

CONF-9506248--1

LA-UR 95-1981

Los Alamos National Laboratory is operated by the University of California for the United States Department of Energy under contract W-7405-ENG-36

**TITLE: COMPACT LOCATION PROBLEMS WITH BUDGET AND
COMMUNICATION CONSTRAINTS**

AUTHOR(S): S. O. Krumke, H. Noltemeier, S. S. Ravi, M. V. Marathe

SUBMITTED TO: 1st International Conference on Computing and Combinatorics
(COCOON '95)
X'ian, China
June 1995

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

By acceptance of this article, the publisher recognizes that the U.S. Government retains a nonexclusive, royalty-free license to publish or reproduce the published form of this contribution or to allow others to do so, for U.S. Government purposes.

The Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy.

Los Alamos

Los Alamos National Laboratory
Los Alamos New Mexico 87545

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

MASTER

DISCLAIMER

**Portions of this document may be illegible
in electronic image products. Images are
produced from the best available original
document.**

Compact Location Problems with Budget and Communication Constraints

S. O. Krumke¹, H. Noltemeier¹, S. S. Ravi² and M. V. Marathe^{3,*}

¹ University of Würzburg, Am Hubland, 97074 Würzburg, Germany. Email:
{krumke, noltemei}@informatik.uni-wuerzburg.de.

² University at Albany - SUNY, Albany, NY 12222, USA. Email:
ravi@cs.albany.edu.

³ Los Alamos Nat. Lab. P.O. Box 1663, MS M986, Los Alamos, NM 87545, USA.
Email: madhav@c3.lanl.gov.

Abstract. We consider the problem of placing a specified number p of facilities on the nodes of a given network with two nonnegative edge-weight functions so as to minimize the diameter of the placement with respect to the first weight function subject to a diameter- or sum-constraint with respect to the second weight function.

Define an (α, β) -approximation algorithm as a polynomial-time algorithm that produces a solution within α times the optimal value with respect to the first weight function, violating the constraint with respect to the second weight function by a factor of at most β .

We show that in general obtaining an (α, β) -approximation for any fixed $\alpha, \beta \geq 1$ is \mathcal{NP} -hard for any of these problems. We also present efficient approximation algorithms for several of the problems studied, when both edge-weight functions obey the triangle inequality.

1 Introduction and Basic Definitions

Several fundamental problems in location theory [HM79, MF90] involve finding a placement obeying certain “covering” constraints. Generally, the goal of such a location problem is to find a placement of minimum cost that satisfies all the specified constraints. The cost of a placement may reflect the price of constructing the network of facilities, or it may reflect the maximum communication cost between any two facilities. Examples of such cost measures are the total edge cost and the diameter respectively.

Finding a placement of sufficient generality minimizing even one of these measures is often \mathcal{NP} -hard [GJ79]. In practice, it is usually the case that a facility location problem involves the minimization of a certain cost measure, subject to budget constraints on other cost measures.

The problems considered in this paper can be termed as *compact location* problems, since we will typically be interested in finding a “compact” placement of facilities. The following is a prototypical compact location problem: Given an undirected edge-weighted complete graph $G = (V, E_c)$, place a specified number

* Research supported by Department of Energy under contract W-7405-ENG-36.

p of facilities on the nodes of G , with at most one facility per node, so as to minimize some measure of the distances between facilities. This problem has been studied for both diameter and sum objectives [RKM⁺93]. Some geometric versions of this problem have also been studied [AI+91].

Consider the following extension of the compact location problem. Suppose we are given *two* weight-functions δ_c, δ_d on the edges of the network. Let the first weight function δ_c represent the cost of constructing an edge, and let the second weight function δ_d represent the actual transportation- or communication-cost over an edge (once it has been constructed). Given such a graph, we can define a general bicriteria problem $(\mathcal{A}, \mathcal{B})$ by identifying two minimization objectives of interest from a set of possible objectives. A budget value is specified on the second objective \mathcal{B} and the goal is to find a placement of facilities having minimum possible value for the first objective \mathcal{A} such that this solution obeys the budget constraint on the second objective. For example, consider the *diameter-bounded minimum diameter compact location problem* denoted by DC-MDP: Given an undirected graph $G = (V, E)$ with two different nonnegative integral edge weight functions δ_c (modeling the building cost) and δ_d (modeling the delay or the communication cost), an integer p denoting the number of facilities to be placed, and an integral bound B (on the total delay), find a placement of p facilities with minimum diameter under the δ_c -cost such that the diameter of the placement under the δ_d -costs (the maximum delay between any pair of nodes) is at most B . We term such problems as *bicriteria compact location problems*.

