Geometric Rounding: A Dependent Randomized Rounding Scheme

Dongdong Ge[*], Simai He[†], Yinyu Ye[*], Jiawei Zhang[‡]

January 5, 2009

Abstract

We develop a new dependent randomized rounding method for approximation of a number of optimization problems with integral assignment constraints. The core of the method is a simple, intuitive, and computationally efficient geometric rounding that simultaneously rounds multiple points in a multi-dimensional simplex to its vertices. Using this method we obtain in a systematic way known as well as new results for the hub location, metric labeling, winner determination and consistent labeling problems. A comprehensive comparison to the dependent randomized rounding method developed by Kleinberg and Tardos [19] and its variants is also conducted, both theoretically and numerically. Overall, our geometric rounding provides a simple and effective alternative for rounding various integer optimization problems.

Keywords: integer programming; linear programming; approximation algorithm; randomized rounding.

1. Introduction

Approximation algorithms based on randomized rounding of fractional solutions have been applied to a variety of optimization problems in recent years. A generic randomized rounding algorithm first formulates a $0-1$ integer programming problem and solves its relaxation in polynomial time to get a (fractional) optimal solution x^*, then rounds each variable x_i to 0 or 1 by a randomization scheme. A polynomial-time *ρ-approximation algorithm* to a minimization(maximization) problem is defined to be an algorithm that finishes in polynomial time and outputs a solution with a value at most ρ(at least $1/\rho$) times the optimal value. $\rho(\geq 1)$ is called *approximation ratio* or *performance guarantee*. Raghavan and Thompson first introduce the idea rounding variable x_i to 1 independently with probability x_i^* in the study of approximation algorithms for covering, packing and routing problems [25, 26]. A variety of randomized independent rounding techniques, which round each variable x_i independently, have been applied to several classes of discrete optimization problems [6, 12, 15, 18, 28]. Meanwhile, dependent rounding schemes have been also developed in parallel [1, 4, 5, 8, 9, 13, 14, 19, 20, 30]. Although

[*]This research is supported by the Boeing Company. Department of Management Science and Engineering, Stanford University, Stanford, CA, USA. Email: {dongdong, yinyu-ye}@stanford.edu. FAX: 650 723 1614.

[†]Department of Systems Engineering and Engineering Management, The Chinese University of Hong Kong, Shatin, Hong Kong. Email: smhe@se.cuhk.edu.hk

[‡]Stern School of Business, IOMS-Operations Management, New York University, New York, NY, 10012. Email: jzhang@stern.nyu.edu

some of these rounding procedures may be deterministically built [1], or be based on a convex programming relaxation [14], all of them underscore the fact that random choices they make highly emphasize the dependencies in the rounding process.

Bertsimas et al. devise dependent rounding schemes in their seminal work for a series of combinatorial optimization problems [4, 5]. Basically, instead of rounding each x_i independently with probability x_i^*, they compare all x_i^*'s with a threshold U that is generated between $[0, 1]$ uniformly at random and round every variable x_i above the threshold to 1. A revision of this rounding generates a $2(1 - \frac{1}{k})$-approximation algorithm for the multiway cut problem with k terminals. Later, Calinescu et al. [9] improve the approximation ratio to $(1.5 - \frac{1}{k})$ by a more sophisticated rounding scheme. Their algorithm treats the LP relaxation solutions as k-dimensional vectors of a standard simplex, and applies the threshold to every vertex of the simplex *sequentially* to create a random partition of this simplex. The sequential order is carefully designed and the partition automatically creates an allocation pattern.

Later, a similar dependent rounding scheme was presented by Kleinberg and Tardos to solve a classification problem [19, 20]. The idea is to uniformly pick a vertex at random and to apply Bertsimas et al.'s dependent rounding scheme on the selected one-variable. The algorithm keeps looping until no points are left. They consider the classification problem arising from metric labeling and Markov random fields, where one needs to assign one of k labels (or classes) to each of n objects. In contrast to the multiway cut problem, the classification problem has a more complex objective function as the penalty for labeling two nodes differently should depend on the identities of the labels they receive. The resulting labeling problem is equivalent to a type of uncapacitated quadratic assignment problem. Their rounding implies an $O(\log k)$-approximation algorithm for the metric labeling problem [2, 19] and a 2-approximation for the uniform metric case. Recently, Krauthgamer and Roughgarden [21] extend the Kleinberg-Tardos rounding to a series of optimization problems in metric space, including graph decomposition, sparse cover and metric triangulation. They cast these metric clustering problems into a unified modeling and algorithm framework, the consistent labeling problem, and make further analyses of the rounding method and necessary revision on some models.

The rounding methods by Kleinberg and Tardos and Calinescu et al. both make multi-dimensional random choices by sequentially implementing Bertsimas et al.'s rounding on one chosen dimension one at a time, though their sequential orders are differently designed. Therefore, it is natural to consider them as generalizations of Bertsimas et al.'s rounding in "time horizon". In this paper we present a new dependent rounding scheme, the geometric rounding, which can be considered as an extension of Bertsimas et al.'s rounding in "space horizon". The key difference of our geometric rounding from the sequential rounding type is that the former uses one randomly generated vector in a standard simplex as a *multi-dimensional threshold* and makes a simultaneous rounding decisions for every point in the simplex by comparing its coordinates with the threshold vector. Our new rounding method possesses a few desired features to be elaborated in the paper: non-sequential or stationary, simple and intuitive, and computationally efficient (linear time in the number of points). By exploring the intrinsic geometric structure and probabilistic distribution in the simplex, we develop in a systematic way known as well as new results for the hub location, metric labeling, winner determination and consistent labeling problems.

Our study also shows that the new geometric rounding and the Kleinberg-Tardos rounding may have some overlap on the scope of the applicability because both are designed for rounding in a simplex with the same cardinality constraints. In particular, algorithms based on the Kleinberg-Tardos rounding and its variants generate solid theoretical bounds for every problem studied in this paper. So we conduct a comprehensive comparison between the new geometric rounding and the Kleinberg-Tardos rounding both theoretically and numerically. An interesting observation emerging from our investigation of dependent rounding schemes is that two rounding methods generate very similar performance guarantees in most studied problems. The Kleinberg-Tardos rounding gains better theoretical bounds in two cases of consistent labeling problems, while the geometric rounding exhibits its advantage on some special cases of the hub location, metric labeling and winner determination problems. Some examples also illustrate that they are essentially different rounding techniques. Our computational tests and simulation results indicate they do not dominate each other while the geometric rounding obtains higher quality solutions more often in comparison to the Kleinberg-Tardos rounding.

The Hub Location and Metric Labeling Problem. The hub location problem [7, 23] is a classical model of hub-and-spoke networks in operations research. In such networks, traffic is routed from cities of origin to specific destinations through hubs. A solution to the problem needs to specify which hubs to open and the allocations of cities to hubs. Demands between two cities have to be routed through hubs to which they are assigned. The cost function includes the linear assignment costs between cities and hubs, the quadratic interhub costs between hubs and the costs of opening hubs. Most published work on the hub location problem mainly focuses on practical heuristics while providing no theoretical bounds [10, 24, 23]. In this paper we first prove that the geometric rounding provides an approximation algorithm with a constant approximation ratio for its special case: the fixed-hub single allocation problem(FHSAP) in which opening a hub is free of charge. The FHSAP is mathematically identical to the well-studied metric labeling problem in computer science community [11, 19, 20] though they are apparently in the context of different applications. Any specific algorithms to the metric labeling problem can naturally transfer to the FHSAP and vice versa. Our further analysis shows that the geometric rounding actually provides a first $\ln n$-approximation for the equilateral-hub case of the hub location problem. It is noteworthy that the hub location problem is quite general and it encompasses many well known optimization problems as its special cases, such as the set covering and nonmetric facility location problems.

The Winner Determination Problem. Consider the assignment problem in a combinatorial auction problem: given k players (or bidders) and m items. Each player is single-minded, i.e., player j is interested precisely in a subset S_j of items. The utility of set T for him is defined to be $v(S_j)$ if $S_j \subseteq T$ and to be zero otherwise. We call S_j the preferred bundle of player j. Provided that the utility function of each player is explicit, the associated *winner determination problem*, also called the social welfare maximization problem, asks for an allocation of the items to the players that maximizes the total valuation. It is well known that this assignment problem is a special case of the packing integer programming problem and that it can not be approximated better than $O(k)$ or $O(\sqrt{m})$ in polynomial time unless P=NP [17, 22]. For the general case where each item has $B(\geq 1)$ copies, Srinivasan [29] proposes an $O(\sqrt[B]{r})$-approximation algorithm where $r(\leq m)$ is the maximal cardinality of the preferred bundles.

In this paper we generalize our geometric rounding to handle multi-assignment constraints and apply it to the multi-unit winner determination problems. We prove that the geometric rounding provides $O(r)$-approximation algorithms for the both single-unit and multi-unit cases, and a new theoretical performance guarantee $\frac{B+k-1}{B}$ for the uniform multi-copy case. Note that the latter is independent of r and is not dominated by the current best bound [29].

The Consistent Labeling Problem. This class of problems in metric spaces includes

Table 1: Approximation factors of two rounding methods and the current best bounds. Ratios provided with unspecified resources are proved in this paper.

Probelm	Geometric Rounding	Kleinberg-Tardos Rounding	Current Best Bound
Metric labeling; FHSAP	$O(\log k)$	$O(\log k)$ [19]	$O(\log k)$ [19]
Metric Labeling; FHSAP: Equilateral	2	2 [19]	2 [19]
Hub Location	$O(\ln n)$	$O(\ln n)$	$O(\ln n)$
Winner Determination	r	r	$\frac{2}{3}r$ [16]
Uniform MU Winner Determination	$\min\{r+1, \frac{B+k-1}{B}\}$	$2r$ [21]	$O(\sqrt[B]{r})$ [29]
Separating Decomposition	2	2 [21]	2 [21]
Padded Decomposition	?	$O(1)$ [21]	$O(1)$ [21]

computing separating and padded decompositions, sparse covers, and metric triangulations. Recently Krauthgamer and Roughgarden [21] find the new applications of the Kleinberg-Tardos rounding in the study of this class of problems. Computing separating and padded decompositions, two fundamental variants of the metric clustering problem, play an important role in metric embedding. Computing sparse covers and metric triangulations can be considered as a second genre of metric clustering problems with the goal to minimize the overlap between clusters subject to covering constraints.

