OPERATIONS   RESEARCH   CENTER

Working Paper

Separable Concave Optimization Approximately Equals
Piecewise Linear Optimization

by
Thomas L. Magnanti
Dan Stratila

MASSACHUSETTS INSTITUTE
OF TECHNOLOGY

Separable Concave Optimization Approximately Equals Piecewise Linear Optimization[*]

Thomas L. Magnanti[†]   Dan Stratila[‡]

March 10, 2006

Abstract

Consider a separable concave minimization problem with nondecreasing costs over a general ground set $X \subseteq \mathbb{R}_+^n$. We show how to efficiently approximate this problem to a factor of $1 + \epsilon$ in optimal cost by a single piecewise linear minimization problem over X. The number of pieces is linear in $1/\epsilon$ and polynomial in the logarithm of certain ground set parameters; in particular, it is independent of the cost functions. Our main result is that when the minimization is over a polyhedron, the number of pieces, and thus the size of the resulting problem, is polynomial in the input size of the polyhedron and linear in $1/\epsilon$. We present generalizations to problems with grounds sets not contained in $\mathbb{R}_+^n$ and concave functions that are not monotone.

Our approach provides a general technique for applying discrete optimization methods to practical concave cost problems with polyhedral ground sets. We exemplify the approach on two problems. For the concave cost multicommodity flow problem, we devise an approximate computational solution procedure using our technique and a primal-dual solution procedure. We are able to solve randomly generated instances significantly larger than previously possible, and obtain solutions within 4% of optimality on average. For the lot-sizing problem with concave production costs, we derive an algorithm with a new polynomial running time that is not dominated by that of previously known algorithms.

1   Introduction

Minimizing a separable concave function over a polyhedron arises frequently in fields such as transportation, logistics, supply chain management, and telecommunications. In a typical setting, the polyhedral ground set arises due to network structure, capacity requirements, and other constraints, while the concave costs arise due to economies of scale, volume discounts, and other economic factors [see e.g. GP90]. The concave functions can be nonlinear, consist of many pieces, or be given by an oracle.

---

[*]An extended abstract of this research has appeared in [MS04].

[†]School of Engineering and Sloan School of Management, Massachusetts Institute of Technology, Room 1-206, 77 Massachusetts Avenue, Cambridge, MA 02139. E-mail: `magnanti@mit.edu`.

[‡]Operations Research Center, Massachusetts Institute of Technology, Room E40-130, 77 Massachusetts Avenue, Cambridge, MA 02139. E-mail: `dstrat@mit.edu`.

A natural approach for solving these problems is to replace the general cost functions by piecewise linear approximations, an idea known at least since the 1950s [see e.g. Dan63]. Problems with piecewise linear concave costs can in turn be reduced to problems with fixed charge cost functions, which consist of a fixed cost plus a per-unit cost [see e.g. NW99]. Researchers have successfully treated the fixed charge problems using combinatorial optimization and integer programming approaches [e.g. BMW89, HH98]. Recently researchers have achieved significant further advances using new techniques in integer programming [e.g. Ata01, OW03] and approximation algorithms [e.g. JMM$^+$03].

The methods for problems with fixed charge and piecewise linear costs would automatically become promising methods for problems with general separable concave costs, if we could approximate the latter by a single piecewise linear problem with few pieces, and provide an approximation guarantee in terms of optimal cost. However, current piecewise linear approximation approaches either yield a large number of pieces, or do not provide a good approximation guarantee. In fact, we are not aware of any non-trivial bounds on the approximation guarantee in terms of the number of pieces for general separable concave functions.

In this paper, we provide improved methods for approximating separable concave cost problems, and thereby reduce the gap between them and solution methods for fixed-charge and piecewise linear cost problems. We provide theoretical results on approximating separable concave functions in the context of a general minimization problem, efficient worst-case bounds for problems with polyhedral ground sets, and computational as well as algorithmic applications to specific problems.

1.1   Previous Work

Clearly, to improve the quality of the approximation, we would increase the number of pieces; however not much is known about the number of pieces required for a single approximation to attain a desired precision in the general case. Rosen and Pardalos [RP86] consider the minimization of a quadratic concave function over a polyhedron. They reduce the problem to a separable quadratic concave minimization problem over a polyhedron, and then study piecewise linear approximations of the resulting univariate concave functions. They interpolate the functions at equally-spaced intervals and obtain an approximation guarantee that is function-dependent. For a fixed ϵ, the size of the resulting problem is not polynomial in the size of the original problem.

Hajiaghayi et al [HMM03] consider the unit-demand concave cost facility location problem, and use the fact that all n facilities have unit demand to obtain an exact reduction by interpolating the concave functions at points $1, 2, \ldots, n$. The size of the resulting problem is polynomial in the size of the original problem, but the approach is limited to unit-demand problems. Meyerson et al [MMP00], in the context of the single-sink concave cost multicommodity flow problem, remark that a "tight" approximation could be computed. Munagala [Mun03] states, in the same context, that an approximation of arbitrary precision could be obtained with a polynomial number of pieces. They do not mention specific bounds, or any details on how to do so.

A significant body of work on approximating separable general objectives with linear pieces has focused on *convex* functions, for which a scale-and-iterate approach is prevalent:

using an equally spaced grid, solve the approximate problem, then iteratively approximate the problem using an increasingly denser grid on a shrinking feasible region. The analysis relies on properties of both the convex problem, and the algorithm. A classical example is the capacity scaling algorithm for the convex cost flow problem [see e.g. AMO93].

Hochbaum and Shanthikumar [HS90] have conducted perhaps the most general study of this approach. They consider separable convex costs over general polyhedra, and use a scale-and-iterate approach to obtain a $(1 + \epsilon)$-approximate solution. Their algorithm is polynomial in the size of the input, and the absolute value $|\Delta|$ of the largest subdeterminant of the constraint matrix. They measure the approximation in terms of the solution vectors themselves, not the objective values. They suggest methods for achieving objective approximation, with a running time dependent on the cost functions, as well as the size of the input and $|\Delta|$.

1.2 Our contribution

In contrast to previous contributions, we consider general nondecreasing separable concave objectives, and obtain polynomial bounds on the size of the resulting problem when the original problem has a polyhedral ground set in $\mathbb{R}^n_+$, and ϵ is fixed. The key idea that enables us to avoid iterations and scaling, and yet obtain polynomial bounds, is to use pieces exponentially increasing in size. Since the notion of objective value approximation is ill-defined when the sign of the costs is unrestricted, we require the objective functions to be nonnegative.

