A := \{v_{S_i} : i = 1, \ldots, n\}$ and $B := \{v_j : j = 1, \ldots, t\}$. For each $v_S \in A$ we create an edge (v_r, v_S) with cost 0. For each $v_S \in A$ and $v_j \in B$ such that $j \in S$, we create an edge (v_j, v_S) with cost β whose value will be set later to achieve desired hardness results. The hub opening cost is one for all vertices in A and zero for all others. See Fig. 1 for an illustration.

Non-metric HST Based on the above construction, we have the following theorem:

Theorem 2.2. *For any $\lambda \in [0, 1)$, the non-metric HST is NP-hard.*

Proof. Set $\beta > \frac{1}{1-\lambda}$ in Reduction 1. We show that the minimum set cover has cardinality k if and only if the optimal HST cost in Reduction 1 is $k + \lambda\beta t$, where t is the number of elements. Notice this proves the theorem.

Without loss of generality, assume an optimal set cover is $\{S_j\}_{j=1}^k$. Our HST is constructed by opening hubs in $\{v_r\} \cup \{v_{S_j}\}_{j=1}^k \cup B$ and selecting the following set of edges: $\{(v_r, v_{S_j}), (v_{S_j}, v_l) : \forall j = 1, \ldots, k, l \in S_j\}$. In the resulting HST, a vertex in set B is connected to only one vertex in set A (breaking ties arbitrarily between a pair of A nodes that have edges to it). The resulting HST has cost $k + \lambda\beta t$. This shows the minimum cost of the HST problem is less than or equal to $k + \lambda\beta t$. Conversely, we claim all edges in the optimal HST between A and B have both endpoints opened as hubs: suppose not, let (v_S, v_l) be an edge violating this property. We can alternatively open hubs on both endpoints which incurs a unit hub opening cost and reduces the cost of (v_S, v_l) by $(1 - \lambda)\beta > 1$ (this might lead to cost decrease in other edges too), which results in a contradiction. Consequently the set $\mathcal{S} := \{v_S : S \in A$ is an opened hub in HST$\}$ is a valid set cover. Observe that the edges between A and B now have cost $\beta\lambda$ and there are t such edges, so we have $|\mathcal{S}| \leq k$. This shows the minimum set cover has cardinality less than or equal to k. □

Similarly we have the following approximation hardness result.

Theorem 2.3. *If there is an α-approximation algorithm for the non-metric HST problem with $\lambda = 0$, then there is an α-approximation algorithm for the set cover problem.*

Proof. Given a set cover instance with minimum cardinality k, we generate the HST instance as in Reduction 1 with $\lambda = 0$ and $\beta > \alpha k$. Let T be an α-approximate solution of this instance. We will prove that every edge in T has both endpoints opened as hubs. Suppose not, there exists one edge that is not between two hubs, incurring a cost of at least $\beta + (t-1)\beta\lambda = \beta$ (since $\lambda = 0$). Since $\beta > \alpha k$, this implies T is not an α-approximate solution, contradicting our assumption.

Let k be the cardinality of a minimum set cover. Next we show the cost of an optimal HST is at most k: we can open hubs at the k vertices that correspond to the optimal set-cover, whose cost is $k + \lambda\beta t = k$. As a result, the cost of T is at most αk. We can therefore obtain an α-approximate set cover solution by selecting those sets opened as hubs. □

We obtain the following corollary by combining Theorems 2.1 and 2.3.

Corollary 2.4. *When $\lambda = 0$, for any $\delta > 0$ it is NP-hard to approximate the non-metric HST problem within a factor of $(1-\delta)\ln n$.*

Metric HST For notational convenience, we shall denote nodes in A as *A-nodes* and nodes in B as *B-nodes*. We call an A-node *selected* if its corresponding set is included in the set cover solution.

Theorem 2.5. *For any $\lambda \in (0,1)$, the metric HST problem with uniform hub opening cost is NP-hard.*

Proof. We modify Reduction 1 for the non-metric HST. We assign a unit hub opening cost for every node. Recall that the edge weight between the root and an A-node is 0, and the edge weight between an A-node and an B-node is β. We take the metric completion of this graph, i.e., we add all the edges in the complete graph where the cost of an edge is defined as the shortest path length between its two endpoints, based on these edge weights. Recall t is the number of elements. We claim that for $\beta > \max\{\frac{1}{\lambda}, \frac{2}{1-\lambda}\}$, the minimum cost of a HST is $k + t + t\lambda\beta$ if and only if the size of the minimum set cover is k.

