Multicommodity Facility Location

R. Ravi[*] A. Sinha[†]

Abstract

Multicommodity facility location refers to the extension of facility location to allow for different clients having demand for different goods, from among a finite set of goods. This leads to several optimization problems, depending on the costs of opening facilities (now a function of the commodities it serves). In this paper, we introduce and study some variants of multicommodity facility location, and provide approximation algorithms and hardness results for them.

1 Introduction

Facility location refers to the class of problems involving the location of facilities to serve clients in a metric space. The classical uncapacitated facility location problem (UFL), for example, is defined in a metric space, a subset of whose points are clients and another subset consists of facilities. Each facility has an opening cost, and the objective is to compute a solution which minimizes the total cost of open facilities and the sum of the distances from each client to its nearest open facility. This problem and its variants, have been studied extensively (See Section 1.3).

In this paper, we study a natural extension of the above model where there is a finite set of commodities, and each client demands a subset of those commodities. We call this *Multicommodity facility location* (MCFL). Facility costs are now a function of the location and the commodities served. This model has many practical applications. For example, a consumer goods manufacturer who produces five different goods might choose to locate three plants to manufacture them, to utilize the economies of scale arising from having one plant producing two commodities instead of a separate plant for each commodity.

1.1 Model Our model and notation extend the well-studied UFL model. There is a set of clients D, and a set of facilities F, with a distance function $c : D \times F \rightarrow \mathbb{R}^+$ which is a metric. There is also a set of commodities S, and each client j has demand for one unit of commodity $d_j \in S$. The number of clients $|D|$, facilities $|F|$ and commodities $|S|$ are abbreviated n, m and k respectively. Facility i can be opened in any *configuration* $\sigma(i) \in 2^S$, specifying which commodities it is serving, at a cost $f_i(\sigma) \in \mathbb{R}^+$. We assume that the facilities have *decreasing marginal costs*, that is, for any facility i and any two configurations σ, σ', we have $f_i(\sigma \cup \sigma') \leq f_i(\sigma) + f_i(\sigma')$. We remark on the implications of relaxing this assumption later.

In a version we dub the t-MCFL, every facility has a subset $T \subseteq S$ of upto t commodities which may be involved in its allowable configurations reducing the input size from 2^k to 2^t per facility. Going one step further, we define a compact class of natural facility cost-functions that are *linear*. In this model, each facility i has an opening cost (fixed cost) f_i^0, and for each commodity s, an incremental cost f_i^s, so that the cost of opening facility i in configuration σ is $f_i(\sigma) = f_i^0 + \sum_{s \in \sigma} f_i^s$. Linear cost functions also satisfy decreasing marginal costs, and in many real situations are a close approximation to general cost functions. Since a linear cost function can be compactly represented by $k + 1$ numbers, there is hope of devising algorithms with running time polynomial in k, n, m instead of $2^k, n, m$.

A feasible solution is specified by a set of facilities F', together with a configuration $\sigma(i)$ of each facility $i \in F'$, such that each commodity which has at least one client demanding it is served by at least one facility in F'. Given F', each client j is assigned to $\phi(j) \in F'$, the nearest open facility which includes commodity d_j in its configuration. The total facility cost is $\sum_{i \in F'} f_i(\sigma(i))$, the sum of the costs incurred in opening each facility in F' in its chosen configuration. The total service cost is $\sum_{j \in D} c_{j,\phi(j)}$, the sum of distances from each client to the facility assigned to it. The total cost is the sum of these two, and the objective is to minimize the total cost.

While the general model allows for clients to specify a vector of demands for the different commodities, we will study the simplified version where each client demands a single unit of a single commodity. This is for ease of exposition; all our algorithms extend

[*]GSIA, Carnegie Mellon University. ravi@cmu.edu. Supported in part by NSF grant CCR-0105548 and ITR grant CCR-0122581 (the ALADDIN project).

[†]GSIA, Carnegie Mellon University. asinha@andrew.cmu.edu. Supported in part by a CBI graduate fellowship, NSF grant CCR-0105548 and ITR grant CCR-0122581 (the ALADDIN project).

to the general case - for example, if a client requires several commodities, we can create several copies of the client each requiring a single commodity. In fact, we can even capture the model when service costs "scale" differently for different commodities, by simply weighting the demand at the clients appropriately.

1.2 Our results and paper outline We begin with a brief survey of existing related work in the next section.

In the next section, we address the main t-MCFL problem. In Section 2.1, we first show that the MCFL problem with linear facility costs is as hard as weighted set cover (Theorem 2.1). We match this with a logarithmic approximation algorithm (Theorem 2.5) in Section 2.3 for the problem with linear facility costs running in time $O(\text{poly}(n, m, k))$. Our hardness result also implies that the t-MCFL problem is MAX-SNP hard even for $t = 3$ via a reduction from weighted set cover with upto t elements per set. Note that the input size to this problem is $\Omega(2^t \text{poly}(n, m))$. Moreover, the best-known performance ratio for the closely related t-set cover problem achievable in running time $O(2^t \text{poly}(n, m))$ is $O(\log t)$ [6]. In Section 2.2, we give an $O(\log t)$-approximation for the general t-MCFL problem running in time $O(2^t \text{poly}(n, m))$ (Theorem 2.3) matching the best-known results for t-set cover.

