Approximating Lower- and Upper-Bounded (Connected) Facility Location via LP Rounding

Zachary Friggstad*, **Mohsen Rezapour, and Mohammad R. Salavatipour**†

Department of Computing Science, University of Alberta,
Edmonton, Alberta, Canada, T6G 2E8
{zacharyf,rezapour,mrs}@ualberta.ca

Abstract

We consider a lower- and upper-bounded generalization of the classical facility location problem, where each facility has a capacity (upper bound) that limits the number of clients it can serve and a lower bound on the number of clients it must serve if it is opened. We develop an LP rounding framework that exploits a Voronoi diagram-based clustering approach to derive the first bicriteria constant approximation algorithm for this problem with non-uniform lower bounds and uniform upper bounds. This naturally leads to the the first LP-based approximation algorithm for the lower bounded facility location problem (with non-uniform lower bounds).

We also demonstrate the versatility of our framework by extending this and presenting the first constant approximation algorithm for some connected variant of the problems in which the facilities are required to be connected as well.

1998 ACM Subject Classification F.2.2 Nonnumerical Algorithms and Problems

Keywords and phrases Facility Location; Approximation Algorithm; LP Rounding

Digital Object Identifier 10.4230/LIPIcs.xxx.yyy.p

1 Introduction

We study the *lower- and upper-bounded facility location* (**LUFL**) problem, a natural generalization of the well-known *capacitated facility location* (**CFL**) and *lower bounded facility location* (**LBFL**) problems. We are given a complete graph $G = (V, E)$, with metric edge lengths $c_e \in \mathbb{Z}_{\geq 0}, e \in E$ containing a set of potential facilities $F \subseteq V$ and a set of demand points (clients) $D \subseteq V$. Each facility $i \in F$ has an opening cost $\mu_i \in \mathbb{Z}_{\geq 0}$ and a capacity (upper bound) $U_i \in \mathbb{Z}_{>0}$, which limits the amount of demand it can serve. Moreover, each facility i has a lower bound $L_i \in \mathbb{Z}_{\geq 0}$ on the amount of demand it must serve if it is opened.

A feasible solution to **LUFL** consists of a set of facilities $I \subseteq F$ to open, and a *valid* assignment $\sigma : D \rightarrow I$ of clients to the open facilities: an assignment is valid if it satisfies the lower and upper bounds

$$L_i \leq |\sigma^{-1}(i)| \leq U_i \qquad \forall i \in I. \tag{1}$$

The goal is to minimize the total cost, i.e., $\sum_{i \in I} \mu_i + \sum_{j \in D} c_{\sigma(j)j}$.

In many real-world applications, particularly in telecommunications, there is an additional requirement to connect the open facilities via high bandwidth *core cables*. This leads to a variant of **LUFL** in which open facilities are connected via a tree-like *core network* that

* Supported by NSERC and funding from the Canada Research Chairs program.
† Supported by NSERC.

consists of infinite capacity cables. We model this variant as a *connected lower- and upper-bounded facility location* (**C-LUFL**) problem. Let us introduce a parameter $M \geq 1$ which reflects the cost per unit length of core cables. A feasible solution to **C-LUFL** is given by a set of facilities $I \subseteq F$, an assignment $\sigma : D \to I$ of clients to the open facilities that is valid, and a Steiner tree of $T \subseteq E$ connecting all facilities I via core cables. The objective of **C-LUFL** is to minimize the total cost, i.e., $\sum_{i \in I} \mu_i + \sum_{j \in D} c_{\sigma(j)j} + M \sum_{e \in T} c_e$.

Both the **CFL** and **LBFL** problems have been well-studied in the literature. However, there is not much work in studying these problems in a complementary way.[1] To address this gap of knowledge, in this paper, we develop a framework that combines LP rounding techniques for facility location problems with a Voronoi diagram-based clustering approach in order to obtain the first (biceriteria) approximation algorithms for several variants of the problems.

▶ **Definition 1.** An (ρ, α, β)-approximation algorithm for LUFL (C-LUFL, resp.) is an algorithm that computes in polynomial time a solution (I, σ) satisfying $\lfloor L_i/\alpha \rfloor \leq |\sigma^{-1}(i)| \leq \lceil \beta U_i \rceil$, $\forall i \in I$, with cost at most $\rho \cdot OPT$, where OPT denotes the minimum cost of a solution to LUFL (C-LUFL, resp.) satisfying (1).

We often loosely refer to a (ρ, α, β)-approximation for **LUFL** or **C-LUFL** when ρ, α, β are constants as a *relaxed constant-factor approximation.*

Related Work. The **CFL** problem is the special case of **LUFL** when $L_i = 0$ for all $i \in F$. There are several approximation algorithms for CFL based on local search techniques. For the case of uniform capacities, Korupolu et al. [13] gave the first constant factor approximation algorithm, with ratio 8. This was later improved to 5.83 [6] and 3 [2]. The first constant factor approximation for the case of non-uniform capacities was proposed by [17] who gave an 9-approximation, which was eventually improve do 5 [5]. An LP-based approach to CFL was employed by Shmoys et al. [18] who gave the first bicriteria approximation for uniform capacities; this was extended to non-uniform capacities [1]. Levi et al. [14] obtained an LP-based 5-approximation algorithm when facilities opening costs are uniform. For a long time it was an open question to prove a constant factor approximation for CFL based on LP-rounding. This was recently answered by An et al. [4] who gave an LP-based 288-approximation algorithm for CFL which works for the general case.

The **LBFL** problem is another special cases of **LUFL** when $U_i = \infty$ for all $i \in F$. This problem was introduced independently by Guha et al. [9] and Karger et al. [12] who gave a bicriteria approximation. The first true approximation algorithm for **LBFL** was given by Svitkina [19] with ratio 448. The factor was then improved to 82.6 by [3] by applying a modified variant of the algorithm of [19], combined with a more careful analysis. We note that the approaches of both papers work only if all lower bounds are uniform. Finding a true approximation for **LBFL** when the lower bounds are non-uniform remains an open problem. To the best of our knowledge, there have been no LP-based approximation (even bicriteria) algorithms for **LBFL** in the literature.

The *Connected Facility Location* (**ConFL**) problem is an obvious special case of **C-LUFL** (when $U_i = \infty$ & $L_i = 0$ for all $i \in F$.) The **ConFL** problem was first introduced by Gupta et al. [10], in the context of reserving bandwidth for *virtual private networks*, where they gave the first constant-factor approximation algorithm for **ConFL**. Using the

---

[1] In an earlier version of [1] there was an attempt to study **LUFL** but there seemed to be an error in the proof. After checking with the authors the claim about **LUFL** is retracted in the current version of [1].

primal-dual technique, the factor was then improved to 8.55 by [20], and to 6.55 by [11]. Applying sampling techniques, the guarantee was later reduced to 4 by [7], and to 3.19 by [8].

