

A Nearly Optimal Deterministic Online Algorithm for Non-Metric Facility Location

Marcin Bienkowski  

Institute of Computer Science, University of Wrocław, Poland

Björn Feldkord  

Heinz Nixdorf Institut & Department of Computer Science, Universität Paderborn, Germany

Paweł Schmidt 

Institute of Computer Science, University of Wrocław, Poland

Abstract

In the online non-metric variant of the facility location problem, there is a given graph consisting of a set F of facilities (each with a certain opening cost), a set C of potential clients, and weighted connections between them. The online part of the input is a sequence of clients from C , and in response to any requested client, an online algorithm may open an additional subset of facilities and must connect the given client to an open facility.

We give an online, polynomial-time deterministic algorithm for this problem, with a competitive ratio of $O(\log |F| \cdot (\log |C| + \log \log |F|))$. The result is optimal up to loglog factors. Our algorithm improves over the $O((\log |C| + \log |F|) \cdot (\log |C| + \log \log |F|))$ -competitive construction that first reduces the facility location instance to a set cover one and then later solves such instance using the deterministic algorithm by Alon et al. [TALG 2006]. This is an asymptotic improvement in a typical scenario where $|F| \ll |C|$.

We achieve this by a more direct approach: we design an algorithm for a fractional relaxation of the non-metric facility location problem with clustered facilities. To handle the constraints of such non-covering LP, we combine the dual fitting and multiplicative weight updates approach. By maintaining certain additional monotonicity properties of the created fractional solution, we can handle the dependencies between facilities and connections in a rounding routine.

Our result, combined with the algorithm by Naor et al. [FOCS 2011] yields the first deterministic algorithm for the online node-weighted Steiner tree problem. The resulting competitive ratio is $O(\log k \cdot \log^2 \ell)$ on graphs of ℓ nodes and k terminals.

2012 ACM Subject Classification Theory of computation \rightarrow Online algorithms; Theory of computation \rightarrow Routing and network design problems

Keywords and phrases Online algorithms, deterministic rounding, linear programming, facility location, set cover

Digital Object Identifier 10.4230/LIPIcs.STACS.2021.14

Funding Supported by Polish National Science Centre grant 2016/22/E/ST6/00499 and by German Research Foundation (DFG) within the Collaborative Research Centre “On-The-Fly Computing” under the project number 160364472 – SFB 901/3.

Acknowledgements We thank Marek Adamczyk and Christine Markarian for helpful discussions. We thank anonymous reviewers of an earlier draft for pointing us to the reduction of Kolen and Tamir [20].

1 Introduction

The facility location (FL) problem [1] is one of the best-known examples of network design problems, extensively studied both in operations research and in computer science. Its simple definition, NP-hardness, and rich combinatorial structure have led to developments of tools and solutions in key areas of approximation algorithms, combinatorial optimization, and linear programming.



© Marcin Bienkowski, Björn Feldkord, and Paweł Schmidt;
licensed under Creative Commons License CC-BY 4.0

38th International Symposium on Theoretical Aspects of Computer Science (STACS 2021).

Editors: Markus Bläser and Benjamin Monmege; Article No. 14; pp. 14:1–14:17

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



An instance of the FL problem consists of a set F of facilities, each with a certain opening cost, and a set C of clients. F and C can be seen as two sides of a bipartite graph. The undirected edges between them have lengths that can either satisfy the triangle inequality (*metric FL*) or be arbitrary (*non-metric FL*). The goal is to open a subset of facilities and connect each client to an open facility. The total cost (the sum of opening and connection costs) is subject to minimization. In the metric scenario, by taking a metric closure, one can assume that each facility is reachable by each client, but it is not the case for the non-metric variant.

Instances and Objectives. In this paper, we focus on an online variant of the non-metric FL problem. We first formalize the offline variant in a way that makes a connection to the online variant more apparent.

A *facility-client graph* $G = (F, C, E, \text{cost})$ is a bipartite graph, whose one side is the set F of facilities and another side is the set of clients C . Set $E \subseteq F \times C$ contains available facility-client connections (edges). We use function cost to denote both costs of opening facilities and connection costs (edge lengths). All costs are non-negative.

An instance of the non-metric FL problem is a pair (G, A) , where $G = (F, C, E, \text{cost})$ is a facility-client graph and $A \subseteq C$ is a subset of *active* clients. A feasible solution to such instance is a set of open (purchased) facilities $F' \subseteq F$ and a subset of purchased edges $E' \subseteq E$, such that any active client $c \in A$ is connected by a purchased edge to an open facility. The cost of such solution is equal to $\sum_{f \in F'} \text{cost}(f) + \sum_{e \in E'} \text{cost}(e)$.

For any facility-client graph G , we define its aspect ratio Δ_G as the ratio of the largest to smallest positive cost in G . These costs include both facilities and connection costs.¹ Note that the aspect ratio is a property of G and is independent of the set of active clients A .

Online Scenario. In an *online variant* of the FL problem, the facility-client graph G is known in advance, but neither elements of A nor its cardinality are known up-front by an online algorithm ALG . The clients from A appear one by one. Upon seeing a new active client, ALG may purchase additional facilities and edges, with the requirement that facilities and edges purchased so far must constitute a feasible solution to all presented active clients. The total cost of ALG is denoted by $\text{ALG}(G, A)$. (We sometimes use $\text{ALG}(G, A)$ to also denote the solution computed by ALG .) Purchase decisions are final and cannot be revoked later. The goal is to minimize *the competitive ratio*, defined as $\sup_{(G,A)} \{\text{ALG}(G, A)/\text{OPT}(G, A)\}$, where OPT is the optimal (offline) algorithm.

1.1 Related Work

Most of the prior work has been devoted to the offline scenario. While the metric variant of the FL problem admits $O(1)$ -approximation algorithms [9, 10, 11, 18, 19, 22, 23, 24, 29], the best approximation ratio for the non-metric one is $O(\log |A|)$ [17], and it cannot be asymptotically improved unless $\text{NP} \subseteq \text{DTIME}[n^{O(\log \log n)}]$ [12]. For a more comprehensive treatment of the offline scenario, including a multitude of variants, we refer the reader to the entry in the Encyclopedia of Algorithms [1] or the survey by Shmoys [28].

For the online metric FL, the problem was resolved over ten years ago by Meyerson [25] and Fotakis [14]: the lower and upper bounds on the competitive ratio are $\Theta(\log |A| / \log \log |A|)$, both for deterministic and randomized algorithms. Simpler deterministic algorithms attaining

¹ In the standard definition of the aspect ratio, only distances are taken into account.

slightly worse competitive ratio of $O(\log |A|)$ were given by Anagnostopoulos et al. [4] and Fotakis [13]. Note that the optimal competitive ratio in the metric case is independent of the set C of potential clients.

1.2 Previous Work on Online Non-Metric Facility Location

For the non-metric FL, the first and currently best online algorithm was a *randomized* algorithm by Alon et al. [2]. It achieves the competitive ratio of $O(\log |F| \cdot \log |A|)$. It is based on solving a natural fractional relaxation of the problem: there is a fractional *opening variable* y_f for each facility f and a *connection variable* $x_{c,f}$ for a client c and a *covering facility* f (facility to which c could be connected). Once a client c arrives, for each covering facility f independently, their algorithm increases either y_f or $x_{c,f}$, whichever is smaller, using multiplicative update method (see, e.g., [5]). The client c is considered fractionally served once the sum of terms $\min\{x_{c,f}, y_f\}$ over all covering facilities is at least 1. The resulting competitive ratio is $O(\log |F|)$.

The computed fractional solution can be then rounded using a random threshold θ_f common for an opening variable y_f and all connection variables involving facility f . Once any variable exceeds its threshold, it is rounded up to 1 and the corresponding object (facility or connection) is purchased. Dynamically adjusting θ_f to have expectation $\Theta(1/\log |A|)$ guarantees that the resulting integral solution is feasible with high probability and the rounding part incurs a factor of $O(\log |A|)$ in the competitive ratio.

