
Naveen Jindal School of Management

2014-01-01

Fixed-Dimensional Stochastic Dynamic Programs: An Approximation Scheme and an Inventory Application

UTD AUTHOR(S): Wei Chen, Milind Dawande and Ganesh Janakiraman

©2014 INFORMS



Operations Research

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Fixed-Dimensional Stochastic Dynamic Programs: An Approximation Scheme and an Inventory Application

Wei Chen, Milind Dawande, Ganesh Janakiraman

To cite this article:

Wei Chen, Milind Dawande, Ganesh Janakiraman (2014) Fixed-Dimensional Stochastic Dynamic Programs: An Approximation Scheme and an Inventory Application. *Operations Research* 62(1):81-103. <http://dx.doi.org/10.1287/opre.2013.1239>

Full terms and conditions of use: <http://pubsonline.informs.org/page/terms-and-conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2014, INFORMS

Please scroll down for article—it is on subsequent pages



INFORMS is the largest professional society in the world for professionals in the fields of operations research, management science, and analytics.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

METHODS

Fixed-Dimensional Stochastic Dynamic Programs: An Approximation Scheme and an Inventory Application

Wei Chen, Milind Dawande, Ganesh Janakiraman

Naveen Jindal School of Management, The University of Texas at Dallas,
 Richardson, Texas 75080 {wei.chen@utdallas.edu, milind@utdallas.edu, ganesh@utdallas.edu}

We study fixed-dimensional stochastic dynamic programs in a discrete setting over a finite horizon. Under the primary assumption that the cost-to-go functions are discrete L^2 -convex, we propose a pseudo-polynomial time approximation scheme that solves this problem to within an arbitrary prespecified additive error of $\epsilon > 0$. The proposed approximation algorithm is a generalization of the explicit-enumeration algorithm and offers us full control in the trade-off between accuracy and running time.

The main technique we develop for obtaining our scheme is approximation of a fixed-dimensional L^2 -convex function on a bounded rectangular set, using only a selected number of points in its domain. Furthermore, we prove that the approximation function preserves L^2 -convexity. Finally, to apply the approximate functions in a dynamic program, we bound the error propagation of the approximation. Our approximation scheme is illustrated on a well-known problem in inventory theory, the single-product problem with lost sales and lead times. We demonstrate the practical value of our scheme by implementing our approximation scheme and the explicit-enumeration algorithm on instances of this inventory problem.

Subject classifications: discrete convexity; multidimensional stochastic dynamic programs; approximation algorithms.

Area of review: Optimization.

History: Received December 2011; revisions received September 2012, March 2013, October 2013; accepted October 2013.

1. Introduction

Consider the following finite-horizon dynamic programming (DP) problem. For period $t = 1, 2, \dots, T$, we have

$$f_t(\mathbf{x}_t) = \min_{\mathbf{a}_t \in \mathcal{A}_t(\mathbf{x}_t)} [G_t(\mathbf{x}_t, \mathbf{a}_t) + \mathbb{E}_{\mathbf{w}_t} f_{t+1}(\mathbb{T}(\mathbf{x}_t, \mathbf{a}_t, \mathbf{w}_t))],$$

where $\mathbf{x}_t \in \mathbb{Z}^{k_1}$, $\mathbf{a}_t \in \mathbb{Z}^{k_2}$, $\mathbf{w}_t \in \mathbb{Z}^{k_3}$. (1)

For period t , f_t is the cost-to-go function starting from that period until the end of the horizon, G_t is the single-period cost function, \mathbf{x}_t is the state vector, \mathbf{a}_t is the action vector, and \mathbf{w}_t is an exogenous random vector. We assume that for all t , \mathbf{x}_t and \mathbf{a}_t belong to bounded rectangular sets, $\mathcal{X}_t \subseteq \mathbb{Z}^{k_1}$ and $\mathcal{A}_t(\mathbf{x}_t) \subseteq \mathbb{Z}^{k_2}$, respectively. We let $\mathbf{w}_t \in \mathbb{Z}^{k_3}$, $t = 1, 2, \dots, T$, be a sequence of independent random vectors. Furthermore, the possible realizations of \mathbf{w}_t and the corresponding probabilities are specified explicitly for each period t . Throughout the paper, we hold k_1 , k_2 , and k_3 fixed. The transformation function \mathbb{T} transforms $(\mathbf{x}_t, \mathbf{a}_t, \mathbf{w}_t)$ to \mathbf{x}_{t+1} ; i.e., $\mathbf{x}_{t+1} = \mathbb{T}(\mathbf{x}_t, \mathbf{a}_t, \mathbf{w}_t)$. The objective is to find the minimum expected cost and an optimal policy for the entire horizon, starting with any given initial state vector \mathbf{x}_1 .

Consider the following backward explicit-enumeration approach to solve the above DP. We start by setting

$f_{T+1}(\mathbf{x}_{T+1}) = 0$ for all \mathbf{x}_{T+1} . For each \mathbf{x}_T , $f_T(\mathbf{x}_T)$ is computed by solving the minimization problem in (1) via straightforward enumeration. By creating a table that stores the values of f_T for every possible value of \mathbf{x}_T , we can similarly compute the value of $f_{T-1}(\mathbf{x}_{T-1})$ for every possible value of \mathbf{x}_{T-1} . Repeating this enumeration through time, we get $f_1(\mathbf{x}_1)$. Assuming $\mathcal{X}_t = \mathcal{S}$ and $\mathcal{A}_t(\mathbf{x}_t) = \mathcal{A}$ for all $t = 1, 2, \dots, T$, the running time of this exact algorithm is $O(T|\mathcal{S}||\mathcal{A}|)$. Note that $|\mathcal{S}|$ and $|\mathcal{A}|$ are usually pseudo-polynomial in the binary size of the input.¹

Halman et al. (2008, 2009) study a special case of this problem when \mathcal{S} , \mathcal{A} , and w_t are one-dimensional (i.e., $k_1 = k_2 = k_3 = 1$) and the functions f_t and G_t are either convex² or monotone. They show that this special case is itself NP-hard. This motivates their consideration of approximation algorithms. They derive an approximation algorithm that can achieve a solution to within an arbitrary specified relative error $\epsilon > 0$ and runs in polynomial time in the binary size of the input and $1/\epsilon$; such an approximation is commonly referred to as a Fully Polynomial-Time Approximation Scheme (FPTAS).

The basic idea of the approach in Halman et al. (2008, 2009) is to reduce the amount of information needed by the explicit-enumeration approach and, instead of the exact solution, obtain one within a relative error guarantee of ϵ .

By doing this, they reduce the running time of the algorithm to a polynomial in the binary size of the input and $1/\varepsilon$. More specifically, in their problem setting, the functions to be minimized in each period, i.e., the functions on the right-hand side of (1), are one-dimensional and convex. This enables them to solve the minimization step using binary search instead of enumeration, yielding a time factor of $O(\log |\mathcal{X}|)$ rather than $O(|\mathcal{X}|)$. Although the values of f_t in period t ($t = 1, 2, \dots, T$) are still represented using a table, the number of entries in the table is reduced from $|\mathcal{S}|$ to $O((1/\varepsilon) \log f^{\max})$, where f^{\max} is the maximum of f_t over all t and the corresponding domains; i.e., $f^{\max} = \max_t \max_{\mathbf{x}_t \in \mathcal{X}_t} f_t(\mathbf{x}_t)$. This is achieved by storing the values of f_t for only a selected number of points in the domain such that the function values at these points alone can be used to construct a *relative* error approximation to the function. The idea used is that of approximating a one-dimensional convex function (f_t) using a piecewise-linear function. Their approximation coincides with f_t on a subset of \mathcal{X}_t . They refer to this subset as a *K-approximation set* for f_t . The use of binary search for minimization and that of the *K-approximation set* of f_t results in an approximation algorithm that offers a relative error guarantee of ε and runs in polynomial time in the binary input size and $1/\varepsilon$ (since $O((1/\varepsilon) \log f^{\max})$ is polynomial in the binary input size and $1/\varepsilon$). The assumption that f_t is one-dimensional and convex is crucial for the success of this approach.

In this paper, we attempt a generalization of the results of Halman et al. (2008, 2011, 2009, 2013) for fixed dimensions k_1, k_2 , and k_3 . By assuming L^b -convexity of the cost-to-go functions, we obtain *additive* error approximations of these functions by storing their values on subsets of their domains. Furthermore, we prove that these approximations are also L^b -convex functions. Finally, we bound the error propagation of the approximation to establish the required guarantee.

For a detailed comparison of our algorithm with the explicit enumeration procedure, we refer the readers to §6.4. Our approximation algorithm can be viewed as a flexible algorithmic framework. On the one hand, as the desired additive-error guarantee ε reduces, the algorithm approaches the explicit-enumeration procedure; on the other hand, as ε increases, the running time decreases fast. To illustrate the usefulness of our algorithm, we consider an inventory application: the single-product stochastic inventory control problem with lost sales and lead times; see, e.g., Morton (1969) and Zipkin (2008b). We show that the DP for this problem satisfies the assumptions of our analysis. Consequently, we obtain a pseudo-polynomial additive-approximation scheme for this problem. We demonstrate the practical value of our scheme by implementing our approximation scheme and the explicit-enumeration algorithm on instances of this inventory problem.

The paper is organized as follows: Section 2 reviews the related literature. In §3, we present the definition of a discrete L^b -convex function and related results. In §4, we

describe our approximation of a discrete L^b -convex function and prove that this approximation is also L^b -convex. In §5, we present our notation and assumptions, discuss the explicit-enumeration procedure for the DP, and establish the nonexistence of a polynomial-time additive-error approximation algorithm, unless $P = \mathcal{NP}$. In §6, we present our approximation scheme, prove its correctness, analyze its running time, and discuss its properties. Section 7 demonstrates the applicability of our approximation scheme to the above-mentioned inventory problem. Section 8 reports our computational results. We conclude in §9 by discussing an important open question.

2. Literature Review, Our Challenges, and Contributions

We review three fields of literature that are most related to our problem: approximation algorithms for dynamic programs, discrete convexity, and approximation algorithms for stochastic inventory control problems. We end this section with comments regarding some unique challenges we face in approximating fixed-dimensional DPs.

2.1. Approximation Schemes for Dynamic Programs

This is a relatively new direction of research. For deterministic single-dimensional dynamic programs, Woeginger (2000) investigates the conditions under which an FPTAS exists. Halman et al. (2008, 2011) extend this investigation to stochastic single-dimensional dynamic programs. When the possible realizations of the random variables are explicitly specified and are independent in all periods, they propose sufficient conditions to guarantee the existence of an FPTAS. Halman et al. (2013) propose a computationally efficient FPTAS for convex stochastic dynamic programs using techniques in Halman et al. (2008, 2011). The motivation of our work is to obtain an effective (not necessarily polynomial-time) approximation scheme that offers an arbitrarily small *additive* guarantee.

2.2. Discrete Convexity

In the DP (1), we are required to solve a minimization problem in every period. The domain of the objective function in that problem is multidimensional and discrete (more precisely, a subset of \mathbb{Z}^{k_2}). Since we seek an efficient solution procedure for the DP, we are interested in identifying a notion of convexity for functions with discrete domains such that the following conditions hold:

1. Positive scaling, addition, translation, uniform variable scaling, and minimization preserve discrete convexity.
2. Efficient minimization algorithms exist.
3. Convex extension over a continuous domain can be efficiently obtained.

Several notions of discrete convexity have been proposed and studied in the literature—for example, Miller’s discrete convexity (Miller 1971), integer convexity (Favati and

Tardella 1990), and L^{\natural} -convexity (Murota 2003). For our purpose, we find that the notion of L^{\natural} -convexity satisfies conditions 1, 2, and 3.

For a general introduction to discrete convexity, we refer the reader to Murota (2003, 2007).

2.3. Approximation Algorithms for Stochastic Inventory Problems

As mentioned earlier, we will illustrate our approximation scheme on a well-known inventory problem. The important results for this problem are reviewed in §7.

We proceed to review the results on approximation algorithms for stochastic inventory control problems. The first set of papers that propose such algorithms includes Levi et al. (2007, 2008a, and 2008c). All three papers study single location, single product problems and together encompass capacitated and uncapacitated models with correlated demand processes, and both the possibilities of backordering of excess demand and lost sales. For each of these problems, a 2-approximation algorithm is presented. Given the generality of the models and that demands and inventories are not restricted to be integral, a rigorous complexity analysis of the approximation algorithms is not provided in these papers. However, these algorithms have the attractive computational feature that they are nonrecursive, as opposed to the recursive enumeration approach for solving dynamic programs.

Another stream of work on approximation algorithms for stochastic inventory control problems is that of Halman et al. (2009, 2008, 2011). For a single-product version of the problem with discrete demands, Halman et al. (2009) investigate the possibility of approximation algorithms with arbitrarily close errors. They show that the model with convex cost functions and backlogging but no setup cost is weakly NP-hard and develop an FPTAS for this problem. This result is generalized in Halman et al. (2008, 2011) for single-dimensional stochastic dynamic programs. Our work is an attempt to generalize this stream of work to a fixed number of dimensions. Recall that we have already discussed the main elements of the analysis in Halman et al. (2008, 2009) in §1.

2.4. Challenges in Approximating Fixed-Dimensional DPs and Our Contributions

We now briefly discuss some of the difficulties in approximating fixed-dimensional dynamic programs.