For the 'if' part, given a set cover of size k, we install hubs on all B-nodes and selected A-nodes. We connect each A-node to the root by edge cost 0. We connect each B-node to a selected A-node which includes that element by paying $\lambda\beta$. This gives an HST with cost $k + t + t\lambda\beta$.

For the 'only if' part, since the root and all A-nodes are connected by the edges of cost 0, no hub is needed at the root to reduce the edge cost between the root and an A-node. We will ensure that opened hubs among A-nodes exactly represent selected sets. To do so, we need to ensure two things in an optimal HST.

(I) Each edge between an A-node and a B-node is an hub-level edge (i.e., hubs are opened on its both endpoints);
(II) No edge exists between two A-nodes, or two B-nodes.

A sufficient condition for (I) is $\beta > \lambda\beta + 2$ where β is the lower-level edge cost for connecting an B-node to an A-node and $\lambda\beta + 2$ is the hub-level edge cost and the hub opening costs of its two end nodes. This condition also implies that, if an edge joining two B-nodes is used in the optimal HST, it is the hub-level edge because lower-level edges joining two B-nodes are of cost at least 2β ($> \beta > \lambda\beta + 2$). For (II), first notice two A-nodes are already connected via the root by two 0-cost edges. Second, under the condition for (I), it is sufficient to have $2\lambda\beta > \beta\lambda + 1$ where $2\lambda\beta$ is the cost of the hub-level edges joining two elements-nodes and $\beta\lambda + 1$ is the cost for connecting a B-node to an A-node by an hub-level edge by opening a hub on the common A-node. To summarize, we need $\beta > \max\{\frac{1}{\lambda}, \frac{2}{1-\lambda}\}$ which are the conditions in the claim. □

3. Non-metric hub Steiner tree

In this section we reduce the HStT problem to the node-weighted Steiner tree problem defined below.

Definition 4 *(Node-weighted Steiner tree (NWST)).* Let G be an undirected graph with nonnegative costs assigned to its nodes and edges. Let $R \subseteq V$ be a set of terminals. A Steiner tree for R in G is a connected subgraph of G containing all the nodes of R. The *node-weighted Steiner tree* (NWST) problem is to find a minimum-cost Steiner tree.

For NWST, Klein and Ravi [14] showed a greedy algorithm which achieves a logarithmic approximation factor.

Theorem 3.1 *([14]). The NWST problem admits a polynomial-time $2\ln k$-approximation algorithm, where k is the number of terminals.*

Below we present our reduction.

Reduction 2. Given an HStT problem as in Definition 1. We create an NWST instance as follows. Let V' and E' be the node and edge set of the NWST instance respectively. Let $c' : E' \to \mathbb{R}^+$ be the edge weight function. For each node $v \in V$ in HStT, we create a pair of nodes v_h, v_l. Let $V_{\text{hub}} := \{v_h : v \in V\}$ and $V_{\text{low}} := \{v_l : v \in V\}$, where V_{hub} stands for the 'hub-level nodes' and V_{low} for the 'lower-level nodes'. Let $V' = V_{\text{hub}} \cup V_{\text{low}}$. Define the set of terminals $R' := \{v_l : v \in R\}$. For each edge $e = (u, v)$ in HStT, we add to E' all possible edges between these vertices: $(u_h, u_l), (v_h, v_l), (u_h, v_h), (u_l, v_l), (u_h, v_l), (u_l, v_h)$. Edge weights are defined as follows:

$$c'(u_h, u_l) = c'(v_h, v_l) := 0.$$

$$c'(u_l, v_l) = c'(u_h, v_l) = c'(u_l, v_h) := c(u, v).$$

$$c'(u_h, v_h) := \lambda c(u, v).$$

For each $v \in V$, the node weight on v_h is defined as f_v and on v_l as zero. See Fig. 2 for an illustration.

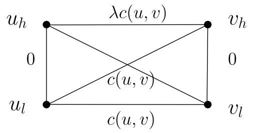

Fig. 2. Edge weights in Reduction 2.

Theorem 3.2. *If we have a γ-approximation algorithm for the NWST problem, then there exists a γ-approximation algorithm for the HStT problem.*

Proof. First, we show that the optimal value of the reduced NWST instance is at most that of the given HStT instance. Let T be a hub Steiner tree T of cost $c(T)$ in the HStT instance. We construct a Steiner tree T' of cost at most $c(T)$ for the reduced NWST instance. For hub-level edges (u,v) in T, we add $(u_h,v_h),(u_h,u_l),(v_h,v_l)$ to T'. For lower-level edges (u,v) in T, we add (u_l,v_l) to T'. It is straightforward to verify that T' has the same cost as T. Next we show T' is indeed a Steiner tree that connects terminals in R'. Consider any pair of nodes (u,v) in R; Since T is a Steiner tree in HStT, there exists a path that connects u and v in T. Call this path P. We will find a path P' in T' that connects u_l and v_l as follows: for any hub-level edge (a,b) in P, add edges $(a_l,a_h),(a_h,b_h),(b_h,b_l)$ to P'. For any lower-level edge (a,b) in P, add an edge (a_l,b_l). It is easy to see that P' indeed connects u_l and v_l.