We study MCFL with additional restrictions in Section 3. First, we study the priority facility location (PFL) problem: Here, each client also has a level-of-service requirement, indicated by an integer between 1 and k. Each facility has a cost function which is non-decreasing in its level-of-service, ie, $f_i(s) \geq f_i(s')$ whenever $s \geq s'$. The objective is to construct a minimum cost solution where client j with priority d_j is served by a facility with level-of-service at least d_j. We may think of each priority level as a commodity, so that priority facility location (PFL) is a special case of MCFL. The priority model occurs very often in real-world scenarios; Charikar, Naor and Scheiber [4] studied priority Steiner tree to model multicast networks with level-of-service requirements. While we can model PFL as a special case of MCFL by associating a commodity with each priority level, this results in facility cost functions which are not linear. Hence, the logarithmic hardness result in Theorem 2.1 does not extend to PFL. We provide a 7-approximation algorithm (Theorem 3.1) for PFL in Section 3.1, by adapting techniques of filtering and rounding fractional LP solutions [8, 13]. Once again, the facility cost functions have compact representations, and our algorithm also has running time $O(\text{poly}(n, m, k))$. Our algorithm extends the ideas of the MCFL algorithms, by using the filtering technique of Lin and Vitter twice, once for service costs as in [13], and once more for facility costs.

Finally, we consider the restriction of MCFL where each commodity can be manufactured at only one facility, for example, due to licensing restrictions, or because it is easier to maintain quality control when all goods of a commodity are being manufactured at the same facility. Our hardness result in Theorem 2.1 extends to this case because in the hardness construction, we only have one client per commodity, and hence only the facility serving this client needs to manufacture its commodity. We can extend the results on general t-MCFL to this case to give an $O(\log t)$-approximation for this extension (Theorem 3.2). In the very special case where each facility cost function $f_i(\sigma) = f_i$ for all $\sigma \neq \emptyset$, we show this improves to a constant factor approximation (Theorem 3.3). We conclude with some open questions in Section 4.

1.3 Related work The uncapacitated (single commodity) facility location problem has received a lot of attention lately. The first constant factor approximation uses LP rounding and is due to Shmoys, Tardos and Aardal [13]; our algorithm builds on their results and those of Lin and Vitter [8]. Other work includes the current best approximation ratio of 1.52 due to Mahdian, Ye and Zhang [9], and an inapproximability threshold of 1.46 due to Guha and Khuller [5].

The facility location model has been extended to include capacities on facilities [11], online facility location [10], restrictions on the number of facilities opened (the k-median problem) [2], combining facility location with network design [12], etc. However, none of these models incorporate multiple commodities; to the best of our knowledge, our algorithms are the first non-trivial approximation algorithms for multicommodity facility location. Note that an $O(\log n)$ approximation can be obtained using a greedy algorithm, such as the following: Repeatedly choose to open a facility and configuration that minimizes the ratio of facility opening cost plus service costs to a subset of clients to the number of clients served by the chosen facility-configuration pair. Also, an $O(k)$-approximation can be obtained by considering each commodity separately and combining the resulting approximate solutions.

Due to the generality and applicability of the multicommodity paradigm, it has been well studied in a lot of contexts, evidenced, for example, by the vast amount of work on multicommodity flow. Also, the multicommodity rent-or-buy problem [7] models a problem where we are required to construct a network which connects all clients of the same commodity, under a special concave cost function on the edges. The priority (or level-of-service) model we study in Section 3.1 is inspired by a recent paper on priority Steiner trees by Charikar, Naor and Scheiber [4], who obtain an $O(\log n)$ approximation for their problem. However, in the context of MCFL, the priority model turns out to be a more tractable special case, yielding a constant-factor approximation.

2 Facility location with multiple commodities

We study the general multicommodity facility location problem (MCFL) as defined in Section 1.1. After showing some hardness results, we provide an $O(\log t)$ approximation algorithm for the t-MCFL problem in Section 2.2; we then extend our algorithm to a faster algorithm when the facility cost function is linear in Section 2.3.

2.1 Hardness We prove the hardness of MCFL follows by reducing from set cover. An instance of (weighted) set cover is specified by a collection C of subsets of S, each with a weight w_c. The objective is to find a minimum weight sub-collection $C' \subseteq C$ such that $\cup_{c \in C'} c = S$. Arora and Sudan [1] showed that it is impossible to approximate set cover better than $\Omega(\log p)$ unless $P = NP$ where p is the maximum size of a set.

THEOREM 2.1. *Any ρ-approximation algorithm for MCFL with linear facility costs yields a ρ-approximation algorithm for weighted set cover.*

Proof. Given an instance of weighted set cover specified by (S, C, w), we transform it into an instance of MCFL. The set of commodities is the set of elements S. We also have one client $d \in D$ for each commodity $s \in S$, and a facility i for every set $c \in C$. The distance between any client and any facility is zero. The cost function of facility i is $f_i(\sigma) = w_c$ if $\sigma \subseteq c$ and $f_i(\sigma) = \infty$ otherwise.