To the best of our knowledge, no non-trivial deterministic algorithm was published so far. In particular, the online network design problems (including the non-metric FL problem) have been listed as unresolved challenges by Buchbinder and Naor [8, Section 1.1]. That said, the non-metric facility location can be reduced to a set cover. A usable reduction (not inducing an exponential blow-up of the input size) was given by Kolen and Tamir [20]: it preserves the solution costs up to constant factors and creates a set cover instance consisting of $m = \Theta(|F| + |C| \cdot \log \Delta_G)$ sets and $n = \Theta(|C| \cdot \log \Delta_G)$ elements. Using doubling techniques described in Section 4, one could assume that $\Delta_G = O(|F| \cdot |C|)$. Applying the deterministic algorithm for the online set cover problem by Alon et al. [3] yields a solution whose competitive ratio is $O(\log m \cdot \log n) = O((\log |C| + \log |F|) \cdot (\log |C| + \log \log |F|))$.

1.3 Our Result

In our paper, we improve the bound above, replacing the first factor of $O(\log |C| + \log |F|)$ by $O(\log |F|)$. This is an asymptotic improvement in a typical scenario where $|F| \ll |C|$.

► **Theorem 1.** *There exists a deterministic polynomial-time $O(\log |F| \cdot (\log |C| + \log \log |F|))$ -competitive algorithm for the online non-metric facility location problem on set F of facilities and set C of clients.*

Our algorithm attains a nearly optimal competitive ratio, as no deterministic algorithm can have a ratio smaller than $\Omega(\log |F| \cdot \log |C| / (\log \log |F| + \log \log |C|))$. This follows by the lower bound for the online set cover problem [2, 3] and holds even for uniform facility costs. If we restrict our attention to the *polynomial-time* deterministic solution, then a stricter lower bound of $\Omega(\log |F| \cdot \log |C|)$ holds (assuming $\text{BPP} \neq \text{NP}$) [21].

Challenges. The description of the randomized algorithm by Alon et al. [2] given above seems deceptively simple, but it hides an important and subtle property, implicitly exploited by the authors. Namely, the threshold θ_f is common for facility f and all connections to it.

This ensures the necessary dependency: once $\min\{x_{c,f}, y_f\} \geq \theta_f$, the rounding purchases *both facility f and a connection from c to f* . (Note that the left-hand side of this inequality is the amount that their fractional solution controls.)

It is unclear how to directly extend this property to deterministic rounding. A straightforward attempt would be to focus on facilities only and round them in a deterministic fashion ensuring the necessary coverage of each client. However, neglecting the connection costs in the rounding process easily leads to a situation, where the facilities are rounded “correctly”, but the cost of connecting a client to the closest open facility in the integral solution is incomparably larger than the corresponding fractional cost.

We note that all known deterministic schemes that round fractional solutions generated by the multiplicative updates operate in rather limited scenarios, where elements have to be covered or packed and all important interactions between elements are handled at the time of constructing the fractional solution. This is the case for the deterministic rounding for the set cover problem [3, 7] and the throughput-competitive virtual circuit routing problem [6, 8]. These methods are based on derandomizing the method of pessimistic estimators [27] in an online manner, by transforming a pessimistic estimator into a potential function [30] that can be controlled by the deterministic rounding process.

Our Techniques. In our solution, we create a new linear relaxation of the problem. We first round the graph distances to powers of 2. For any client, we cluster facilities that have the same distance to this client. (Note that such clusters are client-dependent.) To solve the fractional variant, we run two schemes in parallel: we increase connection variables $x_{c,t}$ corresponding to clusters at distance t , and increase facility variables y_f for all facilities in “reachable” clusters (where the corresponding connection variables are 1). The increases in these variables use two different frameworks: dual fitting for linear increases of connection variables and a primal-dual scheme involving multiplicative updates for facility variables. Ensuring an appropriate balance between these two different types of updates is one of the technical difficulties that we tackle in this paper.

We stop increasing variables once there exists a collection of clusters that are both “fractionally open” (sum of variables y_f within these clusters is $\Omega(1)$) and “reachable” by the considered client. To argue about the existence of such a collection, we use both LP inequalities and structural properties of our fractional algorithm.

Finally, we construct a deterministic rounding routine. We focus on facilities only, neglecting whether particular clients are active or not and how far they are from a given facility. However, we strengthen rounding properties, ensuring, for (some) collections of clusters, that if the sum of opening variables in these collections is $\Omega(1)$, then the integral solution contains an open facility in one of these clusters. This ensures that, for a considered client c , the integral solution contains a facility whose distance from c is asymptotically not larger than the cost invested for connecting c in the fractional solution. Ultimately, this yields the desired dependency between facilities and connections.

Note about Up-Front Knowledge of the Facility-Client Graph. Unlike for the randomized variant, obtaining sub-linear guarantees for a deterministic solution requires knowing a priori the set of potential client-facility connections. To see this, consider a graph of $|F|$ facilities with unit opening costs and the set of $|C| = |F|$ clients. The graph edges are constructed dynamically as clients are activated and all revealed possible connections are of cost 0. The first active client can be connected to all facilities. Each subsequent client can be connected to all facilities but the ones already open by an algorithm. This way an online algorithm

needs to eventually open all facilities, for a total cost of $|F|$. On the other hand, the offline optimal algorithm can open the last facility opened by an online algorithm and connect all clients to this facility paying just 1. Thus, under the unknown-graph assumption, the competitive ratio of any deterministic algorithm would be at least $|F|$.

1.4 Preliminaries and Paper Organization

Let T_G contain all powers of two between the largest and the smallest positive distance (inclusively) and also number 0. In particular, T_G contains all distances in G and $|T_G| \leq 2 + \log \Delta_G$. Whenever G is clear from the context, we drop the G subscript.

We may assume that F contains at least two facilities and C contains at least two clients, as otherwise the problem becomes trivial. For a facility $f \in F$, let $\text{set}(f)$ be the set of clients that may be connected to f . For any client $c \in C$ and distance $t \in T$, cluster $F_{c,t}$ contains all facilities that are incident to c using edges of cost t . Note that for a fixed c , clusters $F_{c,t}$ are disjoint (no client has two connections of different costs to the same facility).

Powers-of-Two Assumption. In the whole paper, we assume that all facilities and connection costs are either equal to 0 or are powers of 2 and are at least 1. This can be easily achieved by initial scaling of positive costs and distances, so that they are at least 1 and rounding positive ones up to the nearest power of two. This transformation changes the competitive ratio at most by a factor of 2.

Paper Overview. Our core approach is to solve a carefully crafted fractional relaxation of the problem (Section 2), and then round it in a deterministic fashion (Section 3). This way, we obtain a deterministic online algorithm INT that on any input $(G = (F, C, E, \text{cost}), A)$ computes a feasible solution of cost

$$\text{INT}(G, A) \leq O(\log |F| \cdot (\log |C| + \log \log \Delta_G)) \cdot \text{OPT}(G, A) + 2 \cdot \max_{f \in F} \text{cost}(f).$$

Moreover INT runs in time $\text{poly}(|G|, |A|, \max_{e \in E} \text{cost}(e), \max_{f \in F} \text{cost}(f))$. In Section 4, we apply doubling and edge pruning techniques, to get rid of dependencies on costs in the running time and on Δ_G in the competitive ratio, achieving guarantees of Theorem 1.

Application to Node-Weighted Steiner Tree. Our result has an immediate application to the online node-weighted Steiner tree (NWST) problem. Namely, when we combine Theorem 1 with the randomized solution for NWST by Naor et al. [26], we obtain the first deterministic algorithm with polylogarithmic competitive ratio (see Section 5).

2 Fractional Solution

We fix an instance $(G = (F, C, E, \text{cost}), A)$ of the online non-metric facility problem. For each facility f , we introduce an *opening* variable $y_f \geq 0$ (fractional opening of f) and for each client c and each distance $t \in T$ a *connection* variable $x_{c,t} \geq 0$. Intuitively, $x_{c,t}$ denotes how much, fractionally, client c invests into connections to facilities from cluster $F_{c,t}$. For any set F' of facilities we use $y(F')$ as a shorthand for $\sum_{f \in F'} y_f$.