In Section 2 we introduce our technique for general grounds sets in $\mathbb{R}^n_+$ and nondecreasing cost functions. We need only $1 + \left\lceil \log_{1+4\epsilon+4\epsilon^2} \frac{u_i}{l_i} \right\rceil$ pieces for each concave component of the objective. In this expression, u_i denotes an upper bound on the value of corresponding variable, and l_i the smallest nonzero feasible value of that component. As $\epsilon \to 0$, the number of pieces as a function of ϵ behaves as $\frac{1}{4\epsilon}$. The number of pieces is the same for any concave function, and depends only on the chosen value of ϵ and the bounds u_i and l_i. Our method requires just one function evaluation per piece. In Section 2.1, we show that, for any fixed ϵ, the number of pieces required by our approach is within a constant factor of the best possible. In Section 2.2, we present several extensions, including to cost functions that are not monotone, and to ground sets not contained in $\mathbb{R}^n_+$.

In Section 3 we show that when the feasible set is a polyhedron, a $1 + \epsilon$ approximation can be achieved with a number of pieces polynomial in the input size of the polyhedron and linear in $1/\epsilon$, with no additional conditions or dependencies. Since the input size of the concave cost problem is always at least the input size of the polyhedron, the size of the resulting piecewise linear instance is always polynomial in the size of the original instance. For general polyhedra and nonnegative concave functions, we show that the number of required pieces is polynomial in the input size and the size of the zeroes of the cost functions. The latter are seldom ill-behaved quantities, and are often present as part of the input, thereby making the bound polynomial in the size of the original problem in this case as well.

These results provide a bridge between concave function minimization and piecewise linear minimization over polyhedra. Since our technique requires only a single, polynomially-sized piecewise linear approximation, we can directly apply any algorithm for optimizing

3

piecewise linear or fixed-charge objectives. In Section 3.1 we show that the resulting piecewise linear optimization problems can be reduced to fixed charge optimization problems while often preserving the underlying structure (for example, network structure). For practical problems, these advantages are amplified by the possibility of establishing significantly lower bounds on the number of pieces.

In Section 4 we illustrate our method on the practical and pervasive uncapacitated concave cost multicommodity flow problem with complete demand. We derive considerably smaller bounds on the number of required pieces than in the general case. Since our method preserves structure, the resulting fixed charge problems are network design problems. Using a primal-dual method [BMW89], we solve large problems with up to 80 nodes, 1,580 edges, 6,320 commodities and 9.9 million flow variables to within 4% of guaranteed optimality, on average. These problems are, to the best of our knowledge, significantly larger than previously solved concave cost multicommodity flow problems with full demand.

In Section 5 we illustrate our method on the lot-sizing problem with general concave production cost functions. We obtain a polynomial $O(n \log n \log \beta + n \log \beta \log \log \beta))$ algorithm; in this setting n denotes the number of periods, and β denotes the sum of demands divided by the smallest demand. According to Aggarwal and Park [AP93], the fastest algorithm for lot-sizing with general concave functions is still the $O(n^2)$ algorithm of Wagner and Whitin [WW58]. Neither our algorithm, nor that of Wagner and Whitin dominates the other in general. For example, our algorithm is faster when n is moderate or large, and the ratio of the largest to the smallest demand is moderate or small.

We chose multicommodity flows and lot-sizing as our examples because of the central role these problems play in the literature. However, the same approach is applicable to a wide variety of problems, such as capacitated multicommodity flows and multi-level inventory problems. In fact, our technique is not limited even to the general optimization framework of Section 2. It is potentially applicable for approximating problems in continuous dynamic programming, continuous optimal control, algorithmic game theory, and other settings where new solutions methods become available when switching from nonlinear to piecewise linear functions.

2   General ground sets

We examine the general concave minimization problem

$$Z_1^* = \min \left\{ \phi(x) : x \in X, x \geq 0 \right\}, \tag{1}$$

defined by a closed ground set $X \subseteq \mathbb{R}^n$ and a separable concave cost function $\phi : \mathbb{R}_+^n \to \mathbb{R}_+$ with $\phi(x_1, \ldots, x_n) = \sum_{i=1}^n \phi_i(x_i)$. The ground set need not be convex or connected (for example, it could be the ground set of an integer program). Let $[n] = \{1, \ldots, n\}$. We impose the following assumption.

Assumption 1. (a) The function ϕ is nondecreasing. (b) The problem has an optimal solution x^* and bounds $0 < l \leq u$ such that $x_i^* \in \{0\} \cup [l, u]$ for $i \in [n]$.

To approximate problem (1) within a factor of $1 + \epsilon$, we approximate each function ϕ_i with a piecewise linear function $\psi_i : \mathbb{R}_+ \to \mathbb{R}_+$. Each function ψ_i consists of $1 + P$ pieces,

with $P := \lceil \log_{1+\epsilon} \frac{u}{l} \rceil$, and is defined by the coefficients

$$c_i^p = \frac{d}{dx} f_i \left(l(1+\epsilon)^p \right), \qquad\qquad p \in \{0, \ldots, P\}, \qquad (2a)$$

$$f_i^p = f_i \left(l(1+\epsilon)^p \right) - l(1+\epsilon)^p c_i^p, \qquad\qquad p \in \{0, \ldots, P\}. \qquad (2b)$$

The symbol $\frac{d\phi_i(x_i')}{dx_i}$ denotes the derivative of ϕ_i at $x_i = x_i'$ if ϕ_i is differentiable at x_i', and an arbitrary supergradient of ϕ_i at x_i' otherwise. Each coefficient pair defines a line with nonnegative slope c_i^p and y-intercept f_i^p, which is tangent to the graph of ϕ_i at the point $l(1+\epsilon)^p$. For $x_i > 0$, the function ψ_i is defined by the lower envelope of these lines:

$$\psi_i(x_i) = \min\{f_i^p + c_i^p x_i : j = 0, \ldots, P\}. \qquad (3)$$

We let $\psi_i(0) = \phi_i(0)$ and $\psi(x) = \sum_{i=1}^n \psi_i(x_i)$. Substituting ψ for ϕ, we obtain the piecewise linear concave minimization problem

$$Z_4^* = \min\{\psi(x) : x \in X, x \geq 0\}. \qquad (4)$$

Lemma 1. $Z_1^* \leq Z_4^* \leq (1+\epsilon)Z_1^*$.

Proof. Let x^* be an optimal solution of problem (4). The graph of any line $f_i^p + c_i^p x_i^*$ lies on or above the graph of ϕ_i, hence $\phi_i(x_i^*) \leq \psi_i(x_i^*)$ for $i \in [n]$. Therefore, $Z_1^* \leq \phi(x^*) \leq \psi(x^*) = Z_4^*$.