The multidimensionality of the domain of $f_{i|}$ in (1) generates additional requirements (relative to those for a single dimension) on our approximation scheme. First, the notion of discrete convexity needs to be chosen carefully. Second, to approximate a fixed-dimensional function $f_{i|}$ to within an additive error, the K -approximation technique in one dimension needs to be adapted appropriately. Third, to prove that our approximation preserves L^{\natural} -convexity (our choice of the notion of discrete convexity),

we need to exploit several properties of L^{\natural} -convex functions and impose additional requirements on our adaptation of the concept of K -approximation sets. Finally, we need some additional developments and assumptions to make our approximation work through time, which is considerably easier for the case of a one-dimensional convex DP.

The main contributions of our work are (i) approximating a fixed-dimensional discrete L^{\natural} -convex function on a bounded rectangular set using only a selected number of points in its domain and (ii) proving that this approximation is also an L^{\natural} -convex function. Furthermore, (iii) we apply this approximation to the DP (1) and develop a pseudo-polynomial time approximation scheme for f_1 with an additive-error guarantee under the primary assumption that all the cost-to-go functions $f_{i|}$ and $G_{i|}$ are L^{\natural} -convex. Our approximation scheme is a generalization of the enumeration algorithm and offers us full control in the trade-off between accuracy and running time; for a more detailed discussion, see §6.4.

Specifically, in our approximation scheme, starting from the last period T , in (1) we replace $f_{T|}$ by its approximation $\tilde{f}_{T|}$. For period $T-1$, this results in an approximation $\tilde{f}_{T-1|}$ of $f_{T-1|}$ on the left-hand side of (1). We then prove that $\tilde{f}_{T-1|}$ is also an L^{\natural} -convex function. This enables us to construct the approximation $\tilde{f}_{T-1|}$ of $\tilde{f}_{T-1|}$ and replace $f_{T-1|}$ by $\tilde{f}_{T-1|}$ in (1) for period $T-2$. Repeating this procedure until the first period, we can obtain an approximation of f_1 . By bounding the error of the approximation in each period, we achieve the desired additive-error guarantee. The usefulness of our approximation scheme is illustrated by applying it to a well-known stochastic inventory control problem, the single-product problem with lost sales and lead times.

To end this section, it is important to clarify that our approach is fundamentally different from that of approximate dynamic programming; see, e.g., Powell (2007). The aim of that approach is to develop practical solutions for high-dimensional dynamic programs using statistical methods for approximating the cost-to-go functions.

3. Discrete L^{\natural} -Convexity and Related Results

This section is organized as follows. In §3.1, we present three different ways of obtaining continuous extensions⁴ of a discrete function—(i) local extension, (ii) global extension, and (iii) convex envelope—and discuss the relationships between them. We show that the convex envelope of a discrete function is the same as its global extension (Lemma 1). In §3.2, we present two different notions of discrete convexity: integer convexity and L^{\natural} -convexity. Several characterizations of integer convexity and L^{\natural} -convexity are also presented. We also present two related definitions: L^{\natural} -convex sets and polyhedral L^{\natural} -convex functions in this section. Finally, in §3.3, we present some properties of discrete L^{\natural} -convex functions that we will use later in our analysis. In particular, we introduce the concept of the Lovász Extension and illustrate its use to obtain a continuous extension

of a discrete L^1 -convex function. A result regarding polyhedral L^1 -convex functions (Lemma 11) is also established in this section.

Throughout the paper, we use $L, U]_{\mathbb{Z}}$ to denote the set of integers between (and including) two given integers L and U , $-\infty < L \leq U < \infty$; i.e., $L, U]_{\mathbb{Z}} = L, U] \cap \mathbb{Z}$. Similarly, for a positive integer n , let $L, U]_{\mathbb{Z}}^n = L, U] \cap \mathbb{Z}^n$. We use $\mathbf{1}$ to denote the vector $(1, 1, \dots, 1)$, where its dimensionality should be clear from the context. Also, we use *increasing* to mean nondecreasing and *decreasing* to mean nonincreasing.

3.1. Continuous Extensions of a Discrete Function

We start this section by defining the convex hull of a set of points in \mathbb{R}^n .

DEFINITION 1. Given a set of m points $H_m = \{\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_m\}$ in \mathbb{R}^n , the convex hull of H_m , denoted by $\text{Conv}(H_m)$, is defined as follows:

$$\text{Conv}(H_m) = \left\{ \mathbf{p} \in \mathbb{R}^n \text{ such that } \mathbf{p} = \sum_{i=1}^m \lambda_i \mathbf{p}_i, \lambda_i \geq 0 \text{ and } \sum_{i=1}^m \lambda_i = 1 \right\}.$$

It is easy to see that $\text{Conv}(L, U]_{\mathbb{Z}}^n) = L, U]_{\mathbb{Z}}^n$. Next, we formally define a piecewise-linear function in n -dimensional real space.

DEFINITION 2 (PIECEWISE-LINEAR FUNCTIONS). Let $D] \subseteq \mathbb{R}^n$ be an n -dimensional polyhedron. A function $f: D] \rightarrow \mathbb{R}$ is referred to as *piecewise-linear* if there exists a set of n -dimensional polyhedra $\mathcal{P} = \{P_1, P_2, \dots, P_i, \dots\}$, indexed by $i \in I$, such that

1. $\bigcup_{i \in I} P_i = D$.
2. For all $i, j \in I$, $i \neq j$, the dimension of $P_i \cap P_j$ is strictly less than n .
3. The function f is linear in \mathbf{x} on P_i for all $i \in I$.

Given a discrete function, there are different ways to obtain a continuous extension. One natural way to do so is via *local* extension: on each unit hypercube in its domain, extend the discrete function to the convex hull of the function graph and finally paste together all these extensions to obtain a continuous function. The next definition, due to Favati and Tardella (1990), formalizes this idea. Given a point $\mathbf{x} \in L, U]_{\mathbb{Z}}^n$, let

$$N(\mathbf{x}) = \{ \mathbf{y} \in L, U]_{\mathbb{Z}}^n : \|\mathbf{x} - \mathbf{y}\|_{\infty} < 1 \},$$

where $\|\mathbf{x}\|_{\infty}$ is the L_{∞} -norm. We refer to $N(\mathbf{x})$ as the discrete neighborhood of \mathbf{x} .

DEFINITION 3 (LOCAL EXTENSION OF A DISCRETE FUNCTION). For $n \in \mathbb{Z}^+$, given a discrete function $f: L, U]_{\mathbb{Z}}^n \rightarrow \mathbb{R}$,

\mathbb{R}^+ , its local extension $f^l: L, U]_{\mathbb{Z}}^n \rightarrow \mathbb{R}^+$ is defined as follows:

$$f^l(\mathbf{x}) = \min \left\{ \sum_{i=1}^{|N(\mathbf{x})|} \alpha_i f(\mathbf{z}_i) : \mathbf{z}_i \in N(\mathbf{x}), \sum_{i=1}^{|N(\mathbf{x})|} \alpha_i \mathbf{z}_i = \mathbf{x}, \sum_{i=1}^{|N(\mathbf{x})|} \alpha_i = 1, \alpha_i \geq 0 \right\}.$$

Observe that for an arbitrary discrete function, its local extension is convex in each unit hypercube and piecewise-linear on the entire domain.

Another natural way to obtain a continuous extension of a discrete function on a bounded domain is the following: instead of the *local* extension, we use the *global* extension, that is, the convex hull of the function graph on the entire domain. The following definition formalizes this idea:

DEFINITION 4 (GLOBAL EXTENSION OF A DISCRETE FUNCTION). For $n \in \mathbb{Z}^+$, given a discrete function $f: L, U]_{\mathbb{Z}}^n \rightarrow \mathbb{R}^+$, its global extension $f^g: L, U]_{\mathbb{Z}}^n \rightarrow \mathbb{R}^+$ is defined as follows:

$$f^g(\mathbf{x}) = \min \left\{ \sum_{i=1}^k \alpha_i f(\mathbf{z}_i) : \mathbf{z}_i \in L, U]_{\mathbb{Z}}^n, \sum_{i=1}^k \alpha_i \mathbf{z}_i = \mathbf{x}, \sum_{i=1}^k \alpha_i = 1, \alpha_i \geq 0 \right\},$$

where $k = (U - L + 1)^n$ is the number of points in the set $L, U]_{\mathbb{Z}}^n$.

Clearly, the global extension of any discrete function is convex.

A third way to extend a discrete function is via its convex envelope. The following definition is adapted from Falk and Hoffman (1976).

DEFINITION 5 (CONVEX ENVELOPE OF A DISCRETE FUNCTION). The convex envelope of function $f: D] \subseteq \mathbb{Z}^n \rightarrow \mathbb{R}$ is a continuous function $\tilde{f}: \text{Conv}(D) \rightarrow \mathbb{R}$ such that

1. $\tilde{f}(\mathbf{x})$ is a convex function.
2. For all $\mathbf{x} \in D$, $\tilde{f}(\mathbf{x}) \leq f(\mathbf{x})$.
3. For any other convex function g defined over $\text{Conv}(D)$, if $g(\mathbf{x}) \leq f(\mathbf{x}), \forall \mathbf{x} \in D$, then $g(\mathbf{x}) \leq \tilde{f}(\mathbf{x}), \forall \mathbf{x} \in \text{Conv}(D)$.

For a function $f: D] \subseteq \mathbb{R}^n \rightarrow \mathbb{R}$ and a set $P] \subseteq D$, let the *restriction* of f to P be defined as $f|_P: P] \rightarrow \mathbb{R}$, where $f|_P(\mathbf{x}) = f(\mathbf{x}), \forall \mathbf{x} \in P$. As observed in Favati and Tardella (1990), for a discrete function $f: L, U]_{\mathbb{Z}}^n \rightarrow \mathbb{R}$, we have the identity $\overline{f|_{N(\mathbf{x})}}(\mathbf{x}) = f^l(\mathbf{x}), \forall \mathbf{x} \in L, U]_{\mathbb{Z}}^n$.

The next lemma characterizes the relationship between the global extension of a discrete function and its convex envelope on its domain. This lemma is a generalization of Theorem 3 in Falk and Hoffman (1976). The proof is simple and, therefore, omitted.

LEMMA 1. For any discrete function $f: L, U]_{\mathbb{Z}}^n \rightarrow \mathbb{R}^+$, its global extension coincides with its convex envelope defined on its domain. That is, for all $\mathbf{x} \in L, U]_{\mathbb{Z}}^n$, $f^g(\mathbf{x}) = \tilde{f}(\mathbf{x})$.

3.2. Definitions of Discrete Convexity

We review two notions of discrete convexity that are used in our analysis. We first define submodularity of a discrete function.

DEFINITION 6 (SUBMODULARITY). For $n \in \mathbb{Z}^+$, a function $f: L, \mathcal{U} \xrightarrow{\mathbb{Z}^+} \mathbb{R}^+$ is submodular if it satisfies the following condition:

$$f(\mathbf{x}) + f(\mathbf{y}) \geq f(\mathbf{x} \vee \mathbf{y}) + f(\mathbf{x} \wedge \mathbf{y}),$$

for all $\mathbf{x}, \mathbf{y} \in L, \mathcal{U}$, where the vector $\mathbf{x} \vee \mathbf{y}$ (respectively, $\mathbf{x} \wedge \mathbf{y}$) denotes the component-wise maximum (respectively, minimum) of \mathbf{x} and \mathbf{y} .

The first notion of discrete convexity we introduce is integer convexity.

DEFINITION 7 (INTEGER CONVEXITY, FAVATI AND TARDELLA 1990). For $n \in \mathbb{Z}^+$, a function $f: L, \mathcal{U} \xrightarrow{\mathbb{Z}^+} \mathbb{R}^+$ is called integrally convex if its local extension $f^\dagger(\mathbf{x}): L, \mathcal{U} \rightarrow \mathbb{R}^+$ is convex.

The following characterizations of integer convexity are due to Favati and Tardella (1990).

LEMMA 2. *The following conditions are equivalent:*

1. Function $f: L, \mathcal{U} \xrightarrow{\mathbb{Z}^+} \mathbb{R}^+$ is integrally convex.
2. The local extension of f coincides with its global extension; i.e.,

$$f^\dagger(\mathbf{x}) = f(\mathbf{x}), \quad \forall \mathbf{x} \in L, \mathcal{U}.$$

In general, addition does not preserve integer convexity; see Favati and Tardella (1990). However, a stronger notion of discrete convexity, namely L^\natural -convexity, is preserved under addition.

DEFINITION 8 (L^\natural -CONVEXITY, MUROTA 2003). For $n \in \mathbb{Z}^+$, a function $f: L, \mathcal{U} \xrightarrow{\mathbb{Z}^+} \mathbb{R}^+$ is called L^\natural -convex if the function $\psi(\mathbf{x}, \mathbf{k}) = f(\mathbf{x} - \mathbf{k}), \mathbf{k} \in \mathbb{Z}, \mathbf{k} \geq 0$, is submodular in (\mathbf{x}, \mathbf{k}) .

The next lemma presents several characterizations of L^\natural -convexity. The second and the third characterization in the following lemma are due to Murota (2003, 2007). The fourth characterization is due to Fujishige and Murota (2000).

LEMMA 3. *The following conditions are equivalent:*

1. Function $f: L, \mathcal{U} \xrightarrow{\mathbb{Z}^+} \mathbb{R}^+$ is L^\natural -convex.
2. Function f satisfies translation submodularity, defined as follows:

$$f(\mathbf{x}) + f(\mathbf{y}) \geq f(\mathbf{x} - \alpha \mathbf{1}) \vee \mathbf{y}) + f(\mathbf{x} \wedge (\mathbf{y} + \alpha \mathbf{1})),$$

for all $\mathbf{x}, \mathbf{y} \in L, \mathcal{U}$, for all $\alpha \in \mathbb{Z}^+$ such that $\mathbf{x} - \alpha \mathbf{1}, \mathbf{y} + \alpha \mathbf{1} \in L, \mathcal{U}$.