Primal Program. After k clients from A arrive (we denote their set by A_k), we consider the following linear program \mathcal{P}_k .

$$\begin{aligned}
 & \text{minimize} && \sum_{f \in F} \text{cost}(f) \cdot y_f + \sum_{c \in A_k} \sum_{t \in T} t \cdot x_{c,t} \\
 & \text{subject to} && x_{c,t} \geq z_{c,t} && \text{for all } c \in A_k, t \in T, \\
 & && y(F_{c,t}) \geq z_{c,t} && \text{for all } c \in A_k, t \in T, \\
 & && \sum_{t \in T} z_{c,t} \geq 1 && \text{for all } c \in A_k,
 \end{aligned}$$

and non-negativity of all variables.

Serving Constraints. The LP constraints combined are equivalent to the set of the following (non-linear) requirements

$$\sum_{t \in T} \min\{x_{c,t}, y(F_{c,t})\} \geq 1 \quad \text{for all } c \in A_k. \quad (1)$$

We call (1) for client c the *serving constraint* for client c . In our description, we omit variables $z_{c,t}$ and the original constraints, ensuring only that the serving constraints hold and implicitly setting $z_{c,t} = \min\{x_{c,t}, y(F_{c,t})\}$.

The LP above is indeed a valid relaxation of the FL problem. To see this, take any feasible integral solution. For any facility f opened in the integral solution, set variable y_f to 1. For each client c connected to facility f , set variable $x_{c,\tau}$ to 1, where $\tau = \text{cost}(f, c)$. This guarantees that $\min\{x_{c,\tau}, y(F_{c,\tau})\} = 1$, and thus the serving constraint (1) is satisfied for each client c .

Dual Program. The program \mathcal{D}_k dual to \mathcal{P}_k is

$$\begin{aligned}
 & \text{maximize} && \sum_{c \in A_k} \gamma_c \\
 & \text{subject to} && \gamma_c \leq \alpha_{c,t} + \beta_{c,t} && \text{for all } c \in A_k, t \in T, \\
 & && \alpha_{c,t} \leq t && \text{for all } c \in A_k, t \in T, \\
 & && \sum_{c \in \text{set}(f) \cap A_k} \beta_{c, \text{cost}(f,c)} \leq \text{cost}(f) && \text{for all } f \in F,
 \end{aligned}$$

and non-negativity of all variables.

2.1 Overview

Our algorithm FRAC creates a solution to \mathcal{P}_k , ensuring that the serving constraint (1) holds for all clients $c \in A_k$. As outlined in the introduction, the computed solution guarantees some additional properties that are useful for the rounding part later.

Whenever a client c arrives, FRAC increases connection variables $x_{c,t}$ one by one starting from the smallest t , at the pace proportional to $1/t$. We ensure that $x_{c,t} \in [0, 1]$, i.e., once any of these variables reaches 1, FRAC stops increasing them. A distance t , for which $x_{c,t} = 1$, is called *saturated*.

In parallel to manipulating variables $x_{c,t}$, FRAC increases all variables y_f for facilities reachable from client c using saturated distances. The variables y_f are increased using the multiplicative update rule [5] (scaled appropriately to take costs of facilities into account).

Together with the solution to \mathcal{P}_k , FRAC also constructs an *almost-feasible* solution to \mathcal{D}_k . That is, its solution to \mathcal{D}_k is feasible when all dual variables are scaled down by a factor of $O(\log |F|)$. By the weak duality, the scaled-down value of this solution serves as a lower-bound for the optimum. Thus, as typical for the primal-dual type of analysis, the dual variables can be thought of as budgets whose increase balances the increase of primal variables.

2.2 Algorithm FRAC

At the very beginning, before any client arrives, FRAC sets all variables y_f to 0 for all positive-cost facilities and to 1 for zero-cost ones. There are no other variables as the set A_0 of active clients is empty. Note that the dual program already contains the last type of constraints, but the sums on their left-hand sides range over empty sets of β variables, and hence these constraints are trivially satisfied.

Whenever a new client c arrives in step k , FRAC updates the primal (dual) programs from \mathcal{P}_{k-1} (\mathcal{D}_{k-1}) to \mathcal{P}_k (\mathcal{D}_k), and then computes a feasible solution to \mathcal{P}_k (based on the already created solution to \mathcal{P}_{k-1}) and a nearly-feasible solution to \mathcal{D}_k .

New variables in primal and dual programs: FRAC sets $x_{c,t} \leftarrow 0$ for all $t \in T \setminus \{0\}$ and sets $x_{c,0} \leftarrow 1$. In the dual solution, it sets $\gamma_c \leftarrow 0$, $\alpha_{c,t} \leftarrow 0$ and $\beta_{c,t} \leftarrow 0$ for all $t \in T$.

Update primal program: A new serving constraint $\sum_{t \in T} \min\{x_{c,t}, y(F_{c,t})\} \geq 1$ appears in the primal program (and is violated unless $y(F_{c,0}) \geq 1$). As we never decrease primal variables, the serving constraints (1) that existed already in \mathcal{P}_{k-1} are satisfied and will not become violated.

Update dual program: New constraints appear in the dual program and new variables $\beta_{c,t}$ appear on the left-hand side of the already existing inequalities. Since the new variables are initialized to 0, the validity of all dual constraints is unaffected.

Update primal and dual solutions: Let $T_c^1 = \{t \in T : x_{c,t} \geq 1\}$ be the set of saturated distances, i.e., initially FRAC sets $T_c^1 \leftarrow \{0\}$. While the serving constraint for c is violated, FRAC executes the *update operation* consisting of the following steps:

1. Set $\gamma_c \leftarrow \gamma_c + 1$.
2. For each $t \in T$, independently, adjust one dual variable: if $t \in T_c^1$, then set $\beta_{c,t} \leftarrow \beta_{c,t} + 1$ and otherwise set $\alpha_{c,t} \leftarrow \alpha_{c,t} + 1$.
3. If $T_c^1 \subsetneq T$, choose *active distance* $t^* \leftarrow \min(T \setminus T_c^1)$ to be the smallest non-saturated distance, and then set $x_{c,t^*} \leftarrow x_{c,t^*} + 1/t^*$. (Note that $0 \in T_c^1$, and thus $t^* > 0$.)
4. For any facility $f \in \bigcup_{t \in T_c^1} F_{c,t}$, independently, perform *augmentation of y_f* , setting

$$y_f \leftarrow \left(1 + \frac{1}{\text{cost}(f)}\right) \cdot y_f + \frac{1}{|F| \cdot \text{cost}(f)}.$$

5. Update the set of saturated distances, setting $T_c^1 \leftarrow \{t \in T : x_{c,t} \geq 1\}$.

We now argue that if variable y_f is augmented in Step 4, then $\text{cost}(f) > 0$ (i.e., Step 4 is well defined). Let $\tau = \text{cost}(c, f)$. As y_f is augmented, the distance τ is saturated ($x_{c,\tau} = 1$). If $\text{cost}(f) = 0$, then y_f would have been initialized to 1, and then $y(F_{c,\tau}) \geq 1$, in which case the serving constraint for c would be already satisfied.

Sidenote about T. For the sake of coherence and more streamlined analysis, FRAC increases also connection variables $x_{c,t}$ to empty sets $F_{c,t}$, i.e., invests into distances to non-existing facilities. Fixing this overspending would not lead to asymptotic improvement of the performance.

2.3 Structural Properties

We focus on a single client c processed by FRAC. We start with a property of connection variables $x_{c,t}$. The distances from T that are neither saturated nor active are called *inactive*. The following claim follows by an immediate induction on update operations performed by FRAC.

► **Lemma 2.** *At all times when a client c is considered, $x_{c,t} \in [0, 1]$ for any $t \in T$. In particular, $x_{c,t} = 1$ for any saturated distance $t \in T_c^1$. Furthermore,*

1. *either all distances are saturated,*
2. *or there exists an active distance $t^* > 0$, such that (i) all smaller distances are saturated, and (ii) all larger distances are inactive and the corresponding variables $x_{c,t}$ are equal to zero.*

Augmentation is performed on variables y_f corresponding to facilities whose distance from c is saturated.

