



Available at
www.ElsevierComputerScience.com
POWERED BY SCIENCE @ DIRECT®
Journal of Algorithms 50 (2004) 194–214

**Journal of
Algorithms**

www.elsevier.com/locate/jalgor

Cooperative facility location games [☆]

Michel X. Goemans ^{a,*,1} and Martin Skutella ^{b,2}

^a *Department of Mathematics, Room 2-351, MIT, Cambridge, MA 02139, USA*

^b *TU Berlin, Fakultät II, Institut für Mathematik, MA 6-1, Straße des 17 Juni 136, 10623 Berlin, Germany*

Received 6 July 2000

Abstract

The location of facilities in order to provide service for customers is a well-studied problem in the operations research literature. In the basic model, there is a predefined cost for opening a facility and also for connecting a customer to a facility, the goal being to minimize the total cost. Often, both in the case of public facilities (such as libraries, municipal swimming pools, fire stations, . . .) and private facilities (such as distribution centers, switching stations, . . .), we may want to find a ‘fair’ allocation of the total cost to the customers—this is known as the cost allocation problem. A central question in cooperative game theory is whether the total cost can be allocated to the customers such that no coalition of customers has any incentive to build their own facility or to ask a competitor to service them. We establish strong connections between fair cost allocations and linear programming relaxations for several variants of the facility location problem. In particular, we show that a fair cost allocation exists if and only if there is no integrality gap for a corresponding linear programming relaxation; this was only known for the simplest unconstrained variant of the facility location problem. Moreover, we introduce a subtle variant of randomized rounding and derive new proofs for the existence of fair cost allocations for several classes of instances. We also show that it is in general NP-complete to decide whether a fair cost allocation exists and whether a given allocation is fair.

© 2003 Elsevier Inc. All rights reserved.

Keywords: Facility location; Cooperative games; LP relaxation; Randomized rounding; Core

[☆] The research was performed when the authors were at C.O.R.E., Louvain-la-Neuve, Belgium.

* Corresponding author.

E-mail addresses: goemans@math.mit.edu (M.X. Goemans), skutella@mpi-sb.mpg.de (M. Skutella).

¹ Supported in part by NSF contract 9623859-CCR.

² Supported by DONET within the frame of the TMR Programme (contract number ERB FMRX-CT98-0202).
Current affiliation: Max-Planck Institut für Informatik, Stuhlsatzenhausweg 85, 66123 Saarbrücken, Germany.

1. Introduction

1.1. Fair cost allocations

In the facility location problem, customers from a given set N are in need of a certain service which can be provided by connecting them to facilities. These facilities could be railway stations, sports complexes, libraries, switching stations, antennas, or supermarkets, to cite a few examples. From a given set F of possible locations for the facilities, one has to decide first which facilities to open (build) and then each customer must be assigned (connected) to an open facility. Opening facility $i \in F$ causes a fixed cost $f_i \geq 0$ and the cost for connecting customer $j \in N$ to this facility is denoted by $c_{ij} \geq 0$. We refer to this problem as the *unconstrained* (or *uncapacitated*) facility location problem; this problem is also sometimes referred to as the uncapacitated plant location problem.

In many situations, further constraints have to be taken into consideration. The facilities can typically handle only a limited number of customers, say at most k_i customers for facility i . Certain customers cannot be assigned to certain facilities (for example, if they are geographically too far apart from each other); this can be handled in the original unconstrained model by simply letting the corresponding cost c_{ij} be very large (or infinite). Other relevant constraints can occur when the set of customers is heterogeneous and some quota must be met: at least some fraction of the customers connected to a facility must belong to a certain subgroup or minority (e.g., at least 40% walloons, 33% socio-democrats, an equal number of women and men, etc.). Or, members of different groups cannot be assigned to the same facility (e.g., each facility represents a factory that can produce only one product and customers request one of the different products).

From a central authority's point of view, it is interesting to ask for a cheapest possible solution, i.e., to minimize the total cost which is made up by the cost to build facilities and to connect the customers to the open facilities.

This facility location problem has attracted much attention in the operations research literature, see, for example, the book of Mirchandani and Francis [26].

1.2. Facility location games

We can also ask whether the total cost can be allocated to the different customers in a fair way. This is known as the cost allocation problem. For example, towns would pay for the building of libraries, or sports complexes, but they do not want to pay more than their fair share of the total cost, whatever that means. In the area of cooperative game theory, see, for example, [28], fairness means that no group of customers, or *coalition*, has any incentive to break apart and obtain the service on their own. In other words, if v_j denotes the price being paid by customer j , we would like that $\sum_{j \in S} v_j \leq \text{OPT}(S)$ where S is any subset of customers and $\text{OPT}(S)$ represents the cost of providing the service only to the customers in S . The core of the cooperative game is then defined as

$$\text{core} = \left\{ v: \sum_{j \in N} v_j = \text{OPT}(N) \right. \\ \left. \sum_{j \in S} v_j \leq \text{OPT}(S) \quad \text{for all } S \subseteq N \right\} \quad (1)$$




Fig. 1. An instance with two facilities (squares) and three customers a , b , and c (circles). The cost for opening a facility is 1 and the connection costs are given by the distances in the drawn graph. Notice that the function $\text{OPT}(\cdot)$ is not submodular since $\text{OPT}(\{a, b, c\}) + \text{OPT}(\{b\}) > \text{OPT}(\{a, b\}) + \text{OPT}(\{b, c\})$. The unique element in the core is given by $v_a = v_c = 2$ and $v_b = 1$.

and a central question in cooperative game theory is whether the core is non-empty, and if so, how to find an allocation vector in the core. Traditionally, the non-vacuity of the core is established by showing that the game is *balanced* (for definitions, see [28]). In linear programming terms, this boils down to showing that any extreme point of the dual to $\max \sum_{j \in N} v_j$ subject to (1) has value at least $\text{OPT}(N)$. It is well known and easy to check that, for a submodular function $\text{OPT}(\cdot)$ (i.e., $\text{OPT}(S) + \text{OPT}(T) \geq \text{OPT}(S \cup T) + \text{OPT}(S \cap T)$ for all $S, T \subseteq N$), the core is non-empty and the Shapley value (cf. [28]) lies in the core. However, it is also known that this condition is in general not satisfied for the facility location game under consideration. In Fig. 1 we give a facility location instance for which the function $\text{OPT}(\cdot)$ is not submodular; nevertheless, the core is non-empty in this example.

1.3. Related LP relaxations

When the core is empty, we would like instead to recover as much as possible and maximize $\sum_{j \in N} v_j$ subject to the constraints (1). Observe that this is a linear program (LP) but the constraints defining it are not only exponential in number but also not well characterized since the right-hand side value $\text{OPT}(S)$ is an NP-hard quantity for general facility location problems. However, the value of this linear program is a lower bound on the optimum value $\text{OPT}(N)$ and thus it can be viewed as a relaxation of the problem.

In this paper, for any kind of constrained facility location problem, we show how to derive an equivalent relaxation in the natural space of variables which contains a variable y_i denoting whether facility i is open and a variable x_{ij} denoting whether customer j is assigned to facility i . In general, our result thus says that the core is non-empty if and only if this *canonical* LP relaxation has no integrality gap for the objective function being considered, i.e., the optimum LP value is equal to $\text{OPT}(N)$. This result and the canonical LP relaxation are described in Section 2.