The previous literature on these metric space problems has focused exclusively on absolute (worst-case) bounds which seek guarantees for every possible metric space. By contrast, Krauthgamer and Roughgarden's work emphasizes relative guarantees that compare the produced solutions to the given input and provides significantly better relative guarantees. In this paper we apply our geometric rounding to a subset of metric clustering problems and show that it provides similar relative guarantees.

The results of our geometric rounding, the results based on the Kleinberg-Tardos rounding by different researchers, as well as the best results on approximating aforementioned problems are summarized in Table 1. It seems that the geometric rounding achieves the same quality approximation ratio as the Kleinberg-Tardos rounding on most studied problems. Given the fact that these problems arise from different fields of research, the geometric rounding provides a simple and intuitive alternative for rounding various optimization problems with cardinality constraints.

The outline of the paper is as follows. In Section 2, we present our geometric rounding and general analyses of the method. In Section 3, 4 and 5, we prove the specific approximation ratios for the hub location, winner determination and consistent labeling problems, respectively. Computational experiments and simulation results to compare our geometric rounding and the Kleinberg-Tardos rounding are also performed. Final remarks and open questions are presented in Section 6.

2. The Geometric Rounding Algorithm

In this section we describe our geometric rounding algorithm and some of its general properties. We illustrate the rounding with a generic integer programming system.

Let the input include a set C of n objects and a set H of k labels. The following constrained

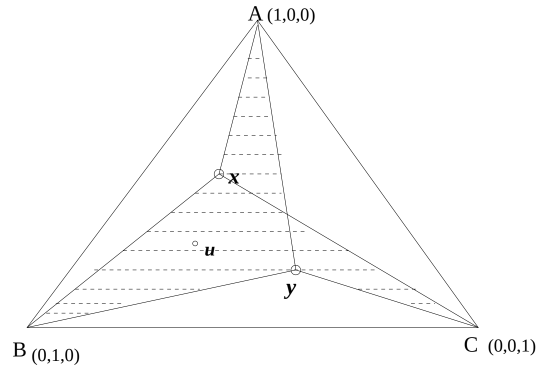

Figure 1: By the geometric rounding, $\hat{x} = (1,0,0)$, $\hat{y} = (0,0,1)$ as the graph indicates.

system encodes that each object is required to be labeled (or assigned) once.

$$\sum_{s \in H} x_{i,s} = 1 \quad \forall i \in C, \tag{1}$$

$$x_{i,s} \in \{0,1\}. \tag{2}$$

In other words, given an optimal solution x^* to its LP relaxation, for each object i, its assignment vector $x_i^* = (x_{i,1}^*, \ldots, x_{i,k}^*)$ is in a standard $(k-1)$ dimensional simplex:

$$\{w \in R^k | w \geq 0, \sum_{i=1}^{k} w_i = 1\}.$$

We denote this simplex by Δ_k.

Therefore, a fractional assignment vector on object i corresponds to a non-vertex point in the Simplex Δ_k. Our goal is to round a fractional vector to a vertex point of Simplex Δ_k, which is of the form:

$$(w \in R^k | w_i \in \{0,1\}, \sum_{i=1}^{k} w_i = 1).$$

It is clear that Δ_k has exactly k vertices. We denote the vertices of Δ_k by $v_1, v_2, \cdots, v_k$, where the ith coordinate of v_i is 1.

For a fractional point $x \in \Delta_k$, connect x with all vertices $v_1, \ldots, v_k$ of Δ_k. Denote the polyhedron which exactly has vertices $\{x, v_1, \ldots, v_{i-1}, v_{i+1}, \ldots, v_k\}$ by $A_{x,i}$. Thus Simplex Δ_k can be partitioned into k polyhedrons $A_{x,1}, \ldots, A_{x,k}$, and any two of these k polyhedrons do not interiorly intersect.

Algorithm: the Geometric Rounding Algorithm (GRA):

1. Solve an LP relaxation of the problem, and get an optimal solution x^*.

2. Generate a random vector u, which follows a uniform distribution in Δ_k.

3. For each $x_i^* = (x_{i,1}^*, \ldots, x_{i,k}^*)$, if u falls into $A_{x_i^*,s}$, let $\hat{x}_{i,s} = 1$; other components $\hat{x}_{i,t} = 0$.

Remark: There are several direct methods to generate a uniform random vector u from the standard simplex Δ_k. We choose the following method: generate k independent unit-exponential random numbers $a_1, ..., a_k$, i.e., $a_i \sim exp(1)$. Then vector u, whose ith coordinate is defined as

$$u_i = \frac{a_i}{\sum_{i=1}^{k} a_i},$$

is uniformly distributed in Δ_k.

Deciding which polyhedron the generated point falls into can be done in linear time by using the following lemma.

Lemma 1. *Given $w = (w_1, w_2, \ldots, w_k) \in \Delta_k$, vector u in Δ_k is in the interior of polyhedron $A_{w,s}$ if and only if $\{s\} = \arg\min_{1\le l\le k}\{\frac{u_l}{w_l}\}$.*

Proof. By symmetry we only need to discuss the case $s = 1$.

If vector u falls into polyhedron $A_{w,1}$, vector u can be written as a convex combination of vertices of $A_{w,1}$. i.e., there exist nonnegative α_i's, such that $\sum_{i=1}^k \alpha_i = 1$ and $u = \alpha_1 w + \sum_{i=2}^k \alpha_i v_i$. It follows that

$$u_1 = \alpha_1 w_1, \quad u_i = \alpha_1 w_i + \alpha_i, \forall i \ge 2.$$

Then, for each $i \ge 2$,

$$\frac{u_i}{w_i} = \frac{\alpha_1 w_i + \alpha_i}{w_i} \ge \alpha_1 = \frac{u_1}{w_1}.$$

If vector u is in the *interior* of polyhedron $A_{w,s}$, then the index set $\arg\min_{1\le l\le k}\{\frac{u_l}{w_l}\}$ is a singleton.

On the other hand, if vector u is in the interior of another polyhedron $A_{w,t}$, then it is easy to see $t \in \arg\min_{1\le l\le k}\{\frac{u_l}{w_l}\}$. This completes the proof. □

Thus the rounding process is "deterministic" once vector u is generated and can be done in $2nk$ operations. The KT rounding can be proved to stop after $O(k \ln n)$ iterations with high probability. So theoretically its time complexity is $O(nk \ln n)$.

2.1 Analysis of the Geometric Rounding

We now present several properties of the geometric rounding. These properties are established from few well known facts of the exponential distribution:

Lemma 2. *The following statements hold.*

- *Assume that $a_1, a_2, \cdots, a_k$ are k independent random variables with $a_i \sim exp(\lambda_i)$. Then for any $1 \le j \le k$,*
$$Pr(a_j = \min_{1\le i\le k} a_i) = \frac{\lambda_j}{\sum_{i=1}^k \lambda_i}.$$
- *If two random variables $Z \sim exp(\mu)$ and $W \sim exp(\lambda)$ are independent, then for any α and β with $0 \le \alpha \le \beta$,*
$$Pr(\alpha Z < W < \beta Z) = \mu(\frac{1}{\mu + \lambda\alpha} - \frac{1}{\mu + \lambda\beta}).$$

The next theorem shows that the geometrically rounded linear cost is optimal in expectation, which is also a fundamental feature of the Kleinberg-Tardos rounding (The sequential rounding by Calinescu et al. does not seem to guarantee this property).

Theorem 3. *For any given $i \in \mathcal{C}, l \in \mathcal{H}$, $E[\hat{x}_{i,l}] = x^*_{i,l}$.*

Proof. For any $\hat{x}_{i,l}$, according to lemma 1,

$$E[\hat{x}_{i,l}] = Pr(\hat{x}_{i,l} = 1) = Pr(u \text{ falls into } A_{x_i^*,l}) = Pr(\frac{u_l}{x_{i,l}^*} = \min_{1\le t\le k}\{\frac{u_t}{x_{i,t}^*}\})$$

Recall that $u_i = \frac{a_i}{\sum_{j=1}^k a_j}$ and $a_i \sim exp(1)$ for any $1 \le i \le k$. This fact implies that $\frac{a_l}{x_{i,l}^*} \sim exp(x_{i,l}^*)$. By Lemma 2, we have

$$Pr(\frac{u_l}{x_{i,l}^*} = \min_{1\le t\le k}\{\frac{u_t}{x_{i,t}^*}\}) = Pr(\frac{a_l}{x_{i,l}^*} = \min_{1\le t\le k}\{\frac{a_t}{x_{i,t}^*}\}) = \frac{x_{i,l}^*}{\sum_{t=1}^k x_{i,t}^*} = x_{i,l}^*.$$

This completes the proof. □

Theorem 3 can also be derived from its geometric structure: the probability that vector u falls into polyhedron $A_{x_i^*,l}$ is equivalent to the volume ratio of polyhedron $A_{x_i^*,l}$ to Simplex Δ_k.

Given a set S of points and a fixed vertex t, the geometric rounding rounds all points in set S to vertex t if random vector u falls into the intersection of $A_{x_i,t}$ for all $i \in S$. The next lemma claims that the polyhedron generated by the intersection of $A_{x_i,t}$ for all $i \in S$, has a single vertex in the interior of Simplex Δ_k. It also constructs an explicit form of the vertex coordinates. Figure 2 depicts such an example of two points.