Conversely, let x^* be an optimal solution of problem (1) satisfying Assumption 1(b). It suffices to show that $\psi_i(x_i^*) \leq (1+\epsilon)\phi_i(x_i^*)$ for $i \in [n]$. If $x_i^* = 0$, then the inequality holds. Otherwise, let $j = \left\lfloor \log_{1+\epsilon} \frac{x_i^*}{l} \right\rfloor \geq 0$, so that $\frac{x_i^*}{l} \in [(1+\epsilon)^p, (1+\epsilon)^{j+1}]$. Because ϕ_i is concave and nondecreasing,

$$\psi_i(x_i^*) \leq f_i^p + c_i^p x_i^* \leq f_i^p + c_i^p l(1+\epsilon)^{j+1} \qquad (5a)$$

$$= f_i^p + c_i^p l(1+\epsilon)(1+\epsilon)^p \leq (1+\epsilon)\left(f_i^p + c_i^p l(1+\epsilon)^p\right) \qquad (5b)$$

$$= (1+\epsilon)\phi_i\left((1+\epsilon)^p\right) \leq (1+\epsilon)\phi_i(x_i^*). \qquad (5c)$$

(See Figure 1 for an illustration.) Therefore, $Z_4^* \leq \psi(x^*) \leq (1+\epsilon)\phi(x^*) = (1+\epsilon)Z_1^*$. □

The previous proof has a simple geometric interpretation, but the approximation ratio of $1 + \epsilon$ is not tight. A tight analysis follows.

Theorem 1. $Z_1^* \leq Z_4^* \leq \frac{1+\sqrt{\epsilon+1}}{2} Z_1^* \leq (1 + \frac{\epsilon}{4})Z_1^*$.

Proof. Without loss of generality, we assume $l = 1$ and $f(0) = 0$, and consider only the segment $[1, 1+\epsilon]$, and the two tangents at $(1, \phi_i(1))$ and $(1+\epsilon, \phi_i(1+\epsilon))$. Suppose these tangents have slopes a and c respectively. The worst case is achievable when ϕ_i consists of 3 linear pieces with slopes $a > b > c$ on $[0, 1]$, $[1, 1+\epsilon]$, and $[1+\epsilon, +\infty]$ respectively. (See Figure 2 for an illustration.)

Let $x_i^* = 1 + \xi \in [1, 1+\epsilon]$. The values yielded by each of the two tangents at x_i^* are $a(1+\xi)$ and $a + b\epsilon - c(\epsilon - \xi)$, and

$$\phi_i(1+\xi) = a + b\xi, \qquad \psi_i(x_i) = \min\{a(1+\xi), a + b\epsilon - c(\epsilon - \xi)\}, \qquad (6)$$

5

Figure 1: Illustration of the proof of Lemma 1. Observe that the height of all points inside the box with the bold lower left and upper right corners exceeds the height of its lower left corner by at most a factor of ϵ.

Since $\xi \le \epsilon$ the worst case is achievable if $c = 0$. Since we seek to find ξ that maximizes

$$\frac{\psi_i(1+\xi)}{\phi_i(1+\xi)} = \min\left\{\frac{a+a\xi}{a+b\xi}, \frac{a+b\epsilon}{a+b\xi}\right\}, \tag{7}$$

we can assume ξ is such that $\frac{a+a\xi}{a+b\xi} = \frac{a+b\epsilon}{a+b\xi}$, which yields $\xi = \frac{b\epsilon}{a}$. Substituting, we now seek to maximize $\frac{1+\epsilon b/a}{1+\epsilon b^2/a^2}$, and letting $d = \frac{b}{a}$, we find that the maximum is achieved at $d = \frac{-1+\sqrt{\epsilon+1}}{\epsilon}$ and equals $\frac{1+\sqrt{\epsilon+1}}{2}$, which is less than $1+\frac{\epsilon}{4}$. □

Equivalently, instead of an approximation ratio of $\frac{1+\sqrt{\epsilon+1}}{2}$ using $1 + \left\lceil \log_{1+\epsilon} \frac{u}{l} \right\rceil$ pieces, we can obtain a ratio of $1+\epsilon$ using only $1 + \left\lceil \log_{1+4\epsilon+4\epsilon^2} \frac{u}{l} \right\rceil$ pieces. We can derive improved bounds on the number of pieces when the functions are known to belong to particular classes (for example, logarithmic functions), and even better bounds when the functions are known.

As a function of ϵ, the number of pieces grows as $\frac{1}{\log(1+4\epsilon+4\epsilon^2)}$. Since $\frac{4\epsilon}{\log(1+4\epsilon+4\epsilon^2)} \to 1$ as $\epsilon \to 0$, the number of pieces behaves as $\frac{1}{4\epsilon}$ as $\epsilon \to 0$. This behavior enables us to apply the approximation technique to practical concave cost problems. In Section 3 we will exploit the logarithmic dependence of our results on $\frac{u}{l}$ to derive polynomial bounds on the number of pieces for a large class of problems.

2.1 A lower bound on the number of pieces

The analysis in the proof of Theorem 1 is tight if we consider a function ϕ_i given by the values a, b, and c at the values obtained in the proof. Therefore, if we introduce the pieces as specified in (2), then $\frac{1+\sqrt{\epsilon+1}}{2}$ is the best approximation ratio that can be achieved. Since

6

Figure 2: Illustration of the proof of Theorem 1.

$\frac{1+\sqrt{\epsilon+1}}{2} \to 1 + \frac{\epsilon}{4}$ and $\frac{d}{d\epsilon}\frac{1+\sqrt{\epsilon+1}}{2} \to \frac{d}{d\epsilon}\left(1 + \frac{\epsilon}{4}\right)$ as $\epsilon \to 0$, $1 + \frac{\epsilon}{4}$ is the best ratio expressible as a linear function of ϵ that can be achieved asymptotically as $\epsilon \to 0$ with our approach.

In the remainder of this section, we establish a lower bound on the number of pieces required by *any* approach. First, we show that by limiting ourselves to tangents, we increase the number of required by at most a constant factor. As before, let $\phi_i(x_i)$ be a concave function, and $\psi_i(x_i)$ a piecewise linear function of $1 + P$ pieces with $\frac{1}{1+\epsilon} \le \frac{\psi_i(x_i)}{\phi_i(x_i)} \le 1 + \epsilon$ for $x_i \in [l, u]$.

Lemma 2. *There is a piecewise linear function $\varphi_i(x_i)$ of at most $2(1 + P)$ pieces such that $\frac{1}{1+\epsilon} \le \frac{\varphi_i(x_i)}{\phi_i(x_i)} \le 1 + \epsilon$ for $x_i \in [l, u]$, and each piece of φ_i is tangent to ψ_i.*

Proof. Fix a piece of ψ_i with intercept f_i^p and slope c_i^p. Since ϕ_i is concave, we can assume that the piece guarantees an $1 + \epsilon$ approximation of ϕ_i for $x_i \in [\xi', \xi'']$ and intersects the graph of ϕ_i at $\xi \in [\xi', \xi'']$. Also assume without loss of generality that the piece lays above the graph for $x_i \in [\xi', \xi)$ and below the graph for $x_i \in (\xi, \xi'']$. We can guarantee an $1 + \epsilon$ approximation on $[\xi', \xi]$ by introducing a tangent at ξ'. The piece with intercept $(1 + \epsilon)f_i^p$ and slope $(1 + \epsilon)c_i^p$ will be above the function on $[\xi, \xi'']$, and will still guarantee a $1 + \epsilon$ approximation on this segment. Therefore, by introducing a tangent at ξ'' we can guarantee a $1 + \epsilon$ approximation on $[\xi, \xi'']$ too. $\qquad\square$

Let $\phi_i(x_i) = \sqrt{x_i}$, and let ψ_i be a piecewise linear function with $\frac{1}{1+\epsilon} \le \frac{\psi_i(x_i)}{\phi_i(x_i)} \le 1 + \epsilon$ for $x_i \in [l, u]$, and each piece of ψ_i is tangent to the graph of ϕ_i. In the following lemma, we compare the number of tangents required by our approach with the minimum number of tangents needed to approximate ϕ_i.