► **Lemma 3.** *On any input $(G = (F, C, E, \text{cost}), A)$, FRAC returns a feasible solution and runs in time $\text{poly}(|G|, |A|, \max_{e \in E} \text{cost}(e), \max_{f \in F} \text{cost}(f))$.*

Proof. Fix any client $c \in A$. By the definition of FRAC, it takes t update operations to increase value $x_{c,t}$ from 0 to 1. Hence, after $\sum_{t \in T} t < 2 \cdot \max_{e \in E} \text{cost}(e)$ update operations, all connection variables are equal to 1. From that point on, all variables y_f for $f \in \biguplus_{t \in T} F_{c,t}$ are augmented in each update operation. Each variable y_f can be augmented at most $|F| \cdot \text{cost}(f)$ times till it reaches or exceeds 1. That is, after at most $2 \cdot \max_{e \in E} \text{cost}(e) + |F| \cdot \max_{f \in F} \text{cost}(f)$ update operations, the serving constraint is satisfied, i.e., the generated solution is feasible. ◀

The following lemma shows the crucial property of FRAC. Namely for any client c , there exist a “good” distance τ , such that the collection of clusters $F_{c,t}$ at distance $t \leq \tau$ is together fractionally half-open and that FRAC invested $\Omega(\tau)$ into connecting client c . For any client c and distance $t \in T$, we define a set $S_{c,t}$ to be a collection of clusters alluded to in the introduction.

$$S_{c,t} = \biguplus_{t' \in T: t' \leq t} F_{c,t'}$$

► **Lemma 4.** *Once FRAC finishes serving client c , there exists a distance $\tau \in T$, such that $y(S_{c,\tau}) \geq 1/2$ and $\sum_{t \in T} t \cdot x_{c,t} \geq \tau/2$.*

Proof. We consider the state of variables once FRAC finishes serving client c . Let $t^* > 0$ be the largest distance from T for which $x_{c,t^*} > 0$. As the serving constraint for client c is satisfied, we have

$$1 \leq \sum_{t \in T} \min\{x_{c,t}, y(F_{c,t})\} = \min\{x_{c,t^*}, y(F_{c,t^*})\} + \sum_{t \in T: t < t^*} \min\{x_{c,t}, y(F_{c,t})\}. \quad (2)$$

We pick τ depending on the value of the last term of (2).

If $\min\{x_{c,t^*}, y(F_{c,t^*})\} \geq 1/2$, we set $\tau = t^*$. Then, $y(S_{c,\tau}) \geq y(F_{c,\tau}) \geq \min\{x_{c,\tau}, y(F_{c,\tau})\} \geq 1/2$, and the first condition of the lemma follows. Furthermore, $\sum_{t \in T} t \cdot x_{c,t} \geq \tau \cdot x_{c,\tau} \geq \tau/2$.

Otherwise, $\min\{x_{c,t^*}, y(F_{c,t^*})\} < 1/2$, and then, by (2), $\sum_{t \in T: t < t^*} \min\{x_{c,t}, y(F_{c,t})\} \geq 1/2$. In such case, we choose τ as the largest distance from T smaller than t^* . Then

$$y(S_{c,\tau}) = \sum_{t \in T: t \leq \tau} y(F_{c,t}) \geq \sum_{t \in T: t \leq \tau} \min\{x_{c,t}, y(F_{c,t})\} \geq 1/2,$$

i.e., the first condition of the lemma holds. By Lemma 2, either t^* is active at the end of processing c or all distances become saturated and t^* is the largest distance from T . In either case, $x_{c,t} = 1$ for any distance $t < t^*$, and thus in particular $x_{c,\tau} = 1$. Hence, the second part of the lemma holds as $\sum_{t \in T} t \cdot x_{c,t} \geq \tau \cdot x_{c,\tau} = \tau$. \blacktriangleleft

2.4 Dual Solution is Almost Feasible

Using primal-dual analysis, we may show that the generated dual solution violates each constraint at most by a factor of $O(\log |F|)$.

► **Lemma 5.** *For any facility f , FRAC augments y_f at most $O(\log |F|) \cdot \text{cost}(f)$ times.*

Proof. First, we observe that variable y_f can be augmented only if prior to augmentation it is smaller than 1. To show that, observe that the augmentation of y_f occurs only when FRAC processes an active client $c \in \text{set}(f)$. Let $\tau = \text{cost}(f, c)$, i.e., $f \in F_{c,\tau}$. As FRAC augments y_f , the distance τ must be saturated, i.e., $x_{c,\tau} = 1$. On the other hand, the serving constraint (1) is not satisfied when y_f is augmented, and thus $\min\{x_{c,\tau}, y(F_{c,\tau})\} < 1$ which implies that y_f must be strictly smaller than 1.

In particular, if $\text{cost}(f) = 0$, then y_f is set to 1 immediately at the beginning, and hence no augmentation of y_f is ever performed, and the lemma follows trivially. As all non-zero costs are at least 1, below we assume $\text{cost}(f) \geq 1$.

During the first $\text{cost}(f)$ augmentations, the value of y_f increases from 0 to at least $1/|F|$ (due to additive increases). Next, during the subsequent $\lceil \log_{1+1/\text{cost}(f)} |F| \rceil$ augmentations, the value of y_f reaches at least 1 (due to multiplicative increases), and hence it will not be augmented further. In total, the number of augmentations is upper-bounded by $\text{cost}(f) + \lceil \log_{1+1/\text{cost}(f)} |F| \rceil = O(\log |F|) \cdot \text{cost}(f)$. In the last relation, we used $\text{cost}(f) \geq 1$. \blacktriangleleft

► **Lemma 6.** *FRAC violates each dual constraint at most by a factor of $O(\log |F|)$.*

Proof. We show the claim for all types of constraints in the dual program.

1. Each dual constraint $\gamma_c \leq \alpha_{c,t} + \beta_{c,t}$ always holds with equality as together with γ_c , for each $t \in T$, FRAC increments either $\alpha_{c,t}$ or $\beta_{c,t}$.
2. Consider a constraint $\alpha_{c,t} \leq t$. Initially $\alpha_{c,t} = 0$ when client c appears, and it is incremented in an update operation only if distance t is not saturated. Distances are processed from the smallest to the largest, and it takes exactly t' update operations for a distance $t' \in T$ to become saturated. Therefore, $\alpha_{c,t}$ can be incremented at most $\sum_{t' \in T: t' \leq t} t'$ times. If $t = 0$, then $\alpha_{c,t} = 0$ trivially. Otherwise, we use the fact that $T \setminus \{0\}$ contains only powers of 2, and hence $\alpha_{c,t} \leq \sum_{t' \in T: t' \leq t} t' < 2 \cdot t$.
3. Finally, fix any facility $f^* \in F$ and consider the constraint $\sum_{c \in \text{set}(f^*) \cap A_k} \beta_{c, \text{cost}(f^*, c)} \leq \text{cost}(f^*)$. We want to show that this constraint is violated at most by a factor of $O(\log |F|)$, i.e., that

$$\sum_{c \in \text{set}(f^*) \cap A_k} \beta_{c, \text{cost}(f^*, c)} \leq O(\log |F|) \cdot \text{cost}(f^*). \quad (3)$$

The left-hand side of (3) is initially 0 and it is incremented only when FRAC processes some active client $c^* \in \text{set}(f^*)$. In a single update operation, FRAC may increment multiple β variables, but only one of them, namely $\beta_{c^*, \text{cost}(f^*, c^*)}$, contributes to the growth of the left-hand side of (3). If variable $\beta_{c^*, \text{cost}(f^*, c^*)}$ is incremented, it means that the distance $\tau = \text{cost}(f^*, c^*)$ is already saturated, i.e., $\tau \in T_{c^*}^1$. Thus, in the same update operation, FRAC augments all variables y_f for $f \in \bigcup_{t \in T_{c^*}^1} F_{c^*, t}$. This set of facilities includes cluster $F_{c^*, \tau}$ and thus also facility f^* . By Lemma 5, the augmentation of y_{f^*} may happen at most $O(\log |F|) \cdot \text{cost}(f^*)$ times, which implies our claim. \blacktriangleleft

2.5 Competitive Ratio of FRAC

Finally, we show that in each update operation the growth of the primal cost is at most constant times the growth of the dual cost. This will imply the competitive ratio of FRAC.