Lemma 4. *For a fixed t, $1 \le t \le k$, define $v_s = \max_{i\in S}\{\frac{x_{i,s}^*}{x_{i,t}^*}\}$, and $z_s = \frac{v_s}{\sum_{l=1}^k v_l}$ for any s, $1 \le s \le k$. We have*

$$\bigcap_{i\in S} A_{x_i^*,t} = A_{z,t}.$$

Proof. If a point $q \in \bigcap_{i\in S} A_{x_i^*,t}$, according to Lemma 1, we have

$$\frac{q_s}{q_t} \ge \frac{x_{i,s}^*}{x_{i,t}^*},\ \forall i \in S, 1 \le s \le k.$$

Thus,

$$\frac{q_s}{q_t} \ge v_s = \frac{v_s}{v_t} = \frac{z_s}{z_t}.$$

So, $q \in A_{z,t}$.

The converse can be easily proved from that $q \in A_{z,t}$ implies $q \in \bigcap_{i\in S} A_{x_i^*,t}$. □

The next theorem estimates the expected distance of two rounded points. For any x and y, define $d(x,y) := \sum_s |x_s - y_s|$.

Theorem 5. *For any $x, y \in \Delta_k$, if we randomly round x and y to vertices $\hat{x}$ and $\hat{y}$ in Δ_k by the geometric rounding, then we have the following observations:*

1. *Assume x and y are collinear with a vertex, then $E[d(\hat{x},\hat{y})] = d(x,y)$.*
2. *In general, $E[d(\hat{x},\hat{y})] \le 2d(x,y)$.*

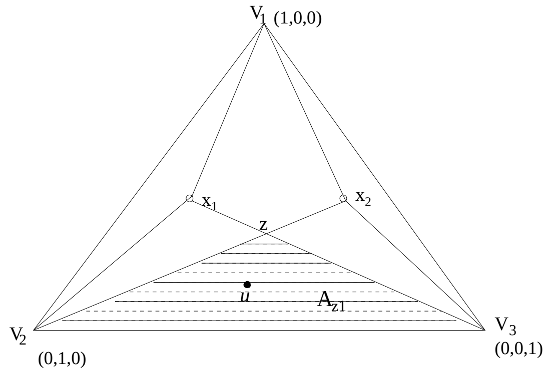

Figure 2: x_1 and x_2 are rounded to V_1 if u falls into $A_{z,1}$.

Proof. (1). Without loss of generality, assume $y = x(s) = (sx_1, sx_2, \ldots, sx_{k-1}, sx_k + (1-s))$, $0 < s < 1$; see Figure 3. We need to prove

$$E[d(\hat{x}, \hat{x}(s))] = d(x, x(s)).$$

By definition, $d(x, x(s)) = 2(1-s)(1-x_k)$. And

$$E[d(\hat{x}, \hat{x}(s))] = 2 * Pr(d(\hat{x}, \hat{x}(s)) \neq 0).$$

Notice that for any i, $1 \leq i \leq k-1$, $\frac{u_i}{x_i} \leq \frac{u_i}{x(s)_i}$. Then, Lemma 1 implies that, given vector u, if $x(s)$ is rounded to vertex v_i, $1 \leq i \leq k-1$, x must be rounded to the same vertex. It follows that the case where $x(s)$ and x are rounded to two different vertices happens only when $x(s)$ is rounded to v_k and x is rounded to a different vertex. In view of Lemma 1, we have

$$\begin{aligned} & Pr(d(\hat{x}, \hat{x}(s)) \neq 0) \\ = \; & Pr(\frac{u_k}{x(s)_k} \leq \min_{1 \leq i \leq k-1}\{\frac{u_i}{x(s)_i}\} \quad \text{and} \quad \frac{u_k}{x_k} > \min_{1 \leq i \leq k-1}\{\frac{u_i}{x_i}\}) \\ = \; & Pr(\frac{a_k}{x(s)_k} \leq \min_{1 \leq i \leq k-1}\{\frac{a_i}{x(s)_i}\} \quad \text{and} \quad \frac{a_k}{x_k} > \min_{1 \leq i \leq k-1}\{\frac{a_i}{x_i}\}), \end{aligned}$$

where the last equality holds because $u_i = \frac{a_i}{\sum_{j=1}^{k} a_j}$ for any $1 \leq i \leq k$. If we define $Z = \min_{1 \leq i \leq k-1}\{\frac{a_i}{x_i}\}$ and $W = \frac{a_k}{x(s)_k}$, then it follows that

$$Pr(d(\hat{x}, \hat{x}(s)) \neq 0) = Pr(\alpha Z < W \leq \beta Z)$$

with $\alpha = \frac{x_k}{sx_k + (1-s)}$ and $\beta = \frac{1}{s}$. Recall that $a_i \sim exp(1)$ for any $1 \leq i \leq k$. Therefore,

$$Z \sim exp(x_1 + x_2 + \ldots + x_{k-1}) = exp(1 - x_k) \quad \text{and} \quad W \sim exp(sx_k + (1-s)).$$

By Lemma 2, we obtain that $Pr(\alpha Z < W \leq \beta Z) = (1-s)(1-x_k)$. Thus, $E[d(\hat{x}, \hat{x}(s))] = 2(1-s)(1-x_k)$.

(2). The claim is a special case of forthcoming Theorem 15 in which each set S has cardinality 2. $\square$

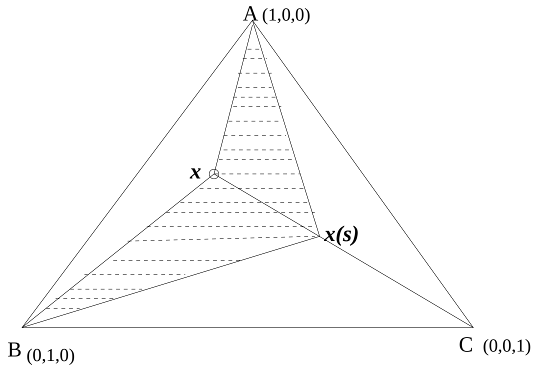

Figure 3: $v_k, x(s), x$ are collinear.

Remark. The bound proved in Theorem 5 is essentially tight. The following observation shows that even rounding two points different only by 2 coordinates may lead to an integrality gap of 2. Define two points x, y in Δ_{k-1}. They only differ on coordinates i and j. Assume $y_i - x_i = x_j - y_j = d > 0$. Then a simple calculation indicates that

$$E[d(\hat{x}, \hat{y})] = \frac{2d}{1+d}(2 - x_i - y_j) = \frac{2 - x_i - y_j}{1+d} d(x, y).$$

3. The Hub Location and Metric Labeling Problems

In this section we first discuss the performance of the geometric rounding on the FHSAP and the metric labeling problem, then extend our discussion to the hub location problem. The discussion illustrates the differences of two rounding methods by a special case and simulation results. The methods we present also provide dependent rounding approaches to the set covering and nonmetric facility location problems.

We first state a quadratic programming formulation for the hub location problem adapted from O'kelly's model [23]. We define a set of potential hubs $H = \{1, 2, \ldots, k\}$ and a set of cities $C = \{1, 2, \ldots, n\}$. Demand d_{ij} to be routed from city i to city j is given. Denote by c_{is} the distance from city i to hub s; by c_{st} the distance from hub s to hub t; and by c_s the cost of opening hub s. $x_{i,s}$ is the assignment variable; and y_s is the decision variable indicating whether hub s is used or not. The formulation is given as follows.

$$\begin{array}{lll}
\text{Minimize} & \displaystyle\sum_{i,j \in C} d_{ij} \left(\sum_{s \in H} c_{is} x_{i,s} + \sum_{t \in H} c_{jt} x_{j,t} + \alpha \sum_{s,t \in H} c_{st} x_{i,s} x_{j,t} \right) + \sum_{s \in H} c_s y_s & \\
\text{Subject to} & \sum_{s \in H} x_{i,s} = 1, & \forall i \in C, \\
 & x_{i,s} \leq y_s, & \forall i \in C, s \in H, \\
 & x_{i,s}, y_s \in \{0, 1\}, & \forall i \in C, s \in H.
\end{array}$$

The first constraints indicate that each city must be assigned to exactly one hub. All coefficients $d_{ij}, c_{is}, c_{jt}, c_{st} \geq 0$, and $c_{st} = c_{ts}$, $c_{ss} = 0$, $\forall i, j \in \mathcal{C}, \forall s, t \in \mathcal{H}$. α is the discount factor and $0 \leq \alpha \leq 1$. Without loss of generality, α can be assumed to be *one*.

3.1 The Fixed-Hub Single Allocation Problem

In the FHSAP, hubs are open and can be used freely, i.e., $c_s = 0$ for every hub s. Consider a special case of the FHSAP in which distances between hubs are uniform. We call it the equilateral-hubs case. Similar to the model in [19], an LP relaxation of the problem can be written as follows.

$$
\begin{aligned}
\text{Minimize } & \sum_{i,j\in C}\sum_{s\in H} c_{is}(d_{ij}+d_{ji})x_{i,s} + \sum_{i,j\in C} d_{ij}y_{i,j} \\
\text{Subject to } & \sum_{s\in H} x_{i,s} = 1, && \forall i \in C, \\
& y_{i,j} = \frac{1}{2}\sum_{s\in H} y_{i,j,s}, && \forall i,j \in C, s\in H, \\
& x_{i,s} - x_{j,s} \le y_{i,j,s}, && \forall i,j \in C, s \in H, \\
& x_{j,s} - x_{i,s} \le y_{i,j,s}, && \forall i,j \in C, s \in H, \\
& x_{i,s}, y_{i,j}, y_{i,j,s} \ge 0, && \forall i \in C, s,t \in H.
\end{aligned}
\tag{3}
$$

This model shows that the FHSAP is mathematically identical to the metric labeling problem defined by Kleinberg and Tardos [19]. In order to solve the original problem, we apply the geometric rounding to the LP relaxation and decide $\hat{y}_{i,j}$ accordingly.

Theorem 6. *The Geometric rounding provides a 2-approximation randomized algorithm for the equilateral-hubs case of the FHSAP.*

Proof. Theorem 3 implies that $E[\hat{x}_{i,l}] = x^*_{i,l}$ for any city i.