Lemma 3. *For fixed ϵ, the minimum number of pieces in ψ_i is within a constant factor of $1 + \left\lceil \log_{1+4\epsilon+4\epsilon^2} \frac{u}{l} \right\rceil$, the number of pieces required by our approach. As $\epsilon \to 0$, the minimum number of pieces behaves as $\sqrt{2\epsilon} \log_{1+4\epsilon+4\epsilon^2} \frac{u}{l}$.*

7

Proof. Fix $\xi_0 \in [l, u]$ and let us determine the segment $[\xi_0 + \delta_1, \xi_0 + \delta_2]$ on which a tangent to the graph of ϕ_i at ξ_0 will guarantee a $1 + \epsilon$ approximation. The values of δ are given by the solutions to the equation

$$\phi_i(\xi_0) + \frac{d\phi_i(\xi_0)}{dx_i}\delta = (1 + \epsilon)\phi_i(\xi_0 + \delta). \tag{8}$$

Solving this quadratic equation yields $\delta = 2\xi_0\left(\epsilon(2 + \epsilon) \pm (1 + \epsilon)\sqrt{\epsilon(2 + \epsilon)}\right)$. Let δ_1 be the negative solution, and δ_2 the positive one; also let $\xi_1 = \xi_0 + \delta_1$. A tangent can provide an approximation on a segment of the form

$$[\xi_1, \gamma(\epsilon)\xi_1] := \left[\xi_1, \frac{-\delta_1 + \delta_2}{1 + \delta_1}\xi_1\right] = \left[\xi_1, \frac{4(1 + \epsilon)\sqrt{\epsilon(2 + \epsilon)}}{(1 + 2\epsilon)^2 - 2(1 + \epsilon)\sqrt{\epsilon(2 + \epsilon)}}\xi_1\right]. \tag{9}$$

Since $\gamma(\epsilon)$ does not depend on ξ_1, it immediately follows that we need $\left\lceil \log_{1+\gamma(\epsilon)} \frac{u}{l} \right\rceil$ pieces to approximate ϕ_i on $[l, u]$. This is within a factor of $1 + \left\lceil \frac{\log(1+\gamma(\epsilon))}{\log(1+4\epsilon+4\epsilon^2)} \right\rceil$ of the number of pieces required by our approach. Since $\lim_{\epsilon \to 0} \frac{\sqrt{\epsilon}\log(1+\gamma(\epsilon))}{\log(1+4\epsilon+4\epsilon^2)} = \sqrt{2}$, the minimum number of pieces behaves as $\left\lceil \sqrt{\epsilon/2} \log_{1+4\epsilon+4\epsilon^2} \frac{u}{l} \right\rceil$ as $\epsilon \to 0$. $\qquad\square$

Therefore, if we do not restrict ourselves to tangents, the minimum number of pieces for approximating $\phi_i(x_i) = \sqrt{x_i}$ behaves as $\sqrt{\epsilon/8} \log_{1+4\epsilon+4\epsilon^2} \frac{u}{l}$ as $\epsilon \to 0$, and is within a constant factor of $1 + \left\lceil \log_{1+4\epsilon+4\epsilon^2} \frac{u}{l} \right\rceil$ for fixed ϵ. The asymptotic behavior and the constant factor are independent of $[l, u]$.

2.2   Extensions

Our approach applies to a broader class of problems. Consider the problem

$$\min\{\phi(x) : x \in X\}, \tag{10}$$

with $\phi : \text{conv}(X) \to \mathbb{R}_+$ a separable and concave function. We relax Assumption 1 as follows.

Assumption 2. Problem (10) has an optimal solution x^* and bounds $0 < l < u$ so that $|x_i^*| \le u$, and either $\phi_i(x_i^*) = 0$ or $\min\{|x_i^* - x_i| : \phi_i(x_i) = 0\} \ge l$, for $i \in [n]$.

The following is a generalization of Theorem 1.

Corollary 1. *Problem (10) can be approximated within a factor of $1 + \epsilon$ by replacing each function ϕ_i with a piecewise linear function ψ_i of $2 + 2\left\lceil \log_{1+4\epsilon+4\epsilon^2} \frac{u}{l} \right\rceil$ pieces, and at most two discontinuity points.*

Proof. We will consider each objective component ϕ_i separately. Any concave function $\phi_i(x_i)$ that is not constant over the projection of $\text{conv}(X)$ to x_i will have at most two zeroes, which we denote by $\zeta_i^L < \zeta_i^R$. Let $\zeta_i' = \max\{-u - l, \zeta_i^L\}$ and $\zeta_i'' = \min\{u + l, \zeta_i^R\}$ and note that we need to approximate ϕ_i only on $[\zeta_i' + l, \zeta_i'' - l]$. Let ζ_i^* be a point where ϕ_i is

8

maximized, and note that ϕ_i is monotonically nondecreasing on $[\zeta_i', \zeta_i^*]$, and monotonically nonincreasing on $[\zeta_i^*, \zeta_i'']$. We will apply Theorem 1 to each of these two segments, by using translation and reflection.

If one of the two segments is empty, the proof is complete. Otherwise, w.l.o.g. consider the segment $[\zeta_i', \zeta_i^*]$. To avoid having to compute ξ_i^*, we simply introduce tangents until the slope is nonpositive. Let the last tangent be at $\zeta_i' + l(1 + 4\epsilon + 4\epsilon^2)^{P_i}$. Since its slope might be negative, Theorem 1 does not guarantee an approximation ratio on the segment $[\zeta_i' + l(1 + 4\epsilon + 4\epsilon^2)^{P_i-1}, \zeta_i' + l(1 + 4\epsilon + 4\epsilon^2)^{P_i}]$. For this reason, we remove the tangent at $\zeta_i' + l(1 + 4\epsilon + 4\epsilon^2)^{P_i}$, and introduce a tangent at $\zeta_i' + l(1 + 4\epsilon + 4\epsilon^2)^{P_i-1}(1 + \epsilon)^p$ for the largest j that yields a positive slope; since $(1+\epsilon)^4 \geq 1+4\epsilon+4\epsilon^2$, $j \leq 3$. The approximation is guaranteed on $[\zeta_i' + l(1 + 4\epsilon + 4\epsilon^2)^{P_i-1}, \zeta_i' + l(1 + 4\epsilon + 4\epsilon^2)^{P_i-1}(1 + \epsilon)^p]$ by Theorem 1, and on $[l(1 + 4\epsilon + 4\epsilon^2)^{P_i-1}(1 + \epsilon)^p], \xi_i^*]$ by Lemma 1.