3. Function f satisfies discrete midpoint convexity, defined as follows:

$$f(\mathbf{x}) + f(\mathbf{y}) \geq f\left(\left\lfloor \frac{\mathbf{x} + \mathbf{y}}{2} \right\rfloor\right) + f\left(\left\lceil \frac{\mathbf{x} + \mathbf{y}}{2} \right\rceil\right),$$

for all $\mathbf{x}, \mathbf{y} \in L, \mathcal{U}$, where for all $\mathbf{v} \in \mathbb{R}^n, \lfloor \mathbf{v} \rfloor = (\lfloor v_1 \rfloor, \lfloor v_2 \rfloor, \dots, \lfloor v_n \rfloor)$ and $\lceil \mathbf{v} \rceil = (\lceil v_1 \rceil, \lceil v_2 \rceil, \dots, \lceil v_n \rceil)$.

4. Function f is both submodular and integrally convex.

Next, we present two concepts that will be useful later in our analysis.

DEFINITION 9 (L^\natural -CONVEX SET, MUROTA 2003). A subset A of \mathbb{Z}^n is an L^\natural -convex set if there exist constants $a_i, b_i, c_{ij} \in \mathbb{Z}$ such that

$$A = \{\mathbf{x} \in \mathbb{Z}^n : a_i \leq x_i \leq b_i, x_i - x_j \leq c_{ij}, \forall i, j \in [n], i \neq j\}.$$

A continuous function is called polyhedral if its epigraph is polyhedral.

DEFINITION 10 (POLYHEDRAL L^\natural -CONVEXITY, MUROTA AND SHIOURA 2000). Let D be a polyhedron. A function $f: D \rightarrow \mathbb{R}$ is polyhedral L^\natural -convex if

1. f is a polyhedral, convex function and
2. f satisfies translation submodularity, i.e.,

$$f(\mathbf{x}) + f(\mathbf{y}) \geq f((\mathbf{x} - \alpha \mathbf{1}) \vee \mathbf{y}) + f(\mathbf{x} \wedge (\mathbf{y} + \alpha \mathbf{1})),$$

for all $\mathbf{x}, \mathbf{y} \in D$, for all $\alpha \in \mathbb{R}^+$ such that $\mathbf{x} - \alpha \mathbf{1}, \mathbf{y} + \alpha \mathbf{1} \in D$.

3.3. Useful Properties of Discrete L^\natural -Convex Functions

As mentioned in §2, discrete L^\natural -convex functions have the following useful properties:

1. Positive scaling, addition, translation, uniform variable scaling, and minimization preserve L^\natural -convexity.
2. Efficient minimization algorithms exist.
3. Global extension over a bounded domain can be efficiently obtained.

In this section, we present these properties formally. Property 1 corresponds to Lemmas 4–7, Property 2 corresponds to Lemma 8, and Property 3 corresponds to Lemma 10. All these results will be used in our analysis in §§4, 6, and 7.

LEMMA 4 (MUROTA 2003). Let $f: \mathbb{Z}^n \rightarrow \mathbb{R}^+$ be a discrete L^\natural -convex function. Then for all $\alpha \in \mathbb{R}^+$, the function $f': \mathbb{Z}^n \rightarrow \mathbb{R}^+$, defined as

$$f'(\mathbf{x}) = \alpha f(\mathbf{x}),$$

is L^\natural -convex.

LEMMA 5 (MUROTA 2003). Let $f_1: \mathbb{Z}^n \rightarrow \mathbb{R}^+$ and $f_2: \mathbb{Z}^n \rightarrow \mathbb{R}^+$ be two discrete L^\natural -convex functions. Then the function $g: \mathbb{Z}^n \rightarrow \mathbb{R}^+$, defined as

$$g(\mathbf{x}) = f_1(\mathbf{x}) + f_2(\mathbf{x}),$$

is L^\natural -convex.

LEMMA 6 (Murota 2003). Let $f: \mathbb{Z}^n \rightarrow \mathbb{R}^+ \cup \{0\}$ be a discrete L^\natural -convex function. Then for all $a \in \mathbb{Z}^n, b \in \mathbb{Z} \setminus \{0\}$, the function $f[\cdot] \oplus a \oplus b \mathbf{x}$, defined as

$$f[\mathbf{x}] \oplus a \oplus b \mathbf{x},$$

is L^\natural -convex.

LEMMA 7. Let $g(\mathbf{x}, \mathbf{y}): C \subseteq \mathbb{Z}^{n+m} \rightarrow \mathbb{R}$ be an L^\natural -convex function, where C is an L^\natural -convex set. Define

$$h(\mathbf{x}) = \min_{\mathbf{y}: (\mathbf{x}, \mathbf{y}) \in C} g(\mathbf{x}, \mathbf{y}).$$

Then $h(\mathbf{x})$ is also L^\natural -convex.⁵

PROOF. To prove the L^\natural -convexity of $h(\mathbf{x})$, we use translation submodularity. That is, for all $\mathbf{x}_1, \mathbf{x}_2$, and $\alpha \in \mathbb{Z}^+$, we prove the following:

$$h(\mathbf{x}_1) \oplus h(\mathbf{x}_2) \leq h((\mathbf{x}_1 - \alpha \mathbf{1}) \vee \mathbf{x}_2) \oplus h(\mathbf{x}_1 \wedge (\mathbf{x}_2 + \alpha \mathbf{1})).$$

Suppose

$$h(\mathbf{x}_1) = g(\mathbf{x}_1, \mathbf{y}_1) \quad \text{and} \quad h(\mathbf{x}_2) = g(\mathbf{x}_2, \mathbf{y}_2).$$

Since $g(\mathbf{x}, \mathbf{y})$ is L^\natural -convex, it satisfies translation submodularity. We have for $\alpha \in \mathbb{Z}^+$:

$$\begin{aligned} h(\mathbf{x}_1) \oplus h(\mathbf{x}_2) &= g(\mathbf{x}_1, \mathbf{y}_1) \oplus g(\mathbf{x}_2, \mathbf{y}_2) \\ &= g\{[(\mathbf{x}_1, \mathbf{y}_1) - \alpha \mathbf{1} \vee (\mathbf{x}_2, \mathbf{y}_2)]\} \\ &\quad + g\{(\mathbf{x}_1, \mathbf{y}_1) \wedge (\mathbf{x}_2, \mathbf{y}_2) \oplus \alpha \mathbf{1}\} \\ &= g\{(\mathbf{x}_1 - \alpha \mathbf{1}) \vee \mathbf{x}_2, (\mathbf{y}_1 - \alpha \mathbf{1}) \vee \mathbf{y}_2\} \\ &\quad + g\{\mathbf{x}_1 \wedge (\mathbf{x}_2 + \alpha \mathbf{1}), \mathbf{y}_1 \wedge (\mathbf{y}_2 + \alpha \mathbf{1})\} \\ &= h((\mathbf{x}_1 - \alpha \mathbf{1}) \vee \mathbf{x}_2) \oplus h(\mathbf{x}_1 \wedge (\mathbf{x}_2 + \alpha \mathbf{1})). \end{aligned}$$

Note that $g\{[(\mathbf{x}_1, \mathbf{y}_1) - \alpha \mathbf{1} \vee (\mathbf{x}_2, \mathbf{y}_2)]\}$ and $g\{(\mathbf{x}_1, \mathbf{y}_1) \wedge (\mathbf{x}_2, \mathbf{y}_2) \oplus \alpha \mathbf{1}\}$ are well defined since $(\mathbf{x}_1, \mathbf{y}_1) - \alpha \mathbf{1} \vee (\mathbf{x}_2, \mathbf{y}_2), (\mathbf{x}_1, \mathbf{y}_1) \wedge (\mathbf{x}_2, \mathbf{y}_2) \oplus \alpha \mathbf{1} \in C$, which, in turn, follows from the characterization of an L^\natural -convex set; see Property (SBS¹Z) of Murota (2003, p. 128). Thus, the result follows. \square

LEMMA 8. Given a discrete L^\natural -convex function $f: L, U \subseteq \mathbb{Z} \rightarrow \mathbb{R}^+$, it takes time $O(\log^k [U - L])$ to find a minimizer of f .

PROOF. We have that

$$\min_{x \in L, U} f(x) = \min_{x_1 \in L, U} \min_{x_2 \in L, U} \cdots \min_{x_k \in L, U} f(\mathbf{x}).$$

Since minimization preserves L^\natural -convexity (Lemma 7), we know that for $i = 1, 2, \dots, k$, the function

$$\min_{x_i \in L, U} \min_{x_{i+1} \in L, U} \cdots \min_{x_k \in L, U} f(\mathbf{x})$$

is again L^\natural -convex in $(x_1, x_2, \dots, x_{i-1})$. Therefore, to find $\arg \min f$ we can use binary search on each dimension recursively. The result follows. \square

The following definition of the Lovász Extension (Lovász 1983) is needed to obtain a continuous extension of an L^\natural -convex function. We define this extension for an arbitrary discrete function (not necessarily L^\natural -convex) on the hypercube $\{L, U\}^n$. Note that $\{L, U\}^n$ refers to the set $\{\mathbf{z} \in \mathbb{Z}^n: z_j \in \{L, U\}, j = 1, 2, \dots, n\}$. This is the set of vertices of L, U^n . It is easy to see that $\text{Conv}(\{L, U\}^n) = \text{Conv}(\{L, U\}^n) = [L, U]^n$.

DEFINITION 11 (LOVÁSZ EXTENSION). Given a discrete function $f: \{L, U\}^n \rightarrow \mathbb{R}$, for any given point $\mathbf{x} \in [L, U]^n$, let σ be a permutation of $\{1, 2, \dots, n\}$ such that $x_{\sigma(1)} \leq x_{\sigma(2)} \leq \dots \leq x_{\sigma(n)}$. For $0 \leq i \leq n$, define $\mathbf{y}(i) \in \{L, U\}^n$ such that $y_{\sigma(1)}(i) = y_{\sigma(2)}(i) = \dots = y_{\sigma(i)}(i) = U$ and $y_{\sigma(i+1)}(i) = y_{\sigma(i+2)}(i) = \dots = y_{\sigma(n)}(i) = L$. Note that for $0 \leq i \leq n$, the point $\mathbf{y}(i)$ depends on \mathbf{x} . Let λ_j be the unique coefficient of $\mathbf{y}(i)$, if we represent \mathbf{x} as a convex combination of $\mathbf{y}(i)$, $i = 0, 1, \dots, n$; i.e.,

$$\mathbf{x} = \sum_{i=0}^n \lambda_i \mathbf{y}(i). \tag{2}$$

The Lovász Extension $f^{\bar{L}}: [L, U]^n \rightarrow \mathbb{R}$ of function f at point \mathbf{x} is defined as follows:

$$f^{\bar{L}}(\mathbf{x}) = \sum_{i=0}^n \lambda_i f(\mathbf{y}(i)).$$

Observe that in Definition 11, for $i = 0, 1, \dots, n$, λ_i can be obtained by solving Equation (2). In closed form, we have the following: $\lambda_0 = (U - x_{\sigma(1)}) / (U - L)$, $\lambda_j = (x_{\sigma(j)} - x_{\sigma(j+1)}) / (U - L)$, $j = 1, 2, \dots, n-1$, and $\lambda_n = (x_{\sigma(n)} - L) / (U - L)$.

To help the readers gain some intuition about the previous definition, we illustrate it using a simple two-dimensional example. Let $f: [0, 1]^2 \rightarrow \mathbb{R}$ be defined as follows: $f(0, 0) = 1$, $f(0, 1) = 2$, $f(1, 0) = 3$, and $f(1, 1) = 4$. Consider the point $\mathbf{x} = (0.4, 0.6)$. We have, $x_{\sigma(1)} = 0.6$ and $x_{\sigma(2)} = 0.4$. Observe that in Definition 11, given any $\mathbf{x} \in [L, U]^n$, \mathbf{x} can be written as a unique convex combination of $\mathbf{y}(0), \mathbf{y}(1), \dots, \mathbf{y}(n)$. By definition, $\mathbf{y}(0) = (0, 0)$, $\mathbf{y}(1) = (0, 1)$, and $\mathbf{y}(2) = (1, 1)$. Since

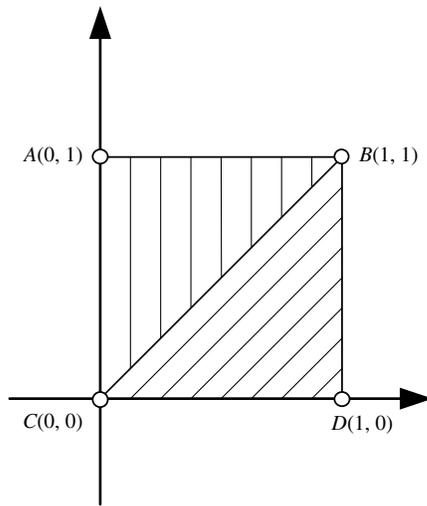
$$\mathbf{x} = 0.4\mathbf{y}(0) \oplus 0.2\mathbf{y}(1) \oplus 0.4\mathbf{y}(2),$$

we conclude that

$$f^{\bar{L}}(\mathbf{x}) = 0.4f(\mathbf{y}(0)) \oplus 0.2f(\mathbf{y}(1)) \oplus 0.4f(\mathbf{y}(2)) = 2.4.$$

It is easy to see that the Lovász Extension of a discrete function on a unit hypercube is piecewise-linear. Furthermore, for an n -dimensional unit hypercube, there are at most $n!$ linear pieces, each corresponding to a domain region such that for any point \mathbf{x} in that region, the order

Figure 1. Illustrating the Lovász Extension of a two-dimensional discrete function defined on the standard unit hypercube.



of its coordinates is fixed. See Figure 1 for an illustration in the two-dimensional case. In this case, the Lovász Extension consists of the two linear pieces on the following triangular regions: $\triangle ABC$ and $\triangle BCD$.