Noticing that variable $y_{i,j}$ is actually the half of $d(i,j)$, Theorem 5 implies that $E[\hat{y}_{ij}] \le 2y^*_{i,j}$. □

The geometric rounding bounds the linear cost optimally and the interhub cost by a factor 2 in expectation, which is the same as the Kleinberg-Tardos rounding. Analogous to Kleinberg and Tardos' idea, we can also handle the general FHSAP by probabilistically embedding the graph into the tree topology and applying the geometric rounding. Naturally it provides a $O(\log k)$-approximation algorithm for the general FHSAP [2, 19].

The next example shows that two rounding methods are essentially different although they provide the same theoretical performance for the FHSAP and the metric labeling problem.

Example 1: Consider a 3-hub case. There are two points inside Simplex Δ_3: $x = (1/3, 1/3, 1/3)$ and $y = (0, 1/2, 1/2)$. Lemma 3 implies that $E[d(\hat{x}, \hat{y})] = 2/3 = d(x, y)$ if applying the geometric rounding. A simple calculation indicates that $E[d(\hat{x}, \hat{y})] = 5/6 > 2/3 = d(x, y)$ with the Kleinberg-Tardos rounding applied.

3.2 The Hub Location problem

In this section we mainly focus on the equilateral-hubs case of the hub location problem. An LP relaxation of this case can be derived from system (3) by adding the opening-hub costs and the corresponding constraints. The hub location problem is significantly harder than the FHSAP due to the addition of the opening costs. The nonmetric facility location problem is a special case of the hub location problem in which the interhub costs are removed. And the set

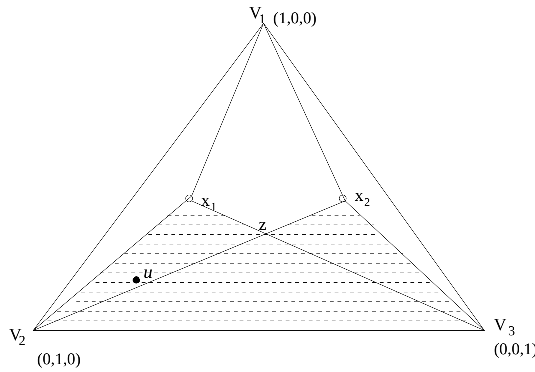

Figure 4: Hub 1 is opened when u falls into the shadow area.

covering problem can be also formulated as a special case with only the opening hub costs as its penalty function. Therefore, even the equilateral-hubs case of the hub location problem cannot be approximated in polynomial time with a factor smaller than $O(\log n)$ unless $P = NP$.

We analyze the performance of the geometric rounding on the opening cost. There always exists an optimal solution to the LP relaxation of the hub location problem satisfying $y_s^* = \max_{i \in C}\{x_{i,s}^*\}$ for all $i \in C$ and $s \in H$. Assume $(x_{i,s}^*, y_s^*)$ is such an optimal solution. Let C_s be the index set including cities who have a positive portion assigned to hub i in the optimal solution above. Thus, $C_s = \{i : i \in C, x_{i,s} > 0\}$.

We estimate the expected opening cost for each hub separately.

Theorem 7. *For each hub s in H, $E[\hat{y}_s] \leq \ln |C_s| \cdot y_s^*$.*

Proof. Without loss of generality, we consider hub 1 and $C_1 = \{1, 2, \cdots, r\}$. Suppose $|C_1| = r$.

Recall the process of generating a uniformly distributed point in the geometric rounding. The algorithm first generates k independent exponentially distributed random variables with parameter 1. Denote them by u_i, $1 \leq u \leq k$.

If hub 1 is opened, equivalently $\hat{y}_1 = 1$. It means that at least one city in C_1 is rounded to vertex 1 in Δ_k (See Figure 4 for an illustration). This case happens if and only if there exists some $i \in C_1$, such that $\frac{u_1}{u_s} \leq \frac{x_{i,1}^*}{x_{i,s}^*}$, $\forall s \in H$.

Equivalently,

$$u_1 \leq \max_{i \in C_1} \min_{s \geq 2, s \in H} u_s \frac{x_{i,1}^*}{x_{i,s}^*}.$$

We also know $u_s \sim exp(1)$, $\forall s \in H$. Thus,

$$
\begin{aligned}
E[\hat{y}_1] &= \int_{R_+^k} I_{(u_1 \le \max_{i\in C_1} \min_{s\ge 2, s\in H} u_s \frac{x_{i,1}^*}{x_{i,s}^*})} dF(u_1, u_2, \ldots, u_k) \\
&= 1 - \int_{R_+^k} I_{(u_1 > \max_{i\in C_1} \min_{s\ge 2, s\in H} u_s \frac{x_{i,1}^*}{x_{i,s}^*})} dF(u_1, u_2, ..., u_k) \\
&= 1 - \int_{R_+^k} I_{(u_1 > \min_{s\ge 2, s\in H} u_s \frac{x_{i,1}^*}{x_{i,s}^*}, \forall i\in C_1)} dF(u_1, u_2, ..., u_k) \\
&= 1 - \int_{R_+} Pr(u_1 > \min_{s\ge 2, s\in H} u_s \frac{x_{i,1}^*}{x_{i,s}^*}, \forall i \in C_1) dF(u_1) \\
&= 1 - \int_{R_+} Pr(u_1 > \min_{s\ge 2, s\in H} u_s \frac{1-x_{i,1}^*}{x_{i,s}^*} \cdot \frac{x_{i,1}^*}{1-x_{i,1}^*},\ \forall i \in C_1) dF(u_1) \\
&\le 1 - \int_{R_+} Pr(v_i < \alpha u_1, \forall i \in C_1) dF(u_1) \\
&\le 1 - \int_{R_+} (Pr(v_1 < \alpha u_1) \cdot Pr(v_2 < \alpha u_1) ... Pr(v_r < \alpha u_1)) dF(u_1) \\
&= 1 - \int_{R_+} e^{-u_1} (1 - e^{-\alpha u_1})^r du_1
\end{aligned}
$$

where

$$
\alpha = \frac{1 - \max_{i\in C_1} x_{i,1}^*}{\max_{i\in C_1} x_{i,1}^*}, \quad v_i = \min_{s\ge 2, s\in H} u_s \frac{1-x_{i,1}^*}{x_{i,s}^*} \sim exp(1).
$$

The first inequality comes from the fact that $\alpha \le \frac{1-x_{i,1}^*}{x_{i,1}^*}$ for any i. The second inequality is proved by Lemma 8 below. Also, because v_i is independent to u_1, we have that

$$
Pr(v_i < \alpha u_1) = 1 - e^{-\alpha u_1}.
$$

Then, the approximation ratio satisfies

$$
\frac{E[\hat{y}_1]}{y_1^*} = (\alpha + 1) E[\hat{y}_1] \le \int_{R_+} (\alpha+1) e^{-u_1} \left(1 - (1 - e^{-\alpha u_1})^r\right) du_1.
$$

By changing variables $y = 1 - e^{-\alpha u_1}$, $\beta = 1/\alpha$, the right side of the above inequality becomes

$$
\begin{aligned}
& \int_{0\le y\le 1} (1+\alpha)(1-y)^\beta (1-y^r) \beta (1-y)^{-1} dy \\
= & \int_{0\le y\le 1} [y^{r-1} + \ldots + 1](1+\beta)(1-y)^\beta dy.
\end{aligned}
$$

Let $z = (1-y)^{1+\beta}$, and $\gamma = 1/(1+\beta)$, the above value is equal to

$$
\int_{0\le z\le 1} [(1 - z^\gamma)^{r-1} + \ldots + 1] dz.
$$

Noticing that this integral is increasing on γ, and the max value is reached at $\gamma = 1$, the bound is

$$\int_{0\le z\le 1}[(1-z)^{r-1}+\ldots+1]dz = \sum_{i=1}^{r} 1/i \approx \ln r.$$

□

Lemma 8. *For each fixed value u_1,*

$$Pr(v_i \le \alpha u_1, \forall i \in C_1) \ge \prod_{i\in C_1} Pr(v_i \le \alpha u_1).$$

Proof. The lemma can be derived from the following recursion:

$$Pr(v_1 \le \alpha u_1, \max_{i\ge 2, i\in C_1} v_i \le \alpha u_1) \ge \quad Pr(v_1 \le \alpha u_1) Pr(\max_{i\ge 2, i\in C_1} v_i \le \alpha u_1).$$

Notice that for any event A_1, A_2 in a probability space, the inequality

$$Pr(A_1 \cap A_2) \ge Pr(A_1)Pr(A_2)$$

is equivalent to

$$1 - Pr(A_1) - Pr(A_2) + Pr(A_1 \cap A_2) \ge 1 - Pr(A_1) - Pr(A_2) + Pr(A_1)Pr(A_2).$$

While the second inequality is equivalent to

$$Pr(A_1^c \cap A_2^c) \ge Pr(A_1^c)Pr(A_2^c).$$

Thus it suffices to prove the inequality:

$$Pr(v_1 \ge \alpha u_1, \max_{i\ge 2, i\in C_1} v_i \ge \alpha u_1) \ge Pr(v_1 \ge \alpha u_1) Pr(\max_{i\ge 2, i\in C_1} v_i \ge \alpha u_1).$$

The above inequality is equivalent to:

$$Pr(\max_{i\ge 2, i\in C_1} v_i \ge \alpha u_1 | \ v_1 \ge \alpha u_1) \ge Pr(\max_{i\ge 2, i\in C_1} v_i \ge \alpha u_1).$$

We prove this inequality by induction. Recall the definition of v_i, we first consider the probability conditioning on one variable u_s for any $s \in H$.