► **Lemma 7.** *For any step k , the value of the solution to \mathcal{P}_k computed by FRAC is at most 3 times the value of its solution to \mathcal{D}_k .*

Proof. As the values of both solutions are initially zero, it suffices to analyze the growth of the primal and dual objectives for a single update operation. The value of the dual solution grows by 1 as γ_c is incremented only for the requested client c . Thus, it is sufficient to show that the primal solution increases at most by 3.

By y_f , $x_{c,t}$ and T_c^1 , we understand the values of these variables before an update operation. Let $F_1 = \biguplus_{t \in T_c^1} F_{c,t}$. As the serving constraint for client c is not satisfied at that point,

$$1 > \sum_{t \in T} \min \{x_{c,t}, y(F_{c,t})\} \geq \sum_{t \in T_c^1} \min \{x_{c,t}, y(F_{c,t})\} \geq \sum_{t \in T_c^1} y(F_{c,t}) = y(F_1). \quad (4)$$

In the last inequality we used that (by Lemma 2), $T_c^1 = \{t \in T : x_{c,t} = 1\}$. The last equality follows as sets $F_{c,t}$ are disjoint for different t .

Within a single update operation, let $\Delta x_{c,t}$ and Δy_f be the increases of variables $x_{c,t}$ and y_f , respectively. By Lemma 2, FRAC increases one connection variable x_{c,t^*} for an active distance t^* (and no connection variable if there is no active distance) and performs augmentations of y_f for all $f \in F_1$. The increase of the primal value is then

$$\begin{aligned} \Delta \mathcal{P} &= \sum_{t \in T} t \cdot \Delta x_{c,t} + \sum_{f \in F_1} \text{cost}(f) \cdot \Delta y_f \leq 1 + \sum_{f \in F_1} \text{cost}(f) \cdot \left(\frac{y_f}{\text{cost}(f)} + \frac{1}{|F| \cdot \text{cost}(f)} \right) \\ &= 1 + y(F_1) + \frac{|F_1|}{|F|} < 3, \end{aligned}$$

where the last inequality follows by (4). ◀

► **Lemma 8.** *For any input $(G = (F, C, E, \text{cost}), A)$, it holds that $\text{FRAC}(G, A) \leq O(\log |F|) \cdot \text{OPT}(G, A)$.*

Proof. Let k be the total number of active clients in A , and let $\text{val}(\mathcal{P}_k)$ and $\text{val}(\mathcal{D}_k)$ be the values of the final primal and dual solutions generated by FRAC. Then,

$$\begin{aligned} \text{FRAC}(G, A) = \text{val}(\mathcal{P}_k) &\leq 3 \cdot \text{val}(\mathcal{D}_k) && \text{(by Lemma 7)} \\ &\leq O(\log |F|) \cdot \text{OPT}(G, A) && \text{(by Lemma 6 and weak duality).} \quad \blacktriangleleft \end{aligned}$$

3 Deterministic Rounding

Now we define our deterministic algorithm INT, which rounds the fractional solution computed by FRAC. For a client $c \in A$, INT observes the actions of FRAC while processing c and on this basis makes its own decisions. First, INT processes augmentations of variables y_f performed by FRAC, and purchases some facilities. Once FRAC finishes handling client c , INT connects c to the closest open facility. (We show below that such facility exists.)

3.1 Purchasing Facilities: Properties of INTFAC

Purchasing facilities by INT is based solely on graph G and on updates of variables y_f produced by FRAC. In particular, it neglects whether a given client is active or not. We use integral variables $\hat{y}_f \in \{0, 1\}$ to denote whether INT opened facility f . Furthermore, for any set F' we use $\hat{y}(F')$ as a shorthand for $\sum_{f \in F'} \hat{y}_f$.

The following lemma is an adaptation of the deterministic rounding routine for the set cover problem by Alon et al. [3] and its proof is postponed to Subsection 3.3.

- **Lemma 9.** *Fix any input $(G = (F, C, E, \text{cost}), A)$. Initially, $\hat{y}_f = y_f = 0$ for any $f \in F$. There exists a deterministic polynomial-time online algorithm INTFAC that transforms increments of fractional variables y_f to increments of integral variables $\hat{y}_f \in \{0, 1\}$, so that*
- *condition $y(S_{c,t}) \geq 1/2$ implies $\hat{y}(S_{c,t}) \geq 1$ for any client $c \in C$ (active or inactive) and any $t \in T$,*
 - $\sum_{f \in F} \text{cost}(f) \cdot \hat{y}_f \leq O(\log |C \times T|) \cdot \sum_{f \in F} \text{cost}(f) \cdot y_f + 2 \cdot \max_{f \in F} \text{cost}(f)$.

3.2 Connecting Clients

Once INT purchases facilities using deterministic routine INTFAC (cf. Lemma 9), it connects client c to the closest open facility. Now we show that such a facility indeed exists and we bound the competitive ratio of INT.

- **Lemma 10.** *On any input (G, A) , the solution generated by INT is feasible and the total cost of connecting clients by INT is at most $2 \cdot \text{FRAC}(G, A)$.*

Proof. Fix any client $c \in A$. By Lemma 4, there exists a distance $\tau \in T$ such that $y(S_{c,\tau}) \geq 1/2$ and $\sum_{t \in T} t \cdot x_{c,t} \geq \tau/2$. By Lemma 9, once INT purchases facilities, it holds that $\hat{y}(S_{c,\tau}) \geq 1$. It means that at least one facility is opened in set $S_{c,\tau}$, i.e., at distance at most τ from c .

Therefore, INT is feasible and by connecting client c to the closest open facility, it ensures that the connection cost is at most $\tau \leq 2 \cdot \sum_{t \in T} t \cdot x_{c,t}$. The proof is concluded by observing that $\sum_{t \in T} t \cdot x_{c,t}$ is the connection cost of FRAC that can be attributed solely to the connection of client c . ◀

- **Lemma 11.** *For any input $(G = (F, C, E, \text{cost}), A)$, it holds that $\text{INT}(G, A) \leq q \cdot \log |F| \cdot (\log |C| + \log \log \Delta_G) \cdot \text{OPT}(G, A) + 2 \cdot \max_{f \in F} \text{cost}(f)$, where q is a universal constant not depending on G or A . Furthermore, INT runs in time polynomial in $|G|$, $|A|$, $\max_{e \in E} \text{cost}(e)$, and $\max_{f \in F} \text{cost}(f)$.*

Proof. Let $\rho = \max_{f \in F} \text{cost}(f)$. Then,

$$\begin{aligned} \text{INT}(G, A) &\leq \sum_{f \in F} \text{cost}(f) \cdot \hat{y}_f + 2 \cdot \text{FRAC}(G, A) && \text{(by Lemma 10)} \\ &\leq O(\log |C \times T|) \cdot \text{FRAC}(G, A) + 2 \cdot \rho && \text{(by Lemma 9)} \\ &= O((\log |C| + \log |T|) \cdot \log |F|) \cdot \text{OPT}(G, A) + 2 \cdot \rho && \text{(by Lemma 8)}. \end{aligned}$$

The bound on the cost of INT is concluded by using $|T| \leq 2 + \log \Delta_G$.

By Lemma 3, FRAC running time is $\text{poly}(|G|, |A|, \max_{e \in E} \text{cost}(e), \max_{f \in F} \text{cost}(f))$. On top of that, INT adds its own computations (in particular the rounding scheme of INTFAC), whose runtime is polynomial in $|G|$ and $|A|$. This implies the second part of the lemma (the running time of INT). ◀

3.3 Purchasing Facilities: Algorithm INTFAC

We start with a technical claim and later we define our rounding procedure INTFAC.