Any solution to this MCFL instance with cost less than ∞ is a feasible solution to the set cover instance. Moreover, any minimal solution selects each open facility in only one configuration (namely, its maximal configuration), so that the total cost of the MCFL solution is the same as the cost of the associated set cover solution. Hence any ρ approximation for MCFL yields a ρ approximation for weighted set cover.

Finally observe that the facility costs are linear and satisfy the decreasing marginal costs property.

We also note that weighted set cover with upto t elements per set and a total of n elements to cover is MAX-SNP-hard even when we allow the algorithm running time $O(2^t \text{poly}(n, m))$ [6], implying the same for t-MCFL via the above reduction.

2.2 Approximation algorithm for t-MCFL We first describe in detail the version of the algorithm when the maximum number of allowable commodities in any facility's configuration t equals the total number of commodities k, for simplicity. We then outline briefly the changes to adapt the result to the t-MCFL problem at the end of the subsection.

2.2.1 Overview Our algorithm begins by solving the LP relaxation of the Balinski integer program formulation extended to MCFL. We then filter the solution so that each client is fractionally assigned only to facilities which are close to it. Following this, we select a set of representative clients, and assign all other clients to these representatives. We then view the fractional solution consisting of the fractionally opened facilities and the representatives as a fractional solution to an instance of k-set cover. (The k-set cover problem is a special case of weighted set cover where each set has no more than k elements.) This k-set cover instance can be rounded to an integer solution within a factor of $O(\log k)$ using standard techniques. Finally, we assign the remaining clients to the facilities who serve their representatives. The rounding process resembles the technique used by Shmoys, Tardos and Aardal [13] for UFL, but with suitable adaptations to incorporate the multiple commodities.

2.2.2 IP formulation The IP formulation extends Balinski's formulation [3] of UFL to MCFL. We maintain an indicator variable y_i^σ for each configuration σ of each facility i. Variable x_{ij}^σ is 1 iff client j is served by facility i in configuration σ.

$$\min \sum_{i,j,\sigma} c_{ij} x_{ij}^\sigma \quad + \quad \sum_{i,\sigma} f_i(\sigma) y_i^\sigma \qquad (IP_{MCFL})$$

$$\sum_i \sum_{\sigma: d_j \in \sigma} x_{ij}^\sigma \quad \geq \quad 1 \qquad \forall j \in D$$

$$x_{ij}^\sigma \quad \leq \quad y_i^\sigma \quad \forall j \in D,\ \forall i \in F,\ \forall \sigma \in 2^S$$

$$x, y \quad \text{non-negative integers}$$

Let LP_{MCFL} denote the linear relaxation of IP_{MCFL}. The size of LP_{MCFL} is polynomial in n, m and 2^k; therefore LP_{MCFL} can be solved to optimality in running time polynomial in its size. We use the optimal solution of LP_{MCFL}, denoted (x, y) with objective function value OPT_{MCFL}, as our lower bound.

2.2.3 Filtering The next step of the algorithm uses the filtering technique of Lin and Vitter [8], as was also done in [13]. We fix a constant $0 < \alpha < 1$. For every client j, we define its optimal fractional service cost to be $c_j^* = \sum_{i,\sigma} c_{ij} x_{ij}^\sigma$. Order the facilities which serve client j according to non-decreasing distance from j. The α point $g_j(\alpha)$ for client j is the smallest distance

c_j^α such that $\sum_{\sigma,i:c_{ij}\le c_j^\alpha} x_{ij}^\sigma \ge \alpha$.

The following theorem, due to Shmoys, Tardos and Aardal [13], can be proved along similar lines because we can treat each facility-configuration pair as a distinct facility.

THEOREM 2.2. *[13] The fractional solution (x, y) can be transformed in time polynomial in $n, m, 2^k$ to a fractional solution $(\overline{x}, \overline{y})$ such that (i) $\overline{x}_{ij}^\sigma > 0 \Rightarrow c_{ij} \le c_j^\alpha$ for all $i \in F, \sigma \in 2^S, j \in D$; (ii) $c_j^\alpha \le \frac{1}{1-\alpha} c_j^*$; (iii) $\overline{y}_i^\sigma \le \min\{1, \frac{y_i^\sigma}{\alpha}\}$ for all $i \in F, \sigma \in 2^S$.*

2.2.4 Selection of representatives

The existence of multiple commodities presents several difficulties if we attempt to round the fractional solution $(\overline{x}, \overline{y})$. Hence we introduce a new step where we select a set of clients as representatives such that no two representatives of the same commodity are fractionally served by the same facility. We do this representative selection independently for each commodity.

