Approximation Algorithms for Buy-at-Bulk Geometric Network Design*

Artur Czumaj[1], Jurek Czyzowicz[2], Leszek Gąsieniec[3], Jesper Jansson[4], Andrzej Lingas[5], and Pawel Zylinski[6]

[1] Centre for Discrete Mathematics and its Applications (DIMAP) and Department of Computer Science, University of Warwick, UK
A.Czumaj@warwick.ac.uk
[2] Departement d'Informatique, Universite du Quebec en Outaouais, Gatineau, Quebec J8X 3X7, Canada
Jurek.Czyzowicz@uqo.ca
[3] Department of Computer Science, University of Liverpool, Peach Street, L69 7ZF, UK
L.A.Gasieniec@liverpool.ac.uk
[4] Ochanomizu University, 2-1-1 Otsuka, Bunkyo-ku, Tokyo-112-8610, Japan
Jesper.Jansson@ocha.ac.jp
[5] Department of Computer Science, Lund University, 22100 Lund, Sweden
Andrzej.Lingas@cs.lth.se
[6] Institute of Computer Science, University of Gdańsk, 80-952 Gdańsk, Poland
Pawel.Zylinski@inf.univ.gda.pl

Abstract. The buy-at-bulk network design problem has been extensively studied in the general graph model. In this paper we consider the *geometric* version of the problem, where all points in a Euclidean space are candidates for network nodes. We present the first general approach for geometric versions of basic variants of the buy-at-bulk network design problem. It enables us to obtain quasi-polynomial-time approximation schemes for basic variants of the buy-at-bulk geometric network design problem with polynomial total demand. Then, for instances with few sinks and low capacity links, we design very fast polynomial-time low-constant approximations algorithms.

1 Introduction

Consider a water heating company that plans to construct a network of pipelines to carry warm water from a number of heating stations to a number of buildings. The company can install several types of pipes of various diameters and prices per unit length. Typically, the prices grow with the diameter while the ratio between the pipe throughput capacity and its unit price decreases. The natural goal of the company is to minimize the total cost of pipes sufficient to construct a network that could carry the warm water to the buildings, assuming a fixed water supply at each source. Similar problems can be faced by oil companies that need to transport oil to refineries or telecommunication companies that need to buy capacities (in bulk) from a phone company.

* Research supported in part by VR grant 621-2005-4085, the Royal Society IJP - 2006/R2, the Centre for Discrete Mathematics and its Applications, the Special Coordination Funds for Promoting Science and Technology (Japan), and the Visby Programme Scholarship 01224/2007.

F. Dehne et al. (Eds.): WADS 2009, LNCS 5664, pp. 168–180, 2009.

The common difficulty of these problems is that only a limited set of types of links (e.g., pipes) is available so the price of installing a link (or, a node respectively) to carry some volume of supply between its endpoints does not grow in linear fashion in the volume but has a discrete character. Even if only one type of link with capacity not less than the total supply is available the problem is NP-hard as it includes the minimum Steiner tree problem. Since the geometric versions of the latter problem are known to be strongly NP-complete [11], these problems cannot admit fully polynomial-time approximations schemes in the geometric setting [11].

In operations research, they are often termed as discrete cost network optimization [4,20] whereas in computer science as minimum cost network (or, link/edge) installation problems [23] or as *buy-at-bulk network design* [3]; we shall use the latter term.

In computer science, the buy-at-bulk network design problem has been introduced by Salman et al. [23], who argued that the case most relevant in practice is when the graph is defined by points in the Euclidean plane. Since then, various variants of buy-at-bulk network design have been extensively studied in the *graph model* [3,5,6,7,10,12,13,14,15,17,19] (rather than in geometric setting). Depending on whether or not the whole supply at each source is required to follow a single path to a sink they are characterized as *non-divisible* or *divisible* [23]. In terms of the warm water supply problem, the divisible graph model means that possible locations of the pipes and their splits or joints are given a priori.

In this paper, we consider the following basic *geometric* divisible variants of the buy-at-bulk network design:

- ▷ *Buy-at-bulk geometric network design (**BGND**)*: for a given set of different edge types and a given set of sources and sinks placed in a Euclidean space construct a minimum cost geometric network sufficient to carry the integral supply at sources to the sinks.
- ▷ *Buy-at-bulk single-sink geometric network design (**BSGND**)*: for a given set of different edge types, a given single-sink and given set of sources construct a minimum cost geometric network sufficient to carry the integral supply at sources to the sink.

Motivated by the practical setting in which the underlying network has to posses some basic structural properties, we distinguish also special versions of both problems where each edge of the network has to be parallel to one of the coordinate system axes, and term them as *buy-at-bulk rectilinear network design (**BRND**)* and *buy-at-bulk single-sink rectilinear network design (**BSRND**)*, respectively.