► **Lemma 12.** *Fix any $q \in [0, 1/2]$ and any $r \geq 0$. Let X be a binary variable being 0 with probability $p > 0$. Then, $\mathbf{E}[\exp(q \cdot X)] \leq \exp(-(3/2) \cdot q \cdot \ln p)$.*

Proof. Using the definition of X , we have

$$\begin{aligned} \mathbf{E}[\exp(q \cdot X)] &= p \cdot e^0 + (1 - p) \cdot e^q = \exp(\ln p) + (1 - \exp(\ln p)) \cdot e^q \\ &\leq 1 + \ln p - e^q \cdot \ln p = 1 - \ln p \cdot (e^q - 1) \\ &\leq 1 - (3/2) \cdot q \cdot \ln p \\ &\leq \exp(-(3/2) \cdot q \cdot \ln p). \end{aligned}$$

In the first inequality, we used that $e^x \cdot 1 + (1 - e^x) \cdot z \leq (1 + x) \cdot 1 + (-x) \cdot z$ for any $x \leq 0$ and $z \geq 1$ and in the second one, we used that $e^x - 1 \leq 3x/2$ for any $x \in [0, 1/2]$. ◀

Algorithm Description. As we mentioned earlier, our routine INTFAC for rounding facilities is an adaptation of the deterministic rounding procedure for the set cover problem by Alon et al. [3]. On the basis of the facility-client graph G , we define the set $C \times T$ of *elements*. Intuitively, our solution FRAC “covers” an element $(c, t) \in C \times T$ by fractionally opening facilities from $S_{c,t}$. The routine INTFAC deterministically rounds these covering choices.

Let $\ell = |C \times T|$, $\rho = \max_{f \in F} \text{cost}(F)$ and $b = 6 \cdot \ln \ell = O(\log |C \times T|)$. We consider the potential function $\Phi = \Phi_1 + \Phi_2$, where

$$\Phi_1 = \sum_{(c,t) : \hat{y}(S_{c,t})=0} \ell^{4 \cdot y(S_{c,t})} \quad \text{and} \quad \Phi_2 = \ell \cdot \exp \left(\sum_{f \in F} \frac{\text{cost}(f)}{2\rho} \cdot (\hat{y}_f - b \cdot y_f) \right).$$

Assume that FRAC augmented variable y_f . Then our algorithm INTFAC chooses whether to set \hat{y}_f to 1 or not (purchase f or not), so that the potential Φ does not increase. (We again emphasize that this choice neglects the current set of active clients.)

Correctness and Performance. In the lemma below, we show that INTFAC is well defined, i.e., it is possible to fix variable \hat{y}_f , so that the potential Φ does not increase. This implies that both Φ_1 and Φ_2 remain upper-bounded, which can be in turn used to show properties of Lemma 9.

► **Lemma 13.** *Assume y_{f^*} is increased by δ . If $\hat{y}_{f^*} = 1$, then Φ does not increase. Otherwise, there is a choice to either set \hat{y}_{f^*} to 1 or not, so that Φ does not increase.*

Proof. By y_f and \hat{y}_f , we mean the values of these variables before an update operation of FRAC.

First, we assume $\hat{y}_{f^*} = 1$. Increasing variable y_{f^*} affects values of $y(S_{c,t})$ for $f^* \in S_{c,t}$: all such $y(S_{c,t})$ increase by δ . However, for any element (c, t) , such that $f^* \in S_{c,t}$, it holds that $\hat{y}(S_{c,t}) \geq \hat{y}_{f^*} = 1$, i.e., element (c, t) is not counted in the sum occurring in Φ_1 . Thus, increasing variable y_{f^*} does not affect Φ_1 . Furthermore, increasing y_{f^*} and keeping \hat{y}_{f^*} unchanged can only decrease Φ_2 . Thus, $\Phi = \Phi_1 + \Phi_2$ does not increase when $\hat{y}_{f^*} = 1$.

Second, we consider the case $\hat{y}_{f^*} = 0$. To show that either setting \hat{y}_{f^*} to 1 or leaving it at 0 does not increase the potential, we use the probabilistic method and show that if we pick such action randomly (setting $\hat{y}_{f^*} = 1$ with probability $1 - \ell^{-4\delta}$), then, in expectation, neither Φ_1 nor Φ_2 increases.

- As observed above, only elements (c, t) for which $S_{c,t}$ contain f^* are affected by the increase of y_{f^*} and possible change of \hat{y}_{f^*} . Let $Q = \{(c, t) : f^* \in S_{c,t} \text{ and } \hat{y}(S_{c,t}) = 0\}$ be the set of such elements contributing to Φ_1 .

Fix any element $(c, t) \in Q$. Its initial contribution towards Φ_1 is $\ell^{4 \cdot y(S_{c,t})}$ and when y_{f^*} increases, the contribution grows to $\ell^{4 \cdot (y(S_{c,t}) + \delta)}$. However, with probability $1 - \ell^{-4 \cdot \delta}$, variable \hat{y}_{f^*} is set to 1, thus $\hat{y}(S_{c,t})$ grows from 0 to 1, and in effect element (c, t) stops contributing to Φ_1 . Hence, the expected final contribution of element (c, t) towards Φ_1 is $\ell^{4 \cdot (y(S_{c,t}) + \delta)} \cdot \ell^{-4 \cdot \delta} + 0 \cdot (1 - \ell^{-4 \cdot \delta}) = \ell^{4 \cdot y(S_{c,t})}$, i.e., is equal to its initial contribution. Therefore, in expectation, the value of Φ_1 is unchanged.

- It remains to bound the expected value of Φ_2 . Let \hat{Y} be the random variable equal to the value of \hat{y}_{f^*} after the random choice (i.e., $\hat{Y} = 1$ with probability $1 - \ell^{-4 \cdot \delta}$) and Φ'_2 denote the value of Φ_2 after increasing y_{f^*} and after the random choice. Using $y_{f^*} = 0$, we obtain

$$\begin{aligned} \Phi'_2 &= \ell \cdot \exp \left(\sum_{f \in F} \frac{\text{cost}(f)}{2\rho} \cdot (\hat{y}_f - b \cdot y_f) + \frac{\text{cost}(f^*)}{2\rho} \cdot \hat{Y} - \frac{b \cdot \text{cost}(f^*)}{2\rho} \cdot \delta \right) \\ &= \Phi_2 \cdot \exp \left(\frac{\text{cost}(f^*)}{2\rho} \cdot \hat{Y} \right) \cdot \exp \left(-\frac{b \cdot \text{cost}(f^*)}{2\rho} \cdot \delta \right). \end{aligned}$$

To estimate $\mathbf{E}[\Phi'_2]$, we upper-bound the expected value of expression $\exp(\hat{Y} \cdot \text{cost}(f^*) / (2\rho))$, using Lemma 12 with $q = \text{cost}(f^*) / (2\rho) \leq 1/2$ and $p = \ell^{-4 \cdot \delta}$, obtaining that

$$\mathbf{E} \left[\exp \left(\frac{\text{cost}(f^*)}{2\rho} \cdot \hat{Y} \right) \right] \leq \exp \left(-\frac{(3/2) \cdot \text{cost}(f^*)}{2\rho} \cdot \ln p \right) = \exp \left(\frac{6 \cdot \ln \ell \cdot \text{cost}(f^*)}{2\rho} \cdot \delta \right).$$

Therefore, $\mathbf{E}[\Phi'_2] \leq \Phi_2$ and the lemma follows. \blacktriangleleft

Proof of Lemma 9. Initially, all variables y_f and \hat{y}_f are zero, and thus $\Phi = \sum_{(c,t) \in C \times T} \ell^0 + \ell \cdot \exp(0) = 2 \cdot \ell$. By Lemma 13, the potential never increases. Since Φ_2 is non-negative, any summand of Φ_1 is always at most $2 \cdot \ell \leq \ell^2$. Therefore, $4 \cdot y(S_{c,t}) \geq 2$ always implies $\hat{y}(S_{c,t}) > 0$, i.e., the first part of the lemma follows.

To show the second part, we again use that $\Phi = \Phi_1 + \Phi_2 \leq 2 \cdot \ell$ at any time. As Φ_1 is non-negative, $\Phi_2 \leq 2 \cdot \ell$. Substituting the definition of Φ_2 , dividing by ℓ , and taking natural logarithm of both sides yields

$$\frac{1}{2\rho} \cdot \sum_{f \in F} (\hat{y}_f \cdot \text{cost}(f) - b \cdot y_f \cdot \text{cost}(f)) \leq \ln(2) < 1.$$

Therefore, $\sum_{f \in F} \hat{y}_f \cdot \text{cost}(f) \leq 2\rho + b \cdot \sum_{f \in F} y_f \cdot \text{cost}(f)$. \blacktriangleleft

4 Handling Large Aspect Ratios

The guarantee of Lemma 11 has two deficiencies: (i) the bound on the competitive ratio of INT depends on the aspect ratio of G and on the cost of the most expensive facility, (ii) the running time of INT depends on the maximal cost in graph G (which can be exponentially large in the input description). We show how to use cost doubling and edge pruning to handle these issues, creating our final deterministic solution DET and proving the main theorem (restated below).