For a discrete submodular function defined on a hypercube $\{L, U\}^n$, its convex envelope is identical to its Lovász Extension. The following result is due to Murota (2003).

LEMMA 9. For a discrete function f defined on a hypercube $\{L, U\}^n$, its convex envelope is identical to its Lovász Extension if and only if f is submodular.

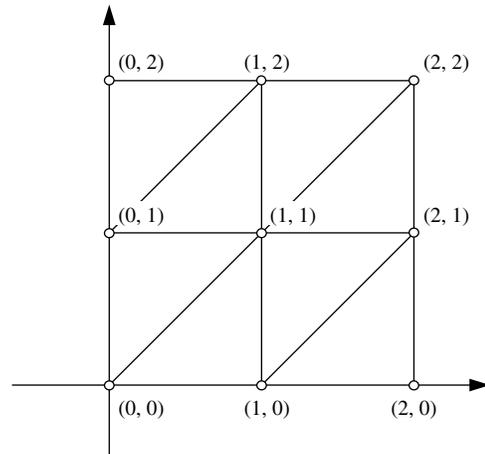
Observe that for a discrete function f , its Lovász Extension is efficiently computable, given oracle access to f . This and the above lemma imply that for a discrete submodular function defined on a hypercube, its convex envelope is efficiently computable. Consequently, for a discrete submodular function on a larger domain, its local extension can be efficiently computed. This is because, as mentioned earlier, for a discrete function f , its local extension restricted to a unit hypercube in its domain is identical to the convex envelope of the function obtained by restricting f to that unit hypercube.

The following result, due to Murota (2003), is about the construction of the global extension of a discrete L^{\natural} -convex function.

LEMMA 10. Given a discrete L^{\natural} -convex function f , its global extension can be obtained as follows: first obtain its Lovász Extension for every unit hypercube in its domain, and then paste all these extensions together.

To help the readers, we explain the proof of this lemma. By the definition of L^{\natural} -convexity, we know that f is both submodular and integrally convex. From the submodularity of f , by Lemma 9, we know that the local extension of f can be constructed by obtaining its Lovász Extension for

Figure 2. Illustrating the convex extension of a two-dimensional discrete L^{\natural} -convex function defined on $[0, 2]^2$.



every unit hypercube in its domain and then pasting all these extensions together. From the integer convexity of f , by Lemma 2, we know that its local extension is identical to its global extension.

From now on, we will refer to the global extension of a discrete L^{\natural} -convex function as its convex extension. Observe that for any discrete L^{\natural} -convex function f , given oracle access to f , the above lemma gives an efficient algorithm to construct its convex extension. See Figure 2 for an illustration of the convex extension of a two-dimensional discrete L^{\natural} -convex function defined on $[0, 2]^2$. The convex extension is linear within each triangular region shown in the figure. Next, we present a result regarding polyhedral L^{\natural} -convex functions that will be useful later in our analysis.

LEMMA 11. Let $f: \mathbb{Z}^n \rightarrow \mathbb{R}^+$ be a discrete L^{\natural} -convex function and f' be its convex extension. Then for $\mathbf{a} \in \mathbb{Z}^n$ and $\beta \in \mathbb{R}^+$, the function $g: \mathbb{Z}^n \rightarrow \mathbb{R}^+$ defined as $g(\mathbf{x}) = f'(\beta) / (\mathbf{x} + \mathbf{a})$ is L^{\natural} -convex.

PROOF.⁶ By Theorem 7.3(1) in Murota (2007),⁷ f' is a polyhedral L^{\natural} -convex function. For any integral \mathbf{a} and positive β , $h(\mathbf{x}) = f'(\beta) / (\mathbf{x} + \mathbf{a})$, defined for a real vector \mathbf{x} , is a polyhedral L^{\natural} -convex function; see Theorem 7.32(1) in Murota (2003). Finally, $g(\mathbf{x}) = h(\mathbf{x})$, defined for an integral vector \mathbf{x} , is a discrete L^{\natural} -convex function; see Theorem 7.1(1) in Murota (2007). \square

4. Approximating L^{\natural} -Convex Functions

Our goal in this section is to approximate an L^{\natural} -convex function by considering its values on only a subset of its domain. To this end, §4.1 defines the notions of an approximation and an ε -approximation set. In §4.2, we present our approximation of an L^{\natural} -convex function and show that this approximation is also an L^{\natural} -convex function. These results will then be used in §6 to develop our approximation algorithm.

4.1. ε -Approximation Sets and Functions

In this subsection, we modify the notion of a K -approximation set in Halman et al. (2009, 2013) for a relative-error approximation to obtain analogous results for an additive-error approximation. We need the following definitions.

DEFINITION 12 (ADDITIVE-ERROR APPROXIMATION). Let $\varepsilon > 0$, $n \in \mathbb{Z}^+$, $D \subseteq \mathbb{R}^n$, and $f: D \rightarrow \mathbb{R}^+$. We say that $\hat{f}: D \rightarrow \mathbb{R}^+$ is an ε -approximation of f , if for all $x \in D$ we have

$$f(x) \leq \hat{f}(x) \leq f(x) + \varepsilon.$$

In the following lemma, we discuss some operations that preserve ε -approximation. The proof is straightforward and, therefore, omitted.

LEMMA 12. Given $\varepsilon_1, \varepsilon_2, \varepsilon_3 > 0$, $m, n \in \mathbb{Z}^+$, $D_1 \subseteq \mathbb{R}^m$, $D_2 \subseteq \mathbb{R}^n$, $f_1: D_1 \rightarrow \mathbb{R}^+$, $f_2: D_2 \rightarrow \mathbb{R}^+$, and $g: D_1 \times D_2 \rightarrow \mathbb{R}^+$. Let $\hat{f}_1: D_1 \rightarrow \mathbb{R}^+$ be an ε_1 -approximation of f_1 , $\hat{f}_2: D_2 \rightarrow \mathbb{R}^+$ be an ε_2 -approximation of f_2 and $\hat{g}: D_1 \times D_2 \rightarrow \mathbb{R}^+$ be an ε_3 -approximation of g . Then the following statements hold:

1. For $\alpha > 0$, $\alpha \hat{f}_1$ is an $\alpha \varepsilon_1$ -approximation of αf_1 .
2. The function $\hat{f}_1 + \hat{f}_2$ is an $(\varepsilon_1 + \varepsilon_2)$ -approximation of $f_1 + f_2$.
3. The function $\min_{\mathbf{y}} \hat{g}(\mathbf{x}, \mathbf{y})$ is an ε_3 -approximation of $\min_{\mathbf{y}} g(\mathbf{x}, \mathbf{y})$.

DEFINITION 13 (ε -APPROXIMATION SET FOR A CONVEX FUNCTION). Let $\varepsilon > 0$ and let $g: L, U \rightarrow \mathbb{R}^+$ be a convex function. An ε -approximation set of g is an ordered set of integers $S = \{i_1 = L < i_2 < \dots < i_r = U\} \subseteq L, U$ satisfying the following properties:

For each $k \in \{1, 2, \dots, r-1\}$ such that $i_{k+1} > i_k + \lambda$, the following holds:

1. If g is monotonically increasing in i_k, i_{k+1} , then

$$g(i_{k+1}) - g(i_k) \leq g(i_k + \lambda) - g(i_k) \leq \varepsilon.$$

2. If g is monotonically decreasing in i_k, i_{k+1} , then

$$g(i_k) - g(i_{k+1}) \leq g(i_{k+1} - \lambda) - g(i_{k+1}) \leq \varepsilon.$$

3. If g first decreases and then increases in i_k, i_{k+1} , let $x^* = \arg \min_{x \in [i_k, i_{k+1}]} g(x)$. Then

$$g(i_k) - g(x^*) \leq \varepsilon \quad \text{and} \quad g(i_{k+1}) - g(x^*) \leq \varepsilon.$$

Note that (i) an ε -approximation set of a convex function always exists, and in the worst case, includes all the points in L, U and (ii) the minimizer $x^* = \arg \min_{x \in L, U} g(x)$ is not necessarily included in an ε -approximation set of g . We now show how an ε -approximation set of a convex function can be used to construct an ε -approximation of that function.

LEMMA 13. Let $\varepsilon > 0$ and let $g: L, U \rightarrow \mathbb{R}^+$ be a convex function. Let $S = \{i_1 = L < i_2 < \dots < i_r = U\} \subseteq L, U$ be an ε -approximation set of g . Then the function $\check{g}: L, U \rightarrow \mathbb{R}^+$, defined as follows, is an ε -approximation of g :

$$\check{g}(x) = \begin{cases} g(x), & \text{if } x \in \{i_1, i_2, \dots, i_r\}, \\ \frac{g(i_{k+1}) - g(i_k)}{i_{k+1} - i_k} x + \frac{g(i_k) i_{k+1} - g(i_{k+1}) i_k}{i_{k+1} - i_k}, & \text{if } i_k < x < i_{k+1}, \text{ for some } k \in \{1, 2, \dots, r-1\}. \end{cases}$$

In words, the function \check{g} is constructed by linearly interpolating the function values defined on any two consecutive points in S . The proof of the lemma follows easily from the convexity of g and is, therefore, omitted.

Note that an ε -approximation set may not be unique. Next, to suit our purpose of approximating a discrete L^1 -convex function, we define the canonical uniform-interval ε -approximation set.

DEFINITION 14 (UNIFORM-INTERVAL ε -APPROXIMATION SET). A uniform-interval ε -approximation set for a convex function $g: L, U \rightarrow \mathbb{R}^+$ is an ε -approximation set for g such that the intervals between two consecutive points in the set have the same length; i.e.,

$$i_r - i_{r-1} = i_{r-1} - i_{r-2} = \dots = i_2 - i_1.$$

The canonical uniform-interval ε -approximation set for g is a uniform-interval ε -approximation set such that the length of the interval is the maximum possible power of 2.

Figure 3 illustrates the above definition. For the function $f: 0, 4 \rightarrow \mathbb{R}^+$ shown in the figure, $\{0, 2, 4\}$ is the canonical uniform-interval 3-approximation set.

The following lemma bounds the size of the canonical uniform-interval ε -approximation set for a monotone convex function g and the time required for its construction.

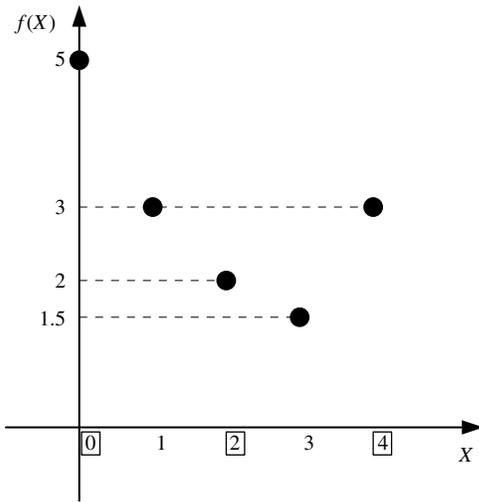
LEMMA 14. Let $g: L, U \rightarrow \mathbb{R}^+$ be a monotone convex function and let $s = \max\{|g(L) - g(L + \lambda)|, |g(U) - g(U - \lambda)|\}$. For any $\varepsilon > 0$, the canonical uniform-interval ε -approximation set S of g has cardinality $O((U - L) \min\{1, s/\varepsilon\})$. Furthermore, it takes time $O((1 + t_g)(U - L) \min\{1, s/\varepsilon\})$ to construct this set, where t_g is the time needed to evaluate g at a point.

PROOF. We prove the statement for the case when g is increasing. The proof for the case when g is decreasing is similar and, therefore, omitted.

For a given $\varepsilon > 0$, let t denote the length of the interval in the canonical uniform-interval ε -approximation set S of g . Since $2t$ is not the interval length, then by definition, there must exist some $x \in [L, L + 2t], [U - 2t, U]$ such that

$$g(x + 2t) - g(x) \leq g(x + \lambda) - g(x) \leq \varepsilon.$$

Figure 3. Illustrating the definition of the canonical uniform-interval ε -approximation set for a convex function f



Then we have

$$g(x+2t) - g(x) \leq g(x+1) - g(x) + 2t \quad (1)$$

$$g(U) - g(U-2t) \leq g(x+1) - g(x) + 2t \quad (2)$$

(convexity of g)

$$g(U) - g(U-2t) \leq g(U) - g(U-t) \quad (g \text{ is increasing})$$

$$2ts \leq g(U) - g(U-t) \quad (3)$$

(convexity of g) and the definition of s)

Thus,

$$2ts \geq \varepsilon. \quad (3)$$

Therefore, the cardinality of S , given by $1 + \lfloor (U-L)/t \rfloor$, is bounded by $O((U-L)(s/\varepsilon))$. Additionally, since the cardinality is also trivially bounded by $O(U-L)$, we conclude that the cardinality is $O((U-L) \min\{1, s/\varepsilon\})$. As a result, the determination of the interval length t and the construction of the ε -approximation set S can be done in time $O(t_g(U-L) \min\{1, s/\varepsilon\})$, by simply checking interval lengths in the sequence of $2^l, 2^{l-1}, \dots, 1$ (assuming $U-L = 2^l$). \square

Note that the upper bound $O((U-L) \min\{1, s/\varepsilon\})$ is not tight because of the strict inequality in (3).