Our contributions and techniques. A classical approach for approximation algorithms for geometric optimization problems builds on the techniques developed for polynomial-time approximation schemes (PTAS) for geometric optimization problems due to Arora [1]. The main difficulty with the application of this method to the general BGND problem lies in the reduction of the number of crossings on the boundaries of the dissection squares. This is because we cannot limit the number of crossings of a boundary of a dissection square below the integral amount of supply it carries into that square. On the other hand, we can significantly limit the number of crossing locations at the expense of a slight increase in the network cost. However with this relaxed approach we cannot achieve polynomial but rather only quasi-polynomial upper bounds on the number of subproblems on the dissection squares in the dynamic programming phase but for very special cases (cf. [2]). Furthermore, the subproblems, in particular the leaf ones, become much more difficult. Nevertheless, we can solve them exactly in the case of BRND with polynomially bounded demands of the sources and nearly-optimally in the case of BGND with polynomially bounded demands of the sources and constant edge capacities, in at most quasi-polynomial time[1].

As the result, we obtain a randomized *quasi-polynomial-time approximation scheme* (QPTAS) for the *divisible buy-at-bulk rectilinear network design problem* in the Euclidean plane with polynomially bounded total supply and a randomized QPTAS for the *divisible buy-at-bulk network design problem* on the plane with polynomially bounded total supply and constant edge capacities. Both results can be derandomized and the rectilinear one can be generalized to include $O(1)$-dimensional Euclidean space. They imply that the two aforementioned variants of buy-at-bulk geometric network design are not *APX*-hard, unless $SAT \in DTIME[n^{\log^{O(1)} n}]$.

These two results are later used to prove our further results about low-constant-factor approximations for more general geometric variants. By using a method based on a novel belt decomposition for the single-sink variant, we obtain a $(2+\varepsilon)$ approximation to the divisible buy-at-bulk rectilinear network design problem in the Euclidean plane, which is fast if there are few sinks and the capacities of links are small; e.g., it runs in $n(\log n)^{O(1)}$ time if the number of sinks and the maximum link capacity are polylogarithmic in n. Similarly, we obtain a $(2+\varepsilon)$ approximation to the corresponding variants of the divisible buy-at-bulk network design problem in the Euclidean plane, which are fast if there are few sinks and the capacities of links are small, e.g., $n(\log n)^{O(1)}$-time if the number of sinks is polylogarithmic in n and maximum link capacity is $O(1)$. For comparison, the best known approximation factor for single-sink divisible buy-at-bulk network design in the graph model is 24.92 [13].

Related work. Salman et al. [23] initiated the algorithmic study of the single-sink buy-at-bulk network design problem. They argued that the problem is especially relevant in practice in the geometric case and they provided a polynomial-time approximation algorithm for the indivisible variant of BSGND on the input Euclidean graph (which differs from our model in that Salman et al. [23] allowed only some points on the plane to be used by the solution, whereas we allow the entire space to be used) with the approximation guarantee of $O(\log D)$, where D is total supply. Salman et al. gave also a constant factor approximation for *general graphs* in case where only one sink and one type of links is available; this approximation ratio has been improved by Hassin et al. [15]. Mansour and Peleg [18] provided an $O(\log n)$ approximation for the multi-sink buy-at-bulk network design problem when only one type of link is available. Awerbuch and Azar [3] were the first who gave a non-trivial (polylogarithmic) approximation for the general graph case for the total of n sinks and sources even in the case where different sources have to communicate with different sinks.

In the *single-sink buy-at-bulk* network design problem for general graphs, Garg et al. [12] designed an $O(K)$ approximation algorithm, where K is the number of edge types,

[1] Our solution method does not work in quasi-polynomial time in the case of the stronger version of BRND and BGND where specified sources must be assigned to specified sinks [3].

and later Guha et al. [14] gave the first constant-factor approximation algorithm for the (non-divisible) variant of the problem. This constant has been reduced in a sequence of papers [10,13,17,24] to reach the approximation ratio of 145.6 for the non-divisible variant and 24.92 for the divisible variant. Recently, further generalizations of the buy-at-bulk network design problem in the graph model have been studied [5,6].

2 Preliminaries

Consider a Euclidean d-dimensional space $\mathbb{E}^d$. Let $s_1, \ldots, s_{n_s}$ be a given set of n_s points in $\mathbb{E}^d$ (*sources*) and $t_1, \ldots, t_{n_t}$ be a given set of n_t points in $\mathbb{E}^d$ (*sinks*). Each source s_i supplies some integral *demand* $d(s_i)$ to the sinks. Each sink t_j is required to receive some integral *demand* $d(t_j)$. The sums $\sum_i d(s_i)$, $\sum_j d(t_j)$ are assumed to be equal and their value is termed as the *total demand* D. There are K types of edges, each type with a fixed cost and capacity. The *capacity* of an edge of type i is c_i and the *cost* of placing an edge e of ith type and length $|e|$ is $|e| \cdot \delta_i$.

The objective of the *buy-at-bulk geometric network design problem* (**BGND**) is to construct a geometric directed multigraph G in $\mathbb{E}^d$ such that:

- each copy of a multi-edge in the network is one of the K types;
- all the sources s_i and the sinks t_j belong to the set of vertices of G (the remaining vertices are called *Steiner vertices*);
- for $\ell = 1, \ldots, D$, there is a supply-demand path (*sd-path* for short) P_ℓ from a source s_i to a sink t_j such that each source s_i is a startpoint of $d(s_i)$ sd-paths, each sink t_j is an endpoint of $d(t_j)$ sd-paths, and for each directed multi-edge of the multigraph the total capacity of the copies of this edge is not less than the total number of sd-paths passing through it;
- the multigraph minimizes the total cost of the copies of its multi-edges.

