The Matroid Median Problem

Ravishankar Krishnaswamy* Amit Kumar† Viswanath Nagarajan‡ Yogish Sabharwal§
Barna Saha¶

Abstract

In the classical k-median problem, we are given a metric space and would like to open k centers so as to minimize the sum (over all the vertices) of the distance of each vertex to its nearest open center. In this paper, we consider the following generalization of the problem: instead of opening at most k centers, what if each center belongs to one of T different types, and we are allowed to open at most k_i centers of type i (for each i=$1, 2, \ldots, T$). The case $T = 1$ is the classical k-median, and the case of $T = 2$ is the *red-blue median problem* for which Hajiaghayi et al. [ESA 2010] recently gave a constant-factor approximation algorithm.

Even more generally, what if the set of open centers had to form an independent set from a matroid? In this paper, we give a constant factor approximation algorithm for such *matroid median* problems. Our algorithm is based on rounding a natural LP relaxation in two stages: in the first step, we sparsify the structure of the fractional solution while increasing the objective function value by only a constant factor. This enables us to write another LP in the second phase, for which the sparsified LP solution is feasible. We then show that this second phase LP is in fact integral; the integrality proof is based on a connection to matroid intersection.

We also consider the penalty version (alternately, the so-called prize collecting version) of the matroid median problem and obtain a constant factor approximation algorithm for it. Finally, we look at the Knapsack Median problem (in which the facilities have costs and the set of open facilities need to fit into a Knapsack) and get a bicriteria approximation algorithm which violates the Knapsack bound by a small additive amount.

*Computer Science Department, Carnegie Mellon University, Pittsburgh, PA 15213, USA. Supported in part by NSF awards CCF-0448095 and CCF-0729022 and an IBM Ph.D Fellowship.

†Computer Science Department, IIT Delhi.

‡IBM T. J. Watson Research Center, Yorktown Heights, NY 10598.

§IBM India Research Laboratory.

¶Computer Science Department, University of Maryland College Park, MD 20742, USA. Supported in part by NSF Award CCF-0728839, ,NSF Award CCF-0937865 and a Google Research Award.

1 Introduction

The k-median problem is an extensively studied location problem. Given an n-vertex metric space (V, d) and a bound k, the goal is to locate/open k centers $C \subseteq V$ so as to minimize the sum of distances of each vertex to its nearest open center. (The distance of a vertex to its closest open center is called its connection cost.) The first constant-factor approximation algorithm for k-median on general metrics was by Charikar et al. [7]. The approximation ratio was later improved in a sequence of papers [16, 15, 6] to the currently best-known guarantee of $3 + \epsilon$ (for any constant $\epsilon > 0$) due to Arya et al. [3]. A number of techniques have been successfully applied to this problem, such as LP rounding, primal-dual and local search algorithms.

Motivated by applications in Content Distribution Networks, Hajiaghayi et al. [12] introduced a generalization of k-median where there are *two types* of vertices (red and blue), and the goal is to locate at most k_r red centers and k_b blue centers so as to minimize the sum of connection costs. For this *red-blue median* problem, [12] gave a constant factor approximation algorithm. In this paper, we consider a substantially more general setting where there are an arbitrary number T of vertex-types with bounds $\{k_i\}_{i=1}^{T}$, and the goal is to locate at most k_i centers of each type-i so as to minimize the sum of connection costs. These vertex-types denote different types of servers in the Content Distribution Networks applications; the result in [12] only holds for $T = 2$.

In fact, we study an even more general problem where the set of open centers have to form an independent set in a given matroid, with the objective of minimizing sum of connection costs. This formulation captures several intricate constraints on the open centers, and contains as special cases: the classic k-median (uniform matroid of rank k), and the CDN applications above (partition matroid with T parts). Our main result is a constant-factor approximation algorithm for this *Matroid Median* problem.

1.1 Our Results and Techniques

In this paper we introduce the Matroid Median problem, which is a natural generalization of k-median, and obtain a 16-approximation algorithm for it. Thus it also gives the first constant approximation for the k-median problem with multiple (more than two) vertex-types, which was introduced in [12].

For the standard k-median problem (and also red-blue median), it is easy to obtain an $O(\log n)$-approximation algorithm using probabilistic tree-embeddings [9], and exactly solving the problem on a tree (via a dynamic program). However, even this type of guarantee is not obvious for the Matroid Median problem, since the problem on a tree-metric does not look particularly easier.

Our algorithm is based on the natural LP-relaxation and is surprisingly simple. Essentially, the main insight is in establishing a connection to matroid intersection. The algorithm computes an optimal LP solution and rounds it in two phases, the key points of which are described below:

- The first phase sparsifies the LP solution while increasing the objective value by a constant factor. This is somewhat similar to the LP-rounding algorithm for k-median in Charikar et al. [7]. However we cannot consolidate fractionally open centers as in [7]; this is because the open centers must additionally satisfy the matroid rank constraints. In spite of this, we show that the vertices and the centers can be clustered into disjoint 'star-like' structures.

- This structure ensured by the first phase of rounding allows us to write (in the second phase) another linear program for which the sparsified LP solution is feasible, and has objective value at most $O(1)$ times the original LP optimum. Then we show that the second phase LP is in fact integral, via a relation to the matroid-intersection polytope. Finally we re-solve the second phase LP to obtain an extreme point solution, which is guaranteed to be integral. This corresponds to a feasible solution to Matroid Median of objective value $O(1)$ times the LP optimum.

We next consider the *Penalty Matroid Median* (a.k.a. prize-collecting matroid median problem), where a vertex could either connect to a center incurring the connection cost, or choose to pay a penalty in the objective function. The prize-collecting version of several well-known optimization problems like TSP, Steiner Tree etc., including k-median and red-blue median have been studied in prior work (See [1, 12] and the references therein). Extending the idea of the Matroid Median algorithm, we also obtain an $O(1)$ approximation algorithm for the Penalty version of the problem.

Finally, we look at the *Knapsack Median* problem (a.k.a. weighted W-median [12]), where the centers have weights and the open centers must satisfy a knapsack constraint; the objective is, like before, to minimize the total connection cost of all the vertices. For this problem we obtain a 16-approximation algorithm that violates the knapsack constraint by an *additive* f_{max} term (where f_{max} is the maximum opening cost of any center). This algorithm is again based on the natural LP relaxation, and follows the same approach as for Matroid Median. However, the second phase LP here is not integral (it contains the knapsack problem as a special case). Instead we obtain the claimed bicriteria approximation by using the iterative rounding framework [10, 14, 20, 17]. It is easy to see that our LP-relaxation for the Knapsack Median problem has unbounded integrality gap, if we do not allow any violation in the knapsack constraint (see eg. [6]). Moreover, we show that the integrality gap remains unbounded even after the addition of *knapsack-cover inequalities* [5] to the basic LP relaxation. We leave open the question of obtaining an $O(1)$-approximation for Knapsack Median without violating the knapsack constraint.

1.2 Related Work The first approximation algorithm for the metric k-median problem was due to Lin and Vitter [18] who gave an algorithm that for any $\epsilon > 0$, produced a solution of objective at most $2(1 + \frac{1}{\epsilon})$ while opening at most $(1 + \epsilon)k$ centers; this was based on the filtering technique for rounding the natural LP relaxation. The first approximation algorithm that opened only k centers was due to Bartal [4], via randomized tree embedding (mentioned earlier). Charikar et al. [7] obtained the first $O(1)$-approximation algorithm for k-median, by rounding the LP relaxation; they obtained an approximation ratio of $6\frac{2}{3}$. The approximation ratio was improved to 6 by Jain and Vazirani [16], using the primal dual technique. Charikar and Guha [6] further improved the primal-dual approach to obtain a 4-approximation. Later Arya et al. [3] analyzed a natural local search algorithm that exchanges up to p centers in each local move, and proved a $3 + \frac{2}{p}$ approximation ratio (for any constant $p \geq 1$). Recently, Gupta and Tangwongsan [11] gave a considerably simplified proof of the Arya et al. [3] result. It is known that the k-median problem on general metrics is hard to approximate to a factor better than $1 + \frac{2}{e}$. On Euclidean metrics, the k-median problem has been shown to admit a PTAS by Arora et al. [2].

Very recently, Hajiaghayi et al. [12] introduced the red-blue median problem — where the vertices are divided into two categories and there are different bounds on the number of open centers of each type — and obtained a constant factor approximation algorithm. Their algorithm uses a local search using single-swaps for each vertex type. The motivation in [12] came from locating servers in Content Distribution Networks, where there are T server-types and strict bounds on the number of servers of each type. The red-blue median problem captured the case $T = 2$. It is unclear whether their approach can be extended to multiple server types, since the local search with single swap for each server-type has neighborhood size $n^{\Omega(T)}$. Furthermore, even a $(T - 1)$-exchange local search has large locality-gap– see Appendix A. Hence it is not clear how to apply local search to MatroidMedian, even in the case of a partition matroid. [12] also discusses the difficulty in applying the Lagrangian relaxation approach (see [16]) to the red-blue median problem; this is further compounded in the MatroidMedian problem since there are exponentially many constraints on the centers.

The most relevant paper to ours with regard to the rounding technique is Charikar et al. [7]: our algorithm builds on many ideas used in their work to obtain our approximation algorithm.

For the penalty k-median problem, the best known bound is a 3-approximation due to Hajiaghayi et al. [12] that improves upon a previous 4-approximation due to Charikar and Guha [6]. Hajiaghayi et al. also consider the penalty version of the red-blue median problem and give a constant factor approximation algorithm.