For the unconstrained facility location problem, this canonical relaxation turns out to simply be a classical LP relaxation of the problem, a result first derived by Kolen [22]. Kolen shows that the dual of this classical LP relaxation for the uncapacitated case can be interpreted as the maximum amount that can be allocated to the customers subject to the constraints (1) and, as a result, the core is non-empty if and only if there is no integrality gap for this relaxation. Our approach is similar and allows to derive (implicitly, and in some cases, explicitly) this canonical relaxation for any facility location problem. Chardaire [8] also considers the classical LP relaxation of the uncapacitated facility location problem and shows that the associated game can be viewed as a linear production game that satisfies

Owen's additivity assumption [31]. As a result, he derives that the core of this linear programming game is non-empty, and hence the core of the original uncapacitated facility location game is non-empty if the LP relaxation has no integrality gap. Furthermore, he uses Kolen's result to derive the converse statement. Chardaire also generalizes some of his results to variants of the capacitated facility location problem but does not characterize the canonical relaxation whose absence of integrality gap is a necessary and sufficient condition for non-emptiness of the core. We refer the reader to Chardaire for a discussion of related work of Sharkey [34] on a similar revenue game and of Göthe-Lundgren, Jörnsten, and Värbrand [17] on a set partitioning game.

Connections between the core and certain LP relaxations have also been found for other cooperative games related to problems in combinatorial optimization. Deng, Ibaraki, and Nagamochi [11] consider general packing and covering games given through integer linear programs (ILP) with a $\{0, 1\}$ -matrix and right-hand side all ones; they show that the core of such a game is non-empty if and only if the natural LP relaxation of the ILP has an integer optimal solution. Moreover, they give applications of this general result for various games on graphs related to well-known combinatorial optimization problems, such as maximum flow, maximum matching, coloring, and others; see also [12]. The result of Kolen [22] on the unconstrained facility location game can also be derived from this more general framework since the unconstrained facility location problem can be formulated as such a covering problem, see, e.g., [1] (cited in [19]) or [23]. Further results in this direction have, for example, been derived by Faigle and Kern [15] and Samet and Zemel [33].

In contrast to the earlier results mentioned above, we develop our results and techniques in a setting which allows us to handle not only the unconstrained facility location problem but also arbitrarily constrained variants of it. Moreover, in comparison to Kolen [22] and Chardaire [8], our proof is more straightforward. For general facility location problems, we can in certain cases give an explicit representation of the canonical relaxation in terms of linear inequalities. However, even if we are unable to completely characterize the relaxation in terms of linear inequalities, we can nevertheless find a fair allocation that maximizes the amount recovered provided that we can find the best feasible assignment of customers to a single facility. This is discussed in Section 3.

1.4. Non-emptiness of the core for special cases

In Section 4, we provide new proofs showing that the core of the unconstrained facility location game is non-empty for two special cases, one in which the facilities can be ordered on a line and the connection costs are unimodal (i.e., first decreasing and then increasing), and the other when the facilities are positioned on a tree and the connection costs are obtained by applying a nondecreasing function (depending on the customer) to the tree metric. These results have been obtained earlier by Trubin [37] (see also [19]) and also by Kolen [22] (see also [10,23]).

Our main contribution in this section is a different proof technique. We use randomized rounding to show that the canonical relaxation has no integrality gap for these special classes of instances. The randomized rounding is performed in a dependent way by assigning subsets of $[0, 1]$ to each facility and to each connection of customers to facilities. This algorithmic proof technique is of independent interest. For example, it can be applied

to simplify a recent result of Bar-Noy et al. [5]. Or it can be used to derive the result of Shmoys and Tardos [35] on the generalized assignment problem without the need for explicitly solving a matching problem, in a way more similar to [25]. For the case of facility location on a line with unimodal connection costs, our randomized rounding approach is closely related to a technique introduced by Bertsimas, Teo, and Vohra [7] to prove the integrality of an LP formulation of the uncapacitated lot-sizing problem.

1.5. Complexity results

In Section 5, we show that even for the unconstrained facility location problem, testing whether there is an integral optimal solution (or no integrality gap) to the canonical LP relaxation is actually NP-complete. Our proof also yields that checking whether a given cost allocation is in the core is an NP-complete problem. On the other hand, if the core is known to be non-empty, the latter problem can be solved in polynomial time and an element of the core can also be computed in polynomial time.

Deng, Ibaraki, and Nagamochi [11] derive similar results for different games on graphs (see Table 1 in their paper); surprisingly, there seems to be no direct correlation between the complexity of those questions for cooperative games and the complexity of the corresponding combinatorial optimization problems. Further results in this direction have been obtained, e.g., by Deng and Papadimitriou [13] and by Faigle et al. [14].

1.6. Known results on the size of the integrality gap

Since for the uncapacitated facility location problem the optimal value of the classical LP relaxation is equal to the maximum amount that can be recovered in the cost allocation problem, results on the worst case ratio of the integrality gap gain a new meaning in the context of cost allocation. It follows for example from the LP-based approximation result of Chudak [9] that, for metric instances (i.e., when the costs c_{ij} arise from a metric on $N \cup F$), there always exists a fair cost allocation that recovers at least a fraction $e/(e+2) \approx 0.576$ of the total cost. On the other hand, Guha and Khuller [21] give a class of metric instances where at most 68% of the total cost can be recovered. This gives worst-case bounds on the amount that the central authority should subsidize in order to ensure the existence of a fair allocation. However, for general cost functions, there exist instances for which the amount one can recover is at most $\text{OPT}(N) 2n/((n+1) \log_2(n+1))$ (using a standard reduction from the set cover problem and using instances with large integrality gaps for the set cover problem, see [38]).

1.7. Extension to a closely related problem

A problem that is closely related to the facility location problem occurs when connections of customers to facilities do not cause costs c_{ij} but produce certain non-negative benefits b_{ij} . Here, a customer can be connected to *at most* one facility and the goal is to maximize the total benefit minus the cost for building facilities. In the corresponding cooperative game we ask for a fair allocation of this amount to the customers, i.e., each

coalition of customers wants to get at least as much as it could gain on its own. The results in this paper can easily be carried forward to this setting.

