On a Leasing Variant of the Online Connected Facility Location Problem

Murilo Santos de Lima*1, Mário César San Felice†2, Orlando Lee‡1

¹Institute of Computing – Campinas State University (Unicamp) Av. Albert Einstein, 1251 – Cidade Universitária – 13083-852 – Campinas – SP – Brazil

²Institute of Mathematics and Statistics – University of São Paulo (USP) Rua do Matão, 1010 – 05508-090 – São Paulo – SP – Brazil

{ra137105, felice, lee}@ic.unicamp.br

Abstract. In the leasing optimization model, resources are leased for K different time periods, instead of being acquired for unlimited duration. The goal is to use these temporary resources to maintain a dynamic infrastructure that serves n requests while minimizing the total cost. We propose and study a leasing variant of the online connected facility location problem, which we call the **online connected facility leasing** problem. In this problem each client that arrives must be connected to a temporary facility, which in turn must be connected to a root facility using permanent edges. We present an algorithm that is $O(K \cdot \lg n)$ -competitive if the scaling factor is M = 1.

1. Introduction

In traditional network design problems, the goal is to acquire resources corresponding to nodes or connections in order to build a minimum-cost network infrastructure that serves client requests. For example, in the **facility location** (FL) problem, we wish to determine where to install servers to satisfy connection requests by a set of clients.

However, in many scenarios, it is not plausible for a small company to install its own servers, and requests may be concentrated in certain time intervals. Therefore, it may be cheaper to **lease** a server in the cloud (from a service like Amazon AWS, Google App Engine or Microsoft Azure) for specific time intervals. This scenario has recently motivated the proposal of a class of problems known as **leasing optimization** [Meyerson 2005, Anthony and Gupta 2007, Nagarajan and Williamson 2013, Abshoff et al. 2015]. In those problems, instead of acquiring resources that last for unlimited duration, each resource may be leased for different time periods (e.g., a day, a week, a month). In this context, it is cheaper to lease a resource for a longer period than for smaller periods totalizing an equivalent time; this represents an economics of scale. This model may be applied to **offline** problems, which are traditional optimization problems in which the whole input is known in advance and there is no causality restriction for a solution, and to **online** problems, in which at every instant a request arrives and must be served without knowledge of future requests, and previous decisions cannot be revoked.

^{*}Supported by FAPESP Proc. 2014/18781-1, and CNPq Proc. 142161/2014-4.

[†]Supported by CAPES PNPD scholarship 1522390.

[‡]Supported by CNPq Proc. 311373/2015-1, and CNPq Proc. 477692/2012-5.

In the **connected facility location** (CFL) problem, we are given a graph with edge costs, a subset F of the vertices (potential facilities), an opening cost for each facility, a subset D of the vertices (clients) and a constant M. The objective is to select a subset of the facilities to open and a subset of the edges that connects the open facilities, so to minimize the cost of the open facilities, plus M times the cost of the edges that connect the facilities, plus the sum of the distances from each client to its closest open facility.

In this paper we propose a leasing variant of this problem, which we call the connected facility leasing (CFLE) problem. In this version, we are given a set F of potential facilities, and clients arriving in different times that we wish to connect to open facilities. However, in the leasing version we no longer buy a facility for an unlimited time, but instead we can lease each facility for one of K different periods of time. Suppose those periods of time have lengths $\delta_1, \dots, \delta_K$, and leasing a facility $f \in F$ for a lease of type k costs $\gamma_f^k \in \mathbb{R}_+$. If we lease a facility f with leasing type k at instant \hat{t} , then we say that facility lease (f, k, \hat{t}) is **active** during interval $[\hat{t}, \hat{t} + \delta_k)$. At every instant $t \in \{0, \dots, T-1\}$ we receive a set D_t of clients which we wish to connect to active facilities. We are given a distance d(f, j) between each potential facility f and each client j. We also have to connect the facilities to a root facility r; however the edges that connect facilities to the root cost M times the cost of an edge connecting a client to a facility, just like a backbone cable costs more than a common cable for a network service provider. The objective is to minimize the cost of leasing the facilities, plus the sum of the distances between each client and the facility to which it is connected, plus M times the cost of the edges connecting the facilities to the root. Note that this problem has a natural temporal requirement, even in an offline setting. Note, also, that we have leasing just on facilities: this models, for example, the problem of a network service provider, who has to install cables between routers to serve clients, but has routers with different limited lifetimes. Our focus will be in online algorithms, in which clients are given in an online manner, facility leases and core edges cannot be removed, and a new facility lease cannot be used to serve an old client request.