The knapsack median problem admits a bicriteria approximation ratio via the filtering technique [18]. The currently best known tradeoff [6] implies for any $\epsilon > 0$, a $\left(1+\frac{2}{\epsilon}\right)$-approximation in the connection costs while violating the knapsack constraint by a *multiplicative* $(1+\epsilon)$ factor. Charikar and Guha [6] also shows that for each $\epsilon > 0$, it is not possible to obtain a trade-off better than $\left(1+\frac{1}{\epsilon}, 1+\epsilon\right)$ relative to the natural LP. On the other hand, our result implies a $(16, 1+\epsilon)$-tradeoff in $n^{O(1/\epsilon)}$ time for each $\epsilon > 0$; this algorithm uses enumeration combined with the natural LP-relaxation. As mentioned in [12], an $O(\log n)$-approximation is achievable for knapsack median (without violation of the knapsack constraint) via a reduction to tree-metrics, since the problem on trees admits a PTAS.

2 Preliminaries

The input to the MatroidMedian problem consists of a finite set of vertices V and a distance function $d : V \times V \to \mathbb{R}_{\geq 0}$ which is symmetric and satisfies the triangle inequality, i.e. $d(u,v) + d(v,w) \geq d(u,w)$ for all $u,v,w \in V$. Such a tuple (V,d) is called a finite metric space. We are also given a matroid $\mathcal{M}$, with ground set V and set of independent sets $\mathcal{I}(\mathcal{M}) \subseteq 2^V$. The goal is to open an independent set $S \in \mathcal{I}(\mathcal{M})$ of centers such that the sum $\sum_{u\in V} d(u,S)$ is minimized; here $d(u,S) = \min_{v\in S} d(u,v)$ is the connection cost of vertex u. We assume some familiarity with matroids, for more details see eg. [19].

In the penalty version, additionally a penalty function $p : V \to \mathbb{R}_{\geq 0}$ is provided and the objective is now modified to minimize $\sum_{u\in V} d(u,S)(1-h(u)) + p(u)h(u)$. Here $h : V \to \{0,1\}$ is an indicator function that is 1 if the corresponding vertex is not assigned to a center and therefore pays a penalty and 0 otherwise.

The KnapsackMedian problem (aka weighted W-median [12]) is similarly defined. We are given a finite metric space (V,d), non-negative weights $\{f_i\}_{i\in V}$ (representing facility costs) and a bound F. The goal is to open centers $S \subseteq V$ such that $\sum_{j\in S} f_j \leq F$ and the objective $\sum_{u\in V} d(u,S)$ is minimized (Section 5).

2.1 An LP Relaxation for MatroidMedian

In the following linear program, y_v is the indicator variable for whether vertex $v \in V$ is opened as a center, and x_{uv} is the indicator variable for whether vertex u is served by center v. Then, the following LP is a valid relaxation for the MatroidMedian problem.

(LP$_1$)

$$\text{minimize} \sum_{u\in V}\sum_{v\in V} d(u,v)x_{uv}$$

$$(2.1)\qquad \text{subject to} \sum_{v\in V} x_{uv} = 1 \qquad \forall\, u \in V$$

$$(2.2)\qquad x_{uv} \leq y_v \qquad \forall\, u\in V, v \in V$$

$$(2.3)\qquad \sum_{v\in S} y_v \leq r_{\mathcal{M}}(S) \qquad \forall\, S \subseteq V$$

$$(2.4)\qquad x_{uv}, y_v \geq 0 \qquad \forall\, u,v \in V$$

If x_{uv} and y_v are restricted to only take values 0 or 1, then this is easily seen to be an exact formulation for MatroidMedian. The first constraint models the requirement that each vertex u must be connected to some center v, and the second one requires that it can do so only if the center v is opened, i.e. $x_{uv} = 1$ only if y_v is also set to 1. The constraints (2.3) are the matroid rank-constraint on the centers: they model the fact that the open centers form an independent set with respect to the matroid $\mathcal{M}$. Here $r_{\mathcal{M}} : 2^V \to \mathbb{Z}_{\geq 0}$ is the *rank-function* of the matroid, which is monotone and submodular. The objective function exactly measures the sum of the connection costs of each vertex. (It is clear that given integrally open centers $y \in \{0,1\}^V$, each vertex $u \in V$ sets $x_{uv} = 1$ for its closest center v with $y_v = 1$.) Let Opt denote an optimal solution of the given MatroidMedian instance, and let LPOpt denote the LP optimum value. From the above discussion, we have that,

LEMMA 2.1. *The LP cost* LPOpt *is at most the cost of an optimal solution* Opt.

2.2 Solving the LP: The Separation Oracle

Even though the LP relaxation has an exponential number of constraints, it can be solved in polynomial time (using the Ellipsoid method) assuming we can, in polynomial time, verify if a candidate solution (x,y) satisfies all the constraints. Indeed, consider any fractional solution (x,y). Constraints (2.1), and (2.2) can easily be verified in $O(n^2)$ time, one by one.

Constraint (2.3) corresponds to checking if the fractional solution $\{y_v : v \in V\}$ lies in the matroid polytope for $\mathcal{M}$. Checking (2.3) is equivalent to seeing whether:

$$\min_{S\subseteq V}\left(r_{\mathcal{M}}(S) - \sum_{v\in S} y_v\right) \geq 0.$$

Since the rank-function $r_{\mathcal{M}}$ is submodular, so is the function $f(S) := r_{\mathcal{M}}(S) - \sum_{v\in S} y_v$. So the above condition (and hence (2.3)) can be checked using submodular function minimization, eg. [19, 13]. There are also more efficient methods for separating over the matroid polytope – refer to [19, 8] for more details on efficiently testing membership in matroid polyhedra. Thus we can obtain an optimal LP solution in polynomial time.

3 The Rounding Algorithm for MatroidMedian

Let (x^*, y^*) denote the optimal LP solution. Our rounding algorithm consists of two stages. In the first stage, we only alter the x^*_{uv} variables such that the modified solution, while still being feasible to the LP, is also very sparse in its structure. In the second stage, we write another LP which exploits the sparse structure, for which the modified fractional solution is feasible, and the objective function has not increased by more than a constant factor. We then proceed to show that the new LP in fact corresponds to an integral polytope. Thus we can obtain an integral solution where the open centers form an independent set of $\mathcal{M}$, and the cost is $O(1)\mathsf{LPOpt}$.

3.1 Stage I: Sparsifying the LP Solution In the first stage, we follow the outline of the algorithm of Charikar et al. [7], but we can not directly employ their procedure because we can't alter/consolidate the y^*_v variables in an arbitrary fashion (since they need to satisfy the matroid polytope constraints). Specifically, step (i) below is identical to the first step (consolidating locations) in [7]. The subsequent steps in [7] do not apply since they consolidate centers; however using some ideas from [7] and with some additional work, we obtain the desired sparsification in steps (ii)-(iii) *without altering the y^*-variables*.

Step (i): Consolidating Clients. We begin with some notation, which will be useful throughout the paper. For each vertex u, let $\mathsf{LP}_u = \sum_{v \in V} d(u,v) x^*_{uv}$ denote the contribution to the objective function LPOpt of vertex u. Also, let $\mathcal{B}(u,r) = \{v \in V \mid d(u,v) \le r\}$ denote the ball of radius r centered at vertex u. For any vertex u, we say that $\mathcal{B}(u, 2\mathsf{LP}_u)$ is the *local ball* centered at u.

Initialize $w_u \leftarrow 1$ for all vertices. Order the vertices according to non-decreasing LP_u values, and let the ordering be $u_1, u_2, \ldots, u_n$. Now consider the vertices in the order $u_1, u_2, \ldots, u_n$. For vertex u_i, if there exists another vertex u_j with $j < i$ such that $d(u_i, u_j) \le 4\mathsf{LP}_{u_i}$, then set $w_{u_j} \leftarrow w_{u_j} + 1$, and $w_{u_i} \leftarrow 0$. Essentially we can think of moving u_i to u_j for the rest of the algorithm (which is why we are increasing the weight of u_j and setting the weight of u_i to be zero).

After the above process, let V' denote the set of locations with positive weight, i.e. $V' = \{v \mid w_v > 0\}$. For the rest of the paper, we will refer to vertices in V' as *clients*. By the way we defined this set, it is clear that the following two observations holds.

OBSERVATION 3.1. *For $u, v \in V'$, we have $d(u,v) > 4\max(\mathsf{LP}_u, \mathsf{LP}_v)$.*

This is true because, otherwise, if (without loss of generality) $\mathsf{LP}_v \ge \mathsf{LP}_u$ and $d(u,v) \le 4\mathsf{LP}_v$, then we would have moved v to u when we were considering v.

OBSERVATION 3.2.

$$\sum_{u \in V'} w_u \sum_{v \in V} d(u,v) x^*_{uv} \le \sum_{u \in V} \sum_{v \in V} d(u,v) x^*_{uv}$$

This is because, when we move vertex u_i to u_j, we replace the term corresponding to LP_{u_i} (in the LHS above) with an additional copy of that corresponding to LP_{u_j}, and we know by the vertex ordering that $\mathsf{LP}_{u_i} \ge \mathsf{LP}_{u_j}$.

Also, the following lemma is a direct consequence of Markov's inequality.

LEMMA 3.1. *For any client $u \in V'$, $\sum_{v \in \mathcal{B}(u, 2\mathsf{LP}_u)} x_{uv} \ge 1/2$. In words, each client is fractionally connected to centers in its local ball to at least an extent of $1/2$.*

Finally, we observe that if we obtain a solution to the new (weighted) instance and incur a cost of C, the cost of the same set of centers with respect to the original instance is then at most $C + 4\mathsf{LPOpt}$ (the additional distance being incurred in moving back each vertex to its original location).

We now assume that we have the weighted instance (with clients V'), and are interested in finding a set $S \subseteq V$ of centers to minimize $\sum_{u \in V'} w_u d(u,S)$. Note that centers may be chosen from the entire vertex-set V, and are not restricted to V'. Consider an LP-solution $\left(x^1, y^*\right)$ to this weighted instance, where $x^1_{uv} = x^*_{uv}$ for all $u \in V'$, $v \in V$. Note that $\left(x^1, y^*\right)$ satisfies constraints (2.1)-(2.2) with u ranging over V', and also constraint (2.3); so it is indeed a feasible fractional solution to the weighted instance. Also, by Observation 3.2, the objective value of $\left(x^1, y^*\right)$ is $\sum_{u \in V'} w_u \sum_{v \in V} d(u,v) x^1_{uv} \le \mathsf{LPOpt}$, i.e. at most the original LP optimum.