For arbitrary positive reals a, b, we want to prove:

$$Pr(\max_{i\ge 2, i\in C_1} v_i \ge a | \ u_s \ge b) \ge Pr(\max_{i\ge 2, i\in C_1} v_i \ge a).$$

By the memoryless property of exponential distribution, the distribution of vector u conditioning on $u_s \ge b$ is the same as $u + be^s$, where e^s is a zero vector except that the sth coordinate is 1. If we view v as function of u, for each j we have that $v_i(u + be^s) \ge v_i(u)$, therefore

$$Pr(\max_{i\ge 2, i\in C_1} v_i \ge a | \ u_s \ge b) = Pr(\max_{i\ge 2, i\in C_1} v_i(u + be^s) \ge a) \ge Pr(\max_{i\ge 2, i\in C_1} v_i \ge a).$$

The inequality above can be easily generalized to prove:

$$Pr(\max_{i\ge 2, i\in C_1} v_i \ge \alpha u_1 | \ v_1 \ge \alpha u_1) \ge Pr(\max_{i\ge 2, i\in C_1} v_i \ge \alpha u_1).$$

□

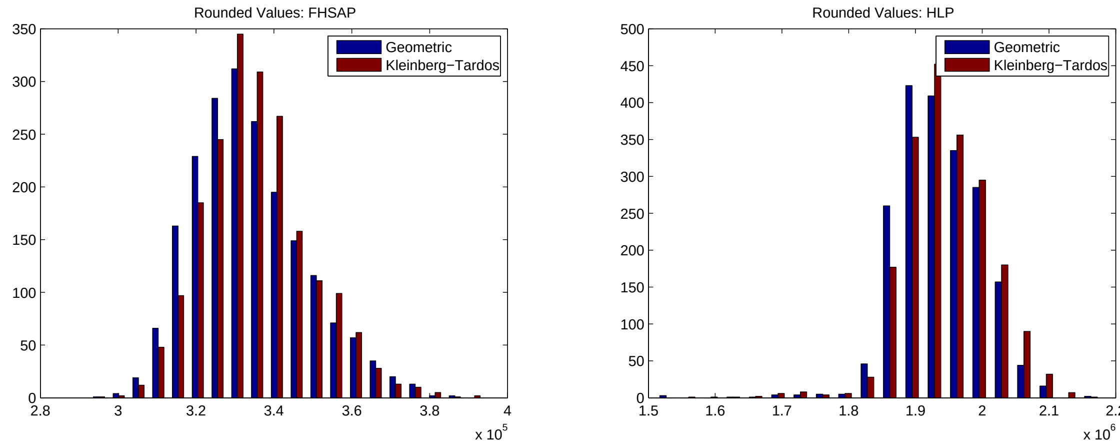

Figure 5: The frequency distribution graph of the rounded solutions for two rounding schemes on the FHSAP and hub location problems.

The allocation and interhub costs of the objective function are up bounded by a factor of 2. So the geometric rounding is actually an $\ln n$-approximation algorithm for equilateral-hubs case. Theorem 7 shows that it is particularly effective for the instances whose LP relaxations have sparse optimal solutions.

We are also interested in the performance of the Kleinberg-Tardos rounding on the hub location problem. The next theorem shows that with high probability it is a $O(\ln n)$-approximation algorithm.

Theorem 9. *Assume that $\hat{y}_s$ is in a rounded feasible solution by the Kleinberg-Tardos rounding. With high probability, the rounding procedure stops after $O(k \ln n)$ rounds, and $E[\hat{y}_s] \leq O(\ln n) \cdot y_s^*$.*

Proof. For any city $i \in C$, the probability that it is allocated to some hub in a round is at most $(x_{i,1} + \cdots + x_{i,n})/k = 1/k$. So the probability that it is not allocated after $ck \ln n$ rounded is at most $(1 - \frac{1}{k})^{ck \ln n} \leq (\frac{1}{e})^{c \ln n} = \frac{1}{n^c}$ for any positive constant c.

Similarly, the probability that hub s is opened in a round is at most y_s/k. So

$$
\begin{aligned}
E[\hat{y}_s] &= Pr(\hat{y}_s = 1) = 1 - Pr(\hat{y}_s = 0) \leq 1 - (1 - \frac{y_s^*}{k})^{ck \ln n} \\
&\leq 1 - (1 - y_s^* \cdot c \ln n) = c \ln n \cdot y_s^*.
\end{aligned}
$$

□

3.3 Simulation Results

Computational experiments are conducted to compare the time efficiency and solution quality of the geometric rounding and the Kleinberg-Tardos rounding. In the experiment we randomly generate 8 medium-sized instances for the FHSAP and hub location problems. The results are presented in Table 2. We test problems of two different sizes: 50 cities and 10 hubs; 30 cities and 30 hubs. The ratio of the shortest edge to the longest edge between hubs is presented in the

Table 2: A comparison of two rounding schemes on the FHSAP and hub location problems.

Type	Size	Ratio	LB	Geometric			Kleinberg-Tardos		
				Best	Avg	Time	Best	Avg	Time
HLP	(50,10)	0.1	1.059	1.151	1.259	4.09	1.156	1.263	5.66
HLP	(50,10)	0.5	1.273	1.372	1.478	4.07	1.380	1.487	5.44
HLP	(30,30)	0.1	2.670	3.378	3.854	1.78	3.423	3.851	2.09
HLP	(30,30)	0.5	3.665	3.790	4.608	1.77	3.783	4.612	2.12
FHSAP	(50,10)	0.1	7.772	9.209	9.704	1.12	9.209	9.749	1.09
FHSAP	(50,10)	0.5	1.163	1.254	1.319	1.14	1.258	1.331	1.10
FHSAP	(30,30)	0.1	2.072	2.589	2.914	0.92	2.603	2.924	1.44
FHSAP	(30,30)	0.5	3.570	4.035	4.340	0.94	4.041	4.384	1.47

table as “Ratio”. For each instance, we run 2000 times of the rounding procedures on the same LP relaxation. Table 2 records the minimum value, the average value of these 2000 rounded solutions and the total running time of the rounding procedure. The lower bound is the optimal solution of the LP relaxation. Table 2 shows that both rounding methods have a very similar performance on all tested instances. The geometric rounding gains a marginal advantage over the Kleinberg-Tardos rounding on both the running time and solution quality on most instances. Figure 5 lists the frequency distribution of the values of 2000 rounded solutions for an instance of the FHSAP and an instance of the hub location problem separately. It can be observed that the solutions generated by the geometric rounding are slightly more concentrated on the “less cost” side.

4. The Winner Determination Problem

In this section we discuss the winner determination problem. Two technical methods are developed to adapt the geometric rounding to the multi-assignment constraints. Particularly, we prove that the geometric rounding for the uniform multi-unit winner determination problem generates a new theoretical performance guarantee that is not dominated by the current best bound [29].

We start the winner determination problem in a single-minded combinatorial auction. In the auction a set of players, $P = \{1, 2, \cdots, k\}$ and a set of items, $I = \{1, 2, \cdots, m\}$, are given. Each player is interested in precisely one subset of items. Each item has only one copy. A feasible assignment allocates each item to at most one player. The problem can be described by a well-known integer program [27].

$$\begin{array}{lll} \text{maximize} & \sum_{j \in P} v(S_j) x_{S_j} & \\ \text{subject to} & \sum_{\forall j: i \in S_j} x_{S_j} \leq 1, & \forall i \in I, \\ & x_{S_j} \in \{0, 1\}, & \forall j \in P. \end{array}$$

By introducing assignment variable $x_{i,j}$ that indicates whether item i is assigned to player

j or not, we derive a new IP formulation:

$$\begin{aligned} \text{maximize} \quad & \sum_{j \in P} v(S_j) x_{S_j} \\ \text{subject to} \quad & \sum_{j \in P} x_{i,j} = 1, \qquad \forall i \in I, \\ & x_{S_j} \le x_{i,j}, \qquad \forall i \in S_j, j \in P, \\ & x_{S_j}, x_{i,j} \in \{0,1\}, \quad \forall i \in I, j \in P. \end{aligned}$$

In general, player j gets his preferred bundle S_j if and only if vector u generated in the rounding falls into the intersection of all $A_{x_i^*,j}$'s for every item i in S_j. Thus, the probability of this event turns out to be the volume of a specific polyhedron in a high dimensional space.

Define r to be the maximal cardinality of S_j, $j \in P$. We have the following theorem.

Theorem 10. *The geometric rounding provides a feasible allocation to the winner determination problem in a single-minded combinatorial auction. Additionally, the solution achieves at least* $\max\{\frac{1}{r}, \frac{1}{k-1}\}$ *of the optimal value.*

Proof. It suffices to prove $E[\hat{x}_{S_j}] \ge \max\{\frac{1}{r}, \frac{1}{k-1}\} x^*_{S_j}$.

Without loss of generality, consider player 1. The event that $\hat{x}_{S_1} = 1$ means that for every $i \in S_1$, vector x_i^* is rounded to vertex 1. So for the geometric rounding algorithm, it is the probability that vector u falls into the intersected polyhedron of all $A_{x_i^*,1}$ for $i \in S_1$.

Lemma 4 shows that this polyhedron only has one vertex z in the interior of Simplex Δ_k. Suppose vector z is the point. We only need to prove

$$z_1 \ge \max\{\frac{1}{k-1}, \frac{1}{r}\} \min_{i \in S_1} x^*_{i,1}.$$

Let $a = \min_{i \in S_1} x^*_{i,1}$. The inequality becomes

$$\frac{v_1}{\sum_{j=1}^n v_j} \ge \max\{\frac{a}{k-1}, \frac{a}{r}\}.$$

It is equivalent to

$$\sum_{j=1}^n v_j \le \min\{\frac{k-1}{a}, \frac{r}{a}\}.$$

For all $j \ge 2$, $x^*_{i,1} \ge a$, so $\frac{x^*_{i,j}}{x^*_{i,1}} \le \frac{1-a}{a}$. It implies $v_j \le \frac{1-a}{a}$.
Recall that $v_1 = 1$, we have

$$\sum_{1 \le j \le n} v_j \quad \le \quad 1 + (k-1)\frac{1-a}{a} = \frac{k-1}{a} - (k-2) \le \frac{k-1}{a}.$$

Also, noticing that $v_j \le \sum_{i \in S_1} \frac{x^*_{i,j}}{x^*_{i,1}}$, we have

$$\sum_{1 \le j \le n} v_j \le \sum_{1 \le j \le n} \sum_{i \in S_1} \frac{x^*_{i,j}}{x^*_{i,1}} = \sum_{i \in S_1} \frac{1}{x^*_{i,1}} \le \frac{r}{a}.$$

□

The best known approximation ratio for this case is $\frac{2}{3}r$ according to [3, 16]. Our rounding provides an LP-based algorithm and a quick effective approach according to the simulation result. We will also need to use this result to tackle the uniform multi-copy case.