Next, we prove the opposite direction. Let T' be a feasible Steiner tree spanning R for the NWST instance. We show that there exists a hub Steiner tree T with cost $c(T)\le c'(T')$ for the HStT instance. For each hub-level node $v_h\in V_{\text{hub}}$ spanned by T', we install a hub on v in T. For each edge (u_h,v_h) in T' where $u_h,v_h\in V_{\text{hub}}$, we add an upper level edge (u,v) to T. For edges of the form (u_h,v_l) or (u_l,v_l), we add a lower-level edge (u,v). For the remaining edges in T', we do nothing. Arbitrarily delete edges to remove cycles in T as necessary. It is easy to verify that T connects all terminals of R and has cost no more than $c'(T')$. $\square$

As a corollary of Theorems 3.1 and 3.2, we obtain the following result.

Corollary 3.3. *There is a polynomial-time $2\ln k$-approximation algorithm for the non-metric HStT problem, where k is the number of terminals.*

4. Metric hub Steiner tree

In the previous section, we reduced the HStT problem to the NWST problem. In this section, we show that, if the edge-costs are metric, the HStT problem can be reduced to the edge-weighted Steiner tree (EWST) problem, the special case of the NWST in which all node costs are zero. The EWST problem admits a number of constant-factor approximations. The currently known best approximation factor is $\rho_{ST}=\ln 4+\epsilon\approx 1.38$ [18,19].

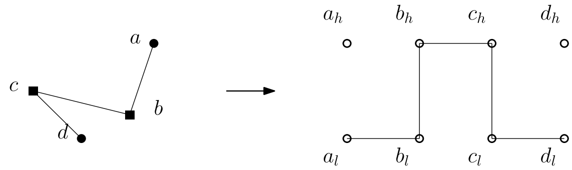

Fig. 3. Convert a HStT to a Steiner tree where the terminal set $R=\{a,b,c,d\}$. On the left, squares (disks, resp.) indicate hubs (non-hubs, resp.). On the right, the corresponding Steiner tree uses two vertical edges, two lower-level edges and one hub-level edge.

Theorem 4.1 *([18,19]). For any constant $\epsilon>0$, there is a polynomial-time $(\ln 4+\epsilon)$-approximation algorithm for the EWST problem.*

Reduction 3. Let V' and E' be the vertex and edge set of the instance we reduce to. Let $c':E'\to\mathbb{R}^+$ be the edge weight function. For each node $v\in V$ in HStT, we create a pair of nodes v_h,v_l. Let V' be the set of all newly created nodes. Define the set of terminals $R':=\{v_l : v\in V\}$. For each edge $e=(u,v)$ in HStT, we add to E' the following edges $(u_h,v_h),(u_l,v_l),(u_h,u_l),(v_h,v_l)$. Edge weights are defined as: $c'(u_h,v_h):=\lambda c(u,v), c'(u_l,v_l):=c(u,v), c'(u_h,u_l):=f_u, c'(v_h,v_l):=f_v$. Call the metric completion of this graph $G'=(V',E')$.

For ease of presentation, we define the following partition of E': $H:=\{(u_h,v_h):u,v\in V\}, L:=\{(u_l,v_l):u,v\in V\}, J:=\{(v_h,v_l),v\in V\}, K:=\{(u_h,v_l),(u_l,v_h):u,v\in V\}$, where H stands for hub-level edges, L for lower-level edges, J for vertical edges and K for cross edges.

Theorem 4.2. *If there exists a γ-approximation algorithm for the EWST problem, then there exists a 2γ-approximation algorithm for the metric HStT problem.*

Proof. First, we show that from a hub Steiner tree T in G, we can construct a Steiner tree T' in G' whose cost is the same as T. Next, we show that for any Steiner tree T' spanning R' in G', we can construct a hub Steiner tree T in G with total cost at most twice the cost of T'.

For the first part, we define a tree T' from T by including all upper-level and lower-level edges in T in addition to each edge of the form (u_l,u_h) that corresponds to installing a hub u in T. Then T' is the required Steiner tree in G'. See Fig. 3 for an illustration.