The number of pieces employed is at most $2 + \left\lceil \log_{1+4\epsilon+4\epsilon^2} \frac{\zeta_i^* - \zeta_i'}{l} \right\rceil + \left\lceil \log_{1+4\epsilon+4\epsilon^2} \frac{\zeta_i'' - \zeta_i^*}{l} \right\rceil \leq 2 + 2 \left\lceil \log_{1+4\epsilon+4\epsilon^2} \frac{u}{l} \right\rceil$, since $\zeta_i'' - \zeta_i' \leq 2u$. Each segment yields at most one discontinuity point. $\qquad\square$

We conclude with a list of further extensions:

1) Since we employ tangents in our method, we require one evaluation of the function and its derivative (or any supergradient) to compute each piece. Our results also hold if we use secants instead of tangents, in which case we only require one function evaluation per piece. The secant approach may be preferable in some computational applications.

2) We can employ separate parameters u_i and l_i for each component. Doing so may lead to fewer required pieces in certain applications.

3) The results in this section, but not in subsequent ones, also apply to concave *maximization* problems, as long as all other assumptions hold.

Our results do not apply to maximization or minimization problems with nonnegative *convex* costs.

3   Polyhedral ground sets

Let $X = \{x : Ax \leq b, x \geq 0\}$ be a rational polyhedron defined by $A \in \mathbb{Q}^{m \times n}$ and $b \in \mathbb{Q}^m$, and let $\phi : X \to \mathbb{R}_+$ be a separable nondecreasing concave function. We consider the problem

$$Z_{11}^* = \min\{\phi(x) : Ax \leq b, x \geq 0\}. \tag{11}$$

We will bound the optimal solution components in terms of input data size. We take the input data size for problem (11) to be the size of A and b alone; omitting the objective functions ϕ_i from the input size computation only strengthens the resulting bounds. Following standard practice [see e.g. KV02], we define the size of rational numbers and matrices as the number of bits needed to represent them:

1) for integers $r \in \mathbb{Z}, \mathrm{size}(r) := 1 + \lceil \log_2(|r| + 1) \rceil$;

2) for rational numbers $r = \frac{r_1}{r_2} \in \mathbb{Q}$ with $\frac{r_1}{r_2}$ irreducible, $\mathrm{size}(r) := \mathrm{size}(r_1) + \mathrm{size}(r_2)$;

3) for vectors or matrices $A \in \mathbb{Q}^{m \times n}, \mathrm{size}(A) := mn + \sum_{i=1}^{m} \sum_{j=1}^{n} \mathrm{size}(a_{ij})$.

The following property is well-known [see e.g. KV02, GLS93].

Lemma 4. *Any vertex x of X has $\text{size}(x) \leq U(A, b) := 4(\text{size}(A) + \text{size}(b) + n^2 + 5n)$.*

To approximate problem (11), we introduce the piecewise linear functions ψ_i as described in equations (2) and (3); each function will have $1 + \left\lceil \frac{2U(A,b)}{\log_2(1+4\epsilon+4\epsilon^2)} \right\rceil$ pieces. Consider the problem

$$Z_{12}^* = \min\{\psi(x) : Ax \leq b, x \geq 0\}. \tag{12}$$

Theorem 2. $Z_{11}^* \leq Z_{12}^* \leq (1+\epsilon)Z_{11}^*$. *Each function ψ_i has a number of pieces polynomial in $\text{size}(A) + \text{size}(b)$, the input size of problem (11).*

Proof. Because $X \subseteq \mathbb{R}_+^n$, it has at least one vertex, and because ϕ is nonnegative, Z_{11} is bounded from below. Therefore, because ϕ is concave, problem (11) has an optimal solution x^* at a vertex of X [Bau58]. Lemma 4 ensures that $x_i^* \in \{0\} \cup \left[2^{-U(A,b)}, \; 2^{U(A,b)}\right]$. The approximation property follows from Theorem 1. $\qquad\square$

Again, a generalization is possible. Consider the problem

$$\min\{\phi(x) : Ax \leq b\}, \tag{13}$$

defined by a polyhedron $X = \{x : Ax \leq b\}$ with at least one vertex, and a separable concave function $\phi : X \to \mathbb{R}_+$. Any concave function $\phi_i(x_i)$ that is not constant over the projection of the feasible polyhedron to x_i will have at most two zeroes; denote them by $\zeta_i' < \zeta_i''$, and assume they are rational.

Corollary 2. *Problem (13) can be approximated within a factor of $1 + \epsilon$ by replacing each function ϕ_i with a piecewise linear function ψ_i of $2 + 2\left\lceil \frac{2U(A,b)+\text{size}(\zeta_i')+\text{size}(\zeta_i'')}{\log_2(1+4\epsilon+4\epsilon^2)} \right\rceil$ pieces, and at most two discontinuity points.*

Proof. Let x^* be a vertex optimal solution of (13). Then $|x_i^*| \leq u := 2^{U(A,b)}$. Moreover, either $x_i^* \in \{\zeta_i', \zeta_i''\}$ or $\min\{|x_i^* - \zeta_i'|, |x_i^* - \zeta_i''|\} \geq l := 2^{-U(A,b)-\text{size}(\zeta_i')-\text{size}(\zeta_i'')}$. Applying Corollary 1 completes the proof. $\qquad\square$

This corollary is motived by the fact that ζ_i' and ζ_i'' are seldom ill-behaved quantities. In many applications, they are included in the input, as part of the description of the concave cost functions. If ζ_i' and ζ_i'' are part of the input for $i \in [n]$, then the number of pieces in each function ψ_i is polynomial in the size of the input.

3.1 Representing the piecewise linear functions

To solve the problems resulting from our approximation technique, we could use several classical methods for representing piecewise linear functions as mixed integer programs. Such methods usually introduce one or more binary variables for each piece and add a coupling constraint that ensures the approximation uses only one piece [see e.g. NW99, CGM03]. However, since the objective function to be minimized is concave, the coupling constraint is

unnecessary, and we can employ the following well-known fixed charge formulation, which is equivalent to formulation (12):

$$\min \sum_{i=1}^{n} \sum_{p=1}^{P} \left(f_i^p z_i^p + c_i^p y_i^p \right), \tag{14a}$$

$$\text{s.t. } Ax \leq b, \tag{14b}$$

$$x_i = \sum_{p=0}^{P} y_i^p, \qquad\qquad\qquad i \in [n], \tag{14c}$$

$$0 \leq y_i^p \leq u_i z_i^p, \qquad\qquad i \in [n], p \in \{0, \dots, P\}, \tag{14d}$$

$$z_i^p \in \{0, 1\}, \qquad\qquad\qquad i \in [n], p \in \{0, \dots, P\}. \tag{14e}$$

We assume without loss of generality that $f(0) = 0$; if $f(0) > 0$ the approximation only becomes tighter. We choose the coefficients u_i so that $x_i \leq u_i$ at any vertex, for example $u_i = 2^{U(A,b)}$.