Our Results and Techniques. We explore LP-based approaches to obtain bicriteria approximations for many combinations of lower/upper/connected facility location. Our first main result is the first constant-factor (bicriteria) approximation algorithm for **LUFL**.

▶ **Theorem 2.** *There is a relaxed constant-factor approximation for instances of* ***LUFL*** *with uniform upper bounds (and non-uniform lower bounds).*

To prove this theorem we start by presenting an LP-based bicriteria approximation for **LBFL** with non-uniform lower bounds. Such approximations were known before, but ours is the first one whose cost can be compared to an LP relaxation. We emphasize that such bounds may be useful to obtain stronger results. For example, the LP-based **CFL** bicritera approximation by [1] was a key component in devising the true LP-based approximation in [4]. Perhaps our result could be used in an analogous result for **LBFL**.

Next, we incorporate connectivity requirement. We obtain the first constant-factor bicriteria approximation for the *connected lower-bounded facility location* problem with non-uniform lower bounds. We then extend this to a relaxed constant-factor approximation for **C-LUFL** when the upper bounds U are uniform and the core cable multiplier M is $O(U)$. Some remarks on the difficulty of extending our approach to the case $M = \omega(U)$ are presented in the conclusion. Our second main result is the following.

▶ **Theorem 3.** *There is a relaxed constant-factor approximation for instances of* ***C-LUFL*** *with uniform upper bounds where $M = O(U)$.*

A key ingredient in our approach is a clustering step to avoid the standard "*filtering*" steps. That is, in classic facility location and **CFL** rounding algorithms a popular approach is to consider a ball around each client j whose radius is roughly the fractional cost of serving j. Values x_{ij} where i lies far outside this ball are set to 0 and the remaining $x_{i'j}$ values are rounded up by a small constant factor in order to get a solution that is "concentrated" around each client. This approach fails when lower bounds are present. We develop a clustering procedure to find a set of cluster *centers* $\mathcal{C}$ using a *Voronoi diagram* which is inspired by approches to *capacitated k-median* problem that was considered in [15, 16].

2 LP Relaxations and Starting steps

We present LP relaxations for **LUFL** as well as **C-LUFL**. For each $i \in F$, y_i indicates if facility i is opened. For each $i \in F$ and $j \in D$, x_{ij} indicates if client j is assigned to facility i.

$$\min \quad \sum_{i \in F} \mu_i y_i + \sum_{j \in D} \sum_{i \in F} c_{ij} x_{ij} \tag{LP-LUFL}$$

$$\sum_{i \in F} x_{ij} = 1 \qquad \forall j \in D \tag{2}$$

$$x_{ij} \leq y_i \qquad \forall i \in F, j \in D \tag{3}$$

$$\sum_{j \in D} x_{ij} \leq U_i y_i \qquad \forall i \in F \tag{4}$$

$$L_i y_i \leq \sum_{j \in D} x_{ij} \qquad \forall i \in F \tag{5}$$

$$x_{ij}, y_i \in [0,1] \qquad \forall i \in F, j \in D$$

Constraints (2) and (3) are standard *facility location constraints* saying that any client has to be assigned to an open facility in an integer solution. Constraints (4) and (5) ensure the lower and upper bounds are satisfied at the open facilities.

Extending LP-LUFL to model a relaxation for **C-LUFL**, we let z_e indicate if edge $e \in E$ is used by the core Steiner tree. We first guess one particular facility r that is open in the optimum solution and we we called r the *root*. LP-C-LUFL is an integer program formulation of **C-LUFL**.

$$\min \quad \sum_{i\in F}\mu_i y_i + \sum_{j\in D}\sum_{i\in F} c_{ij}x_{ij} + M\sum_{e\in E} c_e z_e \qquad \text{(LP-C-LUFL)}$$

$$(2)-(5)$$

$$\sum_{e\in\delta(S)} z_e \geq \sum_{i\in S} x_{ij} \qquad \forall S\subseteq V\setminus\{r\}, j\in D \qquad (6)$$

$$y_r = 1 \qquad (7)$$

$$x_{ij}, y_i, z_e \in [0,1] \qquad \forall i\in F, j\in D, e\in E$$

Constraints (6) guarantee that all open facilities are connected to facility r via core links, where Constraint (7) forces facility r to be opened. Note that while (6) introduces exponentially many constraints, they can easily be separated by an efficient minimum-cut algorithm. Thus we can solve both (LP-LUFL) and (LP-C-LUFL) in polynomial time using the ellipsoid method.

2.1 Reduction Lemmas

In this section we present two lemmas that are used in the algorithms we present. The first lemma is a general clustering step that is applied as a first step of our LP rounding and reduces the facility location problem on hand to solving the problem on a specific cluster of clients facilities. This clustering step has similarities to a Voronoi diagram and for that reason we call it Voronoi clustering (inspired by [15, 16]). The second lemma shows how one can then extend the results obtained via this reduction step to the case where connectivity (with core cables) is required between open facilities.

Let (x, y, z) be a feasible solution to the LP relaxation of (LP-C-LUFL). Let L^j be the connection cost of client j in the LP, i.e. $L^j = \sum_{i\in F} c_{ij}x_{ij}$. The general idea is to select clients in increasing order of their L^j values and selecting them as centers if they are far from all centers so far. We then define a Voronoi cell with center j to be the set of all facilities for which j is the closest center. This Voronoi clustering will be an important tool in decomposition of an LP solution in our rounding algorithms.

The following algorithm finds a set of clients $\mathcal{C}$ that will act as the centers in the Voronoi diagram and a partition $\{\mathcal{P}_j\}_{j\in\mathcal{C}}$ of F where $i \in P_j$ means j is a closest center to i. Here, λ is some parameter that we can specify. Larger values mean the centers are further apart. It also records a cluster center $\delta(j)\in\mathcal{C}$ for each $j\in D$: if $j\in\mathcal{C}$ then $\delta(j)=j$ and if $j\notin\mathcal{C}$ then $\delta(j)$ is the center that caused j to not be included in $\mathcal{C}$ (it may not be the closest center to j).

Algorithm 1. **Voronoi Clustering algorithm** (λ)

$\mathcal{C} \leftarrow \{j^*\}$ where $j^* = \arg\min_j L^j$;
for *each* $j' \in D - \{j^*\}$ *in increasing order of* $L^{j'}$ **do**
  if $c_{jj'} > 2\lambda \cdot L^{j'}$ *for all* $j \in \mathcal{C}$ **then**
    $\mathcal{C} \leftarrow \mathcal{C} \cup \{j'\}$;
    $\delta(j') \leftarrow j'$;
  else
    let $j \in \mathcal{C}$ be some center with $c_{jj'} \leq 2\lambda \cdot L^{j'}$;
    $\delta(j') \leftarrow j$;
  end
end
for *each* $j \in \mathcal{C}$ **do**
  $P_j \leftarrow \{i \in F : c_{ij} \leq c_{ik}$ *for all* $k \in \mathcal{C}, k \neq j\}$;
  Comment: break ties arbitrarily so each $i \in F$ lies in exactly one P_j.
end
return $(\mathcal{C}, \mathcal{P}, \delta)$

Note that by construction of δ we have that $c_{\delta(j)j} \leq 2\lambda L^j$ for each $j \in D$.