2. Integer and linear programming formulations

In order to model the facility location problem we introduce two types of binary variables: for each $i \in F$, the variable y_i is 1 if facility i is opened, and 0 otherwise; for each $i \in F$ and $j \in N$, the variable x_{ij} is 1 if customer j is connected to facility i , and 0 otherwise. A minimum cost solution to the basic version of the facility location problem is then given by the following integer linear program:

$$\begin{aligned} \text{minimize} \quad & \sum_{i \in F} f_i y_i + \sum_{\substack{i \in F \\ j \in N}} c_{ij} x_{ij}, \\ \text{subject to} \quad & \sum_{i \in F} x_{ij} = 1 \quad \text{for all } j \in N, \quad (2) \\ & y_i - x_{ij} \geq 0 \quad \text{for all } i \in F, j \in N, \quad (3) \\ & x_{ij}, y_i \in \{0, 1\} \quad \text{for all } i \in F, j \in N. \end{aligned}$$

Constraints (2) ensure that every customer is connected to exactly one facility. A customer can only be connected to a facility that is open by constraints (3). For the capacitated version of the problem we add the following capacity constraints:

$$k_i y_i - \sum_{j \in N} x_{ij} \geq 0 \quad \text{for all } i \in F. \quad (4)$$

If at least a fraction q_i of the customers connected to facility i have to belong to a subgroup $N' \subseteq N$, we add the quota constraints

$$(1 - q_i) \sum_{j \in N'} x_{ij} - q_i \sum_{j \in N \setminus N'} x_{ij} \geq 0 \quad \text{for all } i \in F. \quad (5)$$

If the set of customers is partitioned into subsets N_p , $p = 1, \dots, l$, and a facility can only serve customers in at most one subset, we add the incompatibility constraints

$$x_{ij} + x_{ik} \leq y_i \quad \text{for all } i \in F, \text{ all } j \in N_p, \text{ and all } k \in N_q, p \neq q. \quad (6)$$

Notice that constraints (4)–(6) do not introduce a coupling between different facilities but can be expressed solely in terms of the variables x_{ij} and y_i for each fixed facility i . We consider a more general class of constraints where, for each facility i , we are given a family of subsets $S \subseteq N$ of customers that can be connected to this facility. In this case we can rewrite the integer program as: minimize $\sum_i f_i y_i + \sum_{i,j} c_{ij} x_{ij}$ subject to $\sum_i x_{ij} = 1$ for all $j \in N$ and $(y_i, x_i) \in P_i$ for all $i \in F$. Here, $x_i := (x_{i1}, \dots, x_{in})$, with $n := |N|$, and $P_i \subseteq \{0, 1\}^{n+1}$ is given by

$$P_i := \{(0, \dots, 0)\} \cup \{(1, \chi_S) : S \subseteq N \text{ feasible for } i\},$$

where $\chi_S \in \{0, 1\}^n$ denotes the characteristic vector of the subset S .

There are several possible ways of deriving a linear programming relaxation for this problem. The most natural would be to try to replace each discrete set P_i by its convex hull $\text{conv}(P_i)$. Notice that the value of the resulting linear program might not be equal to $\text{OPT}(N)$ since the intersection of the convex hulls with the hyperplanes (2) is not necessarily the convex hull of the intersections. A slightly weaker relaxation would be to replace each P_i by its conic hull $\text{cone}(P_i) = \{\sum_{x \in P_i} \lambda_x x : \lambda_x \geq 0\}$. Given the special form of P_i , it is easy to see that $\text{conv}(P_i) = \text{cone}(P_i) \cap \{(y, x) : y \leq 1\}$. This leads to the following relaxation (LP):

$$\begin{aligned}
 \text{(LP)} \quad & \text{minimize} \quad \sum_{i \in F} f_i y_i + \sum_{\substack{i \in F \\ j \in N}} c_{ij} x_{ij}, \\
 & \text{subject to} \quad \sum_{i \in F} x_{ij} = 1 \quad \text{for all } j \in N, \\
 & \quad (y_i, x_i) \in \text{cone}(P_i) \quad \text{for all } i \in F.
 \end{aligned} \tag{7}$$

We now turn to the cost allocation problem. For each coalition $S \subseteq N$, let $\text{OPT}(S)$ denote the minimum cost of the facility location problem restricted to the set of customers S . The maximum cost that can be allocated to the customers is then given by the following linear program (CAP):

$$\begin{aligned}
 \text{(CAP)} \quad & \text{minimize} \quad \sum_{j \in N} v_j, \\
 & \text{subject to} \quad \sum_{j \in S} v_j \leq \text{OPT}(S) \quad \text{for all } S \subseteq N.
 \end{aligned} \tag{8}$$

It is an easy observation that the amount v_j that is paid by customer j in an optimal cost allocation is always nonnegative (since $c_{ij} \geq 0$ implies $\text{OPT}(S) \leq \text{OPT}(S \cup \{j\})$). Although there are exponentially many constraints and although it is in general NP-hard to compute the right-hand side of this linear program, we show that in some cases it can be solved in polynomial time. To obtain this result we develop the following connection to the LP relaxation of the facility location problem introduced above.

Theorem 2.1. *The cost allocation problem (CAP) is equivalent to the dual of the LP relaxation (LP) of the facility location problem. In particular, their values are equal and the core is non-empty if and only if there is no integrality gap for the relaxation (LP) of the facility location problem.*

Proof. We dualize constraints (7) and introduce a vector v of corresponding dual variables v_j for all customers $j \in N$. This leads to the following program of the same value as (LP) by strong duality; see, e.g., [29, Section II.3.6]:

$$\begin{aligned}
 \max_{v, x, y} \quad & \sum_{i \in F} f_i y_i + \sum_{\substack{i \in F \\ j \in N}} (c_{ij} - v_j) x_{ij} + \sum_{j \in N} v_j, \\
 \text{subject to} \quad & (y_i, x_i) \in \text{cone}(P_i) \quad \text{for all } i \in F.
 \end{aligned}$$

For a fixed vector v , the program is decomposed into the sum of $n = |N|$ linear minimization problems over pointed cones. Therefore, the inner minimization problem is either unbounded or an optimum solution is given by $x_{ij} = y_i = 0$, for all $i \in F$, $j \in N$, and has value $\sum_j v_j$. Moreover, since $\text{cone}(P_i)$ is generated by the incidence vectors of feasible assignments to facility i , the inner minimization problem is unbounded if and only if there exists an $i \in F$ and a corresponding feasible coalition $S \subseteq N$, i.e., $(1, \chi_S) \in P_i$, with

$$f_i + \sum_{j \in S} (c_{ij} - v_j) < 0.$$

Thus, we can rewrite the program as

$$\begin{aligned} &\text{minimize} && \sum_{j \in N} v_j, \\ &\text{subject to} && \sum_{j \in S} v_j \leq f_i + \sum_{j \in S} c_{ij} \quad \text{for all } i \in F \text{ and } S \subseteq N \text{ feasible.} \end{aligned} \tag{9}$$

Since the right-hand side of constraints (8) is stronger than the right-hand side of (9), it remains to show that (8) is implicitly contained in (9). Suppose that an optimal solution to the subproblem induced by S is given by

$$\text{OPT}(S) = \sum_{i \in F'} \left(f_i + \sum_{j \in S_i} c_{ij} \right),$$

where $F' \subseteq F$ and the S_i , $i \in F'$, form a partition of S . Now observe that (8) can be obtained by simply aggregating constraints (9) over $i \in F'$ for $S = S_i$. This completes the proof. \square

Let us turn to describing $\text{cone}(P_i)$ for a few special cases. For the unconstrained facility location problem, the conic hull of the set P_i , is given by $\text{cone}(P_i) = \{(y_i, x_i) : 0 \leq x_{ij} \leq y_i \text{ for all } j \in N\}$; this yields the following classical LP relaxation which has been introduced by Balinski [4]:

$$\begin{aligned} &\text{minimize} && \sum_{i \in F} f_i y_i + \sum_{\substack{i \in F \\ j \in N}} c_{ij} x_{ij}, \\ &\text{subject to} && \sum_{i \in F} x_{ij} = 1 \quad \text{for all } j \in N, \\ &&& 0 \leq x_{ij} \leq y_i \quad \text{for all } i \in F, j \in N. \end{aligned} \tag{10}$$

Notice that there exists an optimal solution to this LP relaxation with $y_i \leq 1$ for all $i \in F$. In this case, having replaced P_i by $\text{cone}(P_i)$ rather than $\text{conv}(P_i)$ does not matter. This is no longer true if we turn to the capacitated version of the facility location problem including constraints (4); in this case we get

$$\text{cone}(P_i) = \left\{ (y_i, x_i) : \sum_j x_{ij} \leq k_i y_i \text{ and } 0 \leq x_{ij} \leq y_i \text{ for all } j \in N \right\}$$

such that we only have to add the constraints (4) to the above LP relaxation.