Related Work. [Meyerson 2005] proposed the **parking permit** (PP) problem, which has become the central problem for the leasing optimization model. In the same paper, Meyerson studied the online leasing version for the **Steiner forest** (SF) problem.

[Anthony and Gupta 2007] presented approximation algorithms for offline leasing versions of several NP-hard network design problems. For the **facility leasing** (FLE) problem, they obtained an O(K)-approximation, where K is the number of lease types.

[Nagarajan and Williamson 2013] improved the result for the offline version of problem FLE, obtaining a 3-approximation. For the online version, the **online facility leasing** (OFLE) problem the authors showed an algorithm with competitive factor $O(K \lg n)$, where n is the number of client requests.

[Abshoff et al. 2015] improved the result on problem OFLE, presenting an online $O(\delta_K \lg \delta_K)$ -competitive algorithm, where δ_K is the longest lease duration.

Problem CFL is a generalization of the Steiner tree (ST) problem, so it is NP-hard, and the best known offline algorithm for problem CFL is a 3.19-approximation [Grandoni and Rothvoß 2011]. The **online connected facility location** (OCFL) problem, similarly, has lower bound $\Omega(\lg n)$ for its competi-

tive factor due to the lower bound on the **online Steiner tree (OST)** problem [Imase and Waxman 1991]. There are known $O(\lg n)$ -competitive online algorithms for problem OCFL [Umboh 2015, San Felice et al. 2016].

Our Contribution. We propose a first leasing variant of problem CFL, and give an $O(K \lg n)$ -competitive online algorithm for the case when the scaling factor is M = 1.

Text Organization. In Section 2 we give a formal definition of the problem and present our algorithm. In Section 3 we present some conclusions and future research directions.

2. Online Simple Connected Facility Leasing

We give the following definition of the problem. To simplify our notation, let $\mathcal{D} := \{(j,t): j \in D_t, t \in \{0,\ldots,T-1\}\}$ be the set of all client requests; we denote a client request as a pair (j,t) since the same client can request a connection at different instants.

Problem CFLE $(G,d,r,F,K,\gamma,\delta,M,D_0,\ldots,D_{T-1})$: The input consists of a graph $G=(V_G,E_G)$, a metric distance function $d:V_G\times V_G\mapsto \mathbb{R}_+$ on the vertices of G, a root vertex $r\in V_G$, a set $F\subseteq V_G$ of potential facilities, an integer K>0 that represents the number of lease types, a cost $\gamma_f^k\in\mathbb{R}_+$ for leasing facility $f\in F$ with lease type $k\in [K]$, a function $\delta:[K]\mapsto \mathbb{N}$ that maps each lease type to a time length, a constant $M\in\mathbb{R}_+$, and a sequence $D_0,\ldots,D_{T-1}\subseteq V_G$ of client sets. The goal is to find a set $X\subseteq F\times [K]\times \{0,\ldots,T-1\}$ of leased facilities, a function $a:\mathcal{D}\mapsto X$ that maps each client request (j,t) to a facility leasing $(f,k,\hat{t})\in X$ such that $t\in [\hat{t},\hat{t}+\delta_k)$, and a set $T\subseteq E_G$ of edges that connects the leased facilities in X to r, so to minimize $\sum_{(f,k,\hat{t})\in X}\gamma_f^k+\sum_{(j,t)\in\mathcal{D}}d(a(j,t),j)+M\cdot\sum_{e\in\mathcal{T}}d(e)$.

We suppose that $\delta_1 \leq \delta_2 \leq \ldots \leq \delta_K$, and that $\gamma_f^k/\delta_k \leq \gamma_f^\ell/\delta_\ell$ for every $f \in F$ and $k > \ell$.

Online Algorithm. We propose the following online algorithm, based on the algorithm by [San Felice et al. 2016] for problem OCFL. Our algorithm uses the primal-dual algorithm by [Nagarajan and Williamson 2013] for problem OFLE, which we denote by NWFLE and has offline parameters $(G, d, F, K, \gamma, \delta)$ and an online sequence of client sets. Algorithm NWFLE assumes that facility leases occur only in specific moments; i.e., $X \subseteq \{(f, k, c \cdot \delta_k) : f \in F, k \in [K], c \in \mathbb{Z}_+\}$. This assumption increases the competitive factor of any algorithm for problem CFLE at most by a constant ratio.