In Lemma 4 we show that for each center $j \in \mathcal{C}$, there is a facility i that is close to j whose opening cost can be paid for by the fractional opening cost paid by the LP for facilities near i. Furthermore, this facility i has a small enough lower bound that we can approximately satisfy by assigning to it all fractional client demand that was sent to some facility in $\mathcal{P}_j$. For each client j and positive radius R, we let $B(j,R) = \{v \in V : c_{jv} \leq R\}$ be a ball centered at j.

▶ **Lemma 4.** *Let (x,y) be values satisfying constraints (2), (3), and (5). Suppose $(\mathcal{C}, \mathcal{P}, \delta)$ is returned by calling Algorithm 1 with some given λ. Let $\widehat{X}_j = \sum_{i \in P_j} \sum_{j' \in D} x_{ij'}$ and let $\eta \in (1, \lambda]$. For each $j \in \mathcal{C}$, there exists some $i \in B^j := B(j, \eta L^j)$ fulfilling: (i) $\mu_i \leq \frac{2\eta}{\eta-1} \sum_{i' \in B^j} \mu_{i'} y_{i'}$ and (ii) $L_i \leq \frac{2\eta}{\eta-1} \widehat{X}_j$.*

Proof. First observe that $\sum_{i \in B^j} y_i \geq 1 - \frac{1}{\eta}$. For each $i \in B^j$, let $\overline{y}^j_i = \frac{x_{ij}}{\sum_{i' \in B^j} x_{i'j}}$. Note that $\forall i \in B^j$,

$$\overline{y}^j_i \leq \frac{\eta}{\eta - 1} x_{ij} \leq \frac{\eta}{\eta - 1} y_i, \tag{8}$$

holds by Constraints (3) and the fact that at least $\frac{\eta-1}{\eta}$ portion of j's demand is served within B^j (using Markov's inequality).

Now think of $\overline{y}^j$ as a probability distribution over facilities B^j (note that $\sum_{i \in B^j} \overline{y}^j_i = 1$). Suppose we sample a facility i from this distribution.

▶ Claim 1. $\Pr[\mu_i > \frac{2\eta}{\eta-1} \sum_{i' \in B^j} \mu_{i'} y_{i'}] < 1/2$.

Proof: Observe that $\sum_{i' \in B^j} \mu_{i'} \overline{y}^j_{i'} \leq \frac{\eta}{\eta-1} \sum_{i' \in B^j} \mu_{i'} y_{i'}$. This, with Markov's inequality, implies: $\Pr[\mu_i > \frac{2\eta}{\eta-1} \sum_{i' \in B^j} \mu_{i'} y_{i'}] \leq \Pr[\mu_i > 2 \sum_{i' \in B^j} \mu_{i'} \overline{y}_{i'}] < 1/2$. ◀

▶ Claim 2. $\Pr[L_i > \frac{2\eta}{\eta-1} \widehat{X}_j] < 1/2$.

Proof: Using (5) and (8) and the fact that by choice of η, $B^j \subseteq P_j$ we have

$$\widehat{X}_j \geq \sum_{i \in B^j} \sum_{j' \in D} x_{ij'} \geq \sum_{i \in B^j} y_i L_i \geq \frac{\eta - 1}{\eta} \sum_{i \in B^j} L_i \overline{y}^j_i \tag{9}$$

This implies: $\Pr[L_i > \frac{2\eta}{\eta-1} \widehat{X}_j] \leq \Pr[L_i > 2 \sum_{i \in B^j} \overline{y}^j_i L_i] < 1/2$. ◀

The above two claims immediately imply that with positive probability, there is a facility that satisfies both conditions in inequalities (i) and (ii), respectively. Hence the lemma holds. ◀

Our next lemma demonstrates the utility of our clustering algorithm even in the presence of the connectivity requirements. We show below that if we find a (lower/upper bounded) facility location solution within each cluster and if we can connect those open facilities to the center of the clusters using core cables cheaply then we can connect the centers using core cables cheaply. This helps us to reduce the problem to solving each Voronoi cell separately.

▶ **Lemma 5.** *Let (x, y, z) be values satisfying (2)-(3) and (6)-(7) and $(\mathcal{C}, \mathcal{P}, \delta)$ be returned by Algorithm 1 with x, y, and some given λ. Let $\eta \in (1, \lambda)$. Then we can efficiently find a Steiner tree that connects $\mathcal{C}$ with cost at most $\frac{\lambda}{\lambda-\eta} \cdot \frac{2\eta}{\eta-1} \cdot M \cdot \sum_e c_e z_e$.*

The idea of the proof (which appears in Appendix) is to consider balls $\{B(j, \eta L^j) : j \in \mathcal{C}\}$. These balls are disjoint and one can show that looking at the cut constraints between the balls, the z_e values (from LP) satisfy cut constraints for the Steiner tree problem over these balls.

3 An LP-Based Approximation Algorithm for LUFL

In this section we present a rounding bicriteria approximation algorithm for **LUFL**. We start with the simpler case where we only have lower bounds and then show how to extend the algorithm to work for when there are both upper and lower bounds for facility loads.

3.1 Lower-Bounded Facility Location

We first consider the case where all facilities have infinite capacities. An LP to this case can be written as follows. We let (x, y) and $\mathrm{OPT}_{\mathrm{LP}}$ be the optimal solution and the optimum cost of LP-LFL, respectively.

$$\min \quad \sum_{i \in F} \mu_i y_i + \sum_{j \in D} \sum_{i \in F} c_{ij} x_{ij} \tag{LP-LFL}$$
$$(2),(3),(5)$$
$$x_{ij}, y_i \geq 0$$

It is easy to see that LP-LFL has unbounded integrality gap: Consider a small instance of **LBFL** consisting of $2(L-1)$ clients (with unit demands), two zero-cost facilities each collocated with $L-1$ clients, and an edge of length L between these two facilities. While the optimal cost to IP is $L(L-1)$, LP can manage to pay only $2(L-1)$ by opening both facilities to the extent of $\frac{L-1}{L}$, and thereby only sending $\frac{1}{L}$ demand of each client to its far facility. Hence, the integrality gap can be made arbitrarily large by increasing L. Therefore bicriteria approximation is unavoidable if we use LP.

Let $\eta > 1$ be a parameter we may choose, larger values result in more expensive solutions with smaller violations in the lower bound. Our algorithm for **LBFL** has two steps and works as follows. We first find a Voronoi clustering using Algorithm 1 and then for each cluster center j we open one facility in the cell as guaranteed by Lemma 4. All demand $\widehat{X}_j$ that is fractionally assigned to $\mathcal{P}_j$ is assigned to this open facility. To turn this into an integer assignment of clients to facilities, we then compute the minimum-cost integer flow that satisfies the relaxed lower bounds. The fact that this is cheap is witnessed by the fractional assignment we find in the first part of the algorithm.