Fix a commodity s, and consider all clients which require commodity s in increasing order of c_j^α. Let $D_s = \{j_1, j_2, \ldots, j_{n_s}\}$ be the clients in this order. Iteratively, mark the smallest index (smallest c_j^α) client $j \in D_s$ as a representative. All clients $j' \in D_s$ such that there exists a facility-configuration pair (i, σ) such that $\overline{x}_{ij}^\sigma > 0$ and $\overline{x}_{ij'}^\sigma > 0$ are removed from the list D_s (In this case, note that $d_{j'} = d_j \in \sigma$). We do not consider client j' as a candidate for being a representative, and instead mark client j as the representative for client j'. Let R denote the set of representatives over all commodities, and $\overline{x}|_R$ denote the set of $\overline{x}_{ij}^\sigma$ variables for $j \in R$.

2.2.5 Interpretation as a fractional k-set cover solution

The k-set cover problem is a special case of the set cover problem when each set has cardinality no more than k. We use our fractional solution to construct a k-set cover instance and show that $\overline{y}$ is a fractional solution of the k-set cover instance. We use the following IP formulation for our instance of k-set cover. We create a set (i, σ) with cost $f_i(\sigma)$ for every facility-configuration pair (i, σ). Let z_i^σ be 1 if set (i, σ) is included in our solution. Our universe consists of all clients in R, and a client $j \in R$ is included in a set (i, σ) if and only if $\overline{x}_{ij}^\sigma > 0$.

$$\min \sum_{i,\sigma} f_i(\sigma) z_i^\sigma \qquad (IP_{k-SC})$$

$$\text{s.t.} \sum_{(i,\sigma):\overline{x}_{ij}^\sigma > 0} z_i^\sigma \ge 1 \quad \forall j \in R$$

LEMMA 2.1. *IP_{k-SC} is an instance of k-set cover, and $z = \overline{y}$ is a feasible solution for its linear relaxation of cost no more than $\sum_{i,\sigma} f_i(\sigma)\overline{y}_i^\sigma$.*

Proof. Since we select the set of representatives R in such a way that no two representatives of the same commodity IP_{k-SC} are served by the same facility, the cardinality of any set in IP_{k-SC} is no more than k. Also, since $(\overline{x}, \overline{y})$ is feasible for LP_{MCFL}, we have $\sum_{i,\sigma} \overline{x}_{ij}^\sigma \ge 1$ for every $j \in R$, and $z_i^\sigma = \overline{y}_i^\sigma \ge \overline{x}_{ij}^\sigma$ guarantees the feasibility of z for the linear relaxation of IP_{k-SC}. This also bounds the cost of the fractional solution.

LEMMA 2.2. *There exists an integer solution $(\hat{x}|_R, \hat{y})$ such that (i) $\hat{x}_{ij}^\sigma \le \hat{y}_i^\sigma$; (ii) $\hat{x}_{ij}^\sigma = 1$ only if $\overline{x}_{ij}^\sigma > 0$; (iii) $\hat{y}_i^\sigma = 1$ only if $\overline{y}_i^\sigma > 0$; (iv) $\sum_{i,\sigma} f_i(\sigma)\hat{y}_i^\sigma \le H_k \sum_{i,\sigma} f_i(\sigma)\overline{y}_i^\sigma$, where $H_k = 1 + \frac{1}{2} + \ldots + \frac{1}{k} \approx \log k$.*

Proof. We know that the integrality gap of k-set cover is no more than $\log k$ (for example, see [14]). Hence there exists an integer solution $\hat{z}$ of IP_{k-SC} with cost no more than $\log k \sum_{i,\sigma} f_i(\sigma)\overline{y}_i^\sigma$. We therefore set $\hat{y} = \hat{z}$, guaranteeing (iii) and (iv) above. Also, using the feasibility of $\hat{z}$ and the construction of IP_{k-SC}, for every client $j \in R$ we must have $\hat{y}_i^\sigma = 1$ for at least one facility-configuration pair (i, σ) such that $\overline{x}_{ij}^\sigma > 0$. Define $\hat{x}_{ij}^\sigma = 1$ for this (i, σ) pair, and zero everywhere else. This guarantees (i) and (ii).

In our final solution, we only open those facility-configuration pairs with $\hat{y}_i^\sigma = 1$. Theorem 2.2 and Lemma 2.2 bound the facility opening cost of by $\frac{H_k}{\alpha}$ of the facility cost of OPT_{MCFL}.

2.2.6 Bounding the service cost

Clients in R are assigned to facility-configuration pairs (i, σ) according to the variables $\hat{x}_{ij}^\sigma$. For a client $j' \notin R$, we assign it to the facility-configuration pair which serves the representative j of client j'. Clearly, this results in a feasible solution of IP_{MCFL}.

LEMMA 2.3. *Each client is assigned to a facility at distance no more than $\frac{3}{1-\alpha} c_j^*$.*

Proof. For each client $j \in R$, Theorem 2.2, parts (i) and (ii) guarantee that $\overline{x}_{ij}^\sigma > 0$ only if $c_{ij} \le \frac{c_j^*}{1-\alpha}$. Consider a client $j' \notin R$, and let $j \in R$ be its representative. We therefore have $c_{j'}^\alpha \ge c_j^\alpha$, so that $c_{jj'} \le 2c_{j'}^\alpha \le \frac{2}{1-\alpha} c_{j'}^*$. Since client j' is assigned to the facility $\phi(j)$ which serves j, we have $c_{\phi(j)j'} \le c_{\phi(j)j} + \frac{2}{1-\alpha} c_{j'}^* \le \frac{3}{1-\alpha} c_{j'}^*$.