Lemma 5. *The input size of problem (14) is polynomial in the input size of problem (12).*

A key advantage of fixed-charge formulation (14) is that, in many cases, it preserves the special structure of the original concave cost problem. Therefore, solution methods for fixed charge problems with special structure can be used to approximately solve general concave cost problems with the same structure. A possible drawback of problem (14) is that it has $1 + p$ times more variables. Although for general polyhedra p could be prohibitively large, for many practical problems, we are able to derive significantly smaller expressions for p.

We make use of both these observations when applying our technique to practical problems in the following two sections.

4   Multicommodity Flows

To illustrate our approach on a practical problem, we consider concave cost uncapacitated multicommodity flows (see [GP90] for a survey and applications). Let $G = (V, E)$ be an undirected graph with $|V| = n, |E| = m$, and let $\phi : \mathbb{R}_+^E \to \mathbb{R}_+$ be a separable nondecreasing concave function. Consider the problem

$$\min \left\{ \sum_{ij \in E} \phi_{ij} \left(\sum_{k=1}^{K} (x_{ij}^k + x_{ji}^k) \right) \; \middle| \; \sum_{ij \in E} x_{ij}^k - \sum_{ji \in E} x_{ji}^k = b_i^k, \, x_{ij}^k \geq 0 \right\}. \tag{15}$$

In this model, K is the number of commodities, x_{ij}^k denotes the flow of commodity k from i to j, and $|b_i^k|$ is the supply ($b_i^k > 0$) or demand ($b_i^k < 0$) of commodity k at node i. Let $b_{\min} = \min\{|b_i^k| : i \in V, k \in [K]\}, B^k = \sum_{i:b_i^k>0} b_i^k$, and $B = \sum_{k=1}^{K} B^k$. Since rational numbers can be scaled to obtain integers, for simplicity we assume that the problem data are integral, and that $\phi_{ij}(0) = 0$ for $ij \in E$.

In this setting, formulation (14) yields the well-known fixed-charge multicommodity flow problem, but now on a network with $(1 + p)m$ edges:

$$\min \sum_{ijp \in E} \sum_{k=1}^{K} (f_{ijp} z_{ijp} + c_{ijp}(x_{ijp}^k + x_{jip}^k)), \tag{16a}$$

$$\text{s.t.} \sum_{ijp \in E} x_{ijp}^k - \sum_{jip \in E} x_{ijp}^k = b_i^k, \qquad i \in V, k \in [K], \tag{16b}$$

$$0 \le x_{ijp}^k, x_{jip}^k \le B^k z_{ijp}, \qquad ijp \in E, k \in [K], \tag{16c}$$

$$z_{ijp} \in \{0, 1\}, \qquad ijp \in E. \tag{16d}$$

For each edge $ijp \in E$, the coefficient f_{ijp} can be interpreted as its installation cost, and c_{ijp} as the cost of routing flow on the edge once installed.

Proposition 1. $Z_{15}^* \le Z_{16}^* \le (1 + \epsilon)Z_{15}^*$. *This ratio can be achieved by introducing* $1 + \left\lceil \log_{1+4\epsilon+4\epsilon^2} \frac{B}{b_{\min}} \right\rceil \le 1 + \left\lceil \log_{1+4\epsilon+4\epsilon^2} B \right\rceil$ *pieces for each edge cost function.*

Proof. As is well-known [Bau58, Sch03], in some optimal solution to problem (15), the flow for each commodity occurs on a tree. Consequently, any nonzero flow on any edge will be at least $b_{\min} \ge 1$. The approximation result now follows from Theorem 2 and Lemma 5. $\square$

The special structure of the problem allows us to increase the number of edges by a factor of only $1 + \left\lceil \log_{1+4\epsilon+4\epsilon^2} B \right\rceil$, which is much less than the factor obtained for general polyhedra.

4.1 Computational results

We present computational results for uncapacitated multicommodity flow problems with complete uniform demand. We have generated the instances as follows. To ensure feasibility, for each problem we first generated a random spanning tree. Then we added the desired number of edges between nodes selected uniformly at random. For each number of nodes, we considered a dense network with $\frac{n^2}{4}$ edges, and a sparse network with $3n$ edges. For each network thus generated, we have considered two cost structures.

The first cost structure models moderate economies of scale. We assigned to each edge $ij \in E$ a cost function of the form $a + b(x_{ij})^c$, with a, b, and c randomly generated from uniform distributions over $[0.1, 10], [0.33, 33.4]$, and $[0.8, 0.99]$. For an average cost function from this family, the marginal cost decreases by approximately 30% as the flow on an edge increases from 25 to 1,000. The second cost structure models strong economies of scale. The cost functions are as in the first case, except that c is sampled from a uniform distribution over $[0.0099, 0.99]$. In this case, for an average cost function, the marginal cost decreases by approximately 84% as the flow on an edge increases from 25 to 1,000. (Note that on a network with n nodes, the flow on an edge can range from 2 to $n(n - 1)$.)

Table 1 specifies the problem sizes. Note that although the individual dimensions of the problems are moderate, the resulting number of variables is large, since a problem with n nodes and m edges yields n^2m flow variables. The largest problems we solved have 80 nodes, 1,580 edges, and 6,320 commodities. To approach them with an MIP solver, these

#	n	m	K	Flow Variables	Pieces
1	10	30	90	8,100	41
2	20	60	380	22,800	77
3	20	95	380	36,100	77
4	30	90	870	78,300	98
5	30	215	870	187,050	98
6	40	120	1,560	187,200	113
7	40	390	1,560	608,400	113
8	50	150	2,450	367,500	124
9	50	610	2,450	1,494,500	124
10	60	180	3,540	637,200	133
11	60	885	3,540	3,132,900	133
12	70	210	4,830	1,014,300	141
13	70	1,205	4,830	5,820,150	141
14	80	240	6,320	1,516,800	148
15	80	1,580	6,320	9,985,600	148

Table 1: Network sizes. The column "Pieces" indicates the number of pieces in each piecewise linear function resulting from the approximation.

problem would require 1,580 binary variables, 9,985,600 continuous variables and 10,491,200 constraints, even if we replaced the concave functions by fixed charge costs.

We chose $\epsilon = 0.01 = 1\%$ for the piecewise linear approximation. After applying our piecewise linear approximation technique, we have reduced the the total number of pieces further by noting that close to 1, our approach introduced tangents on a grid denser than the uniform grid $1, 2, 3, \ldots$ For each problem, we have reduced the number of pieces per cost function by approximately 47 by using the uniform grid close to 1, and the grid generated by our approach elsewhere.