Algorithm 2: LBFL rounding

Step 1: Construct a Voronoi clustering $(\mathcal{C}, \mathcal{P}, \delta)$ by running Algorithm 1 with the given x, y, and $\lambda = \eta$.

Step 2: Let $I = \{i(j) : \ j \in \mathcal{C}\}$, where $i(j) \in P_j$ is the facility described in Lemma 4. Open facilities I and find the cheapest assignment of clients to them such that each open facility i serves at least $\frac{\eta-1}{2\eta} L_i$ demand.

▶ **Theorem 6.** *Algorithm 2 computes in polynomial time a solution to* ***LBFL*** *with the following properties:*

(i) The solution cost is at most $\max\{4(\eta+1), \frac{2\eta}{\eta-1}\} \cdot OPT_{LP}$.

(ii) Each open facility $i \in I$ *is serving at least* $\lfloor \frac{\eta-1}{2\eta} L_i \rfloor$ *clients.*

Proof. We provide a solution as described in Step 2 fulfilling the claimed properties. Consider $(\mathcal{C}, \mathcal{P}, \delta)$ costructed at Step 1. Recall $\widehat{X}_j = \sum_{i \in P_j} \sum_{j' \in D} x_{ij'}$.

By Lemma 4 and the fact that P_j cells are disjoint, the total opening cost is bounded as follows.

$$\mu(I) \leq \sum_{j \in \mathcal{C}} \mu_{i(j)} \leq \frac{2\eta}{\eta - 1} \sum_{j \in \mathcal{C}} \sum_{i \in P^j} \mu_i y_i \leq \frac{2\eta}{\eta - 1} \sum_{i \in F} \mu_i y_i. \tag{10}$$

Assigning the fractional demands $\widehat{X}_j$ aggregated at j to each $i(j) \in I$ guarantees the second property; see Lemma 4. Hence, we only need to show this assignment is cheap and how to turn it into an integer assignment of no more cost.

Consider some client $j' \in D$ and some facility $i \in F$. In what follows we show that $x_{ij'}$ units of demand travel a distance of at most $4(\eta L^{j'} + c_{ij'})$. Say that $i \in P_j$ and let $B^j := B(j, \eta L^j)$. Thus, in this assignment the $x_{ij'}$-fraction of demand travels distance $c_{j'i(j)}$. We consider two cases:

Case $j' \notin B^j$: We have

$$\begin{array}{rcll}
c_{j'i(j)} & \leq & c_{i(j)j} + c_{jj'} & \text{(by the triangle inequality)} \\
& \leq & \eta L^j + c_{jj'} & \text{(using the fact that } i(j) \in B^j) \\
& \leq & 2c_{jj'} & \text{(using the fact that } j' \notin B^j) \\
& \leq & 2(c_{ij} + c_{ij'}) & \text{(by the triangle inequality)} \\
& \leq & 2(c_{i\delta(j')} + c_{ij'}) & \text{(using the fact that } i \in P_j) \\
& \leq & 2(c_{j'\delta(j')} + 2c_{ij'}) & \text{(by the triangle inequality)} \\
& \leq & 2(2\eta L^{j'} + 2c_{ij'}) & \text{(from the clustering procedure)}
\end{array}$$

Case $j' \in B^j$: In this case $c_{j'i(j)} \leq 2\eta L^j$ (by the triangle inequality). Below we show that $L^j \leq 2L^{j'}$, which immediately implies $c_{j'i(j)} \leq 4\eta L^{j'}$.

▶ Claim 3. $L^j \leq 2L^{j'}$.

Proof: Assume, for the sake of contradiction, that $L^j > 2L^{j'}$. First observe that by the ordering clients are selected as centers in $\mathcal{C}$, $j' \notin \mathcal{C}$: note that $j \in \mathcal{C}$, and since we assumed $L^{j'} < L^j$, and because $c_{jj'} \leq 2\eta L^j$ (recall $j' \in B^j$), if $j' \in \mathcal{C}$ then it would have prevented j from being added to $\mathcal{C}$ in the first step. Now, consider $\delta(j') \in \mathcal{C}$. Note that $L^{\delta(j')} \leq L^{j'}$ and $c_{j'\delta(j')} \leq 2\eta L^{j'}$. This implies

$$\begin{array}{rcll}
c_{j\delta(j')} & \leq & c_{j'j} + c_{\delta(j')j'} & \text{(by the triangle inequality)} \\
& \leq & \eta L^j + 2\eta L^{j'} & \text{(by } j' \in B^j \text{ and clustering procedure)} \\
& \leq & \eta L^j + \eta L^j & \text{(using the assumption that } L^j > 2L^{j'}) \\
& \leq & 2\eta L^j &
\end{array}$$

which is a contradiction because then $\delta(j')$ would also have blocked j from being added to $\mathcal{C}$. The claim follows. ◀

This completes the proof of this case that $c_{j'i(j)} \leq 4\eta L^{j'}$.

In either case, $x_{ij'}$ travels a distance of at most $4(\eta L^{j'} + c_{ij'})$. Thus, the total assignment cost of this fractional solution is bounded by

$$\begin{aligned} 4\sum_{i\in F}\sum_{j'\in D} x_{ij'}\left(\eta L^{j'} + c_{ij'}\right) &= 4\eta \sum_{j'\in D} L^{j'} \sum_{i\in F} x_{ij'} + 4\sum_{i\in F}\sum_{j'\in D} c_{ij'}x_{ij'} \\ &= 4\eta \sum_{j'\in D} L^{j'} + 4\sum_{i\in F}\sum_{j'\in D} c_{ij'}x_{ij'} \quad \text{using (2)} \\ &= 4(\eta+1)\sum_{i\in F}\sum_{j'\in D} x_{ij'}c_{ij'} \quad \text{(by def. of } L^{j'}) \end{aligned}$$

Together with (10), this implies the claimed bound.

Finally, because of the integrality of flows with integer lower bounds and because we have explicitly described a cheap fractional flow from the clients to the open facilities that satisfies the integer lower bounds $\lfloor \frac{\eta-1}{2\eta} L_i \rfloor$, then there is an integer assignment $\sigma : D \to I$ that also satisfies these lower bounds with no greater cost. Note it could be that $\lfloor \frac{\eta-1}{2\eta} L_i \rfloor = 0$. It is still possible to recover a an $O(1)$-approximation that ensures each open facility $i \in I$ is assigned at least $L_i/3$ clients by choosing η to be sufficiently large and closing any $i \in I$ that is not assigned any clients. ◀

For example, by choosing $\eta = 1.28$ we get a solution of cost at most $9.12 OPT_{LP}$ and the load of each open facility i is at least $\lfloor \frac{L_i}{9.12} \rfloor$.

3.2 The general case with lower and upper bounds

We now consider the case where each facility has capacity U (uniform across all facilities) as well as a given lower bound L_i. Let (x, y) be an optimal solution to (LP-LUFL).