Interestingly, we can prove that the Kleinberg-Tardos rounding also provides a solution with a similar performance guarantee.

Theorem 11. *The solution generated by Kleinberg-Tardos rounding also achieves at least* $\max\{\frac{1}{r}, \frac{1}{k}\}$ *of the optimal value.*

Proof. We need to prove that $E[\hat{x}_{S_j}] \geq \max\{\frac{1}{r}, \frac{1}{k}\}x^*_{S_j}$.

First, the probability that all points i in S are rounded to vertex j at the first step by the Kleinberg-Tardos rounding is at least $\geq \frac{1}{k}\min_{i\in S} x^*_{i,j}$. So $E[\hat{x}_{S_j}] \geq \frac{1}{k}x^*_{S_j}$.

Second, we know that $E[\hat{x}_{S_j}]$ is equal to the probability that all points in set S are rounded to vertex j by the Kleinberg-Tardos rounding. So

$$E[\hat{x}_{S_j}] \geq \frac{1}{k}\min_{i\in S_j} x^*_{i,j} + \sum_{l\in P}\frac{1}{k}(1-\max_{i\in S_j} x^*_{i,l})P_j = \frac{1}{k}x_{S_j} + (1-\frac{1}{k}\sum_{l\in P}\max_{i\in S_j} x^*_{i,l})E[\hat{x}_{S_j}].$$

It is easy to observe that

$$\sum_{l\in P}\max_{i\in S_j} x^*_{i,l} \leq r.$$

Therefore,

$$E[\hat{x}_{S_j}] \geq \frac{1}{r}x^*_{S_j}.$$

□

Remark: Although both rounding methods present the similar theoretical bounds again, they may differ in some cases. Still, consider the example we give for the FHSAP. There are two points inside Δ_3: $x = (1/3, 1/3, 1/3)$ and $y = (0, 1/2, 1/2)$. Assume $S_3 = \{x, y\}$. Then $E[\hat{x}_{S_3}] = 1/3$ by the geometric rounding. If we apply the Kleinberg-Tardos rounding, it is easy to prove that $E[\hat{x}_{S_3}] = 7/24 < 1/3$.

Next we examine the multi-unit case. In this case each item may have multiple copies, that is, $B_i \geq 1$ for each i in I. And these copy numbers may be different. Each player needs at most one copy of an item. We show how to combine the bin-packing technique to handle the multi-assignment problem though the theoretical bound it gains is weaker than the one by Srinivasan [29].

This problem can be formulated as follows.

$$\begin{aligned}
\text{maximize} \quad & \sum_{j\in P} v(S_j)x_{S_j} \\
\text{subject to} \quad & \sum_{j\in P} x_{i,j} = B_i, && \forall i \in I, \\
& x_{S_j} \leq x_{i,j}, && \forall i \in S_j, j \in P, \\
& x_{i,j} \leq 1, && \forall i \in I, j \in P, \\
& x_{S_j}, x_{i,j} \in \{0,1\}, && \forall i \in I, j \in P.
\end{aligned}$$

If $B_i > 1$, the multi-assignment constraint, $\sum_{j\in P} x^*_{i,j} = B_i$, restricts the implementation of the geometric rounding. In order to work around this issue, we revise the rounding by integrating the bin packing technique. The basic idea is to pack all $x^*_{i,j}$'s into B_i unit-volume bins. Each unit-volume bin corresponds to a point in Δ_k whose jth coordinate is $x^*_{i,j}$ if $x^*_{i,j}$ is in this bin and is 0 if not. Next, run the geometric rounding to allocate items to players by rounding these newly created points in Δ_k.

The bin packing idea may not be feasible for some i if $x^*_{i,j}$'s are indivisible. One remedy is to shrink each $x^*_{i,j}$ by half. Denote this shrinked $x^*_{i,j}$ by $x'_{i,j}$. A simple greedy algorithm will guarantee that these $x'_{i,j}$'s can be encapsulated into at most B_i unit-volume bins. Thus we can make a partition of P, $P = P_{i1} \cup P_{i2} \cup \ldots \cup P_{iB_i}$, such that

1. P_{ik}'s are mutually disjoint.

2. $\sum_{j \in P_{ik}} x'_{i,j} \leq 1, \forall 1 \leq k \leq B_i$.

If $\sum_{j \in P_{ik}} x'_{i,j} < 1$, we can stretch any nonzero $x^*_{i,j}$ in P_{ik} to increase the sum to 1. It is easy to see that this modification only increases the quality of the rounded solution. Thus this partition generates B_i points in simplex Δ_k for item i.

Now we apply the geometric rounding to the multi-assignment case with the bin packing idea.

Algorithm: The multi-assignment GRA.

1. Solve the LP relaxation of the multi-assignment case to get an optimal fractional solution x^*.

2. For item i, make the bin packing partition and map $P_{i1}, P_{i2}, \cdots, P_{iB_i}$ to B_i points in Δ_k as described above.

3. Generate a random vector u, which follows a uniform distribution on Δ_k.

4. Run the rounding in the same fashion as the Geometric rounding for all newly created points in Δ_k.

This algorithm provides a general approach to allocation problems with multi-assignment constraints. By following a similar proof of Theorem 10, we have the following theorem.

Theorem 12. *The multi-assignment geometric rounding algorithm yields a feasible assignment to the multi-unit case of the winner determination problem. Moreover, the solution recovers at least* $\max \frac{1}{2}\{\frac{1}{r}, \frac{1}{k-1}\}$ *of the optimal value.*

Factor $\frac{1}{2}$ comes from the fact we shrink the size of each $x^*_{i,j}$ by half during the bin packing and it is not hard to improve. Theorem 10 and 12 show that the geometric rounding generates a quality performance guarantee for the sparse combinatorial allocation problem in which each player is only interested in a small set of items. The analysis for maximum consistent labeling in Lemma 3.1 in Krauthgamer and Roughgarden's work [21] also implies that Kleinberg-Tardos rounding has the same theoretical performance for the multi-unit case.

4.1 The Uniform Multi-Unit Winner Determination Problem

In this section we present a sequential geometric rounding scheme to solve the uniform multi-unit winner determination problem. The solution that our algorithm generates recovers at least a $\max\{\frac{B}{B+k-1}, \frac{1}{1+r}\}$ fraction of the optimal value.

In a uniform multi-unit case, $\sum_{j\in P} x^*_{i,j} = B$ implies $(x^*_{i,1}/B, x^*_{i,2}/B, \ldots, x^*_{i,k}/B) \in \Delta_k, \quad \forall i \in I$. Our algorithm sequentially run B rounds of the geometric rounding to allocate items to players. The allocation is decided by rounding x^*/B with a new generated threshold in each round. The algorithm is stated as follows.
Algorithm: The uniform multi-assignment GRA

1. Solve the LP relaxation of the uniform multi-unit winner determination problem to get an optimal fractional solution x^*. Define $x' = x^*/B$.

2. Choose a sequence of independent random vectors $u^1, u^2, \ldots, u^B$, each of which is uniformly distributed on Δ_k. Run B rounds of the geometric rounding for x' by using u^l at round l.

Intuitively, when B gets large, the possibility that a player gets his preferred bundle will increase in our algorithm. Define binary variable $\hat{x}^l_{i,j}$ to indicate whether the ith item is allocated to player j in round l. Naturally the assignment variable $\hat{x}_{i,j} = 1$ if and only if there exists some l, $1 \le l \le B$, such that $\hat{x}^l_{i,j} = 1$. Let $\hat{x}_{S_j} = \min_{i\in S_j} \hat{x}_{i,j}$. The following lemma gives the estimation of the approximation ratio for each round.

Lemma 13.

$$Pr(\min_{i\in S_j} \hat{x}^l_{i,j} = 1) \ge \max\{\frac{1}{rB}, \frac{1}{k-1}\} \min_{i\in S_j} x^*_{i,j}$$

Proof. Theorem 10 directly implies that $Pr(\min_{i\in S_j} \hat{x}^l_{i,j} = 1) \ge \frac{1}{r} \min_{i\in S_j}\{\frac{x^*_{i,j}}{B}\}$.

For the other part of the bound, we use the same concepts in Theorem 10. Noticing that $v_j \le \frac{1-a}{a}$ for any $j \ge 2$ and $v_1 = 1$, the proof is essentially the same as the one for Theorem 10. □

Theorem 14. *The uniform multi-assignment GRA provides a feasible assignment to the uniform multi-unit WDP. Additionally, the solution recovers at least* $\max\{\frac{B}{B+k-1}, \frac{1}{1+r}\}$ *of the optimal value.*

Proof. At each round one copy of each item is assigned to some player, so the uniform multi-assignment GRA always makes a feasible assignment after B rounds.