Using Lemmas 2.2 and 2.3, we can bound the total cost of our solution.

THEOREM 2.3. *Our algorithm produces a solution with total facility cost no more than $\frac{H_k}{\alpha}$ times the optimal facility cost and service cost no more than $\frac{3}{1-\alpha}$ times the optimal service cost for any $0 < \alpha < 1$.*

Notice that in the above result, we might be able to trade off α with respect to k to get improved guarantees for specific instances depending on the relative contributions of facility costs and service costs in an LP relaxation.

2.2.7 Consequences of decreasing marginal costs The algorithm above essentially treats each facility-configuration pair as a distinct facility. Hence it may end up opening the same facility in more than one configuration. However, if we have decreasing marginal costs ($f_i(\sigma \cup \sigma') \leq f_i(\sigma) + f_i(\sigma')$), we can open the facility in exactly one configuration, which is the union of all the configurations chosen for the facility. Decreasing marginal costs guarantees that this only improves the cost while maintaining feasibility.

If we did not have decreasing marginal costs, our algorithm might require us to open multiple copies of facilities, in different configurations. The model where each facility is restricted to be opened in a single configuration in the absence of decreasing marginal costs can be used to model facility location with hard capacities [11] by setting the cost of opening a facility of capacity u to be infinite for any set of more than u commodities. For this problem, the only known (constant-factor) approximation algorithm uses local search [11], and no IP formulation with a bounded integrality gap is known.

2.2.8 Extension to t-MCFL The main modification is that the IP formulation now only involves upto 2^t configutations per facility i which is the power set of the set of commodities $T_i \subseteq S$ that the facility can produce. this implies that the transformed fractional solution can be interpreted as a fractional t-set cover problem. The rest of the modifications are straightforward and lead to the following result.

THEOREM 2.4. *The above algorithm can be adapted to output a solution with total facility cost no more than $\frac{H_t}{\alpha}$ times the optimal facility cost and service cost no more than $\frac{3}{1-\alpha}$ times the optimal service cost for any $0 < \alpha < 1$.*

2.3 Linear cost functions The running time of the above algorithm is $O(2^k \text{poly}(n, m))$ because if we have no other assumptions on facility costs, we are forced to tabulate the cost of each facility in every configuration. However, if the facility cost functions have compact representations, it is possible to improve the running time. In this section, we study a class of facility cost functions which are compact and yet useful, and show that our algorithm runs in time $O(\text{poly}(n, m, k))$ for these instances.

Suppose each facility i has an opening cost (fixed cost) f_i^0, and for each commodity s, an incremental cost f_i^s, such that the cost of opening facility i in configuration σ is $f_i(\sigma) = f_i^0 + \sum_{s\in\sigma} f_i^s$. We call this class of facility cost functions *linear*. Linear cost functions also satisfy decreasing marginal costs, while in many real situations linear cost functions are a close approximation to general cost functions. Since a linear cost function can be compactly represented by $k+1$ numbers, one would expect an algorithm with running time polynomial in k, n, m instead of $2^k, n, m$. We show that a slight modification of our algorithm satisfies this.

2.3.1 IP formulation We use a slightly different IP formulation for linear MCFL. Variable y_i^0 indicates whether or not facility i is opened, and variable y_i^s indicates whether or not facility i is serving commodity s. Recall that d_j is the commodity demanded by client j.

$$\begin{array}{rcll}
\min \sum_{i,j} c_{ij} x_{ij} & + & \sum_{i\in F} \sum_{s=0}^{k} f_i^s y_i^s & (IP_{L-MCFL}) \\
\sum_i x_{ij} & \geq & 1 & \forall j \in D \\
x_{ij} & \leq & y_i^{d_j} & \forall j \in D,\ \forall i \in F \\
y_i^s & \leq & y_i^0 & \forall i \in D,\ \forall s \in S \\
x, y & & & \text{non-negative integers}
\end{array}$$

Let LP_{L-MCFL} denote the linear relaxation of IP_{L-MCFL}. The size of LP_{L-MCFL} is $\text{poly}(n, m, k)$, and hence we can solve it optimally in time $O(\text{poly}(n, m, k))$. Let (x, y) denote an optimal solution of LP_{L-MCFL}; as before, we use its value as our lower bound.

2.3.2 Filtering and selection of representatives The first two steps of our algorithm are identical to the algorithm for general MCFL. We fix our constant α and obtain an α-close solution $(\overline{x}, \overline{y})$, and we select a set of representatives R such that no two clients with the same commodity are served fractionally by the same facility.