As before, we first use Algorithm 1 to obtain a Voronoi clustering. We then decide to open a number of facilities in each cell to route the clients demand to be served at them while satisfying the upper and lower bounds on the facility loads (approximately). The algorithm consists of two steps and works as follow.

Algorithm 3: LUFL rounding

Step 1: Construct a Voronoi clustering $(\mathcal{C}, \mathcal{P}, \delta)$ by running Algorithm 1 with the given x, y, and $\lambda = \eta$.

Step 2: For each $j \in \mathcal{C}$, we open a subset $I_j \subseteq P_j$ of facilities and send the demand served by facilities in $\mathcal{P}_j$ (namely $\widehat{X}_j = \sum_{j'} \sum_{i\in\mathcal{P}_j} x_{ij'}$) to those facilities as described below, depending on the value of $\widehat{X}_j$:

Case 1. $\widehat{X}_j \geq U$**:** In this case, inspired by ideas from [16], we formulate the described subproblem as another (simpler) facility location (inside the cell) using a simpler (sparse) LP.

We firstly move demand $\widehat{X}_j$ to center j as follows. For each client $j' \in D$ and each facility $i \in \mathcal{P}_j$, we send $x_{ij'}$ demand from j' to i (this is what the LP is doing). Let $\hat{d}^i = \sum_{j'\in D} x_{ij'}$ be the demand sent to i. Next, for each facility $i \in P_j$, we send $\hat{d}^i$ demand from i to j. Obviously, the total cost of this moving is bounded by $\sum_{i\in P_j}\sum_{j'\in D} x_{ij'}\left(c_{ij'} + c_{ij}\right)$.

We now ignore the facility lower bounds and write an LP to solve the subproblem. Solving and then rounding this LP helps us to decide which facilities in P_j to open and how to assign the $\widehat{X}_j$ demand (already aggregated at j) to them. We shall show how the cost of this LP can be bounded by the cost of the original LP restricted to this cell and an optimum solution to this LP satisfies the lower bounds on almost all facilities.

In this LP, we have a variable ω_i for each $i \in \mathcal{P}_j$ indicating how much of the $\widehat{X}_j$ is assigned to i.

$$\min \sum_{i \in P_j} \omega_i \left(\frac{\mu_i}{U} + c_{ij} \right)$$

$$\sum_{i \in P_j} \omega_i = \widehat{X}_j$$

$$0 \leq \omega_i \leq U \qquad \forall i \in P_j$$

Note that setting $\omega_i := \sum_{j' \in D} x_{ij'}$ is a feasible solution with cost at most $\sum_{i \in P_j} \mu_i y_i + \sum_{i \in P_j} \sum_{j' \in D} x_{ij'} c_{ij}$ because $\sum_j x_{ij} \leq U y_i$.

Note that there is only one constraint apart from constraints $0 \leq \omega_i \leq U$. Thus, for all but one $i \in P_j$ we have $\omega_i^* \in \{0, U\}$, where ω^* indicates an optimum extreme point solution to this LP.

To round this solution ω_i^*, let $\zeta > 1$ be a parameter we get to choose. Let $I_j = \{i \in P_j : \omega_i^* = U\}$. If there is some $i' \in P_j$ such $0 < \omega_{i'}^* < U$ then add i' to I_j if $\omega_{i'}^* \geq \frac{U}{\zeta}$. In this case, the upper bound is satisfied for every $i \in I_j$ and the lower bound is violated by no more than a ζ-factor. Assign precisely ω_i^* units of demand to each $i \in I_j$. The cost of this assignment plus the cost of opening I_j is at most $\sum_{i \in P_j} \zeta \mu_i y_i + \sum_{i \in P_j} \sum_{j' \in D} x_{ij'} c_{ij}$.

Otherwise, if $\omega_{i'}^* < \frac{U}{\zeta}$ then let i'' be the facility in I_j closest to j and increase $\omega_{i''}^*$ by $\omega_{i'}^*$. Note that such a facility i'' exists because we are assuming $\widehat{X}_j \geq U$. In this case, no lower bounds are violated at any $i \in I_j$ and the upper bound is violated by at most a $\left(1 + \frac{1}{\zeta}\right)$-factor. The assignment and opening cost in this case are bounded by $\sum_{i \in P_j} \mu_i y_i + \sum_{i \in P_j} \sum_{j' \in D} \frac{\zeta+1}{\zeta} x_{ij'} c_{ij}$.

In either case, we have opened I_j and assigned demand to each $i \in I_j$ to satisfy the relaxed lower bounds L_i/ζ and the relaxed upper bounds $\frac{\zeta+1}{\zeta} U$. The cost of assigning $\widehat{X}_j$ units of demand from j to I_j in this manner is at most $\zeta \sum_{i \in P_j} \mu_i y_i + \frac{\zeta+1}{\zeta} \sum_{i \in P_j} \sum_{j' \in D} x_{ij'} c_{ij}$. Altogether, the total cost (of moving the $\widehat{X}_j$ demand to center j plus the cost of assigning it from j to facilities I_j) is bounded by

$$\zeta \sum_{i \in P_j} \mu_i y_i + \frac{\zeta + 1}{\zeta} \sum_{i \in P_j} \sum_{j' \in D} x_{ij'} c_{ij} + \sum_{i \in P_j} \sum_{j' \in D} x_{ij'} \left(c_{ij'} + c_{ij} \right) \qquad =$$

$$\zeta \sum_{i \in P_j} \mu_i y_i + \frac{2\zeta + 1}{\zeta} \sum_{i \in P_j} \sum_{j' \in D} x_{ij'} c_{ij} + \sum_{i \in P_j} \sum_{j' \in D} x_{ij'} c_{ij'}.$$

▶ **Lemma 7.** *The total cost (over all cells of Voronoi clustering) incurred due to Case 1 of Step 2 of the algorithm is at most* $\zeta \sum_{i \in F} \mu_i y_i + \frac{(2\zeta+1)(2\eta+1)+\zeta}{\zeta} \sum_{i \in F} \sum_{j \in D} x_{ij} c_{ij}$. *(see the appendix for proof.)*

Case 2. $\widehat{X}_j < U$**:** Observe that in this case we can simply ignore the upper bound. So (similar to that for **LBFL**) we open facility $i(j)$ described in Lemma 4 and send the demand to that facility as follows: For each client $j' \in D$ and each facility $i \in \mathcal{P}_j$, we send $x_{ij'}$ demand from j' (directly) to $i(j)$. Let $I_j = \{i(j)\}$ in this case. Note that facility $i(j)$ serves at least $\frac{\eta-1}{2\eta} L_i$.

The following bound can be obtained using the exact same arguments used to bound that in the proof of Theorem 6.

▶ **Lemma 8.** *The total cost incurred due to Case 2 of Step 2 is at most* $\frac{2\eta}{\eta-1} \sum_{i \in F} \mu_i y_i + 4(\eta + 1) \sum_{i \in F} \sum_{j \in D} x_{ij} c_{ij}$.