After this step, even though we have made sure that the clients are well-separated, a client $u \in V'$ may be fractionally dependent on several partially open centers, as governed by the x_{uv} variables. More specifically, it may be served by centers which are contained in the ball $\mathcal{B}(u, 2\mathsf{LP}_u)$, or by centers which are contained in another ball $\mathcal{B}(u', 2\mathsf{LP}_{u'})$, or some centers which do not lie in any of the balls around the clients. The subsequent steps further simplify the structure of these connections.

Remark: To illustrate the high-level intuition behind our algorithm, suppose it is the case that for all $u \in V'$, client u is completely served by centers inside $\mathcal{B}(u, 2\mathsf{LP}_u)$. Then, we can infer that it is sufficient to open a center inside each of these balls, while respecting the matroid polytope constraints. Since we are guaranteed that for $u, v \in V'$, $\mathcal{B}(u, 2\mathsf{LP}_u) \cap \mathcal{B}(v, 2\mathsf{LP}_v) = \emptyset$ (from Observation 3.1), this problem reduces to that of finding an independent set in the intersection of *matroid* $\mathcal{M}$ and the *partition matroid* defined by the balls $\{\mathcal{B}(u, 2\mathsf{LP}_u) \mid u \in V'\}$! Furthermore, the fractional solution (x^*, y^*) is feasible for the natural LP-relaxation of the matroid intersection problem. Now, because the matroid intersection polytope is integral, we can obtain an integer solution of low cost (relative to LPOpt).

However, the vertices may not in general be fully served by centers inside their corresponding local balls, as mentioned earlier. Nevertheless, we establish some additional structure (in the next three steps) which enables us to reduce to a problem (in Stage II) of intersecting matroid $\mathcal{M}$ with some *laminar constraints* (instead of just partition constraints as in the above example).

Step (ii): Making the objective function uniform & centers private. We now simplify connections that any vertex participates outsides its local ball. We start with the LP-solution $\left(x^1, y^*\right)$ and modify it to another solution $\left(x^2, y^*\right)$. Initially set $x^2 \leftarrow x^1$.

(A). For any client u that depends on a center v which is contained in another client u''s local ball, we change the coefficient of x_{uv} in the objective function from $d(u, v)$ to $d(u, u')$. Because the clients are well-separated, this changes the total cost only by a small factor. Formally,

$$\begin{aligned} d(u,v) &\geq d(u,u') - 2\mathsf{LP}_{u'} && \text{Since } v \in \mathcal{B}(u', 2\mathsf{LP}_{u'}) \\ &\geq d(u,u') - d(u,u')/2 && \text{From Obs 3.1} \\ &\geq (1/2) d(u,u') \end{aligned}$$

Thus we can write:

$$\sum_{u \in V'} w_u \left[\sum_{u' \in V' \setminus u} d(u,u') \sum_{v \in \mathcal{B}(u', 2\mathsf{LP}_{u'})} x^2_{uv} \right] \leq 2 \sum_{u \in V'} w_u \sum_{u' \in V' \setminus u} \sum_{v \in \mathcal{B}(u', 2\mathsf{LP}_{u'})} d(u,v) x^1_{uv} \tag{3.5}$$

(B). We now simplify centers that are not contained in any local ball, and ensure that each such center has only one client dependent on it. Consider any vertex $v \in V$ which does not lie in any local ball, and has at least two clients dependent on it. Let these clients be $u_0, u_1, \dots, u_k$ ordered such that $d(u_0, v) \leq d(u_1, v) \leq \ldots \leq d(u_k, v)$. The following claim will be useful for re-assignment.

CLAIM 3.1. *For all $i \in \{1, \dots, k\}$, $d(u_i, u_0) \leq 2d(u_i, v)$. Furthermore, for any vertex $v' \in \mathcal{B}(u_0, 2LP_{u_0})$, $d(u_i, v') \leq 3d(u_i, v)$.*

Proof. From the way we have ordered the clients, we know that $d(u_i, v) \geq d(u_0, v)$; so $d(u_i, u_0) \leq d(u_i, v) + d(u_0, v) \leq 2d(u_i, v)$ for all $i \in \{1, \cdots, k\}$. Also, if v' is some center in $\mathcal{B}(u_0, 2\mathsf{LP}_{u_0})$, then we have $d(u_i, v') \leq d(u_i, u_0) + 2\mathsf{LP}_{u_0} \leq (3/2) d(u_i, u_0)$, where in the final inequality we have used Observation 3.1. Therefore, we have $d(u_i, v') \leq 3d(u_i, v)$ for any $v' \in \mathcal{B}(u_0, 2\mathsf{LP}_{u_0})$, which proves the claim. ■

Now, for each $1 \leq i \leq k$, we remove the connection (u_i, v) (ie. $x^2_{u_i v} \leftarrow 0$) and arbitrarily increase connections (for a total extent $x^1_{u_i v}$) to edges (u_i, v') for $v' \in \mathcal{B}(u_0, 2\mathsf{LP}_{u_0})$ while maintaining feasibility (i.e $x^2_{u_i v'} \leq y^*_{v'}$). But we are ensured that a feasible re-assignment exists because for every client u_i, the extent to which it is connected outside its ball is at most $1/2$, and we are guaranteed that the total extent to which centers are opened in $\mathcal{B}(u_0, 2\mathsf{LP}_{u_0})$ is at least $1/2$ (Lemma 3.1). Therefore, we can completely remove any connection u_i might have to v and re-assign it to centers in $\mathcal{B}(u_0, 2\mathsf{LP}_{u_0})$ and for each of these reassignments, we use $d(u_i, u_0)$ as the distance coefficient. From Claim 3.1 and observing that the approximation on cost is performed on *disjoint* set of edges in **(A)** and **(B)**, we obtain that:

$$\sum_{u \in V'} w_u \sum_{v \in V} d(u,v) x^2_{uv} \leq 2 \cdot \sum_{u \in V'} w_u \sum_{v \in V} d(u,v) x^1_{uv}. \tag{3.6}$$

After this step, we have that for each center v not contained in any ball around the clients, there is only one client, say u, which depends on it. In this case, we say that v is a *private center* to client u. Let $\mathcal{P}(u)$ denote the set of all vertices that are either contained in $\mathcal{B}(u, 2\mathsf{LP}_u)$, or are private to client u. Notice that $\mathcal{P}(u) \cap \mathcal{P}(u') = \emptyset$ for any two clients $u, u' \in V'$. Also denote $\mathcal{P}^c(u) := V \setminus \mathcal{P}(u)$ for any $u \in V'$.

We further change the LP-solution from $\left(x^2, y^*\right)$ to $\left(x^3, y^*\right)$ as follows. In x^3 we ensure that any client which depends on centers in other clients' local balls, will in fact depend *only* on centers in the local ball of its nearest other client. For any client u, we reassign all connections (in x^2) to $\mathcal{P}^c(u)$ to centers of $\mathcal{B}(u', 2\mathsf{LP}_{u'})$ (in x^3) where u' is the closest other client to u. This is possible because the total reassignment for each client is at most half and every local-ball has at least half unit of centers. Clearly the value of $\left(x^3, y^*\right)$ under the new objective is at most that of $\left(x^2, y^*\right)$, by the way we have altered the objective function.

Now, for each $u \in V'$, if we let $\eta(u) \in V' \setminus \{u\}$ denote the closest other client to u, then u depends only on centers in $\mathcal{P}(u)$ and $\mathcal{B}\left(\eta(u), 2\mathsf{LP}_{\eta(u)}\right)$. Thus, the new objective value of $\left(x^3, y^*\right)$ is exactly:

$$\sum_{u \in V'} w_u \left(\sum_{v \in \mathcal{P}(u)} d(u,v) x^3_{uv} + d(u, \eta(u))(1 - \sum_{v \in \mathcal{P}(u)} x^3_{uv}) \right) \leq 2 \cdot \mathsf{LPOpt} \tag{3.7}$$

Observe that we retained for each $u \in V'$ only the x_{uv}-variables with $v \in \mathcal{P}(u)$; this suffices because all other

x_{uw}-variables (with $w \in \mathcal{P}^c(u)$) pay the same coefficient $d(u, \eta(u))$ in the objective (due to the changes made in **(A)** and **(B)**). Since the cost of the altered solution is at most that of (x^2, y^*), we get the same bound of $2\mathsf{LPOpt}$.

Furthermore, for any client u which depends on a private center $v \in \mathcal{P}(u) \setminus \mathcal{B}(u, 2\mathsf{LP}_u)$, it must be that $d(u, v) \leq d(u, \eta(u))$; otherwise, we can re-assign this (uv) connection to a center $v' \in \mathcal{B}(\eta(u), 2\mathsf{LP}_{\eta(u)})$ and improve in the (altered) objective function; again we use the fact that u might depend on $\mathcal{P}(u) \setminus \mathcal{B}(u, 2\mathsf{LP}_u)$ to total extent at most half and $\mathcal{B}(\eta(u), 2\mathsf{LP}_{\eta(u)})$ has at least half unit of open centers.

To summarize, the above modifications ensure that fractional solution (x^3, y^*) satisfies the following:

(i) For any two clients $u, u' \in V'$, we have $d(u, u') > 4\max(\mathsf{LP}_u, \mathsf{LP}_{u'})$. In words, this means that all clients are well-separated.

(ii) For each center v that does not belong to any ball $\mathcal{B}(u, 2\mathsf{LP}_u)$, we have only one client that depends on it.

(iii) Each client u depends only on centers in its ball, its private centers, and centers in the ball of its nearest client. The extent to which it depends on centers of the latter two kinds is at most $1/2$.

(iv) If client u depends on a private center v, $d(u, v) \leq d(u, u')$ for any other client $u' \in V'$.

(v) The total cost under the modified objective is at most $2 \cdot \mathsf{LPOpt}$.