2.3.3 Interpretation as a fractional k-set cover solution Our k-set cover instance is defined as in IP_{k-SC}, with an element for every client $j \in R$ and a set for every facility-configuration pair. Client j is included in a facility-configuration pair (i, σ) only if $\overline{x}_{ij} > 0$ and $d_j \in \sigma$. This by itself is an instance which is exponential in the number of commodities. However, the following lemma casts our fractional solution $\overline{y}$ as a fractional solution to the k-set cover instance so that only polynomially many facility-configuration pairs have non-zero fractional variables. This allows us to restrict our attention to a k-set cover instance (and fractional solution) of size $O(\text{poly}(n, m, k))$.

LEMMA 2.4. *There is a fractional solution z to the k-set cover instance IP_{k-SC} such that (i) $\sum_{(i,\sigma):\overline{x}^{\sigma}_{ij}>0} z^{\sigma}_i \geq 1$ for all $j \in R$; (ii) for every facility i, there are at most k configurations for which $z^{\sigma}_i > 0$; (iii) the total cost of the fractional solution is no more than $\sum_{i,\sigma} f_i(\sigma)\overline{y}^{\sigma}_i$.*

Proof. Consider facility i, and order the commodities so that $\overline{y}^0_i \geq \overline{y}^1_i \geq \overline{y}^2_i \geq \ldots \geq \overline{y}^k_i$. Let $[s] = \{1, 2, \ldots, s\}$. We open facility i in configuration $[s]$ to extent $z^{[s]}_i = \overline{y}^s_i - \overline{y}^{s+1}_i$ for $s = 1, 2, \ldots k-1$, and $z^{[k]}_i = \overline{y}^k_i$. The fractional solution z can be verified to satisfy the properties stated in the lemma.

The rest of the algorithm proceeds as before, by first rounding the fractional k-set cover solution and then assigning clients in $D \setminus R$ to the facilities which serve their representatives. This yields the following analog of Theorem 2.3.

THEOREM 2.5. *There is an algorithm running in time $O(poly(n, m, k))$ which produces a solution for linear MCFL with facility cost no more than $\frac{H_k}{\alpha}$ times the optimal facility cost and service cost no more than $\frac{3}{1-\alpha}$ times the optimal service cost.*

3 Extensions

3.1 Priority facility location

In this section, we provide a constant-factor approximation algorithm for PFL. Recall that the facility cost functions have compact representations, and hence our algorithm also has running time $O(\text{poly}(n, m, k))$. Our algorithm extends the filter-and-round ideas of the MCFL algorithms, by using it twice, with the second application aimed at satisfying the priority constraints.

3.1.1 IP formulation

We begin with the following IP formulation for PFL. Variable y^s_i indicates whether or not facility i is operating at level-of-service s, and variable x^s_{ij} is 1 if client j is served by facility i operating at level s.

$$\min \sum_{i,j} \sum_{s=1}^{k} c_{ij} x^s_{ij} + \sum_{i \in F} \sum_{s=1}^{k} f_i(s) y^s_i \quad (IP_{PFL})$$

$$\sum_{s=d_j}^{k} \sum_{i \in F} x^s_{ij} \geq 1 \qquad \forall j \in D$$

$$x^s_{ij} \leq y^s_i \qquad \forall j \in D, \ \forall i \in F, \ s = 1, 2, \ldots, k$$

$$x, y \quad \text{non-negative integers}$$

Once again, IP_{PFL} has size polynomial in n, m, k. Let (x, y) denote the optimal solution to the linear relaxation LP_{PFL} of IP_{PFL}.

3.1.2 Two-stage filtering

The first step of our algorithm is unchanged; we obtain an α-close solution $(\overline{x}, \overline{y})$ satisfying the properties mentioned in Theorem 2.2. Our next step is the crucial second-stage filtering which enables us to take care of the priorities. The idea is to obtain a fractional solution where each client is served by facilities which cost no more than $\frac{1}{\beta}$ times the average cost of the facilities serving it, for some fixed constant β.

LEMMA 3.1. *The fractional solution $(\overline{x}, \overline{y})$ can be transformed in time polynomial in (n, m, k) to a fractional solution $(\tilde{x}, \tilde{y})$ such that (i) $(\tilde{x}, \tilde{y})$ is feasible for LP_{PFL}; (ii) $\sum_{s=d_j}^{k} \sum_{i \in F} \tilde{x}^s_{ij} \geq 1$ for all clients $j \in D$; (iii) $\tilde{x}^s_{ij} > 0 \Rightarrow f_i(s) \leq \frac{1}{\beta} \sum_{s=d_j}^{k} \sum_{i \in F} \overline{x}^s_{ij} f_i(s)$ for all $j \in D$; (iv) $\tilde{y}^s_i \leq \min\{1, \frac{1}{1-\beta} \overline{y}^s_i\}$ for all facilities $i \in F$ and all $s = 1, 2, \ldots, k$.*

Proof. For each client j, we do the following. Order the variables x^s_{ij} in non-decreasing order of $f_i(s)$, and define the β-cost to be the smallest cost f^*_j such that $\sum_{(i,s):f_i(s) \leq f^*_j} \overline{x}^s_{ij} \geq \beta$. Define $\tilde{x}^s_{ij} = \min\{1, \frac{1}{1-\beta} \overline{x}^s_{ij}\}$ if $f_i(s) \leq f^*_j$, and 0 otherwise. Finally, for every facility-priority pair (i, s), define $\tilde{y}^s_i = \min\{1, \frac{1}{1-\beta} \overline{y}^s_i\}$. It can easily be verified that $(\tilde{x}, \tilde{y})$ satisfies the properties stated.