If the set of sinks is a singleton then the problem is termed as the *buy-at-bulk single-sink geometric network design problem* (**BSGND** for short). If the multigraph is required to be rectilinear, i.e., only vertical and horizontal edges are allowed, then the problem is termed as the *buy-at-bulk rectilinear network design problem* (**BRND** for short) and its single-sink version is abbreviated as **BSRND**.

We assume, that the types of the edges are ordered $c_1 < \cdots < c_K$, $\delta_1 < \cdots < \delta_K$ and $\frac{\delta_1}{c_1} > \cdots > \frac{\delta_K}{c_K}$, since otherwise we can eliminate some types of the edges [23].

In this paper, we will always assume that the Euclidean space under consideration is a Euclidean plane $\mathbb{E}^2$, even though the majority of our results can be generalized to any Euclidean $O(1)$-dimensional space.

Zachariasen [25] showed that several variants and generalizations of the minimum rectilinear Steiner problem in the Euclidean plane are solvable on the *Hanan grid* of the input points, i.e., on the grid formed by the vertical and horizontal straight-lines passing through these points. The following lemma extends this to BRND.

Lemma 1. *Any optimal solution to BRND in the plane can be converted into a planar multigraph (so the sd-paths do not cross) where all the vertices lie on the Hanan grid.*

3 Approximating Geometric Buy-at-Bulk Network Design

In this section, we present our QPTAS for BRND and BGND. We begin with generalizations of several results from [1,22] about PTAS for TSP and the minimum Steiner tree in the plane. We first state a generalization of the Perturbation Lemma from [1,22].

Lemma 2. [22] *Let $G = (V, E)$ be a geometric graph with vertices in $[0, 1]^2$, and let $U \subseteq V$. Denote by $E(U)$ the set of edges incident to the vertices in U. One can perturb the vertices in U so they have coordinates of the form $(\frac{i}{k}, \frac{j}{k})$, where i, j are natural numbers not greater than a common natural denominator k, and the total length of G increases or decreases by an additive term of at most $\sqrt{2} \cdot |E|/k$.*

Consider an instance of BGND or BRND with sources $s_1 \dots s_{n_s}$ and sinks $t_1 \dots t_{n_t}$. We may assume, w.l.o.g., that the sources and the sinks are in $[0, 1)^2$.

Suppose that the total demand D is $n^{O(1)}$ where $n = n_s + n_t$. It follows that the maximum degree in a minimum cost multigraph solving the BGND or BRND is $n^{O(1)}$. Hence, the total number of copies of edges incident to the sources and sinks in the multigraph is also, w.l.o.g., $n^{O(1)} = n^{O(1)} \times n$. In the case of BRND, we infer that even the total number of copies of edges incident to all vertices, i.e., including the Steiner points, is, w.l.o.g., $n^{O(1)} = n^{O(1)} \times O(n^2)$ by Lemma 1.

Let $\delta > 0$. By using a straightforward extension of Lemma 2 to include a geometric multigraph and rescaling by $L = \frac{n^{O(1)}}{\delta}$ the coordinates of the sources and sinks, we can alter our BGND or BRND with all vertices on the Hanan grid such that:

- the sources and sinks of the BGND and BRND as well as the Steiner vertices of the BRND lie on the integer grid in $[0, L)^2$, and
- for any solution to the BGND with the original sources and sites (or, BRND with all vertices on the Hanan grid) and for any type of edge, the total length of copies of edges of this type in the solution resulting for the BGND with the sources and sinks on the integer grid (or, for BRND with all vertices on the integer grid, respectively) is at most $L(1 + \delta)$ times larger, and
- for any solution to the BGND with the sources and sinks on the integer grid (or, for BRND with all vertices on the integer grid, respectively), the total length of copies of edges of this type in the solution resulting for the BGND with the original sources and sites (or, BRND with all vertices on the Hanan grid) is at most $(1 + \delta)/L$ times larger.

Note the second and the third properties imply that we may assume further that our input instance of BGND has sources and sinks on the integer grid in $[0, L)^2$, since this assumption introduces only an additional $(1 + \delta)$ factor to the final approximation factor. We shall call this assumption the **rounding assumption**. In the case of BRND, we may assume further, w.l.o.g., not only that our input instance has sources and sinks on the integer grid but also that Steiner vertices may be located only on this grid by the second and third property, respectively. This stronger assumption in the case of BRND introduces also only an additional $(1 + \delta)$ factor to the final approximation factor by the aforementioned properties. We shall term it the **strong rounding assumption**.

Now we pick two integers a and b uniformly at random from $[0, L)$ and extend the grid by a vertical grid lines to the left and $L - a$ vertical grid lines to the right. We similarly increase the height of the grid using the random integer b, and denote the obtained grid by $L(a, b)$. Next, we define the recursive decomposition of $L(a, b)$ by dissection squares using quadtree. The dissection quadtree is a 4-ary tree whose root corresponds to the square $L(a, b)$. Each node of the tree corresponding to a dissection square of area greater than 1 is dissected into four child squares of equal side length; the four child squares are called *siblings*. The obtained quadtree decomposition is denoted by $Q(a, b)$.