Step (iii): Building Small Stars. Let us modify mapping η slightly: for each $u \in V'$, if it only depends (under LP solution x^3) on centers in $\mathcal{P}(u)$ (ie. centers in its local ball or its private centers) then reset $\eta(u) \leftarrow u$. Consider a directed dependency graph on just the clients V', having arc-set $\{(u, \eta(u)) | u \in V', \eta(u) \neq u\}$. Each component will almost be a tree, except for the possible existence of one 2-cycle[1] (see Figure 3.1). We will call such 2-cycles *pseudo-roots*. If there is a vertex with no out-arc, that is also called a pseudo-root. Observe that every pseudo-root contains at least a unit of open centers.

The procedure we describe here is similar to the reduction to "3-level trees" in [7]. We break the trees up into a collection of stars, by traversing the trees in a bottom-up fashion, going from the leaves to the root. For any arc (u, u'), we say that u' is the *parent* of u, and u is a *child* of u'. Any client $u \in V'$ with no in-arc is called a *leaf*. Consider any non-leaf vertex u which is not part of a pseudo-root, such that all its children are leaves. Let u^{out} denote the parent of u.

1. Suppose there exists a child u_0 of u such that $d(u_0, u) \leq 2d(u, u^{\mathsf{out}})$, then we make the following modification: let u_1 denote the child of u that is closest to u; we replace the directed arc (u, u^{out}) with (u, u_1), and make the collection $\{u, u_1\}$ (which is now a 2-cycle), a pseudo-root. Observe that $d(u_0, u) \geq d(u, u^{\mathsf{out}})$ because u chose to direct its arc towards u^{out} instead of u_0.

2. If there is no such child u_0 of u, then for every child u^{in} of u, replace arc (u^{in}, u) with a new arc $(u^{\mathsf{in}}, u^{\mathsf{out}})$. In this process, u has its in-degree changed to zero thereby becoming a leaf.

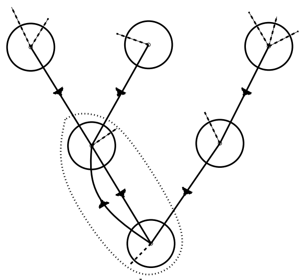

Figure 3.1: The Dependency Tree: Dashed edges represent private centers, circles represent the local balls

Notice that we have maintained the invariant that there are no out-arcs from any pseudo-root, and every node has at most one out-arc. Define mapping $\sigma : V' \to V'$ as follows: for each $u \in V'$, set $\sigma(u)$ to u's parent in the final dependency graph (if it exists); otherwise (if u is itself a pseudo-root) set $\sigma(u) = u$. Note that the final dependency graph is a collection of stars with centers as pseudo-roots.

CLAIM 3.2. *For each $w \in V'$, we have $d(w, \sigma(w)) \leq 2 \cdot d(w, \eta(w))$.*

Proof. Suppose that when w is considered as vertex u in the above procedure, step 1 applies. Then it follows that the out-arc of w is never changed after this, and by definition of step 1, $d(w, \sigma(w)) \leq 2 \cdot d(w, \eta(w))$. The remaining case is that when w is considered as vertex u, step 2 applies.

Then from the definition of steps 1 and 2, we obtain that there is a directed path $\langle w = w_0, w_1, \cdots, w_t \rangle$ in the initial dependency graph such that $\eta(w) = w_1$ and $\sigma(w) = w_t$. Let $d(w, \eta(w)) = d(w_0, w_1) = a$.

We claim by induction on $i \in \{1, \cdots, t\}$ that $d(w_i, w_{i-1}) \leq a/2^{i-1}$. The base case of $i = 1$ is obvious. For any $i < t$, assuming $d(w_i, w_{i-1}) \leq a/2^{i-1}$, we

[1] In general, each component might have one cycle of any length; but since all edges in a cycle will have the same length, we may assume without loss of generality that there are only 2-cycles.

will show that $d(w_{i+1}, w_i) \leq a/2^i$. Consider the point when w's out-arc is changed from (w, w_i) to (w, w_{i+1}); this must be so, since w's out-arc changes from (w, w_1) to (w, w_t) through the procedure. At this point, step 2 must have occurred at node w_i, and w_{i-1} must have been a child of w_i; hence $d(w_{i+1}, w_i) \leq \frac{1}{2} \cdot d(w_i, w_{i-1}) \leq a/2^i$.

Thus we have $d(w, \sigma(w)) \leq \sum_{i=1}^{t} d(w_i, w_{i-1}) \leq a \sum_{i=1}^{t} \frac{1}{2^{i-1}} < 2a = 2 \cdot d(w, \eta(w))$. ■

At this point, we have a fractional solution (x^3, y^*) that satisfies constraints (2.1)-(2.4) and:

$$\sum_{u \in V'} w_u \left[\sum_{v \in \mathcal{P}(u)} d(u,v) x^3_{uv} + d(u, \sigma(u))(1 - \sum_{v \in \mathcal{P}(u)} x^3_{uv}) \right] \leq 4 \cdot \mathsf{LPOpt} \tag{3.8}$$

The inequality follows from (3.7) and Claim 3.2.

3.2 Stage II: Reformulating the LP Based on the star-like structure derived in the previous subsection, we propose another linear program for which the fractional solution (x^3, y^*) is shown to be feasible with objective value as in (3.8). Crucially, we will show that this new LP is integral. Hence we can obtain an integral solution to it of cost at most $4 \cdot \mathsf{LPOpt}$. Finally we show that any integral solution to our reformulated LP also corresponds to an integral solution to the original MatroidMedian instance, at the loss of another constant factor.

Consider the LP described in Figure 3.2.

The reason we have added the constraint 3.10 is the following: In the objective function, each client incurs only a cost of $d(u, \sigma(u))$ to the extent to which a private facility from $\mathcal{P}(u)$ is not assigned to it. This means that in our integral solution, we definitely want a facility to be chosen from the pseudo-root to which u is connected if we do not open a private facility from $\mathcal{P}(u)$; this fact becomes clearer later. Also, this constraint does not increase the optimal value of the LP, as shown below.

CLAIM 3.3. *The linear program* LP_2 *has optimal value at most* $4 \cdot \mathsf{LPOpt}$.

Proof. Consider the solution z defined as: $z_v = \min\{y^*_v, x^3_{uv}\} = x^3_{uv}$ for all $v \in \mathcal{P}(u)$ and $u \in V'$; all other vertices have z-value zero. It is easy to see that constraints (3.9) and (3.11) are satisfied.

Constraint (3.10) is also trivially true for pseudo-roots consisting of only one client. Else, let $\{u_1, u_2\}$ be any pseudo-root consisting of two clients. Recall that each $u \in \{u_1, u_2\}$ is connected to centers in ball $\mathcal{B}(u, 2\mathsf{LP}_u) \subseteq \mathcal{P}(u)$ to extent at least half; hence the total z-value inside $\mathcal{P}(u_1) \cup \mathcal{P}(u_2)$ is at least one. Thus z is feasible for LP_2, and by (3.8) its objective value is at most $4 \cdot \mathsf{LPOpt}$. ■

We show next that LP_2 is in fact, an integral polytope.

LEMMA 3.2. *Any basic feasible solution to* LP_2 *is integral.*

Proof. Consider any basic feasible solution z. Firstly, notice that the characteristic vectors defined by constraints (3.9) and (3.10) define a laminar family, since all the sets $\mathcal{P}(u)$ are disjoint.

Therefore, the subset of these constraints that are tightly satisfied by z define a laminar family (of mostly disjoint sets). Also, by standard uncrossing arguments (see eg. [19]), we can choose the linearly-independent set of tight rank-constraints (3.11) to form a laminar family (in fact even a *chain*).

But then the vector z is defined by a constraint matrix which consists of *two laminar families* on the ground set of vertices. Such matrices are well-known to be *totally unimodular* [19], and this fact is used in proving the integrality of the matroid-intersection polytope. For completeness, we outline a proof of this fact in the Appendix B. This finishes the integrality proof. ■

It is clear that any integral solution feasible for LP_2 is also feasible for MatroidMedian, due to (3.11). We now relate the objective in LP_2 to the original MatroidMedian objective:

LEMMA 3.3. *For any integral solution* $C \subseteq V$ *to* LP_2*, the* MatroidMedian *objective value under* C *is at most* 3 *times that it was paying in the* LP_2 *solution.*

Proof. We show that each client $u \in V'$ pays in MatroidMedian at most 3 times that in LP_2. Suppose that $C \cap \mathcal{P}(u) \neq \emptyset$. Then u's connection cost is identical to its contribution to the LP_2 solution's objective. Therefore, we assume $C \cap \mathcal{P}(u) = \emptyset$.

Suppose that u is not part of a pseudo-center; let $\{u_1, u_2\}$ denote the pseudo-center that u is connected to. By constraint (3.10), there is some $v \in C \bigcap (\mathcal{P}(u_1) \cup \mathcal{P}(u_2))$. The contribution of u is $d(u, \sigma(u))$ in LP_2 and $d(u, v)$ in the actual objective function for MatroidMedian. We will now show that $d(u, v) \leq 3 \cdot d(u, \sigma(u))$.

Without loss of generality let $\sigma(u) = u_1$ and suppose that $v \in \mathcal{P}(u_2)$; the other case of $v \in \mathcal{P}(u_1)$ is easier. From the property of private centers, we know $d(u_2, v) \leq d(u_2, \eta(u_2)) \leq d(u_2, u_1)$. Now if (u_1, u_2) is created as a new pseudo-root in step (iii).1, then we have the property that $d(u_1, u_2) \leq d(u_1, u)$, since we choose the closest leaf to pair up with its parent to form a pseudo-root. Else (u_1, u_2) is the original pseudo-root even before the modifications of step (iii). Thus in that case, by definition $d(u_1, u_2) = d(u_1, \eta(u_1)) \leq d(u_1, u)$. Therefore, $d(u, v) \leq d(u, u_1) + d(u_1, u_2) + d(u_2, v) \leq d(u, u_1) + 2 \cdot d(u_1, u_2) \leq 3 \cdot d(u, u_1) = 3 \cdot d(u, \sigma(u))$.