For the second part, we first partition edges of T' into four sets as follows. Define $E_H:=H\cap T', E_L:=L\cap T', E_J=J\cap T', E_K=K\cap T'$. Recall that for each edge $(u_h,v_l)\in E_K$, there exists a shortest path $P_{u_hv_l}$ from u_h to v_l realizing the distance on this edge which only uses edges from $H\cup L\cup J$. Let $\mathcal{P}$ be the set of such paths, i.e. $\mathcal{P}:=\{P_{u_hv_l}:(u_h,v_l)\in E_K\}$. To construct a hub Steiner tree in G, we install hubs $F:=\{v:(v_h,v_l)\in E_J\cup\{P\cap J:P\in\mathcal{P}\}\}$. We add in all edges (u,v) such that their copies (u_h,v_h) or (u_l,v_l) is in $E_H\cup E_L\cup\{P\backslash J:P\in\mathcal{P}\}$. Let S be the graph constructed as described above. Since T' is a Steiner tree, S guarantees the connectivity for terminals R. We may assume by short-cutting edges that S is a tree. Fig. 4 shows an example of our construction. On the left, solid lines represent edges in the Steiner tree. The dashed path between

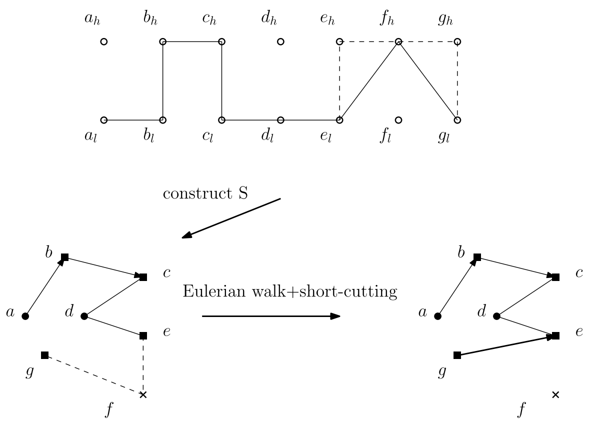

Fig. 4. Convert a Steiner tree to a HStT where the terminal set $R = \{a, b, c, d, e, g\}$. By replacing solid edges with shortest paths, we construct S (bottom left). Its hub level restriction S_H contains two components (subtrees): one containing a single edge (b, c) and the other containing two edges (e, f) and (f, g) where node f is an unhubbed node. By postprocessing (doubling tree edges, taking Eulerian walks and short-cutting on S_H), we obtain a valid HStT (bottom right).

e_l and f_h represents the shortest path between these two nodes. Similarly for the dashed path between f_h and g_l. By definition S contains all solid edges except (e_l, f_h) and (f_h, g_l) which we replace by four dashed edges.

Let S_H be the restriction of S on the hub-level edges (i.e., the edges (u, v) added to S corresponding to an edge $(u_h, v_h) \in H$). S_H may have multiple connected components, each of which may contain *unhubbed nodes* (for which we do not have vertical edges of the form (u_h, u_l) in F). In Fig. 4, the left bottom tree corresponds to S with two components (subtrees): one containing a single edge (b, c) and the other containing two edges (e, f) and (f, g) where node f is an unhubbed node. For each subtree, by doubling the tree, taking an Eulerian walk and short-cutting edges, we can construct a new subtree on only the hubbed nodes with the cost at most doubled w.r.t. the original subtree. The final solution consists of edges from all these new subtrees, as well as edges that are contained in S but not in any of the original subtrees. In Fig. 4, the bottom right shows this solution after postprocessing, which short-cuts the visit to node f.

Notice any two components are connected by our construction, which implies that this solution is connected. Recall S spans the terminal set R. As a result, the solution also spans R. In particular, it means the solution is a valid HStT whose cost is at most twice the cost of the original Steiner tree. The theorem is then proved by combining the two parts. □

We get the following corollary from Theorems 4.1 and 4.2.

Corollary 4.3. *There is a polynomial-time $2\rho_{ST}$-approximation algorithm for the metric HStT problem, where $\rho_{ST} = \ln 4 + \epsilon$ for any constant $\epsilon > 0$.*

5. Conclusion

In this paper, we introduced the HStT problem and presented hardness results and approximation algorithms. Our hardness results rely on reductions from the set cover problem, and approximation algorithms rely on reductions to the node-weighted or edge-weighted instances of Steiner tree.

The generalized network design problem [20] is a well-known generalization of the Steiner tree problem (including the Steiner forest problem e.g.), where the connectivity constraint is specified by a proper set-function over the node set V. The greedy algorithm of Klein and Ravi [14] is known to work for the node-weighted version of this generalized network design problem. It is not hard to verify that Reduction 2 works for this generalization without any modification. As a result, we can obtain an approximation algorithm for the two-level hub version of this generalized network design problem with logarithmic performance ratio. It is also straightforward to extend our reduction of metric HStT to the hub network design version of the generalized network design problem with metric edge costs.