We now generalize the lemma above to convex functions. The proof is similar and, therefore, omitted.

LEMMA 15. Let $g: [L, U]_{\mathbb{Z}} \rightarrow \mathbb{R}^+$ be a convex function, and let $s = \max\{|g(L) - g(L+1)|, |g(U) - g(U-1)|\}$. For every $\varepsilon > 0$, the canonical uniform-interval ε -approximation set S of g has cardinality $O((U-L) \min\{1, s/\varepsilon\})$ and can be constructed in time $O((1 + t_g) \min\{1, s/\varepsilon\})$, where t_g is the time needed to evaluate g at a point.

4.2. Approximating a Discrete L^{\square} -Convex Function

In this subsection, our purpose is to construct an additive-error approximation of a discrete L^{\square} -convex function $f: [L, U]_{\mathbb{Z}}^k \rightarrow \mathbb{R}^+$. Our approximation only uses the values of the original function on a subset of its domain, and is a generalization of the single-dimensional approximation of §4.1. Hereafter, we refer to the domain of f , i.e., $[L, U]_{\mathbb{Z}}^k$ as the domain cube.

Let t be a factor of $U-L$ (i.e., t divides $U-L$ exactly). Let $a_{ij} \in \{L, L+t, \dots, U\}$, $j = 1, 2, \dots, k$. Then the induced gridlines of the hypercube $[L, U]^k$ corresponding to a_{ij} , $j = 1, 2, \dots, k$, are as follows:

$$\{(x_1, a_2, a_3, \dots, a_k) \mid x_1 \in [L, U]_{\mathbb{Z}}\},$$

$$\{(a_1, x_2, a_3, \dots, a_k) \mid x_2 \in [L, U]_{\mathbb{Z}}\},$$

$$\vdots$$

$$\{(a_1, a_2, a_3, \dots, x_k) \mid x_k \in [L, U]_{\mathbb{Z}}\}.$$

The complete set of induced grid-lines of $[L, U]^k$ is obtained by considering all possible choices of $a_1, a_2, \dots, a_k \in \{L, L+t, \dots, U\}$. The following definition generalizes to multiple dimensions our notion of canonical uniform-interval ε -approximation set for a single-dimensional function (Definition 14).

DEFINITION 15 (ε -COARSE CUBE). Let $\varepsilon > 0$ and t be a factor of $U-L$. Let $f: [L, U]_{\mathbb{Z}}^k \rightarrow \mathbb{R}^+$ be an L^{\square} -convex function. We refer to the set $[L, U]^k$ as the ε -coarse cube of f and t as the length of the ε -coarse cube if the following two conditions hold:

1. For each induced gridline E of $[L, U]^k$, the set $[L, U]^k \cap E$ is an ε -approximation set of $f|_E$ where $f|_E$ is the restriction of f to E .
2. The length t is the maximum power of 2 such that condition 1 above holds.

Note that this definition can be easily generalized to the case when the domain is a rectangular set instead of a hypercube.

We use values of f only at points on its ε -coarse cube to construct its approximation \check{f} .

DEFINITION 16 (CONVEX ε -APPROXIMATION).⁸ Given a discrete L^{\square} -convex function $f: [L, U]_{\mathbb{Z}}^k \rightarrow \mathbb{R}^+$, $\varepsilon > 0$, and the ε -coarse cube C of f , we define $\check{f}: [L, U]_{\mathbb{Z}}^k \rightarrow \mathbb{R}^+$ as follows:

$$\check{f}(\mathbf{x}) = \min_{z \in C} f(\mathbf{x}, z) \in \mathcal{C}(\text{Conv}(H)),$$

where H is the set of all points $(\mathbf{p}, f(\mathbf{p}))$ such that $\mathbf{p} \in C$. Thus, $\check{f}(\mathbf{x}) = \overline{f|_C}(\mathbf{x})$, $\forall \mathbf{x} \in [L, U]_{\mathbb{Z}}^k$.

Our next result shows that \check{f} constructed above closely approximates f and preserves L^{\square} -convexity.

LEMMA 16. Given a discrete L^{\square} -convex function $f: L, \mathcal{U} \rightarrow \mathbb{R}^+ / \varepsilon > 0$ and the ε -coarse cube C of f the convex ε -approximation function \check{f} of f (obtained as in Definition 16) is its $k\varepsilon$ -approximation in the additive sense. Furthermore, \check{f} is also L^{\square} -convex.

PROOF. We first prove the L^{\square} -convexity result and then prove the approximation result.

Let t be the length of the ε -coarse cube C of f . Let $f_1(\mathbf{x}) = f(\mathbf{x} + L\mathbf{1})$. By Lemma 6, f_1 is a discrete L^{\square} -convex function defined on $0, \mathcal{U} - L \frac{k}{\mathbb{Z}}$. Let $f_2(\mathbf{x}) = f_1(t\mathbf{x})$. Again, by Lemma 6, f_2 is a discrete L^{\square} -convex function defined on $0, \mathcal{U} - L / t \frac{k}{\mathbb{Z}}$. Let f_2' be the convex extension of f_2 ; i.e., for all $\mathbf{x} \in 0, \mathcal{U} - L / t \frac{k}{\mathbb{Z}}$, we have

$$f_2'(\mathbf{x}) = \min \left\{ \sum_{i=1}^n \alpha_i f_2(\mathbf{z}_i) : \mathbf{z}_i \in \left[0, \frac{\mathcal{U} - L}{t} \right] \frac{k}{\mathbb{Z}}, \sum_{i=1}^n \alpha_i \mathbf{z}_i = \mathbf{x}, \sum_{i=1}^n \alpha_i = 1, \alpha_i \geq 0 \right\}, \quad (4)$$

where $n = \ell(\mathcal{U} - L / t + \lambda) k$ is the number of points in the set $0, \mathcal{U} - L / t \frac{k}{\mathbb{Z}}$.

Let function $g: L, \mathcal{U} \rightarrow \mathbb{R}^+ / \varepsilon$ be defined as

$$g(\mathbf{x}) = f_2' \left(\frac{1}{t}(\mathbf{x} - L\mathbf{1}) \right).$$

Using the above definitions of g, f_2', f_2 , and f_1 , we obtain that for all $\mathbf{x} \in L, \mathcal{U} \frac{k}{\mathbb{Z}}$,

$$g(\mathbf{x}) = \min \left\{ \sum_{i=1}^n \alpha_i f(\mathbf{z}_i) : \mathbf{z}_i \in C, \sum_{i=1}^n \alpha_i \mathbf{z}_i = \mathbf{x}, \sum_{i=1}^n \alpha_i = 1, \alpha_i \geq 0 \right\},$$

where C is the ε -coarse cube of f and $n = |C| = \ell(\mathcal{U} - L) / t + \lambda k$.

Noticing that $\check{f}(\mathbf{x})$ is the restriction of $g(\mathbf{x})$ to the discrete domain $L, \mathcal{U} \frac{k}{\mathbb{Z}}$ from Lemma 11, we conclude that $\check{f}(\mathbf{x})$ is discrete L^{\square} -convex.

Next, we prove that $f(\mathbf{x}) - \check{f}(\mathbf{x}) \leq k\varepsilon$ for all $\mathbf{x} \in L, \mathcal{U} \frac{k}{\mathbb{Z}}$. Pick any point $\mathbf{x} \in L, \mathcal{U} \frac{k}{\mathbb{Z}}$. From the convexity of f , we have $f(\mathbf{x}) \leq \check{f}(\mathbf{x})$. To prove that $\check{f}(\mathbf{x}) - f(\mathbf{x}) \leq k\varepsilon$, consider any hypercube $\mathcal{H} \subseteq C$ such that $\mathbf{x} \in \text{Conv}(\mathcal{H})$ and each edge of \mathcal{H} is of length t . By construction, we have

$$\check{f}(\mathbf{x}) = \sum_{i=0}^k \lambda_i f(\mathbf{y}(i)),$$

where $\mathbf{y}(i), i = 0, \dots, k$, are as defined in Definition 11 and $\mathbf{x} = \sum_{i=0}^k \lambda_i \mathbf{y}(i)$. Notice that for $i = 0, \dots, k - 1$, we have $\mathbf{y}(i) \vee \mathbf{y}(i + 1) = \mathbf{y}(i + 1)$ and $\mathbf{y}(i) \wedge \mathbf{y}(i + 1) = \mathbf{y}(i)$. To proceed, we need to prove

$$\sum_{i=0}^k \lambda_i f(\mathbf{y}(i)) - f(\mathbf{x}) \leq k\varepsilon.$$

To this end, observe that

$$\sum_{i=1}^{k-1} f(\mathbf{y}(i)) - f(\mathbf{x}) + f(\mathbf{x} \wedge \mathbf{y}(1)) + f(\mathbf{x} \vee \mathbf{y}(1) \wedge \mathbf{y}(2)) + f(\mathbf{x} \vee \mathbf{y}(2) \wedge \mathbf{y}(3)) + \dots + f(\mathbf{x} \vee \mathbf{y}(k-1) \wedge \mathbf{y}(k)). \quad (5)$$

The above inequality is obtained by adding the following inequalities, which all hold since f is submodular:

$$\begin{aligned} f(\mathbf{x}) + f(\mathbf{y}(1)) &\geq f(\mathbf{x} \wedge \mathbf{y}(1)) + f(\mathbf{x} \vee \mathbf{y}(1)), \\ f(\mathbf{x} \vee \mathbf{y}(1)) + f(\mathbf{y}(2)) &\geq f(\mathbf{x} \vee \mathbf{y}(1) \wedge \mathbf{y}(2)) + f(\mathbf{x} \vee \mathbf{y}(2)) \\ &\quad (\text{since } \mathbf{y}(1) \vee \mathbf{y}(2) = \mathbf{y}(2)), \\ &\vdots \\ f(\mathbf{x} \vee \mathbf{y}(k-1)) + f(\mathbf{y}(k)) &\geq f(\mathbf{x} \vee \mathbf{y}(k-1) \wedge \mathbf{y}(k)) + f(\mathbf{y}(k)) \\ &\quad (\text{since } \mathbf{x} \vee \mathbf{y}(k) = \mathbf{y}(k)). \end{aligned}$$

Therefore,

$$\begin{aligned} &\sum_{i=0}^k \lambda_i f(\mathbf{y}(i)) - f(\mathbf{x}) \\ &= \sum_{i=0}^k \lambda_i f(\mathbf{y}(i)) + \sum_{i=1}^{k-1} \lambda_i f(\mathbf{y}(i)) - \left[\sum_{i=1}^{k-1} f(\mathbf{y}(i)) + f(\mathbf{x}) \right] \\ &\quad + \sum_{i=0}^k \lambda_i f(\mathbf{y}(i)) + \sum_{i=1}^{k-1} \lambda_i f(\mathbf{y}(i)) - f(\mathbf{x} \wedge \mathbf{y}(1)) \\ &\quad - f(\mathbf{x} \vee \mathbf{y}(1) \wedge \mathbf{y}(2)) - f(\mathbf{x} \vee \mathbf{y}(2) \wedge \mathbf{y}(3)) - \dots / \\ &\quad - f(\mathbf{x} \vee \mathbf{y}(k-1) \wedge \mathbf{y}(k)) \quad (\text{from (5)}) \\ &\quad \check{f}(\mathbf{x} \wedge \mathbf{y}(1)) - f(\mathbf{x} \wedge \mathbf{y}(1)) + \check{f}(\mathbf{x} \vee \mathbf{y}(1) \wedge \mathbf{y}(2)) \\ &\quad - f(\mathbf{x} \vee \mathbf{y}(1) \wedge \mathbf{y}(2)) + \check{f}(\mathbf{x} \vee \mathbf{y}(2) \wedge \mathbf{y}(3)) \\ &\quad - f(\mathbf{x} \vee \mathbf{y}(2) \wedge \mathbf{y}(3)) + \dots + \check{f}(\mathbf{x} \vee \mathbf{y}(k-1) \wedge \mathbf{y}(k)) \\ &\quad - f(\mathbf{x} \vee \mathbf{y}(k-1) \wedge \mathbf{y}(k)) \\ &\leq k\varepsilon \quad (\text{Definition 15 and Lemma 13}). \end{aligned}$$

The second to last inequality above holds since

$$\begin{aligned} \mathbf{x} &= \sum_{i=0}^k \lambda_i \mathbf{y}(i), \\ \mathbf{x} \wedge \mathbf{y}(1) &= \sum_{i=0}^k \lambda_i \mathbf{y}(i) \wedge \mathbf{y}(1) = \lambda_0 \mathbf{y}(0) + \sum_{i=1}^k \lambda_i \mathbf{y}(1), \\ \mathbf{x} \vee \mathbf{y}(1) \wedge \mathbf{y}(2) &= \sum_{i=0}^k \lambda_i \mathbf{y}(i) \vee \mathbf{y}(1) \wedge \mathbf{y}(2) = \sum_{i=0}^k \lambda_i \mathbf{y}(2), \\ &\vdots \\ \mathbf{x} \vee \mathbf{y}(k-1) \wedge \mathbf{y}(k) &= \sum_{i=0}^k \lambda_i \mathbf{y}(i) \vee \mathbf{y}(k-1) \wedge \mathbf{y}(k) = \sum_{i=0}^k \lambda_i \mathbf{y}(k), \end{aligned}$$

Figure 4. Original discrete function f (proof of Lemma 16).