If u is itself (a singleton) pseudo-center then it must be that $C \cap \mathcal{P}(u) \neq \emptyset$ by (3.10), contrary to the above assumption. If u is part of a pseudo-center $\{u, u'\}$. Then it must

$$\text{(LP}_2\text{)} \quad \text{minimize} \sum_{u \in V'} w_u \left[\sum_{v \in \mathcal{P}(u)} d(u,v) z_v + d(u,\sigma(u)) \left(1 - \sum_{v \in \mathcal{P}(u)} z_v\right)\right]$$

$$\text{(3.9)} \quad \text{subject to} \sum_{v \in \mathcal{P}(u)} z_v \le 1 \qquad \forall\, u \in V'$$

$$\text{(3.10)} \quad \sum_{v \in \mathcal{P}(u_1)} z_v + \sum_{v \in \mathcal{P}(u_2)} z_v \ge 1 \qquad \forall \text{ pseudo-roots } \{u_1, u_2\}$$

$$\text{(3.11)} \quad \sum_{v \in S} z_v \le r_{\mathcal{M}}(S) \qquad \forall\, S \subseteq V$$

$$\text{(3.12)} \quad z_v \ge 0 \qquad \forall\, v \in V$$

Figure 3.2: Stage II LP Relaxation

be that there is some $v \in C \cap \mathcal{P}(u')$, by (3.10). The contribution of u in LP_2 is $d(u, \sigma(u))$, and in MatroidMedian is $d(u,v) \le d(u,u') + d(u',v) \le 2 \cdot d(u,u') = d(u,\sigma(u))$ (the second inequality uses property of private centers). ■

To make this result algorithmic, we need to obtain in polynomial-time an extreme point solution to LP_2. Using the Ellipsoid method (as mentioned in Section 2.2) we can indeed obtain *some* fractional optimal solution to LP_2, which may not be an extreme point. However, such a solution can be converted to an extreme point of LP_2, using the method in Jain [14]. (Due to the presence of both "$\le$" and "$\ge$" type constraints in (3.9)-(3.10) it is not clear whether LP_2 can be cast directly as an instance of matroid intersection.)

Altogether, we obtain an integral solution to the weighted instance from **step (i)** of cost $\le 12 \cdot \mathsf{LPOpt}$. Combined with the property of **step (i)**, we obtain:

THEOREM 3.1. *There is a 16-approximation algorithm for the* MatroidMedian *problem.*

We have not tried to optimize the constant via this approach. However, getting the approximation ratio to match that for usual k-median would require additional ideas.

4 MatroidMedian with Penalties

In the MatroidMedian problem with penalties, each client either connects to an open center thereby incurring the connection cost, or pays a penalty. Again, we are given a finite metric space (V, d), a matroid $\mathcal{M}$ with ground set V and a set of independent sets $\mathcal{I}(\mathcal{M}) \subseteq 2^V$; in addition we are also given a penalty function $p : V \to \mathbb{R}_{\ge 0}$. The goal is to open centers $S \in \mathcal{I}(\mathcal{M})$ and identify a set of clients C_1 such that $\sum_{u \in C_1} d(u, S) + \sum_{u \in V \setminus C_1} p(u)$ is minimized. Such objectives are also called "prize-collecting" problems. In this section we give a constant factor approximation algorithm for MatroidMedian with penalties, building on the rounding algorithm from the previous section.

Our LP relaxation is an extension of the one (LP_1) for MatroidMedian. In addition to the variables $\{y_v;\, v \in V\}$ and $\{x_{uv};\, u, v \in V\}$, we define for each client an indicator variable h_u whose value equals 1 if client u pays a penalty and is not connected to any open facility. Then, it is straightforward to see that the following LP is indeed a valid relaxation for the problem.

$$\min \sum_{u \in V} \sum_{v \in V} d(u,v) x_{uv} + \sum_{u \in V} p(u) h_u \qquad (\mathsf{LP}_3)$$

$$\text{(4.13)} \quad \text{s. t} \sum_{v \in V} x_{uv} + h_u = 1 \qquad \forall\, u \in V$$

$$\text{(4.14)} \quad x_{uv} \le y_v \qquad \forall\, u \in V, v \in V$$

$$\text{(4.15)} \quad \sum_{v \in S} y_v \le r_{\mathcal{M}}(S) \qquad \forall\, S \subseteq V$$

$$\text{(4.16)} \quad x_{uv}, y_v, h_u \ge 0 \qquad \forall\, u, v \in V$$

Let Opt denote an optimal solution of the given penalty MatroidMedian instance, and let LPOpt denote the LP_3 optimum value. Since LP_3 is a relaxation of our problem, we have that,

LEMMA 4.1. *The* LP_3 *cost* LPOpt *is at most the cost of an optimal solution* Opt.

Let (x^*, y^*, h^*) denote the optimal LP solution. We round the fractional variables to integers in two stages. In the first stage, only x^*, h^* variables are altered such that the modified solution is still feasible but sparse and has a cost that is $O(1)\mathsf{LPOpt}$. This stage is similar in nature to the first stage rounding for MatroidMedian but we need to be careful when comparing the x_{uv}'s and the y_v's — the primary difficulty is that for any client u, the sum $\sum_v x_{uv}$ is not 1. Typically, this is the case in most LP-based prize-collecting problems, but often, one could argue that if $\sum_v x_{uv} \ge 2/3$, then by scaling we could ensure that it is at least 1; *therefore* the (scaled) LP would be feasible to the original problem (without penalties). However, in our case (and also k-median with penalties) since we also have packing constraints (the matroid rank constraints), simply scaling the fractional solution is not a viable route.

Once we handle these issues, we show that a new LP can be written for which the modified fractional solution is feasible. The constraints of this LP are identical to that for MatroidMedian; but the objective function is different. Since that polytope is integral (Lemma 3.2) we infer that the new LP for MatroidMedian with penalties is integral. This immediately gives us the constant factor approximation for MatroidMedian with penalties. We now get into the details.

4.1 Stage I: Sparsifying the LP Solution Like in the matroid median setting, the goal in this stage is to argue that there exists a *sparse* fractional solution of near optimal cost. This will enable us to write another LP, which will be characterized by an integral polytope.

Step (i): Thresholding Penalties.

Let $\mathcal{C}$ denote the set of clients paying a penalty at most to an extent of $1/4$ in the fractional solution, i.e, $\mathcal{C} = \{u \in V \,|h^*_u < \frac{1}{4}\}$. For each client $u \in V \setminus \mathcal{C}$, we round its h^*_u to one and set x^*_{uv} to zero for all $v \in V$. Let (x^1, h^1, y^*) denote this modified solution. We make the following observations:

OBSERVATION 4.1. *After the above thresholding operation, the following inequalities are satisfied.*

1. $\forall\, u \in \mathcal{C}$, $\sum_{v\in V} x^1_{uv} > \frac{3}{4}$
2. $\forall\, u \in \mathcal{C}$, $h^1_u < \frac{1}{4}$
3. $\sum_{u\in V} \left(\sum_{v\in V} x^1_{uv} d(u,v) + h^1_u p(u)\right) \le 4\,\mathsf{LPOpt}$
4. $\forall u, v \in V$, *if* $x^1_{uv} > 0$ *then* $p(u) \ge d(u,v)$.
5. $\forall u, v \in V$, *if* $x^1_{uv} > 0$ *then* $x^1_{uv} = y^*_v$ *or* $h^1_u = 0$.

Here, the second-to-last property is true because of the following: For any client $u \in V$ and center $v \in V$, if $x^1_{uv} > 0$ and $d(u,v) > p(u)$, then we can increase h^1_u to $h^1_u + x^1_{uv}$ and set $x^1_{uv} = 0$. Such a modification maintains feasibility and only decreases the objective function value.

The final property can be seen from the following argument: because $p(u) \ge d(u,v)$, whenever $x^1_{uv} > 0$, we can increase the connection variable x^1_{uv} and decrease h^1_u equally without increasing the objective function until h^1_u becomes 0 or x^1_{uv} hits y^*_v.

Step (ii): Clustering Clients.

Let $|\mathcal{C}| = n'$. For each client $u \in \mathcal{C}$, let $D_u = \sum_{v\in V} d_{uv} x^1_{uv}$ denote the fractional connection cost of client u and $X^1_u = \sum_{v\in V} x^1_{uv}$ denote the total fractional assignment towards connection. Also let $\mathcal{B}(u,R) = \{v \in V \,|d(u,v) \le R\}$ denote the ball of radius R centered at u. For any vertex u, we say that $\mathcal{B}(u, 4D_u)$ is the 'local ball' centered at u. The following is a direct consequence of Markov's inequality and Observation 4.1-(1).

LEMMA 4.2. *For any client* $u \in \mathcal{C}$, $\sum_{v\in\mathcal{B}(u,4D_u)} x^1_{uv} \ge \frac{1}{2}$. *In words, each client is fractionally connected to centers in its local ball to an extent of at least* $1/2$.

Now order the clients according to non-decreasing D_u values, and let the ordering be $u_1, u_2, \ldots, u_{n'}$. Consider the vertices in the order $u_1, u_2, \ldots, u_{n'}$. For a vertex u_i, if there exists a vertex u_j with $j < i$ such that $d(u_i, u_j) \le 8D_{u_i}$, then we denote this event by $u_i \twoheadrightarrow u_j$ and modify the instance by shifting the location of client u_i to u_j.

For each client u_i, define $\pi(u_i) = u_j$ iff $u_i \twoheadrightarrow u_j$, and $\pi(u_i) = u_i$ if u_i was not shifted. Let $V' = \pi(\mathcal{C})$ denote the set of clients that maintain their own local balls (i.e were not shifted in the above process). For each $u \in V'$, let $C_u := \{u' \in \mathcal{C} | \pi(u') = u\}$. The new instance consists of $|C_u|$ clients located at each vertex $u \in V'$ having respective penalty values $\{p(u') | u' \in C_u\}$.