Consider an example with two facilities of capacity 1 and fixed cost 0 and two customers that are located at the first facility; they can be connected to this facility for free while the connection to the second facility costs 1. In an optimal solution one of the customers has to be connected to the second facility which causes cost 1; however, an optimal LP solution has value 0 since it can open the first facility with $y_1 = 2$ and connect both customers to this facility. If we add the constraint $y_1 \leq 1$ to the LP relaxation, its optimum value increases to 1.

In fact, in this example, none of the customers is willing to pay anything for the service since he can argue that he could connect to the first facility for free. Therefore, the cost allocation problem is equivalent to the dual of the weak LP relaxation but not to the dual of the stronger relaxation including constraints $y_i \leq 1$ for all $i \in F$, which would have been obtained if we had relaxed P_i to $\text{conv}(P_i)$. This answers an open question discussed by Chardaire [8].

If, instead of capacity constraints, we have that each facility must serve the same number of customers from N_1 and from N_2 , then we simply have that

$$\text{cone}(P_i) = \left\{ (y_i, x_i): \sum_{j \in N_1} x_{ij} = \sum_{j \in N_2} x_{ij} \text{ and } 0 \leq x_{ij} \leq y_i \text{ for all } j \in N \right\}.$$

If we have incompatibility constraints $x_{ij} + x_{ik} \leq y_i$ for certain pairs $(j, k) \in E_i$ as in (6), we need to include all inequalities that describe (the conic version of) the stable set polytope for the graph (N, E_i) . For the specific form (6), the corresponding graph is a complete l -partite graph, and therefore the clique constraints

$$\sum_p x_{ij(p)} \leq y_i \quad \text{for all } (j(1), j(2), \dots, j(l)) \in N_1 \times N_2 \times \dots \times N_l$$

are sufficient since the graph is perfect; see [20].

However, if we consider quota constraints (5) or we combine simultaneously say capacity constraints (4) and incompatibility constraints (6), then additional more complicated inequalities are needed to describe the conic hull.

3. Polynomial-time solvability of the cost allocation problem

As a result of Theorem 2.1, the cost allocation problem (CAP) can be solved in polynomial time by linear programming if we know a compact (i.e., with a polynomial number of linear inequalities) representation of the cones generated by the P_i 's. In this case, an element of the core can simply be obtained by solving (LP) and extracting the dual variables on the constraints (7).

Even if we do not know or there does not exist a compact (possibly extended) representation of each cone, we can still solve the cost allocation problem in polynomial time provided we can optimize in polynomial time over each discrete set P_i . This can be done by using the ellipsoid method and exploiting the equivalence between optimization and separation, see [20]. Since we have a polynomial bound on the size of the inequalities needed to describe (CAP), the equivalence is between *strong* (or exact) optimization and

separation. The separation problem associated to (CAP) is given a vector v to decide whether there exists a set $S \subseteq N$ such that $\sum_{j \in S} v_j > \text{OPT}(S)$, and if so to find one such set. By the same reasoning as in the proof of Theorem 2.1, this is equivalent to deciding whether there exists $i \in F$ and a feasible S for i such that

$$f_i + \sum_{j \in S} (c_{ij} - v_j) < 0.$$

And this can be decided by optimizing over P_i .

Similarly, if we have a polynomial approximation scheme for the P_i 's then we can derive a polynomial approximation scheme for (CAP). This is, for example, useful when the customers have different demands, say d_j for customer j , and we have a capacity constraint on each facility

$$\sum_j d_j x_{ij} \leq k_i y_i.$$

4. On the existence of core elements

In the literature, positive and negative results on the existence of elements in the core have been shown for several important classes of instances. In view of the result in Theorem 2.1, we provide a novel variant of *randomized rounding* in order to give new and simple proofs for the existence of integral solutions for the LP relaxation of facility location problems in certain cases. Randomized rounding is a well-known technique in combinatorial optimization for turning a fractional solution into an integer solution making use of the structural information contained in the fractional solution; we refer the reader to [27] for further information.

For the unconstrained facility location problem, we take an optimum solution (x, y) to the LP relaxation discussed in Section 2 and try to round it randomly to a feasible integral solution by interpreting the fractional values x_{ij} and y_i as probabilities. A similar technique was used by Chudak [9] to compute near-optimal solutions for metric instances. However, while Chudak opened facilities randomly with probabilities y_i but established connections by a different routine, the main problem for our approach is to make sure that a variable x_{ij} is only rounded to 1 if facility i is open, i.e., if the variable y_i is also rounded to 1. This condition forces a coupling of the random decisions which makes it necessary to introduce a subtle correlation between the different random variables.

We first ‘color’ all facilities $i \in F$ by subsets I_i of the interval $[0, 1)$ such that the measure $|I_i|$ of I_i is equal to y_i . Later, we will draw a random variable α uniformly distributed from $[0, 1)$ and open all facilities i with $\alpha \in I_i$; in fact, the probability for opening facility i is then equal to y_i . In order to determine the connections of customers j to facilities i , we construct subsets I_{ij} of $[0, 1)$ such that $|I_{ij}| = x_{ij}$ and establish a connection from customer j to facility i if $\alpha \in I_{ij}$. To make sure that each customer is connected to exactly one facility, the subsets I_{ij} , $i \in F$, should form a partition of the interval $[0, 1)$; notice that $\sum_{i \in F} |I_{ij}| = \sum_{i \in F} x_{ij} = 1$ by constraints (10). Moreover, since a customer should only be connected to an open facility, we require $I_{ij} \subseteq I_i$ for all $i \in F$, $j \in N$.

Lemma 4.1. Given an optimum solution (x, y) to the LP relaxation of the unconstrained facility location problem, if we can find subsets I_i and I_{ij} of $[0, 1)$ with the following properties:

- (i) $|I_i| = y_i$, $|I_{ij}| = x_{ij}$, for all $i \in F$, $j \in N$;
- (ii) $\bigcup_{i \in F} I_{ij} = [0, 1)$, for all $j \in N$, and $I_{ij} \cap I_{i'j} = \emptyset$, for all $i \neq i' \in F$, $j \in N$;
- (iii) $I_{ij} \subseteq I_i$, for all $i \in F$, $j \in N$;

then there exists an integral optimal solution to the LP relaxation.