In this paper, we study bicriteria compact location problems motivated by practical problems arising in diverse areas such as statistical clustering, pattern recognition, processor allocation and load-balancing.

2 Preliminaries and Summary of Results

Let $G = (V, E_c)$ be a complete undirected graph with $n = |V|$ nodes and let p ($2 \leq p \leq n$) be the number of facilities to be placed. We call any subset $P \subseteq V$ of cardinality p a *placement*. Given a nonnegative weight- or cost-function $\delta : E_c \rightarrow \mathbb{Q}$, we will use $\mathcal{D}_\delta(P)$ to denote the *diameter* of a placement P with respect to δ ; that is

$$\mathcal{D}_\delta(P) = \max_{\substack{u, v \in P \\ u \neq v}} \delta(u, v).$$

Similarly, we will let $\mathcal{S}_\delta(P)$ denote the *sum of the distances* between facilities in the placement P ; that is

$$\mathcal{S}_\delta(P) = \sum_{\substack{u, v \in P \\ u \neq v}} \delta(u, v).$$

We note that the average length of an edge in a placement P equals $\frac{2}{p(p-1)} \mathcal{S}_\delta(P)$.

As usual, we say that a nonnegative distance δ on the edges of G satisfies the *triangle inequality*, if we have

$$\delta(v, w) \leq \delta(v, u) + \delta(u, w)$$

for all $v, w, u \in V$,

The *Minimum Diameter Placement Problem* (denoted by MDP) is to find a placement P that minimizes $\mathcal{D}_\delta(P)$. Similarly, the *Minimum Average Placement Problem* (denoted by MAP) is to find a placement P such that $\mathcal{S}_\delta(P)$ is minimized. Both problems are known to be \mathcal{NP} -hard, even when the distance δ obeys the triangle inequality [RKM⁺93]. Moreover, if the distances are not required to satisfy the triangle inequality, then as observed in [RKM⁺93], there can be no polynomial time relative approximation algorithm for MDP or MAP unless $\mathcal{P} = \mathcal{NP}$.

In the sequel we will restrict ourselves to those instances of the problems where the weights on the edges obey the triangle inequality. Given a problem Π , we use $\text{TI-}\Pi$ to denote the problem Π restricted to graphs with edge weights satisfying the triangle inequality.

Following [HS86], the *bottleneck graph* $\text{bottleneck}(G, \delta, \Delta)$ of $G = (V, E_c)$ with respect to δ and a bound Δ is defined by

$$\text{bottleneck}(G, \delta, \Delta) := (V, E'), \text{ where } E' := \{e \in E_c : \delta(e) \leq \Delta\}.$$

We now formally define the problems studied in this paper.

Definition 1. [*Diameter Constrained Minimum Diameter Placement Problem (DC-MDP)*]

Input: An undirected complete graph $G = (V, E_c)$ with two nonnegative weight functions $\delta_c, \delta_d : E_c \rightarrow \mathbb{Q}$, an integer $2 \leq p \leq n$ and a number $\Omega \in \mathbb{Q}$.

Output: A set $P \subseteq V$, with $|P| = p$, minimizing the objective

$$\mathcal{D}_{\delta_c}(P) = \max_{\substack{v, w \in P \\ v \neq w}} \delta_c(v, w)$$

subject to the constraint

$$\mathcal{D}_{\delta_d}(P) = \max_{\substack{v, w \in P \\ v \neq w}} \delta_d(v, w) \leq \Omega.$$

Definition 2. [*Sum Constrained Minimum Diameter Placement Problem (SC-MDP)*]

Input: An undirected complete graph $G = (V, E_c)$ with two nonnegative weight functions $\delta_c, \delta_d : E_c \rightarrow \mathbb{Q}$, an integer $2 \leq p \leq n$ and a number $\Omega \in \mathbb{Q}$.