OBSERVATION 4.2. *For* $u, v \in V'$, *we have* $d(u,v) > 8\max(D_u, D_v)$.

We obtain a feasible solution (x^2, h^2, y^*) to this modified instance as follows. For each event $u_i \twoheadrightarrow u_j$, do:

Case (i): If $X^1_{u_i} \le X^1_{u_j}$: Start with $x^2_{u_i v} = 0$ for all v. Then, for each vertex v with $x^1_{u_j v} > 0$, connect u_i to an extent of $x^1_{u_j v}$ to v, i.e. set $x^2_{u_i v} = x^1_{u_j v}$. Finally, set $h^2_{u_i} = h^1_{u_j} \le h^1_{u_i}$.

Case (ii): If $X^1_{u_i} > X^1_{u_j}$: Start with $x^2_{u_i v} = 0$ for all v. For each v with $x^1_{u_j v} > 0$, connect u_i to an extent of $x^1_{u_j v}$ to v, i.e. set $x^2_{u_i v} = x^1_{u_j v}$. Since $X^1_{u_j} < X^1_{u_i}$, we need to further connect u_i to other centers to extent of at least $X^1_{u_i} - X^1_{u_j}$ in order to avoid increasing h_{u_i}. To this end, set $x^2_{u_i,w} = x^1_{u_i,w}$ for all $w \in V$ with $x^1_{u_j,w} = 0$. Observe that client u_i is now connected to extent at least $X^1_{u_i}$; so $h^2_{u_i} \le 1 - X^1_{u_i} = h^1_{u_i}$.

Also if a client is not shifted in the above routine, its x^2, h^2 variables are the same as in x^1, h^1. The following lemma certifies that the objective of the modified instance is not much more than the original.

LEMMA 4.3. $\sum_{u\in\mathcal{C}} \left[\sum_{v\in V} x^2_{uv} \cdot d(\pi(u), v) + h^2_u \cdot p(u)\right] \le \sum_{u\in\mathcal{C}} \left[10 \sum_{v\in V} x^1_{uv} \cdot d(u,v) + h^1_u \cdot p(u)\right]$.

Proof. We prove the inequality term-wise. Any client that is not shifted in the above process maintains its contribution to the objective function. Hence consider a client u_i that is shifted to u_j (ie. $u_i \twoheadrightarrow u_j$). It is clear that $h^2_{u_i} \le h^1_{u_i}$, so the penalty contribution $h^2_u \cdot p(u) \le h^1_u \cdot p(u)$. There are two cases for the connection costs:

Case (i): $X^1_{u_i} \le X^1_{u_j}$.

In this case we have, $\sum_{v\in V} x^2_{u_i v} d(\pi(u_i), v) = \sum_{v\in V} x^2_{u_i v} d(u_j, v) = \sum_{v\in V} x^1_{u_j v} d(u_j, v) = D_{u_j} \le D_{u_i}$.

Case (ii): $X^1_{u_i} > X^1_{u_j}$.

Here, note that $x^2_{u_i v} \le x^1_{u_i v} + x^1_{u_j v}$ for all $v \in V$. So,

$$\begin{aligned} &\sum_{v\in V} x^2_{u_i v} \cdot d(u_j, v) \\ &\le \sum_{v\in V} x^1_{u_j v} \cdot d(u_j, v) + \sum_{v\in V} x^1_{u_i v} \cdot d(u_j, v) \\ &\le D_{u_j} + \sum_{v\in V} x^1_{u_i v} \cdot d(u_i, v) + \left(\sum_{v\in V} x^1_{u_i v}\right) \cdot d(u_j, u_i) \\ &\le D_{u_j} + D_{u_i} + d(u_j, u_i) \le 10 \cdot D_{u_i} \end{aligned}$$

Hence in either case, we can bound the new connection cost by $10 \cdot D_{u_i}$. ■

Thus it follows that (x^2, h^2, y^*) is a feasible LP solution to the modified instance of objective value at most 10 LPOpt. We additionally ensure (by locally changing x^2, h^2) that condition 4 of Observation 4.1 holds, namely:

(4.17)
$$\forall u \in V', u' \in C_u, v \in V, \text{ if } x^2_{u'v} > 0 \text{ then } d(u, v) \le p(u').$$

Note that any feasible integral solution to the modified instance corresponds to one for the original instance, wherein the objective increases by at most an additive term of $8 \cdot \sum_{w\in\mathcal{C}} D_w \le 8 \cdot$ LPOpt. Hence in the rest of the algorithm we work with this modified instance.

Next we modify the connection-variables (leaving penalty variables h^2 unchanged) of clients exactly as in **step (ii)** of the previous section, and also alter the coefficients of some x variables just like in the algorithm for MatroidMedian. This results in a disjoint set of private centers $\mathcal{P}(u)$ for each $u \in V'$ (where $\mathcal{P}(u)$ can be thought of as the collection of all private centers for $u' \in C_u$; notice that these are disjoint for different vertices in V'), and new connection variables $\tilde{x}^3$ such that:

(a) Each client u' depends only on centers $\mathcal{P}(\pi(u'))$ and centers in the local ball nearest to $\pi(u')$. The connection to the latter type of centers is at most half.

(b) For any client $u \in V'$ and center $v \in \mathcal{P}(u)$, we have $d(u, v) \le d(u, w)$ for any other client $w \in V'$.

(c) The total cost under the modified objective is at most $20 \cdot$ LPOpt (the factor 2 loss is incurred due to changing the objective coefficients and rearrangements).

(d) For any $u' \in \mathcal{C}$, $v \in V$ with $\tilde{x}^3_{u'v} > 0$ we have $d(\pi(u'), v) \le 3 \cdot p(u')$. Additionally, for $u' \in \mathcal{C}$, $v \in \mathcal{P}(\pi(u'))$ with $\tilde{x}^3_{u'v} > 0$ we have $d(\pi(u'), v) \le p(u')$.

The first three properties above are immediate from the corresponding properties after step (ii) of Section 3. The last property uses (4.17) and Claim 3.1.

We now modify the penalty variables as follows (starting with $h^3 = h^2$ and $x^3 = \tilde{x}^3$). For each client u', if it is connected to centers in the local-ball of any $w \in V' \setminus \{\pi(u')\}$ then reset $h^3(u') = 0$; and increase the connection-variables $x^3(u', \cdot)$ to centers in the local-ball of w until client u' is connected to extent one. (Such a modification is possible since u' is already connected to extent at least half in the local ball of $\pi(u')$, and there is at least half open centers in any local ball.) Furthermore, using property (d) above, the new objective value of (x^3, h^3, y^*) is at most *thrice* that of $(\tilde{x}^3, h^3, y^*)$, ie. at most $60 \cdot$ LPOpt.

OBSERVATION 4.3. *Any client u' that has $h^3(u') > 0$ is connected* only *to centers in $\mathcal{P}(\pi(u'))$.*

We also apply **step (iii)** from Section 3 to obtain a mapping $\sigma : V' \to V'$ satisfying Claim 3.2 (recall that $\eta : V' \to V'$ maps each client in V' to its closest other client). This increases the objective value by at most factor 2.

Let $M = \{u' \in \mathcal{C} | h^3(u') > 0\}$ denote the clients that have non-zero penalty variable. For each $u' \in M$ let $T(u') \subseteq \mathcal{P}(\pi(u'))$ denote the centers that client u' is connected to (We may assume that $T(u')$ consists of centers in $\mathcal{P}(u)$ closest to u). The objective of (x^3, h^3, y^*) can then be expressed as in the equation in Figure 4.3. From the arguments above, the cost of this solution is at most $120 \cdot$ LPOpt.

4.2 Stage II: Reformulaing the LP

Reducing center variables y^*. For any $u \in V'$, if $\sum_{v\in\mathcal{P}(u)} y^*_v > 1$ then we reduce the y^*-values in $\mathcal{P}(u)$ one center at a time (starting from the farthest center to u) until $y^*(\mathcal{P}(u)) = 1$. Clearly this does not cause the objective to increase. Additionally, y^* still satisfies the matroid independence constraints. Thus we can ensure that $\sum_{v\in\mathcal{P}(u)} y^*_v \le 1$ for all $u \in V'$. Additionally, the following two modifications do not increase the objective.

1. *Client* $u' \in M$. For all $v \in T(u')$ we have $d(\pi(u'), v) \le p(u')$ (property (d) above); set $x^3(u', v) = y^*_v$.

2. *Client* $u \in \mathcal{C} \setminus M$. For all $v \in \mathcal{P}(\pi(u))$ we have $d(\pi(u), v) \le d(\pi(u), \sigma(\pi(u)))$ (property (b) above); again set $x^3(u, v) = y^*_v$.

Thus we can re-write the objective from (4.18) as shown in Figure 4.4 (which is just that in Figure 4.3 with the x variables replaced by the y variables). Notice that there are no x-variables in the above expression. Furthermore, y^* satisfies all the constraints (3.9)-(3.12). We now consider linear program LP$_4$ with the linear objective (4.19) and con-

$$\sum_{u \in M} \left[\sum_{v \in T(u)} d(\pi(u), v) \cdot x^3_{uv} + p_u \cdot \left(1 - \sum_{v \in T(u)} x^3_{uv} \right) \right] +$$

$$\sum_{u \in \mathcal{C} \setminus M} \left[\sum_{v \in \mathcal{P}(\pi(u))} d(\pi(u), v) x^3_{uv} + d(\pi(u), \sigma(\pi(u))) \cdot (1 - \sum_{v \in \mathcal{P}(\pi(u))} x^3_{uv}) \right], \tag{4.18}$$

Figure 4.3: Modified Objective Function of (x^3, h^3, y^*)

$$\sum_{u \in M} \left[\sum_{v \in T(u)} d(\pi(u), v) \cdot y^*_v + p_u \cdot \left(1 - \sum_{v \in T(u)} y^*_v \right) \right]$$

$$+ \sum_{u \in \mathcal{C} \setminus M} \left[\sum_{v \in \mathcal{P}(\pi(u))} d(\pi(u), v) \cdot y^*_v + d(\pi(u), \sigma(\pi(u))) \cdot \left(1 - \sum_{v \in \mathcal{P}(\pi(u))} y^*_v \right) \right]. \tag{4.19}$$

Figure 4.4: Objective Function for Sparse LP

straints (3.9)-(3.12). This can be optimized in polynomial time to obtain an optimal integral solution F (as described in Subsection 3.2). From the reductions in the previous subsection, the objective value of F under (4.19) is at most $120 \cdot \mathsf{LPOpt}$. Finally, using Lemma 3.3 we obtain that F is a feasible solution to MatroidMedian with penalties, of objective at most $360 \cdot \mathsf{LPOpt}$.