Proof. We randomly construct an integral solution (\bar{x}, \bar{y}) . Choose a random variable α uniformly distributed from $[0, 1)$; open all facilities i with $\alpha \in I_i$ (i.e., set $\bar{y}_i = 1$) and establish all connections ij with $\alpha \in I_{ij}$ (i.e., set $\bar{x}_{ij} = 1$). By the properties of the sets I_i and I_{ij} , this gives a feasible integral solution of expected value

$$\mathbb{E} \left[\sum_{i \in F} f_i \bar{y}_i + \sum_{\substack{i \in F \\ j \in N}} c_{ij} \bar{x}_{ij} \right] = \sum_{i \in F} f_i y_i + \sum_{\substack{i \in F \\ j \in N}} c_{ij} x_{ij} = \text{OPT}_{\text{LP}}.$$

This expected value is a convex combination of the values of all integral solutions corresponding to possible choices of α . In particular, there exists an integral solution whose value is upper bounded by the optimum LP value. \square

4.1. Facility location on a line with unimodal connection costs

We apply this proof technique to unconstrained facility location problems where all facilities can be ordered in such a way that, for any customer j , the connection costs c_{ij} are unimodal as a function of i . This means that there exists an ordering $1, \dots, m$ of the facilities and for any customer j , there exists a facility $i(j)$ such that c_{ij} is nonincreasing for $i \leq i(j)$ and nondecreasing for $i \geq i(j)$. This is, for example, the case when all facilities are located on a line in the plane or a higher-dimensional Euclidean space and the connection cost between customer j and facility i is a nondecreasing function of their Euclidean distance (see Fig. 2). As an illustration, this situation occurs when we have a (straight) railway line and the problem is to decide where to build railway stations so as to provide an optimal service to the inhabitants of the region around the railway line.




Fig. 2. The unimodal case with facilities on a line; the rectangles represent facilities while the circles correspond to customers.

Theorem 4.2 [22,37]. *There is no integrality gap for the unconstrained facility location problem with unimodal connection costs; in particular, the core is non-empty in this case.*

Proof. In the following we assume without loss of generality that the optimal LP solution (x, y) fulfills $y_i \leq 1$, for all $i \in F$, and has the following property:

If facility i lies between facility $i(j)$ and facility i' and is distinct from i' (i.e., $i(j) \leq i < i'$ or $i' < i \leq i(j)$) and $x_{i'j} > 0$, then $x_{ij} = y_i$.

Otherwise one can modify the solution accordingly by increasing x_{ij} and simultaneously decreasing $x_{i'j}$ without an increase in cost since $c_{ij} \leq c_{i'j}$. In particular, this implies that $x_{i(j)j} = y_{i(j)}$ for all $j \in N$.

Let $a_0 := 0$ and $a_i := \sum_{k=1}^i y_k$, for $i = 1, \dots, m$, and assign the set $I_i := ([a_{i-1}, a_i) \bmod 1) \subseteq [0, 1)$ to facility i . Notice that by construction $|I_i| = y_i$. We also assign to each pair formed by a facility i and a customer j a subset of measure x_{ij} :

$$I_{ij} := \begin{cases} [a_i - x_{ij}, a_i) \bmod 1 & \text{if } i \leq i(j), \\ [a_{i-1}, a_{i-1} + x_{ij}) \bmod 1 & \text{if } i \geq i(j). \end{cases}$$

Notice that $I_{i(j)j} = [a_{i(j)-1}, a_{i(j)}) = I_{i(j)}$ since $x_{i(j)j} = y_{i(j)}$. The fact that $x_{ij} \leq y_i$ implies that $I_{ij} \subseteq I_i$. Moreover, by our assumption on the LP solution (x, y) , the subsets I_{ij} , $i \in F$, form a partition of the interval $[0, 1)$ for each fixed customer j . Therefore, the sets I_i and I_{ij} fulfill the conditions of Lemma 4.1 and the result follows. \square

Notice that in the unimodal case the cost functions for the connections along the line are not necessarily symmetric around $i(j)$. An important application with non-symmetric connection costs is the lot sizing problem.

The line represents a time axis and facilities and customers correspond to discrete points in time when a product can be produced and has to be delivered, respectively. In particular, if we do not allow backlogging, a customer can only be served by a facility ‘in the past,’ i.e., on its left-hand side on the line. In this case, Theorem 4.2 was proved by Krarup and Bilde [24].

The result in Theorem 4.2 cannot be generalized to the capacitated version of the problem. This follows from the example discussed in Section 2 that can obviously be realized on a line and where the customers are not willing to pay anything although the cost of an optimum solution is positive.

In the case where the facilities are located on a line and each facility must serve the same number of customers from N_1 and from N_2 , the core can also be empty. Consider the example given in Fig. 3 where the same numbers of women and men have to be served by each facility. The fixed cost for opening a facility is 1. The cost for connecting a customer to a facility is equal to the corresponding distance (number of edges on the shortest path) in the graph given by the dotted edges. In an optimal solution, all facilities are opened resulting in total cost 9. However, in the LP relaxation we can open the facility in the middle with fraction 1/2 and also connect the neighboring customers with fraction 1/2; the optimal LP value is therefore only 17/2. It also follows from this example that in the more general case of quota constraints (5) the core can be empty.




Fig. 3. An instance of the facility location problem on a line where each facility has to serve the same number of women and men; the core is empty.




Fig. 4. An instance of the facility location problem on a cycle with empty core.

4.2. Facility location on a cycle

Another possible direction for generalizing the result in Theorem 4.2 is to switch to more complicated topologies than the line. However, it has been observed by Tamir [36, Example 1] (see also [23, Example 6.3]) that already instances defined on a cycle with unimodal connection costs can have an empty core; we say that the connection costs on a cycle are unimodal if, for each customer j , opening the cycle at a facility i with maximum c_{ij} yields an unimodal cost function on the resulting line. In the example given in Fig. 4 the fixed cost for opening each facility is 2 and the connection cost between a customer and a facility is equal to their distance in the cyclic graph. An optimal solution opens two facilities and has total cost 7. However, the optimum solution to the LP relaxation opens each facility with fraction $1/2$ and has value 6. Notice that the sum of the y_i 's is equal to $3/2$ and thus not integral in this case.

Theorem 4.3. *If there exists an optimum solution to the LP relaxation with $\sum_{i \in F} y_i \in \mathbb{Z}$, then there is no integrality gap for the unconstrained facility location problem with unimodal connection costs on a cycle; in particular, the core is non-empty in this case.*

Proof. The proof is a straightforward generalization of the proof of Theorem 4.2. We assume that the facilities $1, \dots, m$ are ordered clockwise along the cycle and define the sets I_i exactly as in the proof of Theorem 4.2. Notice that $a_m = \sum_{i \in F} y_i = (a_0 \bmod 1)$ in this case. In order to assign subsets to the connections ij for a fixed customer j , we ‘open’ the cycle at a facility with maximum c_{ij} yielding a line and construct the sets I_{ij} again as in the proof of Theorem 4.2 (using the same assumption on the optimum LP solution). The resulting sets I_i and I_{ij} fulfill the conditions of Lemma 4.1 and the result follows. \square

The rounding technique in the proof of Theorem 4.3 has been applied before by Bertsimas, Teo, and Vohra [7] to prove the integrality of an LP relaxation of the k -median problem on cycles. The insight from Theorem 4.3 can be used to prove the following result.