Output: A set $P \subseteq V$, with $|P| = p$, minimizing the objective

$$\mathcal{D}_{\delta_d}(P) = \max_{\substack{v, w \in P \\ v \neq w}} \delta_d(v, w)$$

and satisfying the budget-constraint

$$\mathcal{S}_{\delta_c}(P) = \sum_{\substack{v_i, v_j \in P \\ v_i \neq v_j}} \delta_c(v_i, v_j) \leq \Omega.$$

Let $\Pi \in \{\text{TI-DC-MDP, TI-SC-MDP}\}$. Define an (α, β) -approximation algorithm for Π to be a polynomial-time algorithm, which for any instance I of Π does one of the following:

- (a) It produces a solution within α times the optimal value with respect to the first distance function (δ_c), violating the constraint with respect to the second distance function (δ_d) by a factor of at most β .
- (b) It returns the information that no feasible placement exists at all.

Notice that if there is no feasible placement but there is a placement violating the constraint by a factor of at most β , an (α, β) -approximation algorithm has the choice of performing either action (a) or (b).

In this paper we study the complexity and approximability of the problems DC-MDP and SC-MDP. We show that, in general, obtaining an (α, β) -approximation for any fixed $\alpha, \beta \geq 1$ is \mathcal{NP} -hard for any of these problems. We also present efficient approximation algorithms for several of the problems studied, when both edge-weight functions obey the triangle inequality. For TI-DC-MDP problem, we provide a $(2, 2)$ -approximation algorithm. We also show that no polynomial time algorithm can provide an $(\alpha, 2 - \varepsilon)$ - or $(2 - \varepsilon, \beta)$ -approximation for any fixed $\varepsilon > 0$ and $\alpha, \beta \geq 1$, unless $\mathcal{P} = \mathcal{NP}$. This result is proved to remain true, even if one fixes $\varepsilon' > 0$ and allows the algorithm to place only $2p/|V|^{1/6-\varepsilon'}$ facilities. Our techniques can be extended to devise approximation algorithms for TI-SC-MDP. For this problem, our heuristics provide performance guarantees of $(2 - 2/p, 2)$ and $(2, 2 - 2/p)$ respectively. These techniques can also be used to find efficient approximation algorithms for TI-DC-MDP and TI-SC-MDP when there are node and edge weights. Due to lack of space, the discussion on the node-weighted cases is omitted in this version of the paper.

3 Related Work

While there has been much work on finding minimum-cost networks (see for example [DF85, FG88, Go85, IC+86, LV92, Won80]) for each of the cost measures considered in our bicriteria formulations, there has been relatively little work on approximations for multi-objective network-design. In this direction, Bar-Ilan and Peleg [BP91] considered balanced versions of the problem of assigning network centers, where a bound is imposed on the number of nodes that any center can service. Warburton [Wa87] has considered multi-objective shortest path problems. We refer the reader to [MR+95, RMR+93] for a detailed survey of the work done in the area of algorithms for bicriteria network design and location theory problems. Other researchers have addressed multi-objective approximation algorithms for problems arising in areas other than network design. This includes research in the areas of computational geometry [AF+94], numerical analysis, network design [ABP90, KRY93, Fi93] and scheduling [ST93].

Due to lack of space the rest of the paper consists of selected proof sketches.

4 Diameter Constrained Problems

As shown in [RKM⁺93], TI-MDP is \mathcal{NP} -hard. Here we can extend this result to obtain the following non-approximability result.

Proposition 3. *Let $\varepsilon > 0$ and $\varepsilon' > 0$ be arbitrary. Suppose that A is a polynomial time algorithm that, given any instance of TI-DC-MDP, either returns a subset $S \subseteq V$ of at least $\frac{2p}{|V|^{1/6-\varepsilon'}}$ nodes satisfying $\mathcal{D}_{\delta_d}(S) \leq (2-\varepsilon)\Omega$, or provides the information that no placement of p nodes having communication diameter of at most Ω does exist. Then $\mathcal{P} = \mathcal{NP}$. \square*

We can interchange the roles of δ_c and δ_d in the proof of the last proposition to show that the optimal value of the problem cannot be approximated by a factor of $(2 - \varepsilon)$. Moreover, replacing 2 by a suitable function $f \in \Theta(2^{\text{poly}(|V|)})$, which given an input length of $\Theta(|V|)$ is polynomial time computable, it is easy to see that, if the triangle inequality is not required to hold, there can be no polynomial time approximation with performance ratio $O(2^{\text{poly}(|V|)})$ for neither the optimal function value nor the constraint (modulo $\mathcal{P} = \mathcal{NP}$). Thus we obtain:

Lemma 4. *Unless $\mathcal{P} = \mathcal{NP}$, for any fixed $\varepsilon > 0$ and $\varepsilon' > 0$ there can be no polynomial time approximation algorithm for TI-DC-MDP that is required to place at least $2p/|V|^{1/6-\varepsilon'}$ facilities and has a performance guarantee of $(\alpha, 2-\varepsilon)$ or $(2-\varepsilon, \beta)$. If the triangle inequality is not required to hold, then the existence of an $(f(|V|), g(|V|))$ -approximation algorithm for any $f, g \in O(2^{\text{poly}(|V|)})$ implies that $\mathcal{P} = \mathcal{NP}$. \square*

PROCEDURE HEUR-FOR-DIA

1. $G' := \text{bottleneck}(G, \delta_d, \Omega)$
2. $V_{\text{cand}} := \{v \in G' : \deg(v) \geq p - 1\}$
3. IF $V_{\text{cand}} = \emptyset$ THEN RETURN “certificate of failure”
4. Let $\text{best} := +\infty$
5. Let $P_{\text{best}} := \emptyset$
6. FOR each $v \in V_{\text{cand}}$ DO
 - (a) Let $N(v)$ be the set of $p - 1$ nearest neighbors of v in G with respect to δ_c
 - (b) Let $P(v) := N(v) \cup \{v\}$
 - (c) If $\mathcal{M}_{\delta_c}(P(v)) < \text{best}$ THEN $P_{\text{best}} := P(v)$
 $\text{best} := \mathcal{M}_{\delta_c}(P(v))$
7. OUTPUT P_{best}

Fig. 1. Details of the heuristic for TI-DC-MDP and TI-DC-MAP

Using the results in [RKM⁺93] in conjunction with the results in [MR⁺95] we can devise an approximation algorithm with a performance guarantee $(4, 4)$

for TI-DC-MDP. Here we present an improved heuristic HEUR-FOR-DIA for this problem. This heuristic provides a performance guarantee of $(2, 2)$. In view of Lemma 4, this is the best approximation we can expect to obtain in polynomial time. The heuristic is quite simple. The details of the heuristic are shown in Figure 1.

Theorem 5. *Let I be any instance of of TI-DC-MDP such that an optimal solution P^* of diameter cost $OPT(I) = \mathcal{D}_{\delta_c}(P^*)$ exists. Then the algorithm HEUR-FOR-DIA, called with $\mathcal{M}_{\delta_d} := \mathcal{D}_{\delta_d}$, returns a placement P satisfying $\mathcal{D}_{\delta_d}(P) \leq 2\Omega$ and $\mathcal{D}_{\delta_c}(P)/OPT(I) \leq 2$.*

Proof: Consider an optimal solution P^* such that $\mathcal{D}_{\delta_d}(P^*) \leq \Omega$. Then by definition this placement forms a clique of size p in $G' := \text{bottleneck}(G, \delta_d, \Omega)$. Thus in this case V_{cand} is non-empty and the heuristic will not output a “certificate of failure”.

Moreover, any placement $P(v)$ considered by the heuristic will form a clique in $(G')^2$. By the definition of G' as a bottleneck graph with respect to δ_d , the bound Ω and by the assumption that edge weights obey triangle inequality, it follows that no edge e in $(G')^2$ has weight $\delta_d(e)$ more than 2Ω . Thus *every* placement $P(v)$ considered by the heuristic has communication diameter $\mathcal{D}_{\delta_d}(P(v))$ no more than 2Ω .

Consider an arbitrary $v \in P^*$. Clearly $v \in V_{\text{cand}}$. Consider the step of the algorithm HEU-FOR-DIA in which it considers v . For any $w \in N(v)$ we have $\delta_c(v, w) \leq OPT(I)$, by definition of $N(v)$ as the set of nearest neighbors of v and by the fact that every node from the optimal solution is adjacent to v in G' . Thus for $w, w' \in N(v)$ we have $\delta_c(w, w') \leq \delta_c(v, w) + \delta_c(v, w') \leq 2OPT(I)$ by the triangle inequality. Consequently, $\mathcal{D}_{\delta_c}(P(v)) = \mathcal{D}_{\delta_c}(N(v) \cup \{v\}) \leq 2OPT(I)$.