THEOREM 4.1. *There is a constant approximation algorithm for* MatroidMedian *with penalties.*

5 The KnapsackMedian Problem

In this section we consider the KnapsackMedian problem. We are given a finite metric space (V, d), non-negative weights $\{f_i\}_{i \in V}$ and a bound F. The goal is to open centers $S \subseteq V$ such that $\sum_{j \in S} f_j \leq F$ and the objective $\sum_{u \in V} d(u, S)$ is minimized. We can write a LP relaxation (LP_5) of the above problem similar to (LP) in Section 2.1, where we replace the constraint (2.3) with the knapsack constraint $\sum_{v \in V} f_v y_v \leq F$. In addition, we guess the maximum weight facility f_{max} used in an optimum solution, and if $f_v > f_{max}$ we set $y_v = 0$ (and hence $x_{uv} = 0$ as well). This is clearly possible since there are only n many different choices for f_{max}. Unfortunately LP_5 has an unbounded integrality gap if we do not allow any violation in the knapsack constraint. In Subsection 5.1, we show that a similar integrality gap persists even if we add the *knapsack-cover (KC) inequalities* to strengthen LP_5, which have often been useful to overcome the gap of the natural LP [5].

However, in the following section, we show that with an *additive* slack of f_{max} in the budget, we can get a constant factor approximation for the knapsack median problem.

The Rounding Algorithm for KnapsackMedian**.** Let (x^*, y^*) denote the optimal LP solution of LP_5. The rounding algorithm follows similar steps as in MatroidMedian problem. The first stage is identical to Stage I of Section 3 modifying x_{uv} variables until we have a collection of disjoint stars with pseudo-roots. The total connection cost of the modified LP solution is at most a constant factor of the optimum LP cost for LP_5. The sparse instance satisfies the budget constraint since y_u variables are never increased. In Stage II, we start with a new LP (LP_6) by replacing the constraint 3.11 of LP_2 with the knapsack constraint $\sum_{v \in V} f_v z_v \leq F$. However LP_6 is not integral as opposed to LP_2: it contains the knapsack problem as a special case. We now give an iterative-relaxation procedure that rounds the above LP_6 into an integral solution by violating the budget by at most an additive f_{max} and maintaining the optimum connection cost. The following algorithm iteratively creates the set of open centers $\mathcal{C}$.

1. Initialize $\mathcal{C} \leftarrow \emptyset$. While $V \neq \emptyset$ do
   (a) Find an extreme point optimum solution $\hat{z}$ to LP_6.
   (b) If there is a variable $\hat{z_v} = 0$, then remove variable $\hat{z_v}$, set $V = V \setminus \{v\}$.
   (c) If there is a variable $\hat{z_v} = 1$, then $\mathcal{C} \leftarrow \mathcal{C} \cup \{v\}$, $V = V \setminus \{v\}$ and $F = F - f_v$.
   (d) If none of (b), (c) holds, and $|V| = 2$ (say $V = \{x_1, x_2\}$) then:
      - If $x_1, x_2 \in \mathcal{P}(u)$ for some $u \in V'$. If $d(x_1, u) \leq d(x_2, u)$ then $\mathcal{C} \leftarrow \mathcal{C} \cup \{x_1\}$, else $\mathcal{C} \leftarrow \mathcal{C} \cup \{x_2\}$. Break.

- If $x_1, x_2 \in \mathcal{P}(u_1) \cup \mathcal{P}(u_2)$ for some pseudo-root $\{u_1, u_2\}$. Then $\mathcal{C} \leftarrow \mathcal{C} \cup \{x_1, x_2\}$, Break.

2. Return $\mathcal{C}$

The following lemma guarantees that the connection cost is at most the Opt cost of LP_4 and the budget is not exceeded by more than an additive f_{max}.

LEMMA 5.1. *The above algorithm finds a solution for knapsack median problem that has cost at most* Opt *cost of* LP_6 *and that violates the knapsack budget at most by an additive* f_{max}.

Proof. First we show that if the algorithm reaches Step (2), then the solution returned by the algorithm satisfies the guarantee claimed. In Step (c), we always reduce the remaining budget by f_v if we include the center in $\mathcal{C}$. Thus the budget constraint can only be violated at Step (d). In Step (d), in case of tight V'-constraint, we open only one center among the two remaining centers. Thus the budget constraint can be violated by at most $\max\{f_{x_1}, f_{x_2}\} \leq f_{max}$. In case of tight pseudo-root, we have $z(x_1) + z(x_2) = 1$ and thus $f_{x_1} \cdot z(x_1) + f_{x_2} \cdot z(x_2) + \max\{f_{x_1}, f_{x_2}\} \leq f_{x_1} + f_{x_2}$. Hence again the budget constraint can be violated by at most an additive f_{max} term. The total cost of LP_6 never goes up in Step (a)-(c). In Step (d), either the nearer center from $\{x_1, x_2\}$ is chosen (in case of tight V'-constraint), or both the centers $\{x_1, x_2\}$ (in case of tight pseudo-root) are opened. Thus the connection cost is always upper bounded by Opt of LP_6.

To complete the proof we show that the algorithm indeed reaches Step (2). The Steps (1b),(1c) all make progress in the sense that they reduce the number of variables and hence some constraints become vacuous and are removed. Therefore, we want to show whenever we are at an extreme point solution, then either Step (1b),1(c) apply or we have reached (1d) and hence Step (2). Suppose that neither (1b) nor (1c) apply: then there is no $z_v \in \{0, 1\}$. Let the linearly independent tight constraints defining z be: $T \subseteq V'$ from (3.9), and R (pseudo-roots) from (3.10). From the laminar structure of the constraints and all right-hand-sides being 1, it follows that the sets in $T \bigcup R$ are all disjoint. Further, each set in $T \bigcup R$ contains at least two fractional variables. Hence the number of variables is at least $2|T| + 2|R|$. Now count the number of tight linearly independent constraints: There are at most $|T| + |R|$ tight constraints from (3.9)-(3.10), and one global knapsack constraint. Since at an extreme point, the number of variables must equal the number of tight linearly independent constraints, we obtain $|T| + |R| \leq 1$ and that each set in $T \cup R$ contains exactly two vertices. This is possible only when V is some $\{x_1, x_2\}$.

1. $|T| = 1$. Then there must be some $u \in V'$ with $x_1, x_2 \in \mathcal{P}(u)$.
2. $|R| = 1$. Then there must be some pseudo-root $\{u_1, u_2\}$ with $x_1, x_2 \in \mathcal{P}(u_1) \cup \mathcal{P}(u_2)$.

So in either case, Step (1d) applies. ■

Combining Lemma 5.1 with Claim 3.3, Lemma 3.3 and the property of **step (i)**-Stage-I rounding of Section 3, we get the following theorem.

THEOREM 5.1. *There is a 16-approximation algorithm for the* KnapsackMedian *problem that violates the knapsack constraint at most by an additive* f_{max}, *where* f_{max} *is the maximum weight of any center opened in the optimum solution.*

Using enumeration for centers with cost more than ϵF, we can guarantee that we do not exceed the budget by more than ϵF while maintaining a 16-approximation for the connection cost in $n^{O(\frac{1}{\epsilon})}$ time.

5.1 LP Integrality Gap for KnapsackMedian **with Knapsack Cover Inequalities** There is a large integrality gap for LP_5 with a hard constraint for the Knapsack bound, from Charikar and Guha [6].

EXAMPLE 5.1. ([6]) *Consider* $|V| = 2$ *with* $f_1 = N$, $f_2 = 1$, $d(1, 2) = D$ *and* $F = N$ *for any large positive reals* N *and* D. *An optimum solution that does not violate the knapsack constraint can open either center* 1 *or* 2 *but not both and hence must pay a connection cost of* D. LP_5 *can assign* $y_1 = 1 - \frac{1}{N}$ *and* $y_2 = 1$ *and thus pay only* D/N *in the connection cost.*

The above example can be overcome by adding *knapsack covering (KC) inequalities* [5]. We now illustrate the use of KC inequalities in the KnapsackMedian problem. KC-inequalities are used for *covering* knapsack problems. Although KnapsackMedian has a packing constraint (at most F weight of open centers), it can be rephrased as a covering-knapsack by requiring "at least $\sum_{v \in V} f_v - F$ weight of closed centers". Viewed this way, we can strengthen the basic LP as follows.

Define for any subset of centers $S \subseteq V$, $f(S) := \sum_{v \in S} f(v)$. Then to satisfy the knapsack constraint we need to close centers worth of $F' := f(V) - F$. For any subset $S \subseteq V$ of centers with $f(S) < F'$ we can write a KC inequality *assuming* that all the centers in S are closed. Then, the residual covering requirement is:

$$\sum_{v \notin S} \min\{f(v), F' - f(S)\}(1 - y_v) \geq F' - f(S).$$

There are exponential number of such inequalities; however using methods in [5] an FPTAS for the strengthened LP can be obtained. The addition of KC inequalities avoids examples like 5.1; there $F' = 1$ and setting $S = \emptyset$ yields:

$$\min\{1, 1\} \cdot (1 - y_1) + \min\{1, N\} \cdot (1 - y_2) \geq 1,$$

ie. $y_1 + y_2 \leq 1$. Thus the LP optimum also has value D.