We say a graph G is *r-light* if it crosses each boundary between two sibling dissection squares of $Q(a, b)$ at most r times. A multigraph H is *r-fine* if it crosses each boundary between two sibling dissection squares of $Q(a, b)$ in at most r places. For a straight-line segment ℓ and an integer r, an *r-portal* of ℓ is any endpoint of any of the r segments of equal length into which ℓ can be partitioned.

3.1 QPTAS for Buy-at-Bulk Rectilinear Network Design (BRND)

We obtain the following new theorem which can be seen as a generalization of the structure theorem from [1] to include geometric multigraphs, where the guarantee of r-lightness is replaced by the weaker guarantee of r-fineness.

Theorem 1. *For any $\varepsilon > 0$ and any BRND (or BGND, respectively) on the grid $L(a, b)$, there is a multigraph on $L(a, b)$ crossing each boundary between two sibling dissection squares of $Q(a, b)$ only at $O(\log L/\varepsilon)$-portals, being a feasible solution of BRND (BGND, respectively) and having the expected length at most $(1 + \varepsilon)$ times larger than the minimum.*

To obtain a QPTAS for an arbitrary BRND with polynomial total demand in the Euclidean plane it is sufficient to show how to find a minimum cost multigraph for BRND on $L(a, b)$ which crosses each boundary between two sibling dissection squares of $Q(a, b)$ only at r-portals efficiently, where $r = O(\log n/\varepsilon)$.

We specify a subproblem in our dynamic programming method by a dissection square occurring in some level of the quadtree $Q(a, b)$, a choice of crossing points out of the $O(r)$-portals on the sides of the dissection square, and for each of the chosen crossing points p, an integral demand $d(p)$ it should either supply to or receive from the square (instead of the pairing of the distinguished portals [1]). By the upper bound $D \leq n^{O(1)}$, we may assume, w.l.o.g., that $d(p) = n^{O(1)}$. Thus, the total number of such different subproblem specifications is easily seen to be $n^{O(r)}$. The aforementioned subproblem consists of finding a minimum cost r-fine rectilinear multigraph for the BRND within the square, where the sources are the original sources within the square and the crossing points expected to supply some demand whereas the sinks are the original sinks within the square and the crossing points expected to receive some demand.

Each leaf subproblem, where the dissection square is a cell of $L(a, b)$ and the original sources and sinks may be placed only at the corners of the dissection square, and the remaining $O(r)$ ones on the boundary of the cell, can be solved by exhaustive search and dynamic programming as follows. By Lemma 1, we may assume, w.l.o.g., that an optimal solution of the subproblem is placed on the Hanan $O(r) \times O(r)$ grid. We enumerate all directions and total capacity assignments to the edges of the grid in time

$n^{O(r)}$ by using the $n^{O(1)}$ bound on the total demand. For each such grid edge with non-zero total capacity assigned, we find (if possible) the cheapest multi-covering of this capacity with different edge types with capacity bounded by the total demand by using a pseudo-polynomial time algorithm for the integer knapsack problem [11]. Next, we compare the cost of such optimal multi-covering with the cost of using a single copy of the cheapest edge type whose capacity exceeds the total demand (if any) to choose an optimal solution. It follows that all the leaf subproblems can be solved in time $n^{O(r^2)}$.

Then, we can solve subproblems corresponding to consecutive levels of the quadtree $Q(a, b)$ in a bottom up fashion by combining optimal solutions to four compatible subproblems corresponding to the four dissection squares which are children of the dissection square in the subproblem to solve. The compatibility requirement is concerned with the location of the crossing points and their demand requirements. Since there are $n^{O(r)}$ subproblems, solution of a single subproblem also takes $n^{O(r)}$ time.

The bottleneck in the complexity of the dynamic programming are the leaf subproblems. If we could arbitrarily closely approximate their solutions in time $n^{O(r)}$ then we could compute a minimum cost r-fine multigraph for BRND on $L(a, b)$ with polynomially bounded total demand in time $n^{O(r)}$. The following lemma will be helpful.

Lemma 3. *For any $\varepsilon > 0$, one can produce a feasible solution to any leaf subproblem which is within $(1 + \varepsilon)$ from the minimum in time $n^{O(\log^2 r)}$.*

By halving ε both in the dynamic programming for the original problem as well as in Lemma 3 and using the method of this lemma to solve the leaf subproblems, we obtain the following lemma.

Lemma 4. *A feasible r-fine multigraph for BRND on $L(a, b)$ with polynomially bounded total demand and total cost within $1 + \varepsilon$ from the optimum is computable in time $n^{O(r)}$.*

By combining Theorem 1 with Lemma 4 for $r = O(\frac{\log n}{\varepsilon})$ and the fact that the rounding assumption introduces only an additional factor of $(1 + O(\varepsilon))$ to the approximation factor, we obtain our first result.

Theorem 2. *For any $\varepsilon > 0$, there is a randomized $n^{O(\log n/\varepsilon)}$-time algorithm for BRND in the Euclidean plane with a total of n sources and sinks and total demand polynomial in n, which yields a solution whose expected cost is within $(1 + \varepsilon)$ of the optimum.*

Theorem 2 immediately implies the following result for BGND (which will be substantially subsumed in Section 3.2 in the case of constant maximum edge capacity).