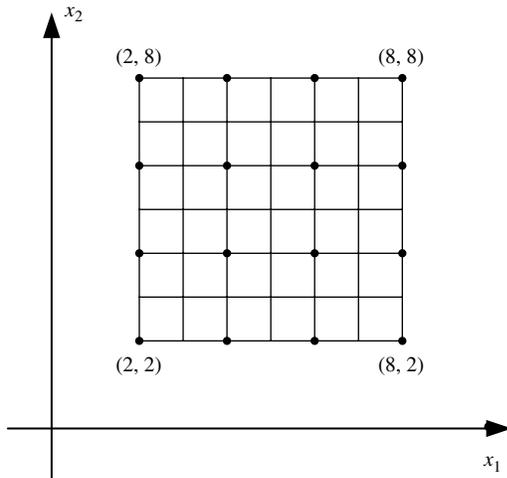


Figure 5. Discrete function f_1 (proof of Lemma 16).

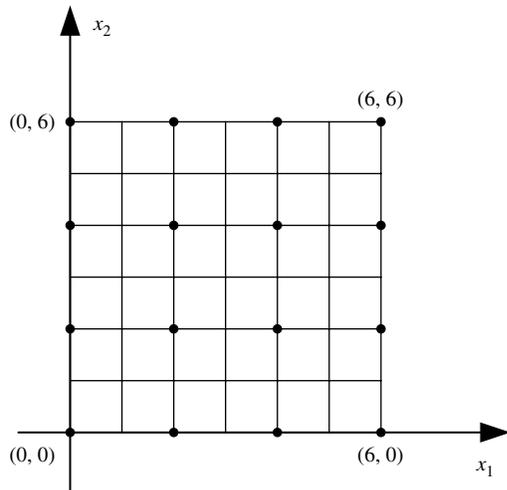


Figure 6. Discrete function f_2 (proof of Lemma 16).

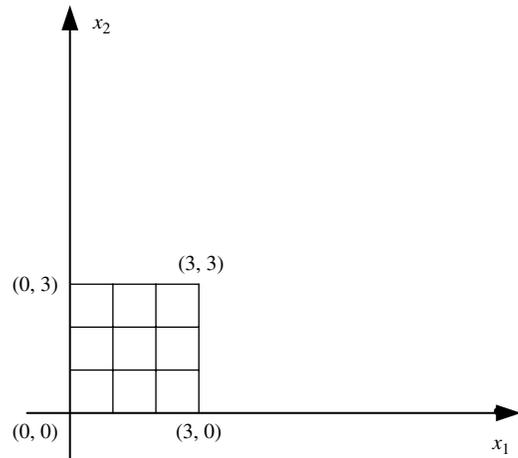
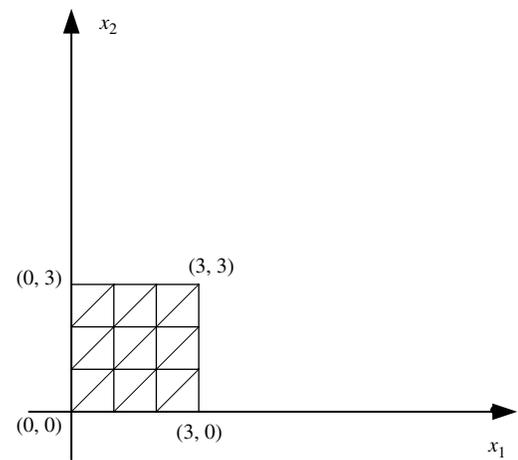


Figure 7. Function f'_2 as the convex extension of f_2 (proof of Lemma 16).



and

$$\begin{aligned} \check{f}(\mathbf{x} \wedge \mathbf{y}(1)) &= \lambda_0 f(\mathbf{y}(0)) + \left[\sum_{i=1}^k \lambda_i f(\mathbf{y}(1)) \right] \\ \check{f}([\mathbf{x} \vee \mathbf{y}(1)] \wedge \mathbf{y}(2)) &= \left[\sum_{i=0}^1 \lambda_i f(\mathbf{y}(1)) + \sum_{i=2}^k \lambda_i f(\mathbf{y}(2)) \right] \\ &\vdots \\ \check{f}([\mathbf{x} \vee \mathbf{y}(k-1)] \wedge \mathbf{y}(k)) &= \left[\sum_{i=0}^{k-1} \lambda_i f(\mathbf{y}(k-1)) + \lambda_k f(\mathbf{y}(k)) \right] \end{aligned}$$

The result follows. \square

Figures 4 through 9 illustrate the proof of Lemma 16. In the figures, the original function f is defined on $2, \mathbb{B}^2_{\mathbb{Z}}$. The length of its ε -coarse cube is 2. The domain points on the ε -coarse cube are depicted using black dots.

Figure 8. Function g (proof of Lemma 16).

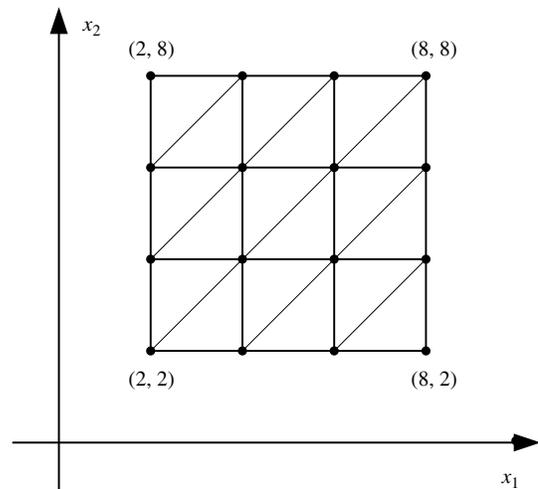
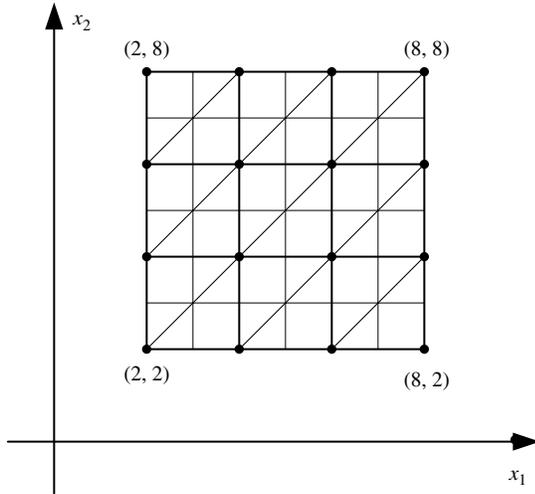


Figure 9. Discrete function \check{f} (proof of Lemma 16).

We now generalize Lemma 15 to higher dimensions. The proof is similar and is, therefore, omitted.

LEMMA 17. Let $f: [L, U]_{\mathbb{Z}}^k \rightarrow \mathbb{R}^+$ be a discrete L^{\square} -convex function, and let s be an upper bound on the absolute value of the slope of the function $f|_{E_1}$ for all induced gridline E of $[L, U]_{\mathbb{Z}}^k$. For every $\varepsilon > 0$, the ε -coarse cube of f has cardinality $O((U-L)^k \min\{1, s/\varepsilon\}^k)$ and can be constructed in time $O((1 + t_f)(U-L)^k \min\{1, s/\varepsilon\}^k)$, where t_f is the time needed to evaluate f at a point.

Given an L^{\square} -convex function $f: [L, U]_{\mathbb{Z}}^k \rightarrow \mathbb{R}^+$ and its ε -coarse cube C , the proof of Lemma 16 suggests an alternative way to construct \check{f} as an approximation of f . First, we obtain function g defined in the proof using the following method: On each smallest hypercube within C , obtain the Lovász Extension of f and then paste these extensions together. Then our desired approximation \check{f} can be obtained by restricting g to the discrete domain. In other words, for any given $\mathbf{x} \in [L, U]_{\mathbb{Z}}^k$, $\check{f}(\mathbf{x})$ is the Lovász Extension of f at point \mathbf{x} in the smallest hypercube in C that contains \mathbf{x} .

Recall that to obtain the Lovász Extension of f for a given point \mathbf{x} , by Definition 11, we need to access the neighboring points of \mathbf{x} on the ε -coarse cube and the corresponding values of f . Later, in our approximation algorithm, for a given $\varepsilon > 0$ and a discrete L^{\square} -convex function f , we will store the points of its ε -coarse cube and the function values at these points and use this information to support future queries of \check{f} . The next result characterizes the query time, given this information. The query time of $\check{f}(\mathbf{x})$ is the time taken to return the value of $\check{f}(\mathbf{x})$ for any \mathbf{x} in the domain.

LEMMA 18. For a given L^{\square} -convex function $f: [L, U]_{\mathbb{Z}}^k \rightarrow \mathbb{Z}^+$ and its ε -coarse cube C , if all points in C and the function values at these points are stored in a sorted list, then for any $\mathbf{x} \in [L, U]_{\mathbb{Z}}^k$, the query time of $\check{f}(\mathbf{x})$ is $O((k + 1)(\log(U-L) + \beta) + \beta + t)$, where 2^l is the length of C

and β is the amount of space required to store $f(\mathbf{x})$ for a given \mathbf{x} in the domain.

PROOF. As discussed earlier, for any $\mathbf{x} \in [L, U]_{\mathbb{Z}}^k$, $\check{f}(\mathbf{x})$ is the Lovász Extension of f at point \mathbf{x} in the smallest hypercube in C that contains \mathbf{x} . To compute this, we first determine the smallest hypercube in C that contains \mathbf{x} and then use Definition 11. Recall that in Definition 11, for any given \mathbf{x} , we find the $k+1$ points $\mathbf{y}(0), \mathbf{y}(1), \dots, \mathbf{y}(k)$ and interpolate the corresponding function values $f(\mathbf{y}(0)), f(\mathbf{y}(1)), \dots, f(\mathbf{y}(k))$. This requires time $O((k+1)(\log(U-L) + \beta))$. Since 2^l is the length of C and \check{f} is constructed using linear interpolation, $2^l \check{f}(\mathbf{x})$ is integer-valued for all $\mathbf{x} \in [L, U]_{\mathbb{Z}}^k$. Thus, outputting the exact value of $\check{f}(\mathbf{x})$ takes time $O(\beta + t)$. The result follows. \square

We will use this lemma in §6 to analyze the running time of our approximation algorithm.

5. Assumptions on the Dynamic Program and Analysis of the Explicit-Enumeration Algorithm

We first define the notation, state the assumptions for our problem, and specify the input in §5.1. Next, in §5.2, we present an explicit-enumeration dynamic programming algorithm and analyze its running time in the binary size of the input. At the end of this section, we show that no polynomial-time approximation algorithm with a fixed additive-error guarantee can exist for the problem unless $P = \text{NP}$.

5.1. Notation and Assumptions

Let T denote the length of the finite horizon. In each period $t \in \{1, 2, \dots, T\}$, the sequence of events is as follows: At the beginning of the period, \mathbf{x}_t is observed. Then an action \mathbf{a}_t is taken. At the end of the period, the random variable \mathbf{w}_t is realized. For each period, the action \mathbf{a}_t belongs to a constrained action space $\mathcal{A}_t(\mathbf{x}_t)$. The state transition equation is $\mathbf{x}_{t+1} = \mathcal{T}(\mathbf{x}_t, \mathbf{a}_t, \mathbf{w}_t)$, where $\mathcal{T}: \mathbb{Z}^{k_1+k_2+k_3} \rightarrow \mathbb{Z}^{k_1}$ is a transformation function. Given $\mathbf{x}_t, \mathbf{a}_t$, and \mathbf{w}_t the cost incurred in period t is $r_t(\mathbf{x}_t, \mathbf{a}_t, \mathbf{w}_t)$. Let $G_t(\mathbf{x}_t, \mathbf{a}_t) = \sum_{\mathbf{w}_t} r_t(\mathbf{x}_t, \mathbf{a}_t, \mathbf{w}_t)$ denote the expected cost incurred in period t for a given \mathbf{x}_t and \mathbf{a}_t . Starting from period t with state \mathbf{x}_t the expected total cost incurred until the end of the horizon is denoted by $f_t(\mathbf{x}_t)$. For convenience, the notation is summarized below.

- \mathbf{x}_t : beginning state vector for period t ;
- \mathbf{a}_t : action vector for period t ;
- \mathbf{w}_t : random vector, realized after the action is decided in period t ;

$\mathcal{T}(\mathbf{x}_t, \mathbf{a}_t, \mathbf{w}_t)$: ending state vector for period t , which is also \mathbf{x}_{t+1} ;

$G_t(\mathbf{x}_t, \mathbf{a}_t)$: the expected cost for period t , when the starting state is \mathbf{x}_t and the action decided is \mathbf{a}_t ;

$f_t(\mathbf{x}_t)$: the expected total cost incurred, starting from period t with state \mathbf{x}_t , until the end of the horizon.