Extending Theorem 2.3 for $\lambda > 0$ is open. Requiring that the hub level network in the solution consists of a single connected component is a practically motivated extension that merits further study in future work.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgement

We thank Patrick Briest of McKinsey & Company who introduced us to the two-level logistic network design problem based on his consulting experience.

References

[1] M. Andrews, L. Zhang, The access network design problem, in: Proceedings. 39th Annual Symposium on the Foundations of Computer Science, IEEE, 1998, pp. 40–49.

[2] A. Balakrishnan, T.L. Magnanti, P. Mirchandani, Modeling and heuristic worst-case performance analysis of the two-level network design problem, Manag. Sci. 40 (7) (1994) 846–867.

[3] T. Carpenter, H. Luss, Telecommunications access network design, in: Handbook of Optimization in Telecommunications, Springer, 2006, pp. 313–339.

[4] M. Andrews, Hardness of buy-at-bulk network design, in: 45th Annual IEEE Symposium on Foundations of Computer Science, 2004. Proceedings, IEEE, 2004, pp. 115–124.

[5] C. Chekuri, M.T. Hajiaghayi, G. Kortsarz, M.R. Salavatipour, Approximation algorithms for nonuniform buy-at-bulk network design, SIAM J. Comput. 39 (5) (2010) 1772–1798.

[6] F.S. Salman, J. Cheriyan, R. Ravi, S. Subramanian, Approximating the single-sink link-installation problem in network design, SIAM J. Control Optim. 11 (3) (2001) 595–610.

[7] K. Talwar, The single-sink buy-at-bulk lp has constant integrality gap, in: International Conference on Integer Programming and Combinatorial Optimization, Springer, 2002, pp. 475–486.

[8] S. Guha, A. Meyerson, K. Munagala, A constant factor approximation for the single sink edge installation problem, SIAM J. Comput. 38 (6) (2009) 2426–2442.

[9] A. Gupta, A. Kumar, T. Roughgarden, Simpler and better approximation algorithms for network design, in: Proceedings of the Thirty-Fifth Annual ACM Symposium on Theory of Computing, ACM, 2003, pp. 365–372.

[10] S. Alumur, B.Y. Kara, Network hub location problems: the state of the art, Eur. J. Oper. Res. 190 (1) (2008) 1–21.

[11] R.Z. Farahani, M. Hekmatfar, A.B. Arabani, E. Nikbakhsh, Hub location problems: a review of models, classification, solution techniques, and applications, Comput. Ind. Eng. 64 (4) (2013) 1096–1109.

[12] G. Nagy, S. Salhi, Location-routing: issues, models and methods, Eur. J. Oper. Res. 177 (2) (2007) 649–672.

[13] C. Prodhon, C. Prins, A survey of recent research on location-routing problems, Eur. J. Oper. Res. 238 (1) (2014) 1–17.

[14] P. Klein, R. Ravi, A nearly best-possible approximation algorithm for node-weighted Steiner trees, J. Algorithms 19 (1) (1995) 104–115.

[15] M. Bentert, R. van Bevern, A. Nichterlein, R. Niedermeier, P.V. Smirnov, Parameterized algorithms for power-efficiently connecting wireless sensor networks: theory and experiments, INFORMS J. Comput. (2021).

[16] A. Erzin, R. Plotnikov, Y.V. Shamardin, On some polynomially solvable cases and approximate algorithms in the optimal communication tree construction problem, J. Appl. Ind. Math. 7 (2) (2013) 142–152.

[17] I. Dinur, D. Steurer, Analytical approach to parallel repetition, in: Proceedings of the Forty-Sixth Annual ACM Symposium on Theory of Computing, ACM, 2014, pp. 624–633.

[18] J. Byrka, F. Grandoni, T. Rothvoß, L. Sanità, Steiner tree approximation via iterative randomized rounding, J. ACM 60 (1) (2013) 6.

[19] M.X. Goemans, N. Olver, T. Rothvoß, R. Zenklusen, Matroids and integrality gaps for hypergraphic Steiner tree relaxations, in: Proceedings of the 44th Symposium on Theory of Computing Conference, STOC 2012, New York, NY, USA, May 19-22, 2012, ACM, 2012, pp. 1161–1176.

[20] M.X. Goemans, D.P. Williamson, A general approximation technique for constrained forest problems, SIAM J. Comput. 24 (2) (1995) 296–317.