Corollary 1. *For any $\varepsilon > 0$, there is a randomized $n^{O(\log n/\varepsilon)}$-time algorithm for BGND in the Euclidean plane with the total of n sources and sinks and with polynomial in n total demand, which yields a solution whose expected cost is within $(\sqrt{2} + \varepsilon)$ from the optimum.*

3.2 QPTAS for the Buy-at-Bulk Geometric Network Design Problem (BGND)

We can arbitrarily closely approximate BGND analogously as BRND if it is possible to solve or very closely approximate the leaf subproblems where all the sources and sinks are placed in $O(\log n/\varepsilon)$ equidistant portals on a boundary of a dissection square, and feasible solutions are restricted to the square area. Note that such a leaf subproblem is logarithmic as for the number of sources and sinks but the total capacity of its sources or sinks might be as large as the total capacity D of all sources. We shall assume D to be polynomial in the number of sinks an sources as in the previous section.

By an *h-square BGND*, we mean BGND restricted to instances where h sources and sinks are placed on a boundary of a square. By a *logarithmic square BGND*, we mean an *h-square BGND* where the total demand of the sources is $O(\log n)$.

Lemma 5. *If there is an $n^{O(\log n)}$-time approximation scheme for a logarithmic square BGND then there is an $n^{O(\log n)}$-time approximation scheme for an $O(\log n)$-square BGND with maximum edge capacity $O(1)$.*

Proof. Let D denote the total capacity of the sources in the h-square BGND, where $h = O(\log n)$. Consider an optimal solution to the h-square BGND. It can be decomposed into D sd-paths, each transporting one unit from a source to a sink. There are $O(h^2)$ types of the sd-paths in one-to-one correspondence with the $O(h^2)$ pairs source-sink. Analogously as in the rectilinear case (see Lemma 1), we may assume, w.l.o.g., that the sd-paths do not intersect and that that the minimum edge capacity is 1. Let M be the maximum edge capacity in the h-square BGND.

For a type t of sd-path, let N_t be the number of sd-paths of type t in the optimal solution. Since these sd-paths do not intersect, we can number them, say, in the cyclic ordering around their common source, with the numbers in the interval $[1, N_t]$. Note that each of these paths whose number is in the sub-interval $[M, N_t - M + 1]$ can use only edges which are solely used by sd-paths of this type in the optimal solution. Let $k = \lfloor \frac{1}{\varepsilon} \rfloor$, and let ϱ be the ratio between the cost δ_1 (per length unit) of an edge of capacity 1 and the cost δ_{max} of an edge of the maximum capacity M divided by M. Suppose that $N_t \geq M + \varrho kM + 2(M-1)$. Let $q = \lceil (N_t - 2(M-1))/M \rceil$.

Consider the following modification of the optimal solution. Group the consecutive bunches of M sd-paths of type t in the sub-interval $[M, qM - 1]$, and direct them through q directed edges of capacity M from the source to the sink corresponding to the type t. Remove all edges in the optimal solution used by these sd-paths in this sub-interval. Note that solely at most $M - 1$ sd-paths of the type t immediately to the left of $[M, N_t - M + 1]$ as well as at most $M - 1$ sd-paths of the type t immediately to the right of this interval can loose their connections to the sink in this way. Independently of whether such a path looses its connection or not, we direct it through a direct edge of capacity 1 from the source to the sink.

The total cost of the directed edges of capacity M from the source to the sink in the distance d is $q\delta_{max}d$. It yields the lowest possible cost per unit, sent from the source to the sink corresponding to t, equal to $\frac{\delta_{max}}{M}d$. Thus the total cost of the removed edges must be at least $q\delta_{max}d$! The additional cost of the $2(M-1)$ direct edges of capacity 1 from the source to the sink is $\leq \varepsilon$ fraction of $q\delta_{max}d$ by our assumption on N_t.

By starting from the optimal solution and performing the aforementioned modification of the current solution for each type t of sd-path satisfying $N_t \geq M + \varrho kM + 2(M-1)$, we obtain a solution which is at most $(1+\varepsilon)$ times more costly than the optimal, and which is decomposed into two following parts. The first, explicitly given part includes all sd-paths of type t satisfying $N_t \geq M + \varrho kM + 2(M-1)$ whereas the second unknown part includes all paths of types t satisfying $N_t < \varrho kM + 2(M-1)$. It follows that it is sufficient to have an $(1+\varepsilon)$-approximation of an optimal solution to the logarithmic square BGND problem solved by the second part in order to obtain an $(1+O(\varepsilon))$-approximation to the original h-square BGND. □

Lemma 6. *For any $\varepsilon > 0$, the logarithmic square BGND problem with the total capacity of the sources D can be $(1+\varepsilon)$-approximated in time $(D/\varepsilon)^{O(D(\log D/\varepsilon))}$ if $c_{max} = O(1)$.*

By combining Lemma 5 with Lemma 6 for $D = O(\log n/\varepsilon)$ and straightforward calculations, we obtain an arbitrarily close to the optimum solutions to the $n^{O(\log n/\varepsilon)}$ leaf problems in total time $n^{O(\log n/\varepsilon^{O(1)})}$. Hence, analogously as in case of BRND, we obtain a QPTAS for BGND with polynomially bounded demand when $c_{max} = O(1)$.