The traditional (first) filtering according to service costs allows each representative client to be served by any facility with $\overline{x}_{ij} > 0$ and other non-representative clients by a nearby representative's facility. The new (second) filtering according to facility costs that we introduce allows each representative client to open the *highest priority facility* with $x_{ij} > 0$ thus serving all other non representative clients depending on it at an adequate priority level.

3.1.3 Rounding

We now round $(\tilde{x}, \tilde{y})$ into an integer solution $(\hat{x}, \hat{y})$ which is feasible for IP_{PFL} and doesn't cost much more than $(\tilde{x}, \tilde{y})$.

We process clients $j \in D$ in increasing order of $c^*_j = \sum_{i,s} c_{ij} x^s_{ij}$. Initially, all clients are "unserved". Pick the unserved client with least c^*_j, and let F_j be the set of all facility-priority pairs (i, s) with $\tilde{x}^s_{ij} > 0$. Let $(i', s') \in F_j$ maximize s' among all facility-priority pairs in F_j. Open facility i' in priority s' (ie, set $\hat{y}^{s'}_{i'} = 1$), and close all others in F_j (ie, set $\hat{y}^s_i = 0$). All clients j' such that $\tilde{x}^s_{ij'} > 0$ for some $(i, s) \in F_j$ are assigned to (i', s') (ie, we set $\hat{x}^{s'}_{i'j'} = 1$ and $\hat{x}^s_{ij'} = 0$ for all other (i, s)).

This is feasible, since $\tilde{x}^s_{ij'} > 0 \Rightarrow d_i \le s \Rightarrow d_i \le s'$, because s' was the maximum priority in F_j. We mark all such clients "served", as well as client j. We repeat this with unserved clients till all clients are served.

The following theorem can be proved along the same lines as Theorem 2.3.

THEOREM 3.1. *The procedure outlined above runs in time $O(poly(n, m, k))$ and produces an integer solution $(\hat{x}, \hat{y})$ such that (i) $(\hat{x}, \hat{y})$ is feasible for IP_{PFL}; (ii) $\sum_{i,s} f_i(s)\hat{y}^s_i \le \frac{1}{\alpha\beta(1-\beta)} \sum_{i,s} f_i(s)y^s_i$; (iii) every client j is assigned to a facility i such that $c_{ij} \le \frac{3}{1-\alpha} c^*_j$.*

Choosing $\alpha = \frac{4}{7}$ and $\beta = \frac{1}{2}$ yields a 7-approximation. This can probably be improved by techniques such as scaling the facility and service costs differently.

3.2 Unique facilities for commodities In this section, we extend the algorithm in Section 2 to the variant where a single facility must serve each commodity.

3.2.1 Selection of representatives Let D_s be the set of clients which have demand for commodity s. Define the star cost of a node $j \in D_s$ to be $\hat{c}_j = \sum_{j' \in D_s} c_{jj'}$. Let j_s be the node with minimum star cost among all nodes in D_s, and select j_s to be the representative for commodity s. Let V_S be the set of all representatives.

3.2.2 Selection of facilities Consider the t-MCFL instance obtained by relaxing the requirement that a single facility serve each commodity. Compute an approximate solution to this (using method APX, say), and let $\phi(j)$ be the facility which serves client j in this solution. Assign all clients in D_s to be served by $\phi(j_s)$. Let $\phi^*(j)$ denote the facility which serves client j in some fixed optimal solution.

LEMMA 3.2. *$\sum_{j \in D} c_{j,\phi(j)}$ is no more than $\rho_{SC} \cdot OPT_{SC}$ plus $4 \cdot \sum_{j \in D} c_{j,\phi^*(j)}$, where OPT_{SC} is the service cost of an optimal solution, and ρ_{SC} is the approximation ratio of APX for the service cost.*

Proof. Consider a commodity s, and let $\phi^*(s)$ be the facility which serves it in the fixed optimal solution. Let $j^*_s \in D_s$ be the client which minimizes $c_{j,\phi^*(s)}$ among all clients in D_s. Summing the inequality $c_{j^*_s,\phi^*(s)} \le c_{j,\phi^*(s)}$ over all $j \in S$ implies that $\hat{c}_{j^*_s} \le 2\sum_{j \in D_s} c_{j,\phi^*(s)}$. Since our representative j_s minimized $\hat{c}_j$, we have $\hat{c}_{j_s} \le \sum_{j \in D_s} 2c_{j,\phi^*(s)}$.