Corollary 4.4. *The unconstrained facility location problem with unimodal connection costs on a cycle can be solved in polynomial time.*

Proof. We add the constraint $\sum_i y_i \in \mathbb{Z}$ to the LP relaxation. It follows from the proof of Theorem 4.3 that an optimal solution to this stronger relaxation can be turned into an integral optimal solution in polynomial time. Moreover, the stronger relaxation can be solved in the following way: let $\gamma := \sum_i y_i^*$ for an optimum solution of the original (LP). Then there exists an optimal solution to the stronger relaxation with $\sum_i y_i \in \{\lceil \gamma \rceil, \lfloor \gamma \rfloor\}$; this follows from the fact that the optimal value of the parametric linear program which we get by adding the constraint $\sum_i y_i = \mu$ to (LP) is a convex function of μ , see, e.g., [32, Section 6.5]. Therefore, we only need to solve this relaxation for the two values $\mu = \lceil \gamma \rceil$ and $\lfloor \gamma \rfloor$ take the better of the two solutions. \square

An alternative proof of Corollary 4.4 is based on the following observation. As we already mentioned in the introduction, the unconstrained facility location problem can be formulated as a covering problem (see [23]). Since the connection costs are unimodal, the rows in the linear programming formulation of this covering problem fulfill the circular 1's property (for definition, see [3, Chapter 9, Exercise 9.9]). If we add the constraint $\sum_i y_i = p$ for some $p \in \mathbb{N}$ (p can, for example, be determined by binary search), it is an easy exercise to show that an integral optimal solution exists since the matrix of the linear program is totally unimodular, see [3, Chapter 9, Exercise 9.9].

The above discussion has also implications for the p -median problem, where the number of facilities to open is a given integral number p . The proof of Theorem 4.2 shows that the corresponding LP relaxation (with the inequality $\sum_{i \in F} y_i = p$) has an integral optimum solution for instances defined on a cycle with unimodal costs; a generalization of a result first derived by Oudjit [30] (for linear costs).

4.3. Facility location on a tree

In the following we consider unconstrained facility location problems on trees where all customers and facilities are located on the vertices of an acyclic connected graph. As for the case of a line, we assume that the connection costs c_{ij} of each customer j are monotone such that $c_{ij} \leq c_{i'j}$ if facility i lies on the unique path between j and facility i' . Unfortunately, this condition is not sufficient to guarantee a non-empty core. The counterexample in Fig. 5 consists of three customers and three facilities. The fixed cost of each facility is 1 and each customer can be connected to the facility at the same vertex of the tree or to its anti-clockwise neighbor for free; a connection to its clockwise neighbor, however, costs a large amount M . An optimum solution opens two of the three facilities and has value 2. In an optimum LP solution, however, each facility is opened with fraction $1/2$ resulting in total cost $3/2$. Thus, only a fraction of $3/4$ of the total cost can be recovered.




Fig. 5. An instance of the unconstrained facility location problem on a tree with monotone connection costs but empty core.

In the following we restrict to a special class of monotone connection costs on a tree: let $d : E \mapsto \mathbb{R}_0^+$ be a distance function on the edges of the tree. We denote the length of the unique path between two vertices i and j of the tree by d_{ij} .

Theorem 4.5 [22,37]. *If for each customer j the costs c_{ij} for connecting j to facilities $i \in F$ is an arbitrary nondecreasing function (possibly dependent on j) of the distances d_{ij} in an underlying tree metric, then there is no integrality gap for the unconstrained facility location problem and the core is non-empty.*

Using the same technique as in the proofs of Theorem 4.2 and Lemma 4.1, one can show that, for the class of instances under consideration, any feasible solution to the LP relaxation can be written as a convex combination of integral solutions. In particular, all vertices of the underlying polyhedron are integral. This result also follows from the work of Trubin [37] (see also [19]), Kolen [22] (see also [23]), or Barany, Edmonds, and Wolsey [6] (see also [10]). They use a simple reduction of the unconstrained facility location problem to the set cover problem. For the class of instances considered in Theorem 4.5, the constraint matrix of the resulting set cover problem is totally balanced which yields the integrality result. Moreover, Trubin and also Kolen give an $O((|N| + |F|)^3)$ algorithm for solving such instances. This result has been improved by Gimadi [16] to running time $O((|N| + |F|)^2)$. Furthermore, Ageev [2] and Granot and Skorin-Karpov [18] gave a polynomial-time algorithm for solving the unconstrained facility location problem on partial k -trees for fixed k .

Tamir [36] considers a variant of the facility location problem on a tree where each customer j has to be connected to a facility within a given distance r_j . This constraint can be modeled by letting $c_{ij} = 0$ if the distance between customer j and facility i is at most r_j and $c_{ij} = \infty$, otherwise; thus, it is a special case of the problem discussed above. Tamir proves the result given in Theorem 4.5 for this special case.

Proof of Theorem 4.5. We construct sets I_i and I_{ij} that fulfill the conditions in Lemma 4.1. To simplify the presentation of the construction, we assume that there is at least one customer at each node of the tree; otherwise we can, without loss of generality, add dummy customers j with $c_{ij} = 0$ for all $i \in F$. For each customer j we order the set of facilities F by nondecreasing distances d_{ij} and break ties according to increasing indices

$i = 1, \dots, m$. Using the same arguments given in the proof of Theorem 4.2, we can restrict to LP solutions with the following property:

If facility i is closer to customer j than facility i' (with respect to the above ordering) and $x_{i'j} > 0$, then $x_{ij} = y_i$.

As an immediate consequence of this we get the following property:

If customer k lies on the unique path from customer j to facility i , then $x_{ij} \leq x_{ik}$. (*)

In particular, if two customers j and j' are located at the same node of the tree, we get $x_{ij} = x_{ij'}$ for all $i \in F$. Therefore, all customers j located at the same node will get the same sets I_{ij} and we will only consider one customer per node in the following construction.

To assign subsets I_i and I_{ij} to the facilities and connections, respectively, we traverse the nodes of the tree in such a way that the first k visited nodes form a subtree for any k . This can, e.g., be done by rooting the tree at some node and traversing the nodes of the tree by depth- or breadth-first search or in any order such that the predecessor of any node is visited before the node itself. During the algorithm we preserve the following invariant:

For all customers j at visited nodes of the tree, the subsets $I_{ij}, i \in F$, form a partition of the interval $[0, 1)$ such that $|I_{ij}| = x_{ij}$ and $I_{ij} \subseteq I_i$. Moreover, (**)
for all $i \in F$, there exists a visited node j such that $I_{ij} = I_i$ for all $i \in F$.

In particular, this implies that $|I_i| \leq y_i$.

For the customer j at the root of the tree, we set $I_i := I_{ij} := [\sum_{k=1}^{i-1} x_{ik}, \sum_{k=1}^i x_{ik}] \subseteq [0, 1)$ for $i = 1, \dots, m$. Notice that the invariant (**) is fulfilled after this step. When we arrive at a node with a customer j , we denote the customer located at the predecessor of the current node in the tree by k . The following observation is crucial for the assignment of the subsets I_i and I_{ij} .

Claim 1. For all $i \in F$, either $x_{ij} \leq x_{ik}$ or $x_{ik} = |I_i|$.

We postpone the proof of this claim after the complete description of the assignment procedure.