We assume that we are given an oracle, as part of the input, which computes functions G_t for $t \in \{1, \dots, T\}$. Note that to encode an oracle that outputs a positive integer-valued function f , we need at least $\Omega(\log f^{\max})$ space since this is the minimum space required to output the value f^{\max} , where f^{\max} is the maximum value of f on its domain. The possible values of the random vector \mathbf{w}_t and their corresponding rational probabilities are explicitly specified. In period t , the possible values of \mathbf{w}_t are $\mathbf{w}_{t,1}, \dots, \mathbf{w}_{t,n_t}$. The corresponding probabilities are specified via positive integer numbers $q_{t,j}$, $j \in \{1, \dots, n_t\}$, with $\text{Prob } \mathbf{w}_t = \mathbf{w}_{t,j} = q_{t,j} / (\sum_{j=1}^{n_t} q_{t,j})$, for $j \in \{1, \dots, n_t\}$. We assume that the sequence of random variables $\{\mathbf{w}_t\}_{t=1}^T$ is independent. We define the following values (when appropriate, for every $t \in \{1, \dots, T\}$ and $i \in \{1, \dots, n_t\}$):

- $p_{t,j} = \text{Prob } \mathbf{w}_t = \mathbf{w}_{t,j}$: Probability that \mathbf{w}_t is realized as $\mathbf{w}_{t,j}$
- $w^* = \max_{t,j} \|\mathbf{w}_{t,j}\|_\infty$: Maximum L_∞ -norm of possible realizations of the random vector \mathbf{w}_t over the entire horizon;
- $n^* = \max_t n_t$: Maximum value of n_t over the entire horizon;
- $Q_t = \sum_{j=1}^{n_t} q_{t,j}$: A common denominator of all the probabilities in period t ;
- $M_t = \prod_{j=1}^t Q_j$: A common denominator of all the probabilities over periods $t \in \{1, \dots, T\}$;
- $M_{T+1} = 1$.

The objective is to find the minimum expected total cost over the entire horizon, starting from any given initial state vector \mathbf{x}_1 . We recall the DP recursion (1): For period $t \in \{1, \dots, T\}$, we have

$$f_t(\mathbf{x}_t) = \min_{\mathbf{a}_t \in \mathcal{A}_t(\mathbf{x}_t)} [G_t(\mathbf{x}_t, \mathbf{a}_t) + \sum_{\mathbf{w}_t} p_{t,j} f_{t+1}(\mathbb{T}(\mathbf{x}_t, \mathbf{a}_t, \mathbf{w}_t))].$$

Our objective is to solve this problem. For all $t \in \{1, \dots, T\}$, define $\mathcal{D}_t \subseteq \mathbb{Z}^{k_1+k_2}$ as

$$\mathcal{D}_t = \{(\mathbf{x}_t, \mathbf{a}_t) : \mathbf{a}_t \in \mathcal{A}_t(\mathbf{x}_t)\}.$$

We make the following assumptions.

ASSUMPTION 1. In (1), for all $t \in \{1, \dots, T\}$, f_t and G_t are L^b -convex and \mathcal{D}_t is an L^b -convex set.

ASSUMPTION 2. The transformation \mathbb{T} satisfies the following condition: If $f(\mathbf{x})$ is L^b -convex, then $f(\mathbb{T}(\mathbf{x}, \mathbf{a}, \mathbf{w}))$ is also L^b -convex in (\mathbf{x}, \mathbf{a}) for every integral \mathbf{w} .

Let $g_t(\mathbf{x}_t, \mathbf{a}_t) = G_t(\mathbf{x}_t, \mathbf{a}_t) + \sum_{\mathbf{w}_t} p_{t,j} f_{t+1}(\mathbb{T}(\mathbf{x}_t, \mathbf{a}_t, \mathbf{w}_t))$. By Assumptions 1 and 2, we have that for all $t \in \{1, \dots, T\}$, g_t is L^b -convex. This is because addition and positive scaling preserve L^b -convexity; see Lemmas 4 and 5.

ASSUMPTION 3.⁹ For fixed positive integers k_1, k_2 , and k_3 , $\mathbf{x}_t \in \mathbb{Z}^{k_1}, \mathbf{a}_t \in \mathbb{Z}^{k_2}, \mathbf{w}_t \in \mathbb{Z}^{k_3}$.

ASSUMPTION 4. For $t \in \{1, \dots, T\}$ the value of G_t at any point in its domain is a positive rational number and can be evaluated in polynomial time in the size of the input using the corresponding oracle (which implies that $\log G_t^{\max}$ is polynomially bounded in the size of the input).

Let t_{G_t} denote the time it takes for a single call of the oracle.

Next, we determine the minimum space required to specify the input of our problem. Let $\mathcal{A}_t = \mathcal{A}_t$, where $u = \arg \max_{\mathcal{A}_t} |u|$. To specify the constrained action set $\mathcal{A}_t(\mathbf{x}_t)$ for all periods, we need $\Omega(T \log |\mathcal{A}|)$ space. To specify the possible realizations of \mathbf{w}_t over $t \in \{1, \dots, T\}$, we need $O(T n^* k_3 \log w^*)$ space. We assume that $O(1)$ space is needed to specify the transformation function \mathbb{T} . To specify the oracle that computes G_t for all $t \in \{1, \dots, T\}$, we require space $\Omega(T \log G^{\max})$, where $G^{\max} = \max_{t,j} G_{t,j}^{\max}$. Therefore, the overall input size is bounded below by

$$\Omega(T + \log |\mathcal{A}| + n^* k_3 \log w^* + \log G^{\max}).$$

We now present an explicit-enumeration algorithm, which provides an exact solution to the problem. Then we analyze its time complexity.

5.2. An Explicit-Enumeration Dynamic Programming Algorithm

Let $F_t(\mathbf{x}_t) = M_t f_t(\mathbf{x}_t)$ be the integer version of f_t . Algorithm 1 below is an explicit-enumeration dynamic program. Note that all the function values encountered in the algorithm are integers. Let $\mathcal{S} = \mathcal{S}$, where $v = \arg \max_{\mathcal{S}} |v|$. Consider the minimization in Step 4. From Assumption 1, we know that the function to be minimized is L^b -convex. Thus, the minimization step can be performed using recursive binary search (see Lemma 8). Consequently, Algorithm 1 has time complexity $O(t_G (\log(M_T T G^{\max})) + \log |\mathcal{S}| T |\mathcal{S}| \log |\mathcal{A}|)$ and space complexity $O((\log(M_T T G^{\max})) + \log |\mathcal{S}| |\mathcal{S}|)$, where $G^{\max} = \max_{t,j} G_{t,j}^{\max}$ is the maximum value of G_t over all feasible points and over all periods $t \in \{1, \dots, T\}$. Note that $O(\log(M_T T G^{\max}))$ is the time required to store a function value in Step 3 and $O(\log |\mathcal{S}|)$ is the time required to store a point in \mathcal{S} . Observe that $|\mathcal{S}|$ typically depends on the input in a super-log manner. Thus, the algorithm is, both in time and space, pseudo-polynomial in the input size.

Algorithm 1 (An Explicit-Enumeration Dynamic Program)

- 1 Let $F_{T+1} = 0$;
- 2 **for** $t = T$ **to** 1 **do**
- 3 For all points \mathbf{x}_t in the domain, calculate and store
- 4 $F_t(\mathbf{x}_t) = \min_{\mathbf{a}_t \in \mathcal{A}_t(\mathbf{x}_t)} [M_t G_t(\mathbf{x}_t, \mathbf{a}_t) + Q_t \sum_{i=1}^{n_t} p_{t,i} \cdot F_{t+1}(\mathbb{T}(\mathbf{x}_t, \mathbf{a}_t, \mathbf{w}_{t,i}))]$;
- 5 Delete the stored values $F_{t+1}(\mathbf{x}_{t+1})$ for all \mathbf{x}_{t+1} ;
- 6 **end**
- 7 Return $F_1(\mathbf{x}_1) / M_1$.

It is easy to see that the problem studied in Halman et al. (2009)—the single-item stochastic inventory control problem with discrete demands—is a special case of our problem. They show that this problem is NP-hard, which implies that our problem is NP-hard as well. Also, the above pseudo-polynomial time algorithm implies that our problem is weakly NP-hard.

Note that, unless $P = NP$, no polynomial-time approximation algorithm exists with a fixed additive-error guarantee for our problem. Following an argument similar to the proof of Theorem 3.1 of Ausiello et al. (1999, p. 89), this statement can be proved by contradiction. Suppose that there exists a polynomial-time approximation algorithm A with an additive error δ for our problem. Then we can construct an exact algorithm that runs in polynomial time in the following manner: For our problem, for any instance, we can scale up all function values to obtain a new instance where all function values are multiples of $\lceil \delta \rceil$. For the new instance, Algorithm A is an exact algorithm with polynomial running time. However, this implies membership in class P , which contradicts the NP-hardness of the problem unless $P = NP$.

In the next section, we will present our approximation scheme with an arbitrary additive error ε , which runs in time pseudo-polynomial in the size of the input and polynomial in $1/\varepsilon$.

6. An Additive-Error Approximation Scheme

In this section, we propose and analyze an approximation scheme for the dynamic program (1). Section 6.1 provides an overview of this algorithm and an additional result that we will use in our analysis. Then §6.2 presents the details of the algorithm and its analysis. Subsequently, we discuss the corresponding near-optimal policy in §6.3. Finally, in §6.4, we summarize the highlights of our analysis.

6.1. Overview of the Scheme and Additional Developments

Recall the dynamic programming recursion (1): For period $t \in \{1, 2, \dots, T\}$, we have

$$f_t(\mathbf{x}_t) = \min_{\mathbf{a}_t \in \mathcal{A}_t(\mathbf{x}_t)} G_t(\mathbf{x}_t, \mathbf{a}_t) + \mathbb{E}_{\mathbf{w}_t} [f_{t+1}(\mathbb{T}(\mathbf{x}_t, \mathbf{a}_t, \mathbf{w}_t))],$$

where $\mathbf{x}_t \in \mathcal{L}_1, \mathcal{U}_1 \subseteq \mathbb{Z}^{k_1}, \mathbf{a}_t \in \mathcal{L}_2, \mathcal{U}_2 \subseteq \mathbb{Z}^{k_2}$, and $\mathbf{w}_t \in \mathcal{L}_3, \mathcal{U}_3 \subseteq \mathbb{Z}^{k_3}$. The procedure we use is a modification of the explicit-enumeration algorithm. Specifically, we start from the last period T and, assuming $f_{T+1} = 0$, obtain f_T via (1). For period $T-1$, note that the calculation of f_{T-1} requires the value of f_T . Rather than storing f_T for all \mathbf{x}_T (as in explicit enumeration), we approximate f_T using the convex ε -approximation \check{f}_T (Definition 16). This results in the following approximation of f_{T-1} :

$$\check{f}_{T-1}(\mathbf{x}_{T-1}) = \min_{\mathbf{a}_{T-1} \in \mathcal{A}_{T-1}(\mathbf{x}_{T-1})} G_{T-1}(\mathbf{x}_{T-1}, \mathbf{a}_{T-1}) + \mathbb{E}_{\mathbf{w}_{T-1}} [\check{f}_T(\mathbb{T}(\mathbf{x}_{T-1}, \mathbf{a}_{T-1}, \mathbf{w}_{T-1}))].$$

To continue using this approximation method for period $T-2$, it is desirable for \check{f}_{T-1} to have the following two properties: (i) \check{f}_{T-1} is L^2 -convex (this will enable us to use our results in §4.2 to obtain \check{f}_{T-1} as an approximation to f_{T-1}^{10}) and (ii) \check{f}_{T-1} closely approximates f_{T-1} . Furthermore, if we can establish these properties for \check{f}_t for all $t \in \{1, 2, \dots, T\}$, then we can repeatedly approximate f_t for all $t \in \{1, 2, \dots, T\}$ using the following recursion:

$$\check{f}_t(\mathbf{x}_t) = \min_{\mathbf{a}_t \in \mathcal{A}_t(\mathbf{x}_t)} G_t(\mathbf{x}_t, \mathbf{a}_t) + \mathbb{E}_{\mathbf{w}_t} [\check{f}_{t+1}(\mathbb{T}(\mathbf{x}_t, \mathbf{a}_t, \mathbf{w}_t))], \quad (6)$$

where \check{f}_t is the convex ε -approximation of f_t . Our next result establishes these properties.

LEMMA 19. For $t \in \{1, 2, \dots, T\}$, consider functions f_t and f_{t+1} in (1). Suppose that $\check{f}_{t+1}: \mathcal{L}, \mathcal{U} \subseteq \mathbb{Z}^k \rightarrow \mathbb{R}^+$ is an ε -approximation of f_{t+1} and is an L^2 -convex function. Then \check{f}_t (defined as in (6) above) is also an L^2 -convex function and is an ε -approximation of f_t .