Now, since the algorithm HEU-FOR-DIA chooses a placement with minimal diameter among all the placements produced, the claimed performance guarantee with respect to the cost diameter \mathcal{D}_{δ_c} follows. \square

5 Sum Constrained Problems

Next, we study bicriteria compact location problems where the objective is to minimize the diameter \mathcal{D}_{δ_d} subject to budget-constraints of sum type.

Again, it is not an easy task to find a placement P satisfying the budget-constraint or to determine that no such placement exists. Using a reduction from CLIQUE [GJ79] one obtains the following.

Proposition 6. *If the distances δ_c, δ_d are not required to satisfy the triangle inequality, there can be no polynomial time (α, β) -approximation algorithm for SC-MDP for any fixed $\alpha, \beta \geq 1$, unless $\mathcal{P} = \mathcal{NP}$. Moreover, if there is a polynomial time $(\alpha, 1)$ -approximation algorithm for TI-SC-MDP for any fixed $\alpha \geq 1$, then $\mathcal{P} = \mathcal{NP}$.* \square

We proceed to present a heuristic for TI-SC-MDP. The main procedure shown in Figure 2 uses the test procedure from Figure 3.

PROCEDURE HEUR-FOR-SUM

1. Sort the edges of G in ascending order with respect to δ_d
2. Assume now that $\delta_d(e_1) \leq \delta_d(e_2) \leq \dots \leq \delta_d(e_{\binom{n}{2}})$
3. Let $P_{best} :=$ "certificate of failure"
4. $i := 1$
5. Do
 - (a) $G_i := \text{bottleneck}(G, \delta_d, \delta_d(e_i))$
 - (b) $P_{best} := \text{test}(G_i, \delta_c|_{G_i}, \Omega)$
 - (c) $i := i + 1$
6. **UNTIL** $P_{best} \neq$ "certificate of failure"
7. **OUTPUT** P_{best}

Fig. 2. Generic bottleneck procedure

PROCEDURE test(G, δ, Ω)

1. $V_{cand} := \{v \in G : \deg(v) \geq p - 1\}$
2. IF $V_{cand} = \emptyset$ THEN RETURN "certificate of failure"
3. Let $best := +\infty$
4. Let $P_{best} := \emptyset$
5. FOR each $v \in V_{cand}$ Do
 - (a) Let $N(v)$ be the set of $p - 1$ nearest neighbors of v in G with respect to δ
 - (b) Let $P(v) := N(v) \cup \{v\}$
 - (c) IF $S_\delta(P(v)) < best$ THEN $P_{best} := P(v)$
 $best := S_\delta(P(v))$
6. IF $best > (2 - 2/p)\Omega$ THEN RETURN "certificate of failure"
- ELSE RETURN P_{best}

Fig. 3. Test procedure used for TI-SC-MDP

Lemma 7. *Let I be an instance of TI-SC-MDP such that there is an optimal placement P^* . If the test procedure $\text{test}(G_i, \delta_c, \Omega)$ returns a "certificate of failure", then we have $OPT(I) > \delta_d(e_i)$. \square*

Now we can establish the result about the performance guarantee of the heuristic:

Theorem 8. *Let I denote any instance of TI-SC-MDP and assume that there is an optimal placement P^* of diameter $OPT(I) = \mathcal{R}_{\delta_d}(P^*)$. Then HEUR-FOR-SUM with the test procedure test returns a placement P with $S_{\delta_c}(P) \leq (2 - 2/p)\Omega$ and $\mathcal{D}_{\delta_1}(I)/OPT(I) \leq 2$.*

Proof: Consider the case when $\delta_d(e_i) = OPT(I)$. Since in G_i we have deleted only edges e having weight $\delta_d(e) > OPT(I)$ and we assume that there is a feasible solution satisfying the budget-constraint, it follows that the bottleneck graph G_i must contain a clique C of size p such that $S_{\delta_c}(C) \leq \Omega$.