► Theorem 1. *There exists a deterministic polynomial-time $O(\log |F| \cdot (\log |C| + \log \log |F|))$ -competitive algorithm for the online non-metric facility location problem on set F of facilities and set C of clients.*

14:14 A Nearly Optimal Deterministic Online Algorithm for Non-Metric Facility Location

Proof. Fix facility-client graph $G = (F, C, E, \text{cost})$ for the non-metric facility location problem. Recall that we assumed that all non-zero costs and distances in G are powers of 2 and are at least 1. Let $R = \log |F| \cdot (\log |C| + \log \log(|F| \cdot |C|))$.

We now construct a deterministic algorithm DET which is $O(R)$ -competitive on an input (G, A) . Let q be the constant from Lemma 11. DET operates in phases, numbered from 0. In phase j , it executes the following operations.

1. DET *pre-purchases* all facilities and edges of G whose cost is smaller than $2^j/(|F| \cdot |C|)$.
2. DET creates an auxiliary facility-client graph \tilde{G}_j applying the following modifications to G .
 - First, DET creates graph G_j containing only edges and facilities from G whose individual cost is at most 2^j . It also removes connections to facilities that have been removed in this process.
 - Second, the costs of all facilities and edges that have been pre-purchased by DET are set to zero in G_j . In a result, G_j is a sub-graph of G with adjusted distances and costs of facilities, has the same set of clients, its set of facilities is a subset of F , and $\Delta_{G_j} \leq |F| \cdot |C|$.
 - Third, \tilde{G}_j is the modified version of G_j , where all costs have been scaled down, so that the smallest positive cost is equal to 1. We denote the scaling factor by $h_j \leq 1$.
3. DET simulates algorithm INT on input (\tilde{G}_j, A) . That is, for a client $c \in A$, DET verifies whether the overall cost of INT (including serving c) remains at most $h_j \cdot (q \cdot R + 2) \cdot 2^j$. In such case, DET outputs the choices of INT for client c as its own. We emphasize that INT is run also on clients that have been already served in the previous phases; in effect, DET may purchase the same facilities or connections multiple times.
4. Eventually, either the sequence A of active clients ends and the total cost of INT on (\tilde{G}_j, A) is at most $h_j \cdot (q \cdot R + 2) \cdot 2^j$ (in which case DET terminates as well) or the purchases made by INT, while handling a client $c \in A$, caused its cost to exceed $h_j \cdot (q \cdot R + 2) \cdot 2^j$. (This includes the special case where c is disconnected from all facilities in \tilde{G}_j , because all edges incident to c in G were either more expensive than 2^j or were leading to facilities more expensive than 2^j .) In the case of exceeded cost, DET disregards the decisions of INT for client c , terminates INT, and starts phase $j + 1$, processing also all clients that were already served in phase j .

We now analyze the performance of DET. Let $k = \lceil \log(\text{OPT}(G, A)) \rceil \geq 0$. We show that DET terminates latest in phase k . Assume that DET has not finished within phases $0, 1, \dots, k-1$. In phase k , DET creates auxiliary graphs G_k and \tilde{G}_k , and runs INT on graph \tilde{G}_k . Graph G_k contains all edges of G of cost at most 2^k ; their cost in G_k is the same or reset to zero. As $\text{OPT}(G, A) \leq 2^k$, $\text{OPT}(G, A)$ purchases only edges that are in G_k , and thus $\text{OPT}(G, A)$ is also a feasible solution to instance (G_k, A) . Thus, $\text{OPT}(G_k, A) \leq \text{OPT}(G, A) \leq 2^k$. As \tilde{G}_k is the scaled-down copy of G_k , $\text{OPT}(\tilde{G}_k, A) = h_k \cdot \text{OPT}(G_k, A) \leq h_k \cdot 2^k$.

Let \tilde{F}_k be the set of facilities of graph \tilde{G}_k and $\tilde{\text{cost}}_k(f)$ is the cost of opening facility f in graph \tilde{G}_k . Clearly, $|\tilde{F}_k| \leq |F|$ and $\tilde{\text{cost}}_k(f) \leq h_k \cdot \text{cost}(f)$ for any $f \in F$. By our construction, $\Delta_{\tilde{G}_k} = \Delta_{G_k} \leq |F| \cdot |C|$. Hence, Lemma 11 implies that

$$\begin{aligned} \text{INT}(\tilde{G}_k, A) &\leq q \cdot \log |F_k| \cdot (\log |C| + \log \log \Delta_{\tilde{G}_k}) \cdot \text{OPT}(\tilde{G}_k, A) + 2 \cdot \max_{f \in \tilde{F}_k} \tilde{\text{cost}}_k(f) \\ &\leq h_k \cdot q \cdot \log |F| \cdot (\log |C| + \log \log(|F| \cdot |C|)) \cdot 2^k + 2 \cdot h_k \cdot 2^k \\ &= h_k \cdot (q \cdot R + 2) \cdot 2^k. \end{aligned}$$

Therefore, INT is not terminated prematurely within phase k because of high cost and it finishes the entire sequence A . This implies the feasibility of INT: it serves all clients latest in phase k .

To bound the total cost of DET, recall that at the beginning of phase j , DET purchases at most $|F| \cdot |C|$ edges and at most $|F|$ facilities, each of cost at most $2^j / (|F| \cdot |C|)$. The associated overall cost is at most $2 \cdot 2^j$. The cost of the subsequent execution of algorithm INT on \tilde{G}_j is, by our termination rule, at most $h_j \cdot (q \cdot R + 2) \cdot 2^j$, and thus the cost incurred by repeating INT's actions on G is at most $(q \cdot R + 2) \cdot 2^j$. The overall cost is then $\text{DET}(G, A) \leq \sum_{j=0}^k (q \cdot R + 4) \cdot 2^j = O(R) \cdot 2^k = O(R) \cdot \text{OPT}(G, A) = O(\log |F| \cdot (\log |C| + \log \log |F|)) \cdot \text{OPT}(G, A)$.

For the running time of DET, we note that in phase j , INT is run on a graph \tilde{G}_j whose smallest cost is 1, and hence the largest cost is at most $\Delta_{\tilde{G}_j} = \Delta_{G_j} \leq |F| \cdot |C|$. Thus, by Lemma 11, the running time of INT in a single phase is polynomial in $|G|$ and $|A|$, and the number of phases is logarithmic in the maximum cost occurring in G , and thus also polynomial in $|G|$. ◀

5 Application to Online Node-Weighted Steiner Tree

Our result for the non-metric FL problem has an immediate application for the online node-weighted Steiner tree (NWST) problem, where the graph consists of ℓ nodes and an online algorithm is given k terminals to be connected. Namely, the randomized solution for the online NWST problem by Naor et al. [26] is in fact a deterministic polynomial-time “wrapper” around randomized routine solving the non-metric FL problem. To solve an instance of the NWST problem, their algorithm constructs a sub-instance of non-metric FL with $O(\ell)$ facilities, $O(\ell)$ potential clients, and $O(k)$ active clients. Such instance can be solved by the randomized algorithm of Alon et al. [2] with the competitive ratio of $O(\log k \cdot \log \ell)$. The wrapper adds another $O(\log k)$ factor in the ratio, resulting in an $O(\log^2 k \cdot \log \ell)$ -competitive algorithm.

Our deterministic algorithm, when applied to this setting would be $O(\log^2 \ell)$ -competitive on the constructed non-metric FL sub-instance. Therefore, by replacing the randomized algorithm by Alon et al. [2] with our deterministic one, we immediately obtain the first online deterministic solution for online NWST.