Let $I = \cup_{j\in\mathcal{C}} I_j$ be the set of facilities opened over all Voronoi cells. Observe that each of the two cases above finds a solution to **LBFL** in which each open facility $i \in I$ serves at least $\min\left\{\frac{1}{\zeta}, \frac{\eta-1}{2\eta}\right\}L_i$ (based on the two cases above) and at most $\frac{\zeta+1}{\zeta}U$ demand.

Summing our bounds on the cost of the solutions found in each Voronoi diagram (see Lemmas 7 and 8), we see the cost of opening I is at most

$$(\zeta + \frac{2\eta}{\eta-1})\sum_{i\in F}\mu_i y_i, \tag{11}$$

and the cost of assigning demands is at most:

$$\Big(\frac{(2\zeta+1)(2\eta+1)+\zeta}{\zeta} + 4(\eta+1)\Big)\sum_{j\in D}\sum_{i\in F} c_{ij}x_{ij} \tag{12}$$

Together, (11) and (12) and using integrality of flows with integer lower and upper bounds, imply the main results of this section.

▶ Theorem 2 (restated). Algorithm 3 is a polynomial time (ρ, α, β)-approximation for instances of **LUFL** with uniform capacities where $\rho = \max\{\frac{(2\zeta+1)(2\eta+1)+\zeta}{\zeta} + 4(\eta+1), \frac{2\eta}{\eta-1} + \zeta\}$, $\alpha = \max\{\zeta, \frac{2\eta}{\eta-1}\}$, $\beta = \frac{\zeta+1}{\zeta}$.

4 An LP-Based Approximation Algorithm for C-LUFL

In this section we show that our rounding framework for **LUFL** extends to connected variants. In the light of Lemma 5, we observe that our framework works for the connected variants too, as long as we can bound the cost of connecting facilities opened in each Voronoi cell to the center it belongs to.

We begin with the case where all facilities have infinite capacities (denoted by **C-LBFL**) and then show how to extend it to the case with lower bounds. We let (x, y, z) and $\mathrm{OPT_{LP}}$ be the optimal solution and the optimum cost of the LP relaxation for this case, respectively.

Let $\lambda > \eta > 1$ be constant parameters. Following the same general ideas of that for **LBFL** and using our observation described in Lemma 5, we present our algorithm for **C-LBFL** which has three stages and works as follows.

Algorithm 4: C-LBFL rounding
Step 1: Construct a Voronoi clustering $(\mathcal{C}, \mathcal{P}, \delta)$ by running Algorithm 1 with the given x, y, λ.
Step 2: Open facilities $I = \{i(j) \colon j \in \mathcal{C}\}$ and assign clients to them as described in Step 2 of Algorithm 2. Connect each facility $i(j) \in I$ to the center it belongs to via core cables.
Step 3: Compute a core Steiner tree over centers $\mathcal{C}$ as described in Lemma 5.

One can simply adapt the proof of Lemma 5 to bound the extra cost of connecting each center j to the facility $i(j)$ by losing a constant factor. Apart from this the proof of the following theorem is analogous to that for Theorem 6.

▶ **Theorem 9.** *Algorithm 4 computes in polynomial time a solution to **C-LBFL** with the following properties:*
(i) The solution cost is at most $\max\{4(\eta+1), \frac{2\eta}{\eta-1}, \frac{2\cdot(\lambda+\eta)\eta}{(\lambda-\eta)(\eta-1)}\}\cdot OPT_{LP}$.
(ii) Each open facility $i \in I$ *is serving at least* $\lfloor\frac{\eta-1}{2\eta}L_i\rfloor$ *clients.*

We now consider the **C-LUFL** problem. First we show that one can convert an optimum solution of **C-LUFL** to an approximate solution in which each open facility (say) i is assigned a sufficiently large number of clients comparable not only to U and L_i but also to M (core cable cost per unit length). This property of a near optimal solution will help use to compute approximate solutions to **C-LUFL**. Let $\Delta = \min\{M, U\}$. Let $\text{OPT}_{\text{C-LU}}$ be the cost of an optimal solution to **C-LUFL**. Observe that when the number of clients is less than $\frac{\Delta}{2}$, selecting only the cheapest facility to be opened and then assigning all clients to that open facility returns the optimal solution. We hence assume that the number of clients is at least $\frac{\Delta}{2}$.

▶ **Theorem 10.** *There is a feasible solution of cost at most $3OPT_{C\text{-}LU}$ to **C-LUFL** in which each open facility i is assigned at least $\max\{\frac{\Delta}{2}, L_i\}$ units of demand.*
(see the appendix for proof.)

In what follows (instead of approximating **C-LUFL**) we approximate the near optimal solution described above whose property is needed for our analysis to work. We write a modification of LP-C-LUFL to model the approximate solution described above.

$$
\begin{aligned}
\min \quad & \sum_{i \in F} \mu_i y_i + \sum_{j \in D} \sum_{i \in F} c_{ij} x_{ij} + M \sum_{e \in E} c_e z_e \\
& \text{(2)-(4), (6)-(7)} \\
& \Delta y_i \leq 2 \sum_{j \in D} x_{ij} && \forall i \in F \qquad (13) \\
& x_{ij}, y_i, z_e \geq 0
\end{aligned}
$$

We let (x, y, z) be the optimal solution of this LP relaxation. Let $\lambda > \eta > 1$ be constant parameters. Following the algorithm for **LBFL** and using Lemma 5, we extend the algorithm for **C-LBFL** to work for the more general case where each facility has a capacity U and $M = O(U)$. Our algorithm has three steps and works as follows.

Algorithm 5: C-LUFL rounding:
Step 1: Construct a Voronoi clustering $(\mathcal{C}, \mathcal{P}, \delta)$ by running Algorithm 1 with the given x, y, λ.
Step 2: Open facilities $I = \cup_{j \in \mathcal{C}} I_j$ and assign clients to them as described in Step 3 of Algorithm 3. Then, connect each facility $i \in I$ to the center of the cell it belongs to using core cables.
Step 3: Compute a core Steiner tree over centers $\mathcal{C}$ as described in Lemma 5.

▶ Theorem 3 (restated). Algorithm 5 computes in polynomial time a $\left(O(1), \max\{\zeta, \frac{2\eta}{\eta-1}\}, \frac{\zeta+1}{\zeta}\right)$-bicriteria approximation for instances of **C-LUFL** with uniform capacities (and non-uniform lower bounds) and with $M = O(U)$.

The proof is deferred to the Appendix.

5 Conclusion

It would be nice to extend our approximations for C-LUFL to include the case $M = \omega(U)$. As M gets larger, the cost of connecting core cables becomes so large that an optimum solution would open the fewest possible facilities, namely $k := \lceil |D|/U \rceil$. This case resembles the well-studied k-MST problem where it is well-known that even getting a constant-factor bicriteria approximation is not possible using the natural cut-based relaxation. So, this case poses additional difficulties.