It suffices to show that for every j,

$$E[\hat{x}_{S_j}] \ge \max\{\frac{B}{B+k-1}, \frac{1}{1+r}\} x^*_{S_j}.$$

First,

$$\begin{aligned} E[\hat{x}_{S_j}] = Pr(\hat{x}_{S_j} = 1) = Pr(\min_{i\in S_j} \hat{x}_{i,j} = 1) &= 1 - Pr(\min_{i\in S_j} \hat{x}_{i,j} = 0) \\ &= 1 - \prod_{l=1}^{B} Pr(\min_{i\in S_j} \hat{x}^l_{i,j} = 0). \end{aligned}$$

The last inequality uses the fact that u^l are independent and

$$\min_{i\in S_j} \hat{x}_{i,j} = 0 \;\Rightarrow\; \min_{i\in S_j} \hat{x}^l_{i,j} = 0, \quad \forall 1 \le l \le B.$$

Furthermore, we have

$$\prod_{l=1}^{B} Pr(\min_{i\in S_j} \hat{x}^l_{i,j} = 0) = \prod_{l=1}^{B}(1 - Pr(\min_{i\in S_j} \hat{x}^l_{i,j} = 1)) = (1 - Pr(\min_{i\in S_j} \hat{x}^l_{i,j} = 1))^B.$$

Here again we use the independence of u^l.
Define $c = \max\{\frac{1}{rB}, \frac{1}{k-1}\}$.

$$Pr(\min_{i\in S_j} \hat{x}_{i,j} = 1) \geq 1 - (1 - Pr(\min_{i\in S_j} \hat{x}^l_{i,j} = 1))^B \geq 1 - (1 - c\cdot \min_{i\in S_j} x^*_{i,j})^B.$$

Therefore, for any j,

$$\frac{E[\hat{x}_{S_j}]}{x^*_{S_j}} \geq \frac{1 - (1 - c\cdot \min_{i\in S_j} x^*_{i,j})^B}{x^*_{S_j}} \geq \frac{1 - (1 - c\cdot \min_{i\in S_j} x^*_{i,j})^B}{\min_{i\in S_j} x^*_{i,j}}.$$

However, for any $t \in [0, 1]$ (with $0 \leq c \leq 1$),

$$\frac{1 - (1 - ct)^B}{t} = \frac{1 - (1 - ct)}{t}\sum_{i=0}^{B-1}(1 - ct)^i \geq c\sum_{i=0}^{B-1}(1 - c)^i = 1 - (1 - c)^B \geq \frac{cB}{1 + cB},$$

where the last inequality follows from the fact that

$$(1 - c)^B(1 + cB) \leq (1 - c)^B(1 + c)^B \leq 1.$$

Recall that $c = \max\{\frac{1}{rB}, \frac{1}{k-1}\}$, the theorem follows. □

4.2 Simulation Results

Figure 6 and Table 3 show the simulated performance for a comparison of the two rounding methods on the winner determination problem. We test 24 randomly generated instances in 8 groups. Each group of problems have the same size but different unit number B. For each case, we run 1000 times of the rounding procedures separately for two methods. Table 3 records the maximum and average ratios of the rounded value to the LP relaxation value and the total running time of the rounding procedure. We can observe from Table 3 that both rounding methods perform similarly for the single unit case. The Kleinberg-Tardos rounding seems to outperform the geometric rounding when the unit number B is between 2 and 4, while the geometric rounding performs better once B is above 5. There is no theory to explain this interesting scenario yet. We also notice that the geometric rounding outperforms the Kleinberg-Tardos rounding for large unit number case at the expense of significantly longer running time.

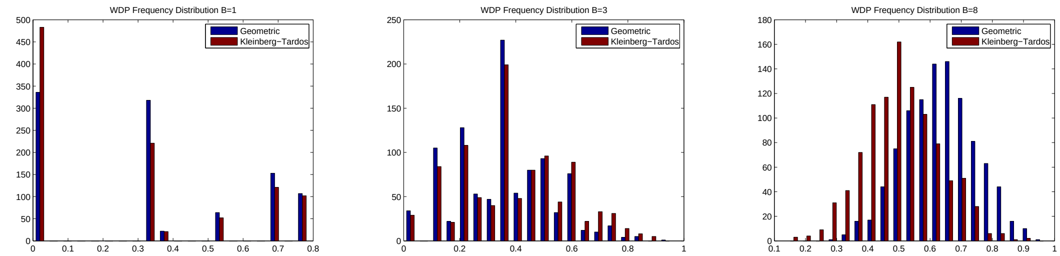

Figure 6: The frequency distribution graph of the rounded solutions for two rounding schemes on the winner determination problem.

Table 3: A comparison of two rounding schemes on the winner determination problem.

size	B=1						B=3						B=8					
(k, m, r)	Geometric			Kleinberg-Tardos			Geometric			Kleinberg-Tardos			Geometric			Kleinberg-Tardos		
	best	avg	time	best	avg	time	best	avg	time	best	avg	time	best	avg	time	best	avg	time
(20, 100, 12)	0.7122	0.3397	2.02	0.7122	0.2636	0.76	0.9083	0.3476	9.94	0.9749	0.3852	0.87	0.9719	0.5771	24.17	0.8517	0.4940	0.85
(50, 50, 12)	0.7190	0.3091	1.22	0.7190	0.2978	0.96	0.8993	0.3310	5.57	0.9104	0.4025	0.70	0.7657	0.4150	13.5	0.7850	0.4609	0.59
(100, 20, 12)	0.7258	0.3981	0.69	0.7258	0.3801	0.97	0.8704	0.4414	2.75	0.9201	0.4991	0.72	0.9053	0.5549	8.12	0.7971	0.5189	0.62
(20, 100, 40)	0.8031	0.3720	1.99	0.8031	0.2817	0.78	0.8034	0.3550	9.8	0.8371	0.4010	0.70	0.9405	0.6317	27.5	0.9017	0.5054	0.75
(50, 50, 24)	0.6924	0.2755	1.2	0.6924	0.2551	1.05	0.6943	0.3171	5.56	0.8079	0.4003	0.70	0.7340	0.4491	15.91	0.7605	0.4852	0.59
(100, 20, 12)	0.8861	0.3735	0.69	0.8861	0.3862	1.01	0.8736	0.4906	3.03	0.8604	0.4955	0.76	0.8581	0.6028	8.13	0.7863	0.5307	0.63
(100, 10, 6)	0.9200	0.8112	0.41	0.9200	0.8110	0.73	1	0.72	1.79	0.9163	0.5886	0.54	0.9593	0.7494	4.90	0.7646	0.5643	0.46
(10, 100, 41)	0.9162	0.4379	1.88	0.9162	0.3478	0.43	0.9721	0.5047	9.71	0.9721	0.4883	0.57	1	0.8673	0.2644	1	0.6879	0.93

5. The Consistent Labeling Problem

This section discusses the relative guarantees of the geometric rounding on metric clustering problems. Krauthgamer and Roughgarden [21] have formulated integer programming models for these problems as variants of the consistent labeling problem and developed LP-based approximation algorithms. Our discussion will mainly focus on the performance analysis of the geometric rounding based on the linear programming relaxations they developed.

The formulations presented in this section are introduced in Krauthgamer and Roughgarden's work. We will follow their system exactly for the convenience of readers. First consider a generic consistent labeling model. Given a set A of objects, a set L_a of allowable labels for each object a (drawn from a ground set L), and a collection $\mathcal{C}$ of subsets of A. Each set $S \in \mathcal{C}$ has a nonnegative weight w_S. A feasible labeling is an assignment of each object a to a subset of L_a. Two main objects here are to minimize the number of labels assigned to each object, and to maximize the number (or total weight) of sets that are consistently labeled, meaning that a common label is assigned to all of the objects in the set. The following constraints are common to all our relaxations:

$$1 \le \sum_{i\in L} x_{ai} \le k, \quad \forall a \in A, \tag{4}$$
$$y_{iS} \le x_{ai}, \quad \forall S \in \mathcal{C}, i \in L, a \in S, \tag{5}$$
$$z_s \le \sum_{i\in L} y_{iS}, \quad \forall S \in \mathcal{C}, \tag{6}$$
$$z_s \le 1, \quad \forall S \in \mathcal{C}, \tag{7}$$
$$x_{ai} = 0, \quad \forall a \in A, i \notin L_a. \tag{8}$$

5.1 Maximum Consistent Labeling

The maximum consistent labeling *(MAX CL)* problem looks for a feasible labeling that assigns at most k labels to every object and maximizes the total weight of the consistently labeled sets. Its LP relaxation model is

$$\text{Maximize} \sum_{S \in \mathcal{C}} w_S z_S \tag{9}$$

subject to (4)–(8).

Another variant, the Maximum Fair Consistent Labeling problem, maximizes the minimum weighted probability that a set S is labeled consistently over $S \in \mathcal{C}$ with the same constraints. In this case, the LP relaxation maximizes a decision variable α subject to (4)–(8) and

$$w_S z_S \geq \alpha \quad \text{for every set } S \in \mathcal{C}. \tag{10}$$

Define $f_{max} = \max_{S \in \mathcal{C}} |S|$. Theorem 3.2 in Krauthgamer and Roughgarden's work proves that the Kleinberg-Tardos rounding generates a $1/2f_{max}$-approximation algorithm for MAX CL and MAX FAIR CL. By an argument similar to Theorem 14, we can easily prove that the geometric rounding generates the same approximation ratio.

The next problem we consider is minimizing the inconsistency probability using the linear program below, which is an LP relaxation for computing separating decomposition.

$$\begin{array}{lrl} \text{Minimize} & \alpha & \\ \text{Subject to} & \sum_{i \in L} x_{ai} = 1, & \forall a \in A, \\ & y_{iS} \geq x_{ai} - x_{a'i}, & \forall S \in \mathcal{C}, a, a' \in S, i \in L, \\ & z_S \geq \frac{1}{|S|} \sum_{i \in L} y_{iS}, & \forall S \in \mathcal{C}, \\ & x_{ai} = 0, & \forall a \in A, i \notin L_a, \\ & \alpha \geq w_S z_S, & \forall S \in \mathcal{C}. \end{array}$$

The next theorem encompasses Lemma 3.3 and Theorem 3.4 in Krauthgamer and Roughgarden's work, while the analysis based on the geometric rounding is quite different.