Our algorithm maintains a set \mathcal{D} of client requests, and a set $\mathcal{D}^{\mathcal{T}}$ of clients that are connected to the core network. At each client request, the algorithm forwards the request to algorithm NWFLE and updates the virtual solution. Then, the client is added to the core network with probability 1/M. In that case, the algorithm adds to the core network the shortest path between the client and the tree, opens the facility indicated by the virtual solution and adds another core edge between the client and the new facility. If the client is not added to the core network, then it is served by the closest active facility lease. Note that we ensure that the root facility is always active at zero cost.

We have a proof that this algorithm is $O(K \cdot \lg n)$ -competitive for the special case when M = 1, where n is the total number of client requests.¹

¹The proof is available at http://www.ime.usp.br/~mslima/appendix-etc2016.pdf.

```
Algorithm OCFLE-ALG(G, d, r, F, K, \gamma, \delta, M)
         set \delta_r^K \leftarrow 0, X \leftarrow \emptyset, \mathcal{D} \leftarrow \emptyset, \mathcal{D}^{\mathcal{T}} \leftarrow \emptyset, \mathcal{T} \leftarrow (\{r\}, \emptyset)
02
        initialize NWFLE with (G, d, F, K, \gamma, \delta); let (X', a') be the virtual solution
03
        upon receipt of D_t do
              if t \mod \delta_K = 0 then X \leftarrow X \cup \{(r, K, t)\}; send (r, t) to NWFLE
04
05
              for each j \in D_t do
                   send (j, t) to NWFLE and update the virtual solution (X', a')
06
                   include j in \mathcal{D}^{\mathcal{T}} with probability 1/M
07
                   if i \in \mathcal{D}^{\mathcal{T}} then
08
                        \mathcal{T} \leftarrow \mathcal{T} \cup \text{path}(j, V(T)), (f, k, \hat{t}) \leftarrow a'(j, t)
09
10
                        if (f, k, \hat{t}) is not opened then
                              X \leftarrow X \cup \{(f, k, \hat{t})\}, \mathcal{T} \leftarrow \mathcal{T} \cup \{(f, j)\}
11
12
                   choose (f, k, \hat{t}) \in X active in t and closest to j
13
                   \mathcal{D} \leftarrow \mathcal{D} \cup \{(j,t)\}, a(j,t) \leftarrow (f,k,\hat{t})
14
        return (X, a, \mathcal{T})
```

3. Conclusion and Future Work

We propose a leasing variant of the online connected facility location problem, and give an algorithm which is $O(K \cdot \lg n)$ -competitive if M=1. We wish to prove the same competitive factor for arbitrary M, maybe by modifying the algorithm so to take into account core edges when deciding to lease a facility; currently we have partial results.

References

- Abshoff, S., Kling, P., Markarian, C., auf der Heide, F. M., and Pietrzyk, P. (2015). Towards the price of leasing online. *Journal of Combinatorial Optimization*, pages 1–20.
- Anthony, B. M. and Gupta, A. (2007). Infrastructure leasing problems. In *Integer Programming and Combinatorial Optimization*, pages 424–438. Springer.
- Grandoni, F. and Rothvoß, T. (2011). Approximation algorithms for single and multi-commodity connected facility location. In Günlük, O. and Woeginger, G. J., editors, *Integer Programming and Combinatoral Optimization*, volume 6655 of *Lecture Notes in Computer Science*, pages 248–260. Springer Berlin Heidelberg.
- Imase, M. and Waxman, B. M. (1991). Dynamic Steiner tree problem. *SIAM Journal on Discrete Mathematics*, 4(3):369–384.
- Meyerson, A. (2005). The parking permit problem. In *Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science*, pages 274–282.
- Nagarajan, C. and Williamson, D. P. (2013). Offline and online facility leasing. *Discrete Optimization*, 10(4):361–370.
- San Felice, M. C., Williamson, D. P., and Lee, O. (2016). A randomized $O(\log n)$ -competitive algorithm for the online connected facility location. *Algorithmica*, pages 1–19.
- Umboh, S. (2015). Online network design algorithms via hierarchical decompositions. In *Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1373–1387.