Theorem 3. *BGND with polynomially bounded demand of the sources and constant maximum edge capacity admits an $n^{O(\log n)}$-time approximation scheme.*

4 Fast Low-Constant Approximation for BRND and BGND

In this section, we present another method for BGND and BRND which runs in polynomial time, gives a low-constant approximation guarantee, and does not require a polynomial bound on the total demand. The method is especially efficient if the edge capacities are small and there are few sinks.

We start with the following two simple lemmas. The first lemma is analogous to the so-called routing lower bound from [18,23] and the second follows standard arguments.

Lemma 7. *Let S be the set of sources in an instance of BGND (BRND), and for each $s \in S$, let $t(s)$ be the closest sink in this instance. The cost of an optimal solution to the BGND (BRND, respectively) is at least $\sum_{s\in S} dist(s,t(s)) \frac{\delta_K}{c_K} d(s)$, where $dist(s,t(s))$ is the Euclidean distance (the L_1 distance, respectively).*

Lemma 8. *Let S be a set of k points within a square of side length ℓ. One can find in time $O(k)$ a Steiner tree of S with length $O(\ell\sqrt{k})$.*

The following lemma describes a simple reduction procedure which yields an *almost* feasible solution to BSGND or BSRND with cost arbitrarily close to the optimum.

Lemma 9. *For any $\varepsilon > 0$, there is a reduction procedure for BSGND (or BSRND, respectively), with one sink and $n-1$ sources and the ratio between the maximum and minimum distances of a source from the sink equal to m, which returns a multigraph yielding a partial solution to the BSGND (or BSRND, respectively) satisfying the following conditions:*

- *all but $O((\frac{1}{\varepsilon})^2 c_K^2 \log m)$ sources can ship their whole demand to the sink;*
- *for each source s there are at most $c_K - 1$ units of its whole demand $d(s)$ which cannot be shipped to the sink.*

The reduction runs in time $O(\frac{c_K}{\varepsilon} \log m \log n + c_K n)$, which is $O(n/\varepsilon^2)$ if $c_K = O(1)$.

Proof. Form a rectilinear $2\lceil m\rceil \times 2\lceil m\rceil$ grid F with unit distance equal to the minimum distance between the only sink t and a source, centered around t. Let μ be a positive constant to be set later.

We divide F into the square R of size $2\lceil \mu\sqrt{c_K}\rceil$ centered in t and for $i = 0, 1, \ldots,$ the *belts* B_i of squares of size 2^i within the L_∞ distance at least $2^i\lceil \mu\sqrt{c_K}\rceil$ and at most $2^{i+1}\lceil \mu\sqrt{c_K}\rceil$ from t. Note that the number of squares in the belt B_i is at most $(4\lceil \mu\sqrt{c_K}\rceil)^2 = O(\mu^2 c_K)$, hence the total number of squares in all the belts is $O(\mu^2 c_K \log m)$ by the definition of the grid.

The reduction procedure consists of two phases. In the first phase, we connect each source s by a multi-path composed of $\lfloor d(s)/c_K\rfloor$ copies of a shortest path from s to t implemented with the K-th type of edges. Observe that the average cost of such a connection per each of the $c_K\lfloor d(s)/c_K\rfloor$ demand units u shipped from s to t is $dist(s,t)\frac{\delta_K}{c_K}$ which is optimal by Lemma 7. Note that after the first phase the remaining demand for each source is at most $c_K - 1$ units.

In the second phase, for each of the squares Q in each of the belts B_i, we sum the remaining demands of the sources contained in it, and for each complete c_K-tuple of demand units in Q, we find a minimum Steiner multi-tree of their sources and connect its vertex v closest to t by a shortest path to t. The total length of the resulting multi-tree is easily seen to be $dist(v,t) + O(2^i\sqrt{c_K}) \leq (1 + O(\frac{1}{\mu}))dist(v,t)$ by the definition of the squares and Lemma 8. Hence, for each unit u in the c_K-tuple originating from its source $s(u)$, we can assign the average cost of connection to t by the multi-tree implemented with the K-th type of edges not greater than $(1 + O(\frac{1}{\mu}))dist(s(u),t)\frac{\delta_K}{c_K}$.

It follows by Lemma 7 that the total cost of the constructed network is within $(1+O(\frac{1}{\mu}))$ from the minimum cost of a multigraph for the input BSGND. By choosing μ appropriately large, we obtain the required $1 + \varepsilon$-approximation.

Since the total number of squares different from R is $O(\mu^2 c_K \log m)$, the total number of their sources with a non-zero remaining demand (at most $c_K - 1$ units) to ship is $O(\mu^2 c_K^2 \log m)$. Furthermore, since the square R can include at most $O(\mu^2 c_K)$ sources, the number of sources with a non-zero remaining demand (at most $c_K - 1$ units) in R is only $O(\mu^2 c_K)$.