Theorem 15. *For every set $S \in \mathcal{C}$, Pr(S is not consistently labeled) $\leq |S| z_S^*$. Therefore, the geometric rounding provides a 2-approximation algorithm for computing a separating decomposition.*

Proof. We first define $r_i = max_{a \in S} x_{ai}^*$; $s_i = min_{a \in S} x_{ai}^*$; and $P_i = Pr$(all objects in S are rounded to label i).

The inequality we need to prove is equivalent to

$$1 - \sum_{i \in L} P_i \leq \sum_{i \in L} y_{iS}^* = \sum_{i \in L} (r_i - s_i).$$

If $u_j / r_j > u_i / s_i$, $\forall j \neq i, j \in L$, then every object in S is rounded to label i.

This implies

$$\begin{aligned} P_i &\geq Pr(u_j/r_j > u_i/s_i, \forall j \neq i) = Pr(\min_{j \in L} \frac{u_j}{r_j} s_i \geq u_i) \\ &= \int_{R^+} e^{-\frac{u_i \sum_{j \in L} r_j}{s_i}} du_i = s_i / \sum_{j \in L} r_j. \end{aligned}$$

Thus $1 - \sum_{i \in L} P_i \leq 1 - \frac{\sum_{i \in L} s_i}{\sum_{i \in L} r_i} \leq \sum_{i \in L} (r_i - s_i)$. □

5.2 Padded Decomposition

Computing a padded decomposition can be modeled as a consistent labeling problem in the same way as for a separating decomposition. The difference is the collection $\mathcal{C}$ of consistency sets is not all pairs of points, but rather all balls of radius Δ/β. Krauthgamer and Roughgarden design an LP-based approximation algorithm in which they apply the iterations of the Kleinberg-Tardos rounding only n times. In order to apply the geometric rounding, one direct approach is to keep the structure of their algorithm and replace the rounding part by the geometric rounding. It is straightforward to get an $O(n)$ approximation. In order to improve the approximation ratio, the algorithm needs to leave some points un-labeled during the rounding. One possible remedy is to leave some points un-labeled in the rounding process. This needs the geometric rounding to handle the case where cardinality constraints are inequalities. Check the padded decomposition algorithm designed by Krauthgamer and Roughgarden. We will keep its structure and replace the rounding part by the geometric rounding, i.e., we redesign Step $3-5$ as follows.

1. Define x'_{ij}, y'_{ij} as the half value of the original x_{ij}, y_{ij}.
2. Consider the geometric rounding in Simplex Δ_{n+2} by adding an artificial node 0. Define $x'_{0j} = 1 - \sum_{i \in X} x'_{ij}$ $(\geq 1/2)$. Similarly define $y'_{0j} = 1 - \sum_{i \in X} y'_{ij}$.
3. Round each j by the geometric rounding on Simplex Δ_{n+2} with probability distribution $(y'_{0,j}, y'_{1,j}, \cdots, y'_{n,j})$. A point is rounded to vertex 0 means that it is not assigned to any cluster.

The revised algorithm obviously improves the chance a point is un-labeled in the rounding process. We conjecture that it gives a constant approximation for the padded decomposition problem.

Krauthgamer and Roughgarden also discuss minimum consistent labeling problems. Those problems seek solutions that consistently label a prescribed fraction of the sets while using as few labels as possible. Since the geometric rounding labels each item once in each round while the Kleinberg-Tardos rounding "consumes" labels in a more efficient way, it seems hard for the geometric rounding to match the theoretical bounds developed by the Kleinberg-Tardos rounding.

6. Remarks and Open Questions

In this paper we have developed a non-sequential, intuitive, and computationally efficient geometric rounding method that simultaneously rounds multiple points in a multi-dimensional simplex to its vertices, and established in a systematic way many known as well as new results for various optimization problems with integral assignment constraints. A comprehensive comparison with the well known dependent randomized rounding method developed by Kleinberg and Tardos [19] and its variants is also conducted, both in theory and in numerical simulation. The results show that the two methods perform similarly on most studied problems, although they outperform each other in few different cases. Overall, the geometric rounding provides a simple and effective alternative for rounding various integer optimization problems.

There are some intriguing and challenging questions left open to explore. One is to prove the constant-approximation conjecture for the padded decomposition. Another is derandomization of the geometric rounding. Any derandomization method of the geometric rounding would be likely linked to computing volumes of high-dimensional polyhedrons, which only admits PTAS algorithms. Finally, since the geometric rounding method is based on a simple geometry, the higher moment correlations among assignment variables may be estimable, which may make the method applicable to approximating optimization problems with high order polynomial objective functions.

References

[1] A. A. Ageev and M. I. Sviridenko. Pipage rounding: a new method of constructing algorithms with proven performance guarantee. *J. Comb. Optim.*, 8, 2004.

[2] 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, Washington, DC, USA, 1996. IEEE Computer Society.

[3] P. Berman and P. Krysta. Optimizing misdirection. In *SODA '03: Proceedings of the fourteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 192–201, Philadelphia, PA, USA, 2003. Society for Industrial and Applied Mathematics.

[4] D. Bertsimas, C. Teo, and R. Vohra. Analysis of lp relaxations for multiway and multicut problems. *Networks*, 34(2):102–114, 1999.

[5] D. Bertsimas, C. Teo, and R. Vohra. On dependent randomized rounding algorithms. *Oper. Res. Lett.*, 24(3):105–114, 1999.

[6] D. Bertsimas and R. Vohra. Rounding algorithms for covering problems. *Math. Prog.*, 80(1):63–89, 1998.

[7] D. Bryan and M. E. O'Kelly. Hub-and-spoke networks in air transportation: An analytical review. *J. Reg. Sci.*, 39(2):275–295, May 1999.

[8] G. Calinescu, R. Chekuri, and J. Vondrk. Maximizing a submodular set function subject to a matroid constraint. Technical report, Proc. of 12 th IPCO, 2006.

[9] Gruia Calinescu, Howard Karloff, and Yuval Rabani. An improved approximation algorithm for multiway cut. In *STOC '98: Proceedings of the thirtieth annual ACM symposium on Theory of computing*, pages 48–52, New York, NY, USA, 1998. ACM.

[10] J. F. Campbell. Hub location and the p-hub median problem. *Oper. Res.*, 44, 1996.

[11] C. Chekuri, S. Khanna, J. Naor, and L. Zosin. Approximation algorithms for the metric labeling problem via a new linear programming formulation. In *Proceedings of the twelfth annual ACM-SIAM symposium on Discrete algorithms(SODA)*, 2001.

[12] W. Fernandez de la Vega and G. S. Lueker. Bin packing can be solved within $1+\epsilon$ in linear time. *Combinatorica*, 1:349–355, 1981.

[13] R. Gandhi, S. Khuller, S. Parthasarathy, and A. Srinivasan. Dependent rounding and its applications to approximation algorithms. *J. ACM*, 53(3):324–360, 2006.

[14] M. X. Goemans and D. P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *J. ACM*, 42:1115–1145, 1995.

[15] M. X. Goemans and D. P. Williamson. New $\frac{3}{4}$-approximation algorithms for the maximum satisfiability problem. *SIAM J. Discrete Math.*, 7(4), 1995.

[16] M. Halldórsson. Approximations of independent sets in graphs. In *APPROX '98: Proceedings of the International Workshop on Approximation Algorithms for Combinatorial Optimization*, pages 1–13, London, UK, 1998. Springer-Verlag.

[17] J. Hastad. Clique is hard to approximate within $n^{1-\epsilon}$. *Acta Math.*, 182, 1999.

[18] D. S. Hochbaum. Tapproximation algorithms for the set covering and vertex cover problems. *SIAM J. Comput.*, 11(3):555–556, 1982.

[19] J. Kleinberg and É. Tardos. Approximation algorithms for classification problems with pairwise relationships: Metric labeling and markov random fields. In *FOCS '99: Proceedings of the 40th Annual Symposium on Foundations of Computer Science*, page 14, Washington, DC, USA, 1999. IEEE Computer Society.

[20] J. Kleinberg and É. Tardos. Approximation algorithms for classification problems with pairwise relationships: metric labeling and markov random fields. *J. ACM*, 49(5):616–639, 2002.

[21] R. Krauthgamer and T. Roughgarden. Metric clustering via consistent labeling. In *SODA '08: Proceedings of the nineteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 809–818, Philadelphia, PA, USA, 2008. Society for Industrial and Applied Mathematics.

[22] D. J. Lehmann, L. I. O'Callaghan, and Y. Shoham. Truth revelation in approximately efficient combinatorial auctions. In *ACM Conference on Electronic Commerce*, pages 96–102, 1999.

[23] M. E. O'Kelly. A quadratic integer program for the location of interacting hub facilities. *Eur. J. Oper. Res.*, 33, 1987.

[24] M. E. O'Kelly and J. Skorin-Kapov D. Skorin-Kapov. Lower bounds for the hub location problem. *Manage. Sci.*, 41, 1995.

[25] P. Raghavan. Probabilistic construction of deterministic algorithms: approximating packing integer programs. *J. Comput. Syst. Sci.*, 37(2):130–143, 1988.

[26] P. Raghavan and C.D. Thompson. Randomized rounding: a technique for provably good algorithms and algorithmic proofs. *Combinatorica*, 7, 1987.

[27] S.Dobzinski, N. Nisan, and M. Schapira. Approximation algorithms for combinatorial auctions with complement free bidders. In *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing(STOC)*, pages 610–618, 2005.

[28] D. B. Shmoys and É. Tardos. An approximation algorithm for the generalized assignment problem. *Math. Prog.*, 62(3):461–474, 1993.

[29] A. Srinivasan. An extension of the lovász local lemma, and its applications to integer programming. In *SODA '96: Proceedings of the seventh annual ACM-SIAM symposium on Discrete algorithms*, pages 6–15, Philadelphia, PA, USA, 1996. Society for Industrial and Applied Mathematics.

[30] A. Srinivasan. Distributions on level-sets with applications to approximation algorithms. In *FOCS '01: Proceedings of the 42nd IEEE symposium on Foundations of Computer Science*, page 588, Washington, DC, USA, 2001. IEEE Computer Society.