We first consider all facilities $i \in F$ with $x_{ij} \leq x_{ik}$ and choose an arbitrary subset $I_{ij} \subseteq I_{ik}$ of measure x_{ij} . For all the other facilities $i \in F$, namely those for which $x_{ij} > x_{ik} = |I_i|$ by Claim 1, choose arbitrary sets $I_{ij} \supset I_{ik}$ of measure x_{ij} such that the sets $I_{ij}, i \in F$, form a partition of $[0, 1)$ (this is possible since the sets $I_{ik}, i \in F$, form a partition) and redefine $I_i := I_{ij}$. Notice that the invariant (**) is still fulfilled after this step.

Thus, the final sets I_i and I_{ij} fulfill the conditions of Lemma 4.1 and the result follows.

Proof of Claim 1. Assume that $x_{ik} < x_{ij}$ and $x_{ik} < |I_i|$ (since k has been visited). By the invariant (**), there exists a customer $h \neq k$ with $I_{ih} = I_i$, implying that $x_{ik} < x_{ih}$. By property (*), k can neither lie on the path from j to i nor on the path from h to i . However, since k is the predecessor of j in the tree and h was visited, k lies on the path between j and h —a contradiction. \square

5. On the complexity of core computations

In this section we prove the following theorem which confirms a conjecture by Chardaire [8].

Theorem 5.1. *For general instances of the unconstrained facility location problem it is NP-complete to decide whether the core is non-empty.*

Proof. By Theorem 2.1 it suffices to show that it is NP-complete to decide whether or not the LP relaxation given in Section 2 has an integral optimal solution. The problem is obviously in NP. To prove that it is NP-hard, we use a reduction from 3SAT. We restrict to instances where each clause contains exactly three (not necessarily different) literals.

Given an instance of 3SAT we construct the following facility location problem. For each variable X we introduce one customer j_X and two facilities i_X and $i_{\bar{X}}$ corresponding to the two literals X and its negation \bar{X} . The cost for connecting customer j_X to these two facilities is 0, the cost for opening one of the facilities is equal to 1 plus the number of occurrences of the corresponding literal in the instance.

For each clause C we introduce one customer j_C and for each of its three literals L a facility i_{CL} and a customer j_{CL} . Customer j_C can be connected to the three facilities at cost 0, the cost for opening each facility is 1; each customer j_{CL} can be connected either to i_{CL} or to i_L at cost 0. Finally, we introduce one additional dummy facility i_0 with fixed cost 0. The customers j_{CL} of all clauses can connect at cost 1 to this facility. An illustrating example of this construction is given in Fig. 6. All connections that are not depicted in the figure have infinite cost.




Fig. 6. A facility location instance corresponding to an instance of 3SAT containing variables X , Y , and Z and a clause $C = (\bar{X} \vee Y \vee Z)$.

We claim that the LP relaxation of this facility location instance has an integral optimal solution if and only if the underlying 3SAT instance can be satisfied. We first show that the optimal LP value is $n + 3m$ where n denotes the number of variables and m the number of clauses of the 3SAT instance. A feasible solution of value $n + 3m$ is given by

$$\left. \begin{aligned} y_{i_L} &= 1/2 && \text{for each literal } L, \\ x_{i_X j_X} &= x_{i_{\bar{X}} j_X} = 1/2 && \text{for each variable } X, \\ y_{i_{CL}} &= x_{i_{CL} j_C} = x_{i_{CL} j_{CL}} = 1/3 \\ & \quad x_{i_L j_{CL}} = 1/2 \\ & \quad y_{i_0} = x_{i_0 j_{CL}} = 1/6 \end{aligned} \right\} \text{for each clause } C \text{ and each of its literals } L.$$

In order to show that this solution is optimal, we construct a dual solution of value $n + 3m$. The dual of the LP relaxation is given by

$$\begin{aligned} &\text{maximize} && \sum_{j \in N} v_j, \\ &\text{subject to} && \sum_{j \in N} w_{ij} = f_i \quad \text{for all } i \in F, \\ & && v_j - w_{ij} \leq c_{ij} \quad \text{for all } i \in F, j \in N, \\ & && w_{ij} \geq 0 \quad \text{for all } i \in F, j \in N. \end{aligned}$$

A feasible dual solution of value $n + 3m$ is given by

$$\left. \begin{aligned} v_{j_X} &= w_{i_X j_X} = w_{i_{\bar{X}} j_X} = 1 && \text{for each variable } X, \\ v_{j_C} &= w_{i_{CL} j_C} = w_{i_0 j_{CL}} = 0 \\ v_{j_{CL}} &= w_{i_{CL} j_{CL}} = w_{i_L j_{CL}} = 1 \end{aligned} \right\} \text{for each clause } C \text{ and each of its literals } L.$$

Since the given dual solution is optimal, every primal optimal solution has to fulfill the following complementary slackness conditions:

$$x_{i_X j_X} = y_{i_X}, \quad x_{i_{\bar{X}} j_X} = y_{i_{\bar{X}}} \quad \text{for each variable } X, \tag{11}$$

$$x_{i_{CL} j_{CL}} = y_{i_{CL}}, \quad x_{i_L j_{CL}} = y_{i_L} \quad \text{for each clause } C \text{ and each of its literals } L. \tag{12}$$

Using these conditions we can show that any primal integral optimal solution (x, y) corresponds to a satisfying truth-assignment of the underlying 3SAT instance. Condition (11) and constraint (10) yield that for each variable X exactly one of the facilities i_X and $i_{\bar{X}}$ has to be open. We set the variable X to the value true if $y_{i_X} = 0$ and to the value false if $y_{i_{\bar{X}}} = 0$. For an arbitrary clause C , at least one of the three facilities i_{CL} corresponding to the three literals of C must be open to serve customer j_C . Condition (12) yields that the customer j_{CL} is connected to this facility. It follows again from (12) that the facility i_L is closed and the literal L is thus set to the value true such that the clause C is fulfilled.

Using the same interpretation one can easily show that any satisfying truth-assignment yields an integral optimal solution of the LP relaxation. This completes the proof. \square

We can state the following interesting corollary of the result in Theorem 5.1.

Corollary 5.2. *Given an instance of the unconstrained facility location problem and a cost allocation vector $v \in \mathbb{R}^N$, it is NP-complete to decide whether or not v is in the core of the corresponding game.*

Proof. The problem is in NP since a positive answer to the question can be proven in the following way: give an integral optimal solution to the LP relaxation and a completion w of the given vector v to an optimal dual solution (v, w) . \square

Quite interestingly, given the information that the core is non-empty, it is easy to compute an element of the core and to decide whether a given cost allocation v belongs to the core. Both problems reduce to solving the dual of the LP relaxation of the problem.

Corollary 5.3. *If the core of an unconstrained facility location game is non-empty, an element of the core can be computed in polynomial time and it can be checked in polynomial time whether a given cost allocation $v \in \mathbb{R}^N$ belongs to the core.*

It follows from the considerations in Section 3 that the result in Corollary 5.3 can be generalized to constrained variants of the facility location problem for which we can optimize over the discrete sets P_i in polynomial time.

Acknowledgments

The authors are much indebted to Michel Le Breton for motivating them to study facility location problems from a game-theoretic point of view and for numerous interesting and helpful discussions on this subject. They also thank Alexander Ageev and Maxim Sviridenko for pointing out the connection to totally balanced matrices mentioned in Section 4.3. Further helpful comments and pointers to the literature were given by two anonymous referees.