The first phase can be implemented in time linear in the number of sources. The second phase requires $O(\mu^2 c_K \log m)$ range queries for disjoint squares and $O(c_K n/c_K)$ constructions of Steiner trees on c_K vertices using the method of Lemma 8. Thus it needs $O(\mu^2 c_K \log m \log n + c_K n)$ time by [21] and Lemma 8. Since, w.l.o.g, $\mu = O(\frac{1}{\varepsilon})$, we conclude that the whole procedure takes $O(\frac{c_K}{\varepsilon^2}\log m \log n + c_K n)$ time. □

Extension to BRND and BGND. We can generalize our reduction to include n_t sinks by finding the Voronoi diagram in the L_2 (or L_1 for BGND) metric on the grid, locating each source in the region of the closest sink, and then running the reduction procedure separately on each set of sources contained in a single region of the Voronoi diagram. The construction of the Voronoi diagram and the location of the sources takes time $O(n\log n)$ (see [16,21]). The n_t runs of the reduction procedure on disjoint sets of sources takes time $O((\frac{1}{\varepsilon})^2 n_t c_K \log m \log n + c_K n)$. The union of the n_t resulting multigraphs may miss to ship the whole demand only from $O((\frac{1}{\varepsilon})^2 n_t c_K^2 \log m)$ sources. This gives the following generalization of Lemma 9.

Lemma 10. *For any $\varepsilon > 0$, there is a reduction procedure for BGND (or BRND, resp.), with n_t sinks and $n - n_t$ sources and the ratio between the maximum and minimum distances of a source from the sink equal to m, which returns a multigraph yielding a partial solution to the BGND (or BRND, resp.) satisfying the following conditions:*

- *all but $O((\frac{1}{\varepsilon})^2 n_t c_K^2 \log m)$ sources can ship their whole demand to the sink;*
- *for each source s there are at most $c_K - 1$ units of its whole demand $d(s)$ which cannot be shipped to the sink.*

The reduction procedure runs in time $O((\frac{1}{\varepsilon})^2 n_t c_K \log m \log n + n(c_K + \log n))$. In particular, if $c_K = (\log n)^{O(1)}$ then the running time is $(\frac{1}{\varepsilon})^2 n \log m (\log n)^{O(1)}$.

Now, we are ready to derive our main results in this section.

Theorem 4. *For any $\varepsilon > 0$, there is a $(2+\varepsilon)$-approximation algorithm for BRND with n_t sinks and $n - n_t$ sources in the Euclidean plane, running in time $O((\frac{1}{\varepsilon})^2 n_t c_K \log^2 n + n(\log n + c_K)) + (n_t c_K^2 \log n)^{O(\frac{\log n_t + \log c_K}{\varepsilon^2})}$, in particular in time $n(\log n)^{O(1)} + (\log n)^{O(\frac{\log\log n}{\varepsilon^2})}$ if $n_t = (\log n)^{O(1)}$ and $c_K = (\log n)^{O(1)}$.*

Proof. By the rounding assumption discussed in Section 3 we can perturb the sinks and the sources so they lie on an integer grid of polynomial size introducing only an additional $(1 + O(\varepsilon))$ factor to the final approximation factor. The perturbation can be easily done in linear time. Next, we apply the reduction procedure from Lemma 10 to obtain an almost feasible solution of total cost not exceeding $(1 + O(\varepsilon))$ of that for the optimal solution to the BSRND on the grid. Note that $m \le n^{O(1)}$ and hence $\log m = O(\log n)$ in this application of the reduction by the polynomiality of the grid. It remains to solve the BRND subproblem for the $O((\frac{1}{\varepsilon})^2 n_t c_K \log n)$ remaining sources with total remaining demand polynomial in their number. This subproblem can be solved with the randomized $(1+O(\varepsilon))$-approximation algorithm of Theorem 2. In fact, we can use here also its derandomized version which will run in time $(n_t c_K^2 \log n)^{O(\frac{\log n_t + \log c_K}{\varepsilon^2})}$. □

As an immediate corollary from Theorem 4, we obtain a $(\sqrt{8} + \varepsilon)$-approximation algorithm for BGND with n_t sinks and $n - n_t$ sources in the Euclidean plane, running in time $O((\frac{1}{\varepsilon})^2 n_t c_K \log^2 n + n(\log n + c_K)) + (n_t c_K^2 \log n)^{O(\frac{\log n_t + \log c_K}{\varepsilon^2})}$. However, the direct method analogous to that of Theorem 4 yields a better approximation, in particular also an $(2 + \varepsilon)$-approximation if $c_K = O(1)$.

Theorem 5. *For any $\varepsilon > 0$, there is a $(1 + \sqrt{2} + \varepsilon)$-approximation algorithm for BGND with n_t sinks and $n - n_t$ sources in the Euclidean plane, running in time $O((\frac{1}{\varepsilon})^2 n_t c_K \log^2 n + n(\log n + c_K)) + (n_t c_K^2 \log n)^{O(\frac{\log n_t + \log c_K}{\varepsilon^2})}$; the running time is $n(\log n)^{O(1)} + (\log n)^{O(\frac{\log\log n}{\varepsilon^2})}$ if $n_t = (\log n)^{O(1)}$ and $c_K = (\log n)^{O(1)}$. Furthermore, if $c_K = O(1)$ then the approximation factor of the algorithm is $2 + \varepsilon$.*

5 Final Remarks

We have demonstrated that BRND and BGND in a Euclidean space admit close approximation under the assumption that the total demand is polynomially bounded. By running the first phase of the reduction procedure from Lemma 9 as a preprocessing, we could get rid of the latter assumption at the expense of worsening the approximation factors by the additive term 1.