For a node $v \in C$ let

$$S_v := \sum_{\substack{w \in C \\ w \neq v}} \delta_c(v, w).$$

Then we have

$$\mathcal{S}_{\delta_c}(C) = \sum_{v \in C} S_v.$$

Now let $v \in C$ be so that S_v is a minimum among all nodes in C . Then clearly

$$\mathcal{S}_{\delta_c}(C) \geq pS_v. \quad (1)$$

By definition of the bottleneck graph G_i and the clique C , the node v must have degree at least $p - 1$ in G_i . Thus v is one of the nodes considered by the test procedure. Let $N(v)$ be the set of $p - 1$ nearest neighbors of v in G_i . Then we have

$$\sum_{\substack{w \in N(v) \\ w \neq v}} \delta_c(v, w) \leq S_v, \quad (2)$$

by definition of $N(v)$ as the set of nearest neighbors, $P(v) := N(v) \cup \{v\}$. Let $w \in N(v)$ be arbitrary. Then

$$\begin{aligned} \sum_{u \in N(v) \cup \{v\} \setminus \{w\}} \delta_c(w, u) &= \delta_c(w, v) + \sum_{u \in N(v) \setminus \{w\}} \delta_c(w, u) \\ &\leq \delta_c(w, v) + \sum_{u \in N(v) \setminus \{w\}} (\delta_c(w, v) + \delta_c(v, u)) \\ &= (p - 1)\delta_c(w, v) + \sum_{u \in N(v) \setminus \{w\}} \delta_c(v, u) \\ &= (p - 2)\delta_c(v, w) + \sum_{u \in N(v)} \delta_c(v, u) \\ &\stackrel{(2)}{\leq} (p - 2)\delta_c(v, w) + S_v. \end{aligned} \quad (3)$$

Now using (3) and again (2), we obtain

$$\begin{aligned} \mathcal{S}_{\delta_c}(P(v)) &= \mathcal{S}_{\delta_c}(N(v) \cup \{v\}) \\ &= \sum_{u \in N(v)} \delta_c(v, u) + \sum_{w \in N(v)} \sum_{u \in N(v) \cup \{v\} \setminus \{w\}} \delta_c(w, u) \\ &\stackrel{(2)}{\leq} S_v + \sum_{w \in N(v)} \sum_{u \in N(v) \cup \{v\} \setminus \{w\}} \delta_c(w, u) \\ &\stackrel{(3)}{\leq} S_v + \sum_{w \in N(v)} ((p - 2)\delta_c(v, w) + S_v) \\ &= S_v + (p - 2)S_v + (p - 1)S_v \\ &= (2p - 2)S_v \\ &\stackrel{(1)}{\leq} (2 - 2/p)OPT(I). \end{aligned}$$

Thus the placement $P(v)$ violates the budget-constraint by a factor of at most $2 - 2/p$. Consequently, as the algorithm chooses the placement with P_{best} with the least constraint-violation, it follows that the test-procedure called with $G_i = \text{bottleneck}(G, \delta_d, OPT(I))$ will not return a “certificate of failure”.

The placement P_{best} that is produced by the algorithm turns into a clique in G_i^2 . Thus the longest edge in the placement with respect to δ_d is at most $2OPT(I)$. \square