► **Corollary 14.** *There exists a polynomial-time deterministic online algorithm for the node-weighted Steiner tree problem, which is $O(\log k \cdot \log^2 \ell)$ -competitive on graphs with ℓ nodes and k terminals.*

We note that the currently best solution for the node-weighted Steiner tree is randomized and achieves the ratio of $O(\log^2 \ell)$ [16, 15] and the best known lower bound for deterministic algorithms is $\Omega(\log \ell \cdot \log k / (\log \log \ell + \log \log k))$ [26, 3].

6 Final Remarks

We presented a deterministic solution to the non-metric facility location problem, whose performance nearly matches that of the best randomized one. By clustering facilities, we encoded dependencies between facilities and clients, which allowed us later to apply the rounding scheme to facilities only, neglecting the actual active clients. It would be however interesting and useful to have an online deterministic rounding routine able to handle such dependencies internally (e.g., by creating a pessimistic estimator that can be computed and handled in an online manner), as it is the case for the set cover problem or throughput-competitive virtual circuit routing [8].

That said, we believe that our distance clustering techniques can be extended to other network design problems for which only randomized algorithms existed so far, e.g., online multicast problems on trees [2], online group Steiner problem on trees [2], or variants of the facility location problem that are used as building blocks for solutions to other node-weighted Steiner problems [15, 16]. (For these problems there are no known direct reductions to the set cover problem). Finally, another open problem is whether these techniques could be also applied more directly for the node-weighted Steiner tree, resulting in a better deterministic competitive ratio.

References

- 1 Karen Aardal, Jaroslaw Byrka, and Mohammad Mahdian. Facility location. In *Encyclopedia of Algorithms*, pages 717–724. Springer, 2016. doi:10.1007/978-1-4939-2864-4_139.
- 2 Noga Alon, Baruch Awerbuch, Yossi Azar, Niv Buchbinder, and Joseph Naor. A general approach to online network optimization problems. *ACM Transactions on Algorithms*, 2(4):640–660, 2006. doi:10.1145/1198513.1198522.
- 3 Noga Alon, Baruch Awerbuch, Yossi Azar, Niv Buchbinder, and Joseph Naor. The online set cover problem. *SIAM Journal on Computing*, 39(2):361–370, 2009. doi:10.1137/060661946.
- 4 Aris Anagnostopoulos, Russell Bent, Eli Upfal, and Pascal Van Hentenryck. A simple and deterministic competitive algorithm for online facility location. *Information and Computation*, 194(2):175–202, 2004. doi:10.1016/j.ic.2004.06.002.
- 5 Sanjeev Arora, Elad Hazan, and Satyen Kale. The multiplicative weights update method: a meta-algorithm and applications. *Theory of Computing Systems*, 8(1):121–164, 2012. doi:10.4086/toc.2012.v008a006.
- 6 Baruch Awerbuch, Yossi Azar, and Serge A. Plotkin. Throughput-competitive on-line routing. In *Proc. 34th IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 32–40, 1993. doi:10.1109/SFCS.1993.366884.
- 7 Niv Buchbinder and Joseph 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. doi:10.1561/04000000024.
- 8 Niv Buchbinder and Joseph Naor. Online primal-dual algorithms for covering and packing. *Mathematics of Operations Research*, 34(2):270–286, 2009. doi:10.1287/moor.1080.0363.
- 9 Jaroslaw Byrka and Karen Aardal. An optimal bifactor approximation algorithm for the metric uncapacitated facility location problem. *SIAM Journal on Computing*, 39(6):2212–2231, 2010. doi:10.1137/070708901.
- 10 Moses Charikar and Sudipto Guha. Improved combinatorial algorithms for facility location problems. *SIAM Journal on Computing*, 34(4):803–824, 2005. doi:10.1137/S0097539701398594.
- 11 Fabián A. Chudak and David B. Shmoys. Improved approximation algorithms for the uncapacitated facility location problem. *SIAM Journal on Computing*, 33(1):1–25, 2003. doi:10.1137/S0097539703405754.
- 12 Uriel Feige. A threshold of $\ln n$ for approximating set cover. *Journal of the ACM*, 45(4):634–652, 1998. doi:10.1145/285055.285059.
- 13 Dimitris Fotakis. A primal-dual algorithm for online non-uniform facility location. *Journal of Discrete Algorithms*, 5(1):141–148, 2007. doi:10.1016/j.jda.2006.03.001.
- 14 Dimitris Fotakis. On the competitive ratio for online facility location. *Algorithmica*, 50(1):1–57, 2008. doi:10.1007/s00453-007-9049-y.
- 15 MohammadTaghi Hajiaghayi, Vahid Liaghat, and Debmalya Panigrahi. Near-optimal online algorithms for prize-collecting steiner problems. In *Proc. 41st Int. Colloq. on Automata, Languages and Programming (ICALP)*, pages 576–587, 2014. doi:10.1007/978-3-662-43948-7_48.

- 16 MohammadTaghi Hajiaghayi, Vahid Liaghat, and Debmalya Panigrahi. Online node-weighted steiner forest and extensions via disk paintings. *SIAM Journal on Computing*, 46(3):911–935, 2017. doi:10.1137/14098692X.
- 17 Dorit S. Hochbaum. Heuristics for the fixed cost median problem. *Mathematical Programming*, 22(1):148–162, 1982. doi:10.1007/BF01581035.
- 18 Kamal Jain, Mohammad Mahdian, Evangelos Markakis, Amin Saberi, and Vijay V. Vazirani. Greedy facility location algorithms analyzed using dual fitting with factor-revealing LP. *Journal of the ACM*, 50(6):795–824, 2003. doi:10.1145/950620.950621.
- 19 Kamal Jain, Mohammad Mahdian, and Amin Saberi. A new greedy approach for facility location problems. In *Proc. 34th ACM Symp. on Theory of Computing (STOC)*, pages 731–740, 2002. doi:10.1145/509907.510012.
- 20 Antoon Kolen and Arie Tamir. Covering problems. In P.B. Mirchandani and R.L. Francis, editors, *Discrete Location Theory*, Wiley Series in Discrete Mathematics and Optimization. Wiley, 1990.
- 21 Simon Korman. On the use of randomization in the online set cover problem. Master’s thesis, The Weizmann Institute of Science, 2004.
- 22 Madhukar R. Korupolu, C. Greg Plaxton, and Rajmohan Rajaraman. Analysis of a local search heuristic for facility location problems. *Journal of Algorithms*, 37(1):146–188, 2000. doi:10.1006/jagm.2000.1100.
- 23 Shi Li. A 1.488 approximation algorithm for the uncapacitated facility location problem. *Information and Computation*, 222:45–58, 2013. doi:10.1016/j.ic.2012.01.007.
- 24 Mohammad Mahdian, Yinyu Ye, and Jiawei Zhang. Approximation algorithms for metric facility location problems. *SIAM Journal on Computing*, 36(2):411–432, 2006. doi:10.1137/S0097539703435716.
- 25 Adam Meyerson. Online facility location. In *Proc. 42nd IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 426–431, 2001. doi:10.1109/SFCS.2001.959917.
- 26 Joseph Naor, Debmalya Panigrahi, and Mohit Singh. Online node-weighted steiner tree and related problems. In *Proc. 52nd IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 210–219, 2011. doi:10.1109/FOCS.2011.65.
- 27 Prabhakar Raghavan. Probabilistic construction of deterministic algorithms: Approximating packing integer programs. *Journal of Computer and System Sciences*, 37(2):130–143, 1988. doi:10.1016/0022-0000(88)90003-7.
- 28 David B. Shmoys. Approximation algorithms for facility location problems. In *Proc. 3rd Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM)*, pages 27–33, 2000. doi:10.1007/3-540-44436-X_4.
- 29 David B. Shmoys, Éva Tardos, and Karen Aardal. Approximation algorithms for facility location problems (extended abstract). In *Proc. 29th ACM Symp. on Theory of Computing (STOC)*, pages 265–274, 1997. doi:10.1145/258533.258600.
- 30 Neal E. Young. Randomized rounding without solving the linear program. In *Proc. 6th ACM-SIAM Symp. on Discrete Algorithms (SODA)*, pages 170–178, 1995.