We used an improved version of the dual ascent method described by Balakrishnan et al. [BMW89] (also known as the primal-dual method [GW97]) to solve the resulting problems. The method produces a feasible solution, whose objective we denote by Z_{16}^{DA}, to problem (16) and a lower bound Z_{16}^{LB} on the optimal value. This allows us to compute, a posteriori, a problem-dependent error bound $\epsilon_{\mathrm{DA}} := \frac{Z_{16}^{\mathrm{DA}}}{Z_{16}^{\mathrm{LB}}} - 1$ with respect to the piecewise linear approximation, and an overall error bound $\epsilon_{\mathrm{ALL}} := (1 + \epsilon)(1 + \epsilon_{\mathrm{DA}})$ with respect to the original problem.

Table 2 summarizes the computational results. We performed all the computations on a Pentium Xeon 2.8 GHz. All the error bounds, times, and edge numbers in the table are averaged over 3 instances of the respective size and cost structure.

We obtained average error bounds of 3.62% for problems with moderate economies of scale, and 4.20% for problems with strong economies of scale. This difference in average error bound is consistent with previous reports in the literature for fixed-charge functions, in which problems with higher fixed to variable cost ratios, and thus stronger economies of scale, have been found harder to solve [BMW89, HS89]. Note that the solutions to problems with moderate economies of scale have more edges than those to problems with

#	Moderate economies of scale				Strong economies of scale			
	Time	Sol. Edges	$\epsilon_{DA}\%$	$\epsilon_{ALL}\%$	Time	Sol. Edges	$\epsilon_{DA}\%$	$\epsilon_{ALL}\%$
1	0.26s	14	0.41	1.41	0.4s	9	0.35	1.35
2	7.57s	31	1.45	2.46	10.6s	19	1.06	2.07
3	8.77s	25.3	1.20	2.21	18.3s	19	3.38	4.42
4	48s	44	1.95	2.96	43.1s	29	1.18	2.20
5	1m33s	43.6	2.16	3.19	1m40s	29	3.50	4.54
6	3m29s	61.6	2.47	3.49	1m46s	39	2.20	3.22
7	6m49s	59	3.24	4.28	4m23s	39	3.17	4.21
8	9m16s	79	2.22	3.24	4m20s	49	3.42	4.46
9	20m51s	74.6	3.10	4.13	8m35s	49	4.22	5.26
10	21m10s	95	2.58	3.61	6m42s	59	3.27	4.30
11	56m58s	95.6	3.64	4.68	16m31s	59	4.25	5.29
12	40m42s	101.6	2.85	3.87	8m32s	69	3.77	4.81
13	1h47m	115.6	4.19	5.24	25m34s	69	4.98	6.03
14	1h18m	127.6	2.82	3.84	13m43s	79	4.10	5.14
15	3h3m	129.3	4.59	5.64	36m2s	79	4.68	5.73
Average			2.59	**3.62**			3.17	**4.20**

Table 2: Computational results. The values in column "Sol. Edges" represent the number of edges in the obtained solutions.

strong economies of scale; in fact, in the latter case, the edges always form a tree.

To the best of our knowledge, the literature does not contains exact or approximate computational results for concave cost multicommodity network flow problems of this size. Bell and Lamar [BL97] propose an exact branch-and-bound approach for *single-commodity* flows, and present computational results on networks with at most 20 nodes and 96 arcs. Fontes et al [FHC03] propose a heuristic approach for single-commodity flows, and present computational results on networks with up to 50 nodes and 200 edges. They obtain average error bounds of less than 13.8%, and conjecture that the actual gap between the obtained solutions and the optimal ones is much smaller.

5 Economic Lot-Sizing

The economic lot-sizing model [WW58] is one of the most celebrated inventory planning models. Although it is a discrete deterministic model, researchers use it in conjunction with safety stock provisions in settings with uncertain demand, and as an approximation for continuous-time models; algorithms for this model are used as a subroutine for material requirement planning systems, and as a solution method for subproblems resulting from Lagrangean relaxation of more complex models. See [FT91] for a brief survey.

Consider a facility facing deterministic demand of a single commodity over n periods. Let b_i denote the demand in period $i \in [n]$, and h_i be a linear per-unit holding for the inventory x_i carried from period i to $i + 1$. Assume the cost of ordering y_i units of the commodity in period i is specified by a nonnegative concave function $\phi_i(y_i)$. Demand must be satisfied at the time it occurs (i.e. backlogging is prohibited). The objective is

to determine a production and holding plan that will minimize total cost and satisfy all demands. The problem can be formulated as follows:

$$\min\left\{\sum_{i=1}^{n}\left(\phi_i(y_i)+h_ix_i\right) \;\middle|\; x_i+y_i-y_{i+1}=b_i,\, x_i\geq 0,\, y_i\geq 0\right\}. \tag{17}$$

We assume, without loss of generality, that $x_0 = 0$ is a constant, which signifies that the initial inventory is 0.

Let $B = \sum_{i=1}^{n} b_i$ be the total demand, $b_{\min} = \min_i b_i$, and $\beta = \frac{B}{b_{\min}}$. We approximate the concave functions ϕ_i with piecewise linear functions ψ_i of $1+P := 1 + \left\lceil \log_{1+4\epsilon+4\epsilon^2} \beta \right\rceil$ pieces, as described in equation (2). The resulting problem becomes a lot-sizing problem with fixed charge functions, and $n(1+P)$ periods:

$$\min\left\{\sum_{i=1}^{n(1+P)}\left(f_iz_i+c_iy_i+h'_ix_i\right) \;\middle|\; x_i+y_i-y_{i+1}=b'_i,\, x_i\geq 0,\, 0\leq y_i\leq Bz_i,\, z_i\in\{0,1\}\right\}. \tag{18}$$

In this model, f_i represents the fixed cost of ordering in period i, and c_i represents the incremental cost. The new demands b'_i (holding costs h'_i) equal $b_{i/(1+P)}$ $(h_{i/(1+P)})$ for i divisible by $1+P$ and 0 otherwise.

As is well-known, if the objective is concave, then $x_i, y_i \in \{0\} \cup \{b_{\min}, B\}$. Therefore, the following proposition follows immediately from Theorem 1.

Proposition 2. *Problem (17) can be approximated within a factor of $1+\epsilon$ by problem (18) with $1 + \left\lceil \log_{1+4\epsilon+4\epsilon^2} \beta \right\rceil$ as many periods.*

Since b_i are part of the input, the resulting instances are polynomially sized with respect to the original problem. By employing one of the $O(n \log n)$ algorithm for lot-sizing with fixed charge production costs [FT91, WvHK92, AP93] on the resulting instances, we obtain a $O(n \log \beta \log(n \log \beta)) = O(n \log n \log \beta + n \log \beta \log \log \beta)$ polynomial algorithm for lot-sizing with general concave production cost functions.

According to Aggarwal and Park [AP93], the fastest algorithm for lot-sizing with general concave functions is the $O(n^2)$ algorithm of Wagner and Whitin [WW58]. Neither our algorithm, nor that of Wagner and Whitin in general. For example, our algorithm is faster when n is moderate or large, and the ratio $\frac{\max_i b_i}{\min_i b_i}$ of the largest to the smallest demand is moderate or small.