Next, consider a client $j \in D_s$, and let $\phi'(j)$ be the facility which serves it in the relaxation where we ignored the constraint that each commodity be served by a single facility. By triangle inequality, we must have $c_{j_s,\phi'(j)} \le c_{j,j_s} + c_{j,\phi'(j)}$, implying that $c_{j,\phi(j)} \le 2c_{j,j_s} + c_{j,\phi'(j)}$. Observing that $\sum_{j \in D_s} c_{j,j_s} \le \hat{c}_{j_s}$ we can sum the inequality bounding $c_{j,\phi(j)}$ over all clients j to get the bound claimed in the lemma.

THEOREM 3.2. *There is a polynomial time algorithm for the restricted version of t-MCFL with approximation ratio $O(\log t)$.*

Proof. The service cost of our solution is bounded by a constant by Lemma 3.2 and Theorem 2.4, while the facility cost of our solution is bounded within $O(\log t)$ of the optimal since we only open facilities opened by the relaxation where we ignore the constraint that every commodity be served at a single facility. This yields the theorem.

When the facility cost is the same for all non-null configurations of a facility ($f_i(\sigma) = f_i \ \forall \sigma \neq \emptyset$ for all facilities i), we can apply any existing constant-factor approximation for UFL to get the following result.

THEOREM 3.3. *There is a constant-factor approximation algorithm for the UFL problem when clients are partitioned into classes (commodities) and all clients in the same class are required to be served by the same facility.*

4 Conclusion

We have introduced and provided approximation algorithms for various models of multicommodity facility location. All of these algorithms match the best known ratios for the appopriate set covering variants they generalize. Since the number of commodities t involved in a facility is likely to be much smaller than the number of clients n, the $O(\log t)$-approximation ratio represents a significant improvement in the state of the art from the naive $O(\min(k, \log n)$ result.

An intriguing questions is whether the running time dependence on 2^k can be removed for general MCFL. In particular, given a cost oracle which could answer questions of the form "Is the cost of serving configuration σ no more than x?" can we use such an oracle to obtain an approximation algorithm with running time polynomial in n, m, k with only polynomially many calls to such an oracle? There is hope for such a solution under decreasing marginal facility costs.

References

[1] S. Arora and M. Sudan. Improved low degree testing and its applications. In *Proceedings of the 29^{th} Annual ACM Symposium on Theory of Computing*, 485-495, 1997.

[2] V. Arya, N. Garg, R. Khandekar, V. Pandit, A. Meyerson and K. Munagala. Local search heuristic for k-median and facility location problems. In *Proceedings of the* 33^{rd} *Annual ACM Symposium on Theory of Computing*, 21-29, 2001.

[3] M.L. Balinski. On finding integer solutions to linear programs. In *Proc. IBM Scientific Computing Symposium on Combinatorial Problems*, 225-248, 1966.

[4] M. Charikar, J. Naor and B. Scheiber. Resource optimization in QoS multicast routing of real-time multimedia. In *Proceedings of IEEE INFOCOM*, 1518-1527, 2000.

[5] S. Guha and S. Khuller. Greedy strikes back: Improved facility location algorithms. In *Proceedings of the* 9^{th} *ACM-SIAM Symposium on Discrete Algorithms*, 649-657, 1998.

[6] M. Halldórsson. Approximating k-set cover and complementary graph coloring. In *Proceedings of the* 5^{th} *International Conference on Integer Programming and Combinatorial Optimization*, 118-131, 1996.

[7] A. Kumar, A. Gupta and T. Roughgarden. A constant-factor approximation algorithm for the multicommodity rent-or-buy problem. In *Proceedings of the* 43^{rd} *Annual IEEE Symposium on Foundations of Computer Science*, 333-342, 2002.

[8] J-H. Lin and J. Vitter. ϵ-approximations with minimum packing constraint violation. In *Proceedings of the* 24^{th} *Annual ACM Symposium on Theory of Computing*, 771-782, 1992.

[9] M. Mahdian, Y. Ye and J. Zhang. A 1.52 approximation algorithm for the uncapacitated facility location problem. In *Approximation Algorithms for Combinatorial Optimization*, 229-242, 2002.

[10] A. Meyerson. Online facility location. In *Proceedings of the* 42^{nd} *Annual IEEE Symposium on Foundations of Computer Science*, 426-431, 2001.

[11] M. Pal, E. Tardos and T. Wexler. Facility location with nonuniform hard capacities. In *Proceedings of the* 42^{nd} *Annual IEEE Symposium on Foundations of Computer Science*, 329-338, 2001.

[12] R. Ravi and A. Sinha. Integrated logistics: Approximation algorithms combining facility location and network design. In *Proceedings of the* 9^{th} *International Conference on Integer Programming and Combinatorial Optimization*, 212-229, 2002.

[13] D. Shmoys, E. Tardos and K. Aardal. Approximation algorithms for the facility location problem. In *Proceedings of the* 29^{th} *Annual ACM Symposium on Theory of Computing*, 265-274, 1997.

[14] V. Vazirani. *Approximation Algorithms*, Springer-Verlag, Berlin, 2001.