All our approximation results for different variants of BRND in Euclidean plane derived in this paper can be generalized to include corresponding variants of BRND in a Euclidean space of fixed dimension. All our approximation schemes are randomized but they can be derandomized similarly as those in [1,8,9,22].

References

1. Arora, S.: Polynomial time approximation schemes for Euclidean traveling salesman and other geometric problems. Journal of the ACM 45(5), 753–782 (1998)
2. Arora, S., Raghavan, P., Rao, S.: Approximation schemes for Euclidean k-medians and related problems. In: Proc. 30th ACM STOC, pp. 106–113 (1998)
3. Awerbuch, B., Azar, Y.: Buy-at-bulk network design. In: Proc 38th IEEE FOCS, pp. 542–547 (1997)
4. Bienstock, D., Chopra, S., Günlük, O., Tsai, C.-Y.: Minimum cost capacity installation for multicommodity network flows. Mathematical Programming 81(2), 177–199 (1998)
5. Chekuri, C., Hajiaghayi, M.T., Kortsarz, G., Salavatipour, M.R.: Polylogarithmic approximation algorithms for non-uniform multicommodity buy-at-bulk network design. In: Proc. 47th IEEE FOCS, pp. 677–686 (2006)
6. Chekuri, C., Hajiaghayi, M.T., Kortsarz, G., Salavatipour, M.R.: Approximation algorithms for node-weighted buy-at-bulk network design. In: SODA, pp. 1265–1274 (2007)
7. Chopra, S., Gilboa, I., Sastry, S.T.: Source sink flows with capacity installation in batches. Discrete Applied Mathematics 85(3), 165–192 (1998)
8. Czumaj, A., Lingas, A.: On approximability of the minimum-cost k-connected spanning subgraph problem. In: Proc. 10th IEEE-SIAM SODA, pp. 281–290 (1999)
9. Czumaj, A., Lingas, A.: Fast approximation schemes for euclidean multi-connectivity problems. In: Welzl, E., Montanari, U., Rolim, J.D.P. (eds.) ICALP 2000. LNCS, vol. 1853, pp. 856–868. Springer, Heidelberg (2000)
10. Gupta, A., Kumar, A., Roughgarden, T.: Simpler and better approximation algorithms for network design. In: Proc. 35th ACM STOC, pp. 365–372 (2003)
11. Garey, M.R., Johnson, D.S.: Computers and Intractability. A Guide to the Theory of NP-completeness. W.H. Freeman and Company, New York (1979)
12. Garg., N., Khandekar, R., Konjevod, G., Ravi, R., Salman, F.S., Sinha, A.: On the integrality gap of a natural formulation of the single-sink buy-at-bulk network design problem. In: Proc. 8th IPCO, pp. 170–184 (2001)
13. Grandoni, F., Italiano, G.F.: Improved approximation for single-sink buy-at-bulk. In: Asano, T. (ed.) ISAAC 2006. LNCS, vol. 4288, pp. 111–120. Springer, Heidelberg (2006)
14. Guha, S., Meyerson, A., Munagala, K.: A constant factor approximation for the single sink edge installation problem. In: Proc. 33rd ACM STOC, pp. 383–399 (2001)
15. Hassin, R., Ravi, R., Salman, F.S.: Approximation algorithms for a capacitated network design problem. In: Jansen, K., Khuller, S. (eds.) APPROX 2000. LNCS, vol. 1913, pp. 167–176. Springer, Heidelberg (2000)
16. Hwang, F.K.: An O(nlogn) algorithm for rectilinear minimal spanning trees. JACM 26(2), 177–182 (1979)
17. Jothi, R., Raghavachari, B.: Improved approximation algorithms for the single-sink buy-at-bulk network design problems. In: Hagerup, T., Katajainen, J. (eds.) SWAT 2004. LNCS, vol. 3111, pp. 336–348. Springer, Heidelberg (2004)

18. Mansour, Y., Peleg, D.: An approximation algorithm for minimum-cost network design. Technical report, The Weizman Institute of Science, Revohot, Israel, CS94-22 (1994)
19. Meyerson, A., Munagala, K., Plotkin, S.: COST-DISTANCE: Two metric network design. In: Proc. 41st IEEE FOCS, pp. 624–630 (2000)
20. Minoux, M.: Discrete cost multicommodity network optimization problems and exact solution methods. Annals of Operations Research 106, 19–46 (2001)
21. Preparata, F., Shamos, M.: Computational Geometry. Springer, New York (1985)
22. Rao, S.B., Smith, W.D.: Approximating geometrical graphs via "spanners" and "banyans". In: Proc. ACM STOC, pp. 540–550 (1998); Full version appeared as TR, NEC (1998)
23. Salman, F.S., Cheriyan, J., Ravi, R., Subramanian, S.: Approximating the single-sink link-installation problem in network design. SIAM J. Optimization 11(3), 595–610 (2000)
24. Talwar, K.: Single-sink buy-at-bulk LP has constant integrality gap. In: Cook, W.J., Schulz, A.S. (eds.) IPCO 2002. LNCS, vol. 2337. Springer, Heidelberg (2002)
25. Zachariasen, M.: A catalog of Hanan grid problems. Networks 38(2), 76–83 (2001)