However the following example shows that the integrality gap remains high even with KC inequalities.

EXAMPLE 5.2. *$V = \{a_i\}_{i=1}^R \bigcup \{b_i\}_{i=1}^R \bigcup \{p,q,u,v\}$ with metric distances d as follows: vertices $\{a_i\}_{i=1}^R$ (resp. $\{b_i\}_{i=1}^R$) are at zero distance from each other, $d(a_1,b_1) = d(p,q) = d(u,v) = D$ and $d(a_1,p) = d(p,u) = d(u,a_1) = \infty$. The facility-costs are $f(a_i) = 1$ and $f(b_i) = N$ for all $i \in [R]$, and $f(p) = f(q) = f(u) = f(v) = N$. The knapsack bound is $F = 3N$. Moreover $N > R \gg 1$.*

An optimum integral solution must open exactly one center from each of $\{a_i\}_{i=1}^R \bigcup \{b_i\}_{i=1}^R$, $\{p,q\}$ and $\{u,v\}$ and hence has connection cost of $(R+2)D$.

On the other hand, we show that the KnapsackMedian LP with KC inequalities has a feasible solution x with much smaller cost. Define $x(a_i) = 1/R$ and $x(b_i) = \frac{N-1}{RN}$ for all $i \in [N]$, and $x_p = x_q = x_u = x_v = \frac{1}{2}$. Observe that the connection cost is $\left(\frac{R}{N}+2\right) D < 3D$. Below we show that x is feasible; hence the integrality gap is $\Omega(R)$.

x clearly satisfies the constraint $\sum_{w\in V} f_w \cdot x_w \leq F$. We now show that x satisfies all KC-inequalities. Recall that $F' = f(V) - F = (R+1)N + R$ for this instance. Note that KC-inequalities are written only for subsets S with $F' - f(S) > 0$. Also, KC-inequalities corresponding to subsets S with $F' - f(S) \geq N = \max_{w\in V} f_w$ reduce to $\sum_{w\notin S} f_w \cdot y_w \leq F$, which is clearly satisfied by x. Thus the only remaining KC-inequalities are from subsets S with $0 < F' - f(S) < N$, ie. $f(S) \in [F' - N + 1, F' - 1] = [RN + R + 1, (R+1)N + R - 1]$. Since all facility-costs are in $\{1, N\}$ and $R < N$, subset S must have *exactly* $R+1$ cost N facilities. Thus there are exactly three cost-N facilities H in S^c. Since $x_w \leq \frac{1}{2}$ for all $w \in V$, we have $\sum_{w\in H}(1 - x_w) \geq \frac{3}{2}$. The KC-inequality from S is hence:

$$\begin{aligned} & \sum_{w\in S^c} \min\{f(w), F' - f(S)\}(1 - x_w) \\ \geq \;& \sum_{w\in H} \min\{f(w), F' - f(S)\}(1 - x_w) \\ = \;& (F' - f(S)) \cdot \sum_{w\in H} (1 - x_w) > F' - f(S). \end{aligned}$$

The equality uses $F' - f(S) < N$ and that each facility-cost in H is N, and the last inequality is by $\sum_{w\in H}(1 - x_w) \geq \frac{3}{2}$ which was shown above.

References

[1] A. Archer, M. H. Bateni, M. T. Hajiaghayi, and H. Karloff. Improved Approximation Algorithms for prize-collecting Steiner tree and TSP. In *FOCS*, pages 427–436, 2009.

[2] S. Arora, P. Raghavan, and S. Rao. Approximation schemes for euclidean k-medians and related problems. In *STOC '98: Proceedings of the thirtieth annual ACM symposium on Theory of computing*, pages 106–113, 1998.

[3] V. Arya, N. Garg, R. Khandekar, A. Meyerson, K. Munagala, and V. Pandit. Local search heuristics for k-median and facility location problems. *SIAM J. Comput.*, 33(3):544–562, 2004.

[4] Y. Bartal. Probabilistic approximation of metric spaces and its algorithmic applications. In *FOCS '96: Proceedings of the 37th Annual Symposium on Foundations of Computer Science*, page 184, 1996.

[5] R. D. Carr, L. Fleischer, V. J. Leung, and C. A. Phillips. Strengthening integrality gaps for capacitated network design and covering problems. In *SODA '00: Proceedings of the eleventh Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 106–115, 2000.

[6] M. Charikar and S. Guha. Improved combinatorial algorithms for facility location problems. *SIAM J. Comput.*, 34(4):803–824, 2005.

[7] M. Charikar, S. Guha, E. Tardos, and D. B. Shmoys. A constant-factor approximation algorithm for the k-median problem. *J. Comput. Syst. Sci.*, 65(1):129–149, 2002.

[8] W. H. Cunningham. Testing membership in matroid polyhedra. *J. Comb. Theory, Ser. B*, 36(2):161–188, 1984.

[9] J. Fakcharoenphol, S. Rao, and K. Talwar. A tight bound on approximating arbitrary metrics by tree metrics. *J. Comput. System Sci.*, 69(3):485–497, 2004.

[10] L. Fleischer, K. Jain, and D. P. Williamson. An iterative rounding 2-approximation algorithm for the element connectivity problem. In *In 42nd Annual IEEE Symposium on Foundations of Computer Science*, pages 339–347, 2001.

[11] A. Gupta and K. Tangwongsan. Simpler analyses of local search algorithms for facility location. *CoRR*, abs/0809.2554, 2008.

[12] M. Hajiaghayi, R. Khandekar, and G. Kortsarz. The red-blue median problem and its generalization. *European Symposium on Algorithms*, 2010.

[13] S. Iwata and J. B. Orlin. A simple combinatorial algorithm for submodular function minimization. In *SODA '09: Proceedings of the twentieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1230–1237, 2009.

[14] K. Jain. A factor 2 approximation algorithm for the generalized steiner network problem. In *FOCS '98: Proceedings of the 39th Annual Symposium on Foundations of Computer Science*, page 448, 1998.

[15] K. Jain, M. Mahdian, and A. Saberi. A new greedy approach for facility location problems. In *STOC '02: Proceedings of the thiry-fourth annual ACM symposium on Theory of computing*, pages 731–740, 2002.

[16] K. Jain and V. V. Vazirani. Approximation algorithms for metric facility location and k-median problems using the primal-dual schema and lagrangian relaxation. *J. ACM*, 48(2):274–296, 2001.

[17] L. C. Lau, J. S. Naor, M. R. Salavatipour, and M. Singh. Survivable network design with degree or order constraints. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing (STOC)*, pages 651–660, 2007.

[18] J.-H. Lin and J. S. Vitter. Approximation algorithms for geometric median problems. *Inf. Process. Lett.*, 44(5):245–249, 1992.
[19] A. Schrijver. *Combinatorial optimization*. Springer, 2003.
[20] M. Singh and L. C. Lau. Approximating minimum bounded degree spanning trees to within one of optimal. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing (STOC)*, pages 661–670, 2007.

A Bad Example for Local Search with Multiple Swaps

Here we give an example showing that any local search algorithm for the T-server type problem (ie. MatroidMedian under partition matroid of T parts) that uses at most $T-1$ swaps cannot give an approximation factor better than $\Omega(\frac{n}{T})$; here n is the number of vertices.

The metric is uniform on $T+1$ locations. There are two servers of each type: Each location $\{2, 3, \ldots, T\}$ contains two servers; locations 1 and $T+1$ contain a single server each. For each $i \in [1, T]$, the two copies of server i are located at locations i (first copy) and $i+1$ (second copy). There are $m \gg 1$ clients at each location $i \in [1, T]$ and just one client at location $T+1$; hence $n = 2T + mT + 1$. The bounds on server-types are $k_i = 1$ for all $i \in [1, T]$. The optimum solution is to pick the *first copy* of each server type and thus pay a connection cost of 1 (the client at location $T+1$). However, it can be seen that the solution consisting of the *second copy* of each server type is locally optimal, and its connection cost is m (clients at location 1). Thus the locality gap is $m = \Omega(n/T)$.

B Proof of TU-ness of Double Laminar Family

We now show that such a matrix is *totally unimodular*. For this we use the following classical characterization: A matrix A is totally unimodular if, for each submatrix A', its rows can be labeled $+1$ or -1 such that every column sum (when restricted to the rows of A' is either $+1$, 0, or -1. Consider such a submatrix A'. Clearly, we have chosen some constraints out of two laminar families, so the chosen rows also correspond to some two laminar families.

Consider one of these laminar families $\mathcal{L}$. We can define a forest by the following rules: We have a node for each set/tight-constraint in $\mathcal{L}$. Nodes S and T are connected by a directed edge from S to T, iff $T \subseteq S$, and there exists no tight constraint $T' \in \mathcal{L} \setminus \{S, T\}$ such that $T \subseteq T' \subseteq S$. Then, we can label each set of $\mathcal{L}$ in the following manner: each node of an odd level gets a label $+1$ and labels of an even levels are -1 (say roots has level of 1, and its children have level 2, and so on). By the laminarity, we know that a variable z_v appears in all the tight constraints which correspond to nodes on a path from some root to some other node. By the way we have labeled these constraints, we know that any such sum is either $+1$, or 0 (see Figure B.5).

Similarly, we can label each set of the second laminar family $\mathcal{L}'$ in the opposite fashion: each node of an odd level gets label -1 and nodes of even levels get $+1$. Again, z_v appears in all the tight constraints corresponding to nodes on a path from some root to some node, and the sum of these labels is either -1 or 0. Therefore, the total sum corresponding to the column for z_v is either $+1$, 0, or -1, which completes the proof.

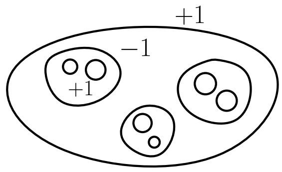

Figure B.5: Labeling the laminar family