Also open is the problem of getting constant-factor biceriteria approximations for **LUFL** when both lower and upper bounds are not necessarily uniform.

References

1 Zoë Abrams, Adam Meyerson, Kamesh Munagala, and Serge Plotkin. The integrality gap of capacitated facility location. *Technical Report CMU-CS-02-199, Carnegie Mellon University*, 2002.

2 Ankit Aggarwal, Anand Louis, Manisha Bansal, Naveen Garg, Neelima Gupta, Shubham Gupta, and Surabhi Jain. A 3-approximation algorithm for the facility location problem with uniform capacities. *Mathematical Programming*, 141, 2013.

3 Sara Ahmadian and Chaitanya Swamy. Improved approximation guarantees for lower-bounded facility location. *In proceedings of WAOA 2012*, pages 257–271.

4 Hyung-Chan An, Monika Singh, and Ola Svensson. LP-based algorithms for capacitated facility location. *In proceedings of FOCS 2014*, pages 256–265.

5 Manisha Bansal, Naveen Garg, and Neelima Gupta. A 5-approximation for capacitated facility location. *In proceedings of ESA 2012*, pages 133–144.

6 Fabián A Chudak and David P Williamson. Improved approximation algorithms for capacitated facility location problems. *Mathematical programming*, 102, 2005.

7 Friedrich Eisenbrand, Fabrizio Grandoni, Thomas Rothvoß, and Guido Schäfer. Connected facility location via random facility sampling and core detouring. *Journal of Computer and System Sciences*, 76(8):709–726, 2010.

8 Fabrizio Grandoni and Thomas Rothvoß. Approximation algorithms for single and multi-commodity connected facility location. *In proceedings of IPCO 2011*.

9 Sudipto Guha, Adam Meyerson, and Kamesh Munagala. Hierarchical placement and network design problems. *In proceedings of FOCS 2000*, pages 603–612, 2000.

10 Anupam Gupta, Jon Kleinberg, Amit Kumar, Rajeev Rastogi, and Bulent Yener. Provisioning a virtual private network: a network design problem for multicommodity flow. *In proceedings of STOC 2001*, pages 389–398.

11 Hyunwoo Jung, Mohammad Khairul Hasan, and Kyung-Yong Chwa. A 6.55 factor primal-dual approximation algorithm for the connected facility location problem. *Journal of combinatorial optimization*, 18(3):258–271, 2009.

12 DR Karget and Maria Minkoff. Building steiner trees with incomplete global knowledge. *In proceedings of FOCS 2000*, pages 613–623.

13 Madhukar R Korupolu, C Greg Plaxton, and Rajmohan Rajaraman. Analysis of a local search heuristic for facility location problems. *Journal of algorithms*, 2000.

14 Retsef Levi, David B Shmoys, and Chaitanya Swamy. LP-based approximation algorithms for capacitated facility location. *Mathematical programming*, 131, 2012.

15 Shi Li. On uniform capacitated k-median beyond the natural lp relaxation. *In proceedings of SODA 2015*, 2015.

16 Shi Li. Approximating capacitated k-median with $(1+\epsilon)$k open facilities. *In proceedings of SODA 2016*, 2016.

17 Martin Pal, Eva Tardos, and Tom Wexler. Facility location with nonuniform hard capacities. *In proceedings of FOCS 2001*, pages 329–338.

18 David B Shmoys, Éva Tardos, and Karen Aardal. Approximation algorithms for facility location problems. *In proceedings of STOC 1997*, pages 265–274.

19 Zoya Svitkina. Lower-bounded facility location. *ACM Transactions on Algorithms (TALG)*, 6(4):69, 2010.

20 Chaitanya Swamy and Amit Kumar. Primal–dual algorithms for connected facility location problems. *Algorithmica*, 40(4):245–269, 2004.

A Missing proofs.

Poof of Lemma 5: Note that we require $\eta < \lambda$. We assume that facility $r \in B(j, \eta L^j)$ for some $j \in \mathcal{C}$. The other case where $r \notin B(j, \eta L^j)$ for any $j \in \mathcal{C}$ is nearly identical and results in the same bound. We also observe that $\{B(j, \eta L^j) : j \in \mathcal{C}\}$ consists of disjoint sets: if $B(j, \eta L^j) \cap B(j', \eta L^{j'}) \neq \emptyset$ for distinct $j, j' \in \mathcal{C}$ then $c_{jj'} \leq 2\lambda \cdot \max\{L^j, L^{j'}\}$ so both j and j' could not be cluster centers.

Note that $\sum_{i \in B(j, \eta L^j)} x_{ij} \geq \frac{\eta-1}{\eta}$ holds for any $j \in \mathcal{C}$, using of Markov's inequality. This, together with (6), implies that vector $\frac{\eta}{\eta-1} z$ is a feasible fractional solution to the standard cut based LP relaxation of the Steiner tree problem with terminals being balls $B(j, \eta L^j)$ contracted at their centers. Thus, we can efficiently find a Steiner tree $\hat{T}$ over these contracted balls (on the resulting graph after contracting balls) with cost at most $\frac{2\eta}{\eta-1} \sum_e c_e z_e$.

Now we have to convert this tree $\hat{T}$ into a Steiner tree over centers $\mathcal{C}$. When we uncontract the balls, each edge of $\hat{T}$ between two balls around centers j, j' can be replaced with the edge between two closest nodes, say $u \in B(j, \eta L^j)$ and $v \in B(j', \eta L^{j'})$. We add edges ju and vj' for each such $uv \in \hat{T}$ to complete the Steiner tree. To bound the cost of these new edges, observe that $\eta < \lambda$ and not only balls $B(j, \eta L^j)$ and $B(j', \eta L^{j'})$ are disjoint, but also balls $B(j, \lambda L^j)$ and $B(j', \lambda L^{j'})$ are disjoint as well by the same argument. So we can "charge" the cost of two new edges ju and vj' to the section of edge uv that falls between the two nested balls as follows. Let $\alpha = \max\{L^j, L^{j'}\}$. Since $u \in B(j, \eta L^j)$ and $v \in B(j, \eta L^{j'})$ then $c_{uj} + c_{j'v} \leq 2\eta\alpha$. Furthermore, $2\lambda\alpha \leq c_{jj'} \leq c_{uj} + c_{uv} + c_{vj'} \leq c_{uv} + 2\eta\alpha$. Therefore,

$$c_{ju} + c_{vj'} \leq 2\eta\alpha = \frac{2\eta}{\lambda - \eta} \cdot (\lambda - \eta) \cdot \alpha \leq \frac{\eta}{\lambda - \eta} c_{uv}.$$

Thus, the total cost of this tree is at most $\left(1 + \frac{\eta}{\lambda-\eta}\right) \cdot \frac{2\eta}{\eta-1} \cdot M \cdot \sum_e c_e z_e$. ◀

Proof of Lemma 7: The total cost is bounded by

$$\begin{aligned}
&\sum_{j \in \mathcal{C}} \Big(\zeta \sum_{i \in P_j} \mu_i y_i + \frac{2\zeta + 1}{\zeta} \sum_{i \in P_j} \sum_{j' \in D} x_{ij'} c_{ij} + \sum_{i \in P_j} \sum_{j' \in D} x_{ij'} c_{ij'} \Big) \\
&= \zeta \sum_{i \in F} \mu_i y_i + \frac{2\zeta + 1}{\zeta} \sum_{j \in \mathcal{C}} \sum_{i \in P_j} \sum_{j' \in D} x_{ij'} c_{ij} + \sum_{i \in F} \sum_{j' \in D} x_{ij'} c_{ij'}, \qquad (14)
\end{aligned}$$

using the fact that P_j cells are disjoint.