References

- [1] A.A. Ageev, On the complexity of minimization problems for polynomials in boolean variables, *Upravlyaemye Sistemy* 23 (1983) 3–11 (in Russian).
- [2] A.A. Ageev, A criterion of polynomial-time solvability for the network location problem, in: E. Balas, G. Cornuéjols, R. Kannan (Eds.), *Integer Programming and Combinatorial Optimization*, Carnegie Mellon Univ., Pittsburgh, 1992, pp. 237–245.
- [3] R.K. Ahuja, T.L. Magnanti, J.B. Orlin, *Network Flows. Theory, Algorithms, and Applications*, Prentice-Hall, Englewood Cliffs, NJ, 1993.
- [4] M.L. Balinski, Integer programming: Methods, uses, computation, *Manag. Sci.* 12 (1965) 253–313.
- [5] A. Bar-Noy, S. Guha, J. Naor, B. Schieber, Approximating the throughput of multiple machines under real-time scheduling, in: *Proceedings of the 31st Annual ACM Symposium on Theory of Computing*, 1999, pp. 622–631.
- [6] I. Bárány, J. Edmonds, L.A. Wolsey, Packing and covering a tree by subtrees, *Combinatorica* 6 (1986) 221–233.
- [7] D. Bertsimas, C. Teo, R. Vohra, On dependent randomized rounding algorithms, *Oper. Res. Lett.* 24 (1999) 105–114.
- [8] P. Chardaire, *Facility location optimization and cooperative games*, PhD thesis, University of East Anglia, Norwich, UK, 1998.
- [9] F.A. Chudak, Improved approximation algorithms for uncapacitated facility location, in: R.E. Bixby, E.A. Boyd, R.Z. Ríos-Mercado (Eds.), *Integer Programming and Combinatorial Optimization*, in: *Lecture Notes Comput. Sci.*, vol. 1412, Springer-Verlag, Berlin, 1998, pp. 180–194.

- [10] G. Cornuéjols, G.L. Nemhauser, L.A. Wolsey, The uncapacitated facility location problem, in: P. Mirchandani, R. Francis (Eds.), *Discrete Location Theory*, Wiley, New York, 1990, pp. 119–171.
- [11] X. Deng, T. Ibaraki, H. Nagamochi, Algorithmic aspects of the core of combinatorial optimization games, *Math. Oper. Res.* 24 (1999) 751–766; also in: *Proceedings of the 8th Annual ACM–SIAM Symposium on Discrete Algorithms*, 1997.
- [12] X. Deng, T. Ibaraki, H. Nagamochi, W. Zang, Totally balanced combinatorial optimization games, *Math. Program.* 87 (2000) 441–452.
- [13] X. Deng, C.H. Papadimitriou, On the complexity of cooperative solution concepts, *Math. Oper. Res.* 19 (1994) 257–266.
- [14] U. Faigle, S.P. Fekete, W. Hochstättler, W. Kern, On the complexity of testing membership in the core of min-cost spanning tree games, *Int. J. Game Theory* 26 (1997) 361–366.
- [15] U. Faigle, W. Kern, On the core of ordered submodular cost games, *Math. Program.* 87 (2000) 483–499.
- [16] E.K. Gimadi, An effective algorithm for the solution of the location problem with service domains connected with respect to an acyclic network, *Upravlyaemye Sistemy* 23 (1983) 12–23 (in Russian).
- [17] M. Göthe-Lundgren, K. Jörnsten, P. Värbrand, On the nucleolus of the basic vehicle routing game, *Math. Program.* 72 (1996) 83–100.
- [18] D. Granot, D. Skorin-Karpov, On some optimization problems on k -trees and partial k -trees, *Discrete Appl. Math.* 48 (1994) 129–145.
- [19] V.P. Grishukhin, On polynomial solvability conditions for the simplest plant location problem, *Trans. Amer. Math. Soc. Ser. 2* 158 (1994) 37–46; translation from *Economico-Mathematical Modelling of Discrete Processes*, Moscow, 1988, pp. 52–65.
- [20] M. Grötschel, L. Lovász, A. Schrijver, *Geometric Algorithms and Combinatorial Optimization*, in: *Algorithms Combin.*, vol. 2, Springer-Verlag, Berlin, 1988.
- [21] S. Guha, S. Khuller, Greedy strikes back: Improved facility location algorithms, in: *Proceedings of the 9th Annual ACM–SIAM Symposium on Discrete Algorithms*, 1998, pp. 649–657.
- [22] A. Kolen, Solving covering problems and the uncapacitated plant location problem on trees, *Eur. J. Oper. Res.* 12 (1983) 266–278.
- [23] A. Kolen, A. Tamir, Covering problems, in: P. Mirchandani, R. Francis (Eds.), *Discrete Location Theory*, Wiley, New York, 1990, pp. 263–304.
- [24] J. Krarup, O. Bilde, Plant location, set covering and economic lot size: an $O(mn)$ -algorithm for structured problems, in: L. Collatz, W. Wetterling (Eds.), *Numerische Methoden bei Optimierungsaufgaben, Band 3 (Optimierung bei graphentheoretischen und ganzzahligen Problemen)*, in: *Int. Ser. Numer. Math.*, vol. 36, Birkhäuser, 1977, pp. 155–180.
- [25] J. Lenstra, D. Shmoys, E. Tardos, Approximation algorithms for scheduling unrelated parallel machines, *Math. Program.* 46 (1990) 259–271.
- [26] P.B. Mirchandani, P.L. Francis (Eds.), *Discrete Location Theory*, Wiley, New York, 1990.
- [27] R. Motwani, P. Raghavan, *Randomized Algorithms*, Cambridge Univ. Press, Cambridge, UK, 1995.
- [28] H. Moulin, *Cooperative Microeconomics: A Game-Theoretic Introduction*, Princeton Univ. Press, Princeton, NJ, 1995.
- [29] G.L. Nemhauser, L.A. Wolsey, *Integer and Combinatorial Optimization*, Wiley, New York, 1988.
- [30] A. Oudjit, Median locations on deterministic and probabilistic multidimensional networks, PhD thesis, Rensselaer Polytechnic Institute, Troy, NY, 1981.
- [31] G. Owen, *Game Theory*, Academic Press, 1995.
- [32] M. Padberg, *Linear Optimization and Extensions*, Springer-Verlag, Berlin, 1995.
- [33] D. Samet, E. Zemel, On the core and dual set of linear programming games, *Math. Oper. Res.* 9 (1984) 309–316.
- [34] W.W. Sharkey, Cores of games with fixed costs and shared facilities, *Internat. Econom. Rev.* 31 (1990) 245–262.
- [35] D.B. Shmoys, E. Tardos, An approximation algorithm for the generalized assignment problem, *Math. Program.* 62 (1993) 461–474.
- [36] A. Tamir, On the core of cost allocation games defined on location problems, *Transp. Sci.* 27 (1992) 81–86; also in: *Proceedings of the International Symposium on Locational Decisions*, Skodsborg, Denmark, 1981.

- [37] V.A. Trubin, An effective algorithm for solving the distribution problem in a network in the form of a tree, *Soviet Math. Dokl.* 17 (1976) 1604–1608; translation from *Dokl. Akad. Nauk SSSR* 231 (1976) 547–550.
- [38] V.V. Vazirani, *Approximation Algorithms*, Springer-Verlag, Berlin, 2001.