References

- [ABP90] B. Awerbuch, A. Baratz, and D. Peleg. Cost-sensitive analysis of communication protocols. In *Proceedings of the 9th Symposium on the Principles of Distributed Computing (PODC)*, pages 177–187, 1990.
- [AF+94] E.M. Arkin, S.P. Fekete, J.S.B. Mitchell, and C.D. Piatko. Optimal Covering Tour Problems. In *Proceedings of the Fifth International Symposium on Algorithms and Computation*, 1994.
- [AI+91] A. Aggarwal, H. Imai, N. Katoh, and S. Suri. Finding k points with Minimum Diameter and Related Problems. *J. Algorithms*, 12(1):38–56, March 1991.
- [BP91] J. Bar-Ilan and D. Peleg. Approximation Algorithms for Selecting Network Centers. In *Proc. 2nd Workshop on Algorithms and Data Structures (WADS)*, pages 343–354, Ottawa, Canada, August 1991. Springer Verlag, LNCS vol. 519.
- [BS94] M. Bellare and M. Sudan. Improved Non-Approximability Results. in *Proceedings of the 26th annual ACM Symposium on the Theory of Computing (STOC)*, May 1994.
- [DF85] M.E. Dyer and A.M. Frieze. A Simple Heuristic for the p -Center Problem. *Operations Research Letters*, 3(6):285–288, Feb. 1985.
- [EN89] E. Erkut and S. Neuman. Analytical Models for Locating Undesirable Facilities. *European J. Operations Research*, 40:275–291, 1989.
- [FG88] T. Feder and D. Greene. Optimal Algorithms for Approximate Clustering. In *ACM Symposium on Theory of Computing (STOC)*, pages 434–444, 1988.
- [Fi93] T. Fischer. Optimizing the Degree of Minimum Weight Spanning Trees. Technical report, Department of Computer Science, Cornell University, Ithaca, New York, April 1993.
- [GJ79] M.R. Garey and D.S. Johnson. *Computers and Intractability*. W.H. Freeman, 1979.
- [Go85] T.F. Gonzalez. Clustering to Minimize the Maximum Intercluster Distance. *Theoretical Computer Science*, 38:293–306, 1985.
- [Ha92] R. Hassin. Approximation schemes for the restricted shortest path problem. *Mathematics of Operations Research*, 17(1):36–42, 1992.
- [HM79] G.Y. Handler and P.B. Mirchandani. *Location on Networks: Theory and Algorithms*. MIT Press, Cambridge, MA, 1979.
- [HS86] D. S. Hochbaum and D. B. Shmoys. A Unified Approach to Approximation Algorithms for Bottleneck Problems. *Journal of the ACM*, 33(3):533–550, July 1986.
- [IC+86] A. Iwainsky, E. Canuto, O. Taraszow, and A. Villa. Network decomposition for the optimization of connection structures. *Networks*, 16:205–235, 1986.

[KRY93] S. Khuller, B. Raghavachari, and N. Young. Balancing Minimum Spanning and Shortest Path Trees. In *Proceedings of the Fourth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 243–250, 1993.

[Lee82] D.T. Lee. On k -nearest neighbor Voronoi diagrams in the plane. *IEEE Trans. Comput.*, C-31:478–487, 1982.

[LV92] J.H. Lin and J. S. Vitter. ε -Approximations with Minimum Packing Constraint Violation. In *ACM Symposium on Theory of Computing (STOC)*, pages 771–781, May 1992.

[MR+95] M.V. Marathe, R. Ravi, R. Sundaram, S.S. Ravi, D.J. Rosenkrantz, and H.B. Hunt III. Bicriteria Network Design Problems. To appear in *Proceedings of the 22nd International Colloquium on Automata Languages and Programming (ICALP)*, 1995.

[MF90] P.B. Mirchandani and R.L. Francis. *Discrete Location Theory*. Wiley-Interscience, New York, NY, 1990.

[PS85] F.P. Preparata and M.I. Shamos. *Computational Geometry: An Introduction*. Springer-Verlag Inc., New York, NY, 1985.

[RKM⁺93] V. Radhakrishnan, S.O. Krumke, M.V. Marathe, D.J. Rosenkrantz, and S.S. Ravi. Compact Location Problems. In *13th Conference on the Foundations of Software Technology and Theoretical Computer Science (FST-TCS)*, volume 761 of *LNCS*, pages 238–247, December 1993.

[RMR⁺93] R. Ravi, M.V. Marathe, S.S. Ravi, D.J. Rosenkrantz, and H.B. Hunt III. Many birds with one stone: Multi-objective approximation algorithms. In *Proceedings of the 25th Annual ACM Symposium on the Theory of Computing (STOC)*, pages 438–447, 1993.

[RRT91] S.S. Ravi, D.J. Rosenkrantz, and G.K. Tayi. Facility Dispersion Problems: Heuristics and Special Cases. In *Proc. 2nd Workshop on Algorithms and Data Structures (WADS)*, pages 355–366, Ottawa, Canada, August 1991. Springer Verlag, LNCS vol. 519. (Journal version: *Operations Research*, 42(2):299–310, March-April 1994.)

[ST93] D.B. Shmoys and E. Tardos. Scheduling unrelated parallel machines with costs. In *Proceedings of the 4th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 438–447, 1993.

[Wa87] A. Warburton. Approximation of Pareto optima in multiple-objective, shortest path problems. *Operations Research*, 35:70–79, 1987.

[Won80] R.T. Wong. Worst case analysis of network design problem heuristics. *SIAM J. Disc. Math.*, 1:51–63, 1980.