Note that for any $i \in P_j$ and any $j' \in D$ we have

$$\begin{aligned}
c_{ij} &\leq c_{i\delta(j')} && \text{(using the fact that } i \in P_j) \\
&\leq c_{ij'} + c_{j'\delta(j')} && \text{(by the triangle inequality)} \\
&\leq c_{ij'} + 2\eta L^{j'} && \text{(from Step 1)}
\end{aligned}$$

Hence, we have:

$$\begin{aligned}
\frac{2\zeta + 1}{\zeta} \sum_{j \in \mathcal{C}} \sum_{i \in P_j} \sum_{j' \in D} x_{ij'} c_{ij} &\leq \frac{2\zeta + 1}{\zeta} \sum_{j \in \mathcal{C}} \sum_{i \in P_j} \sum_{j' \in D} x_{ij'} (c_{ij'} + 2\eta L^{j'}) \\
&= \frac{2\zeta + 1}{\zeta} \sum_{i \in F} \sum_{j' \in D} c_{ij'} x_{ij'} + \frac{(2\zeta + 1)(2\eta)}{\zeta} \sum_{j' \in D} L^{j'} \quad \text{(by (2))} \\
&= \frac{(2\zeta + 1)(2\eta + 1)}{\zeta} \sum_{i \in F} \sum_{j' \in D} c_{ij'} x_{ij'}.
\end{aligned}$$

This, together with (14), implies the claimed bound. ◀

Proof of Theorem 10: Consider an optimum solution of **C-LUFL**. Let I^*, σ^*, and T^* be the optimal set of open facilities, assignment function of clients, and the core network tree, respectively. In the following we construct a feasible solution to **C-LUFL** which costs at most three times the cost of the optimal solution, while respecting the claimed lower bounds for facilities.

We first transform T^* into a TSP tour over I^*, call it Γ (by doubling and short-cutting over repeated facilities). Note that $\sum_{e\in\Gamma} c_e \leq 2\sum_{e\in T^*} c_e$. Our idea is to move around this tour and close some facilities that serve few clients and move the assigned clients to them to the next facility and show that this can be done without increasing the cost too much. Let w_i be the demand that is currently assigned to facility $i \in I^*$; let $w_i = |{\sigma^*}^{-1}(i)|$ initially. For the sake of presentation, let us assume that there exists some facility $\hat{i}_0 \in I^*$ for which $w_{\hat{i}_0} = \frac{\Delta}{2} \geq L_{\hat{i}_0}$ (this can be simply enforced either by replacing some facility with two collocated copies of that facility or just by considering a chain of successive facilities of the tour as the claimed facility). Now, we traverse the tour Γ in (say) clockwise direction and adjust the loads of the facilities one by one. Let $\hat{i}_t \in I^*$ (for $t < |I^*|$) be the t-th facility that is reached in the tour (after leaving $\hat{i}_0$). Starting from $\hat{i}_1$, if $w_{\hat{i}_1} < \Delta$, we send $w_{\hat{i}_1}$ units of demand from $\hat{i}_1$ to $\hat{i}_2$ via edge $\{i_1, i_2\} \in \Gamma$; for otherwise, we send $\max\{w_{\hat{i}_1} - U, 0\}$ units of demand from $\hat{i}_1$ to $\hat{i}_2$. We update $w_{\hat{i}_1}$ and $w_{\hat{i}_2}$ (as well as the assignment mapping), accordingly. This guarantees that $\max\{\Delta, L_{\hat{i}_1}\} \leq w_{\hat{i}_1}$ as far as $w_{\hat{i}_1} > 0$. We perform this procedure for facilities $\hat{i}_2, \hat{i}_3, \cdots, \hat{i}_{|I^*|-1}$ in turn, however for facility $\hat{i}_{|I^*|-1}$ we send its current load to the next facility, namely facility $\hat{i}_0$, only if $w_{\hat{i}_{|I^*|-1}} < \frac{\Delta}{2}$. After this, we simply close the facilities which are now assigned no demand. Clearly, this results a feasible solution satisfying the claimed property. It only remains then to bound the cost of this solution.

Obviously the facility opening and the core connection costs of the new solution is no more than that of the initial solution. Observe that the total demand crossing along each edge of the tour (due to the clients reassignments described above) is no more than Δ. Hence, using the triangle inequalities, one can simply show that the total reassignments cost is bounded by $\Delta \cdot \sum_{e\in\Gamma} c_e + \sum_{j\in D} c_{\sigma^*(j)j}$. Note that that $\Delta \leq M$ and that $\sum_{e\in\Gamma} c_e \leq 2\sum_{e\in T^*} c_e$. Altogether, we deduce that the total reassignments cost is bounded by $2M\sum_{e\in T^*} c_e + \sum_{j\in D} c_{\sigma^*(j)j} \leq 2OPT$. This concludes the proof. ◀

Proof (sketch) of Theorem 3:

▶ **Lemma 11.** *The core connection cost incurred by Step 2 of the algorithm is at most $O(1)\sum_{i\in F}\sum_{j\in D} x_{ij}c_{ij}$.*

Sketch. By (13), one can simply verify that $\widehat{X}_j \geq \frac{\eta-1}{2\eta}\Delta$. This implies

$$\omega_i \geq \max\{\frac{\eta-1}{2\eta}, \frac{1}{\zeta}\}\Delta \tag{15}$$

where ω_i indicates the demand which is sent from center j to facility $i \in I_j$ at the end of this step. (see Step 2 of Algorithm 3 for details).

The core connection cost incurred by Step 2 can be bounded as follows:

$$\sum_{j\in\mathcal{C}} M \sum_{i\in I_j} c_{ij} \leq O(1)\sum_{j\in\mathcal{C}}\sum_{i\in I_j} c_{ij}w_i \leq O(1)\sum_{i\in F}\sum_{j\in D} c_{ij}x_{ij}$$

using (15), the assumption that $M \leq O(1)\Delta$, and the exact same arguments used to bound the cost incurred by Step 2 of Algorithm 3. ◀

Together, Lemma 5, Lemma 11, and the results of Theorem 2 conclude the proof. ◀