The same approach of combining algorithms for problems with fixed-charge costs with our reduction is applicable to various extensions of the lot-sizing problem. For example, we obtain the same running time of $O(n \log n \log \beta + n \log \beta \log \log \beta)$ for the lot-sizing problem with backlogging. The fastest algorithm for this problem with general concave production costs is due to Zangwill [Zan69], and runs in $O(n^3)$ [AP93].

Acknowledgments

References

[AMO93] Ravindra K. Ahuja, Thomas L. Magnanti, and James B. Orlin. *Network flows.* Prentice Hall Inc., Englewood Cliffs, NJ, 1993. Theory, algorithms, and applications.

[AP93] Alok Aggarwal and James K. Park. Improved algorithms for economic lot size problems. *Oper. Res.*, 41(3):549–571, 1993.

[Ata01] Alper Atamtürk. Flow pack facets of the single node fixed-charge flow polytope. *Oper. Res. Lett.*, 29(3):107–114, 2001.

[Bau58] Heinz Bauer. Minimalstellen von Funktionen und Extremalpunkte. *Arch. Math.*, 9:389–393, 1958.

[BL97] Gavin J. Bell and Bruce W. Lamar. Solution methods for nonconvex network flow problems. In *Network optimization (Gainesville, FL, 1996)*, volume 450 of *Lecture Notes in Econom. and Math. Systems*, pages 32–50. Springer, Berlin, 1997.

[BMW89] A. Balakrishnan, T. L. Magnanti, and R. T. Wong. A dual-ascent procedure for large-scale uncapacitated network design. *Oper. Res.*, 37(5):716–740, 1989.

[CGM03] Keely L. Croxton, Bernard Gendron, and Thomas L. Magnanti. A comparison of mixed-integer programming models for nonconvex piecewise linear cost minimization problems. *Management Science*, 49(9):1268–1273, 2003.

[Dan63] George B. Dantzig. *Linear programming and extensions.* Princeton University Press, Princeton, N.J., 1963.

[FHC03] Dalila B. M. M. Fontes, Eleni Hadjiconstantinou, and Nicos Christofides. Upper bounds for single-source uncapacitated concave minimum-cost network flow problems. *Networks*, 41(4):221–228, 2003. Special issue in memory of Ernesto Q. V. Martins.

[FT91] A. Federgruen and M. Tzur. A simple forward algorithm to solve general dynamic lot sizing models with n periods in $O(n \log n)$ or $O(n)$ time. *Management Sci.*, 37:909–925, 1991.

[GLS93] Martin Grötschel, László Lovász, and Alexander Schrijver. *Geometric algorithms and combinatorial optimization*, volume 2 of *Algorithms and Combinatorics*. Springer-Verlag, Berlin, second edition, 1993.

[GP90] G. M. Guisewite and P. M. Pardalos. Minimum concave-cost network flow problems: applications, complexity, and algorithms. *Ann. Oper. Res.*, 25(1-4):75–99, 1990. Computational methods in global optimization.

[GW97] M.X. Goemans and D.P. Williamson. The primal-dual method for approximation algorithms and its application to network design problems. In Dorit S.

Hochbaum, editor, *Approximation algorithms for NP-hard problems*, chapter 4, pages 144–191. PWS Pub. Co., Boston, 1997.

[HH98] Kaj Holmberg and Johan Hellstrand. Solving the uncapacitated network design problem by a Lagrangean heuristic and branch-and-bound. *Oper. Res.*, 46(2):247–259, 1998.

[HMM03] M. T. Hajiaghayi, M. Mahdian, and V. S. Mirrokni. The facility location problem with general cost functions. *Networks*, 42(1):42–47, 2003.

[HS89] Dorit S. Hochbaum and Arie Segev. Analysis of a flow problem with fixed charges. *Networks*, 19(3):291–312, 1989.

[HS90] Dorit S. Hochbaum and J. George Shanthikumar. Convex separable optimization is not much harder than linear optimization. *J. Assoc. Comput. Mach.*, 37(4):843–862, 1990.

[JMM+03] Kamal Jain, Mohammad Mahdian, Evangelos Markakis, Amin Saberi, and Vijay V. Vazirani. Greedy facility location algorithms analyzed using dual fitting with factor-revealing LP. *J. ACM*, 50(6):795–824, 2003.

[KV02] Bernhard Korte and Jens Vygen. *Combinatorial optimization*, volume 21 of *Algorithms and Combinatorics*. Springer-Verlag, Berlin, second edition, 2002. Theory and algorithms.

[MMP00] Adam Meyerson, Kamesh Munagala, and Serge Plotkin. Cost-distance: two metric network design. In *41st Annual Symposium on Foundations of Computer Science (Redondo Beach, CA, 2000)*, pages 624–630. IEEE Comput. Soc. Press, Los Alamitos, CA, 2000.

[MS04] Thomas L. Magnanti and Dan Stratila. Separable concave optimization approximately equals piecewise linear optimization. In *Integer Programming and Combinatorial Optimization X (New York, NY, 2004)*. Springer, New York, NY, 2004.

[Mun03] Kamesh Munagala. *Approximation algorithms for concave cost network flow problems*. PhD thesis, Stanford University, Department of Computer Science, March 2003.

[NW99] George Nemhauser and Laurence Wolsey. *Integer and combinatorial optimization*. Wiley-Interscience Series in Discrete Mathematics and Optimization. John Wiley & Sons Inc., New York, 1999. Reprint of the 1988 original, A Wiley-Interscience Publication.

[OW03] Francisco Ortega and Laurence A. Wolsey. A branch-and-cut algorithm for the single-commodity, uncapacitated, fixed-charge network flow problem. *Networks*, 41(3):143–158, 2003.

[RP86]   J. B. Rosen and P. M. Pardalos. Global minimization of large-scale constrained concave quadratic problems by separable programming. *Math. Programming*, 34(2):163–174, 1986.

[Sch03]   Alexander Schrijver. *Combinatorial optimization. Polyhedra and efficiency. Vol. A*, volume 24 of *Algorithms and Combinatorics*. Springer-Verlag, Berlin, 2003. Paths, flows, matchings, Chapters 1–38.

[WvHK92]   Albert Wagelmans, Stan van Hoesel, and Antoon Kolen. Economic lot sizing: an $O(n \log n)$ algorithm that runs in linear time in the Wagner-Whitin case. *Oper. Res.*, 40(suppl. 1):S145–S156, 1992.

[WW58]   Harvey M. Wagner and Thomson M. Whitin. Dynamic version of the economic lot size model. *Management Sci.*, 5:89–96, 1958.

[Zan69]   W.I. Zangwill. A backlogging model and a multiechelon model of a dynamic economic lot size production systems—a network approach. *Management Sci.*, 15:506–527, 1969.