PROOF. Using Lemma 12, it is easy to show that \check{f}_t is an ε -approximation of f_t . Next, we prove that \check{f}_t is L^2 -convex. For notational convenience, denote

$$\check{g}_t(\mathbf{x}_t, \mathbf{a}_t) = G_t(\mathbf{x}_t, \mathbf{a}_t) + \mathbb{E}_{\mathbf{w}_t} [\check{f}_{t+1}(\mathbb{T}(\mathbf{x}_t, \mathbf{a}_t, \mathbf{w}_t))]. \quad (7)$$

We now have

$$\check{f}_t(\mathbf{x}_t) = \min_{\mathbf{a}_t \in \mathcal{A}_t(\mathbf{x}_t)} \check{g}_t(\mathbf{x}_t, \mathbf{a}_t). \quad (8)$$

We know that $\check{f}_{t+1}(\mathbf{x}_{t+1})$ is L^2 -convex. Next, we prove that $\check{g}_t(\mathbf{x}_t, \mathbf{a}_t)$ is L^2 -convex. We can write $\mathbb{E}_{\mathbf{w}_t} [\check{f}_{t+1}(\mathbb{T}(\mathbf{x}_t, \mathbf{a}_t, \mathbf{w}_t))] = \sum_{i=1}^{n_t} p_{i,t} \check{f}_{t+1}(\mathbb{T}(\mathbf{x}_t, \mathbf{a}_t, \mathbf{w}_{i,t}))$. By Assumption 2, $\check{f}_{t+1}(\mathbb{T}(\mathbf{x}_t, \mathbf{a}_t, \mathbf{w}_{i,t}))$ is L^2 -convex in $(\mathbf{x}_t, \mathbf{a}_t)$ for all $\mathbf{w}_{i,t}$. Observe from (7) that $\check{g}_t(\mathbf{x}_t, \mathbf{a}_t)$ is a positive linear combination of L^2 -convex functions and is, therefore, L^2 -convex; see Lemmas 4 and 5. By Lemma 7, minimization over an L^2 -convex set¹¹ preserves L^2 -convexity. Therefore, we conclude that $\check{f}_t(\mathbf{x}_t)$ is again L^2 -convex. The result follows. \square

Next, we present our approximation algorithm formally.

6.2. Our Approximation Scheme

Our approximation algorithm is a modification of Algorithm 1. As in Algorithm 1, we work with the integer versions of the cost-to-go functions: $F_t(\mathbf{x}_t) = M_t f_t(\mathbf{x}_t)$. In each period $t \in \{1, 2, \dots, T\}$ instead of storing F_t for every possible value of \mathbf{x}_t in the domain, we construct \check{F}_t as an approximation of F_t using less space.¹² Rewriting (6) by multiplying both sides of the equation by M_t we have for $t \in \{1, 2, \dots, T\}$

$$\check{F}_t(\mathbf{x}_t) = \min_{\mathbf{a}_t \in \mathcal{A}_t(\mathbf{x}_t)} M_t G_t(\mathbf{x}_t, \mathbf{a}_t) + \mathbb{E}_{\mathbf{w}_t} [\check{F}_{t+1}(\mathbb{T}(\mathbf{x}_t, \mathbf{a}_t, \mathbf{w}_t))]. \quad (9)$$

all induced gridlines E] of the domain cube $L_1, \mathcal{U}_1 \stackrel{k_1}{\mathbb{Z}^E}$. The time to store the $\varepsilon' M_t$ -coarse cube in Step 4 and the corresponding function values in Step 5 is

$$O \left[t_{\tilde{F}_t} (U_1 - L_1)^{k_1} \left[\min \left\{ 1, \frac{S_t}{\varepsilon'} \right\} \right]^{k_1} (T \log(U_1 - L_1) \log F^{\max}) \right].$$

Note that the term $O(T \log(U_1 - L_1) \log F^{\max})$ accrues from the need to store the k_1 coordinates of a point on the ε -cube and store the corresponding function value (to be precise, this term is $O(k_1 \log(U_1 - L_1) + T \log(U_1 - L_1) \log F^{\max})$, but we only use the dominant factor).

Thus, for any period t , the running time of Algorithm 2 (omitting the constant terms) is

$$O \left[\log |\mathcal{A}| n^3 T^2 \log^2(U_1 - L_1) \log^2 F^{\max}(U_1 - L_1)^{k_1} \cdot \left[\min \left\{ 1, \frac{S_t}{\varepsilon'} \right\} \right]^{k_1} \right].$$

Observe that $O(\log F^{\max}) = O(\log(M_1 T G^{\max}))$. The overall running time is

$$O \left[\log |\mathcal{A}| n^3 T^2 \log^2(U_1 - L_1) \log^2(M_1 T G^{\max})(U_1 - L_1)^{k_1} \cdot \left[\min \left\{ 1, \frac{S_t}{\varepsilon'} \right\} \right]^{k_1} \right].$$

Since k_1 is fixed and $\varepsilon' = \varepsilon / (k_1 T)$, the only nonpolynomial factors in the above expression, with respect to the binary size of the input and $1/\varepsilon$, are $O((U_1 - L_1)^{k_1} / \min\{1, \lfloor sk_1 T \rfloor / \varepsilon\}^{k_1})$. The result follows. \square

6.3. A Corresponding Near-Optimal Policy

Theorem 1 only guarantees a near-optimal *value* for our problem. It is also important for our approximation to generate a corresponding near-optimal *policy*. We first define a policy and its expected cost and then present our result in Theorem 2. For our problem, a policy π] is defined by a set of functions

$$\{a_t^\pi(\mathbf{x}) : \mathcal{X}^E \rightarrow \mathcal{A}_t\}, \quad t = \lambda, \mu, \dots, T]$$

The expected cost of policy π] is given by $f_1^\pi(\mathbf{x}_1)$, where

$$f_t^\pi(\mathbf{x}_t) = G_t(\mathbf{x}_t, \mathbf{h}_t^\pi(\mathbf{x}_t)) + \omega_t f_{t+1}^\pi(\mathbb{T}(\mathbf{x}_t, \mathbf{h}_t^\pi(\mathbf{x}_t), \mathbf{w}_t))$$

for $t = \lambda, \mu, \dots, T]$ and $f_{T+1}^\pi(\mathbf{x}) = \emptyset, \forall \mathbf{x} \in L_1, \mathcal{U}_1 \stackrel{k_1}{\mathbb{Z}^E}$.

Given $\varepsilon > 0$, let π^ε] denote the policy defined as follows:

$$\forall (t, \mathbf{x}_t), a_t^{\pi^\varepsilon}(\mathbf{x}_t) = \underset{\mathbf{a}_t \in \mathcal{A}_t(\mathbf{x}_t)}{\text{arg min}} \left[M_t G_t(\mathbf{x}_t, \mathbf{h}_t) + \mathcal{Q}_t \sum_{i=1}^{n_t} P_{t,i} \cdot \tilde{F}_{t+1}(\mathbb{T}(\mathbf{x}_t, \mathbf{h}_t, \mathbf{w}_t)) \right],$$

where $\tilde{F}_t(\cdot)$, $t = \lambda, \mu, \dots, T]$ + λ , are as defined in Algorithm 2. We have the following result. The proof is straightforward and, therefore, omitted.

THEOREM 2 (ε -APPROXIMATE POLICY). *Under Assumptions 1 through 4, the policy π^ε] defined above is an ε -approximate policy for our problem. That is, for any initial state vector \mathbf{x}_1 , $f_1^{\pi^\varepsilon}(\mathbf{x}_1) - f_1(\mathbf{x}_1) \leq \varepsilon$.*

Note that by the definition of π^ε], we need to obtain $a_t^{\pi^\varepsilon}(\mathbf{x}_t)$] for every t] and \mathbf{x}_t . If we choose to store $a_t^{\pi^\varepsilon}(\mathbf{x}_t)$] for all t] and \mathbf{x}_t , then the space required is linear in $O(|\mathcal{S}|)$, which may be undesirable. Alternatively, for any given t] and \mathbf{x}_t , we can query the corresponding $a_t^{\pi^\varepsilon}(\mathbf{x}_t)$] in the following way: After executing Algorithm 2, we can obtain $a_t^{\pi^\varepsilon}(\mathbf{x}_t)$] by solving

$$a_t^{\pi^\varepsilon}(\mathbf{x}_t) = \underset{\mathbf{a}_t \in \mathcal{A}_t(\mathbf{x}_t)}{\text{arg min}} \left[M_t G_t(\mathbf{x}_t, \mathbf{h}_t) + \mathcal{Q}_t \sum_{i=1}^{n_t} P_{t,i} \cdot \tilde{F}_{t+1}(\mathbb{T}(\mathbf{x}_t, \mathbf{h}_t, \mathbf{w}_t)) \right]$$

in time $O(\log |\mathcal{A}| (t_{G_t} + n_t T) \log(U_1 - L_1)) \log(T M_1 G^{\max})$.

6.4. Highlights of Our Analysis

Generalization of the approximation technique used in Halman et al., (2009, 2008, 2011, 2013)

Although the aim in Halman et al. (2009, 2008, 2011, 2013) is a relative-error approximation and our goal is an additive-error approximation, our approach can be viewed as a generalization of theirs, in the following sense:

1. Recall that in Halman et al. (2009, 2008, 2011) and (2013), the authors use the following idea to approximate a single-dimensional discrete convex function $f] : L, \mathcal{U} \stackrel{k_1}{\mathbb{Z}^E} \rightarrow \mathbb{R}^+$:

- (a) Select only a subset S] of $L, \mathcal{U} \stackrel{k_1}{\mathbb{Z}^E}$ by using the concept of a K -approximation set.
- (b) Let $g]$ be the piecewise-linear extension¹³ of $f]$ induced by S .
- (c) Restrict $g]$ to $L, \mathcal{U} \stackrel{k_1}{\mathbb{Z}^E}$ to approximate¹⁴ $f]$.

2. Our approach to approximate a fixed-dimensional discrete L^k -convex function $f] : L, \mathcal{U} \stackrel{k_1}{\mathbb{Z}^E} \rightarrow \mathbb{R}$] generalizes the above idea:

- (a) Select only a subset S] of $L, \mathcal{U} \stackrel{k_1}{\mathbb{Z}^E}$ by using the concept of an ε -coarse cube.
- (b) Let $g]$ be the global extension (see Definition 4) of the function obtained by restricting $f]$ to S .
- (c) Restrict $g]$ to $L, \mathcal{U} \stackrel{k_1}{\mathbb{Z}^E}$ to approximate $f]$.

6.4.1. Properties of Algorithm 2 and Comparison with Algorithm 1

1. Algorithm 2 preserves the structure of the problem. That is, in each period, our approximation of the cost-to-go function is also L^k -convex, regardless of the choice of the error-guarantee ε .

2. Algorithm 2 can be viewed as a domain-reduction technique. Specifically, in each period, given the desired additive error ε , let 2^l be the length of the corresponding ε -coarse cube. Then, the total number of domain points is $(1 + U] - L)^{k_1}$, whereas the number of points on the ε -coarse cube is only $(1 + \lfloor U] - L \rfloor / 2^l)^{k_1}$.

3. To achieve this reduction, the overhead that our algorithm incurs is the increase in the space required to store the function values of \tilde{F}_t on the points of the ε -coarse cube (see Claim 1 on page 31). However, this increase is only a logarithmic factor. For example, if the lengths of the ε -coarse cubes for period $t \in [T]$ are respectively $2^{l_1}, 2^{l_2}, \dots, 2^{l_T}$, then to store a function value of \tilde{F}_t it takes $\sum_{i=t+1}^T l_i$ additional bits compared to the space required by Algorithm 1.

4. The running time of Algorithm 2 is decided by items 2 and 3 above. For each period, on the one hand, item 2 allows us to reduce the number of points that we store and hence decreases the running time. On the other hand, for each point, item 3 increases the storage space required and hence increases the running time. However, the increase is logarithmic compared to the decrease.

6.4.2. Practical Algorithmic Framework. From the discussion above, Algorithm 2 can be viewed as a practical algorithmic framework that offers us full control in the trade-off between accuracy and running time. To this end, a possible use of Algorithm 2 may take the following steps:

1. For a given problem, we start with a relatively large additive error ε so that Algorithm 2 provides a quick solution. For the justification of considering a relatively large additive error ε , we refer the readers to item 3 below.

2. For any initial state \mathbf{x}_1 , Algorithm 2 outputs $\check{f}_1(\mathbf{x}_1)$ such that $f_1(\mathbf{x}_1) \leq \check{f}_1(\mathbf{x}_1) \leq f_1(\mathbf{x}_1) + \varepsilon$, which implies $\check{f}_1(\mathbf{x}_1) - \varepsilon \leq f_1(\mathbf{x}_1) \leq \check{f}_1(\mathbf{x}_1)$. Thus, Algorithm 2 provides us both an upper bound of f_1 and a lower bound of $f_1 - \varepsilon$ of the original function value f_1 .

3. The bounds in item 2 above imply that the relative error is $(\check{f}_1(\mathbf{x}_1) - f_1(\mathbf{x}_1)) / f_1(\mathbf{x}_1) \leq \varepsilon / (f_1(\mathbf{x}_1) - \varepsilon)$ for any initial state \mathbf{x}_1 . In words, an additive-error guarantee of ε implies a relative-error guarantee of $\varepsilon / (f_1(\mathbf{x}_1) - \varepsilon)$. Therefore, choosing a large value of ε can still give us an acceptable relative error (assuming $f_1(\mathbf{x}_1)$ is large enough). This motivates the choice of a large additive error ε .

4. If the obtained bounds are unsatisfactory, then the additive error ε can be reduced. The running time will increase accordingly. This step can be repeated multiple times.

7. Application to an Inventory Problem

Our purpose in this section is to illustrate the use of our approximation scheme for the discrete version of the following well-known problem: the single-product stochastic inventory control problem with lost sales and lead times, studied, among others, in Morton (1969) and Zipkin (2008b). Hereafter, we will refer to this problem as INV-LS.

For INV-LS, Morton (1969) proves certain monotonicity properties of the optimal policy with respect to the state vector. Morton's proof involves an inductive analysis of the second-order derivatives of the cost-to-go functions. Recently, Zipkin (2008b) rederives Morton's results using

the following technique: First, he transforms the state and action space to create a new dynamic program. He then shows that the cost-to-go functions of this new DP are L^2 -convex. This implies certain monotonicity properties of the optimal policy in the new DP that, in turn, imply the desired properties of the optimal policy in the original DP. We will use Zipkin's results to show that INV-LS satisfies the assumptions of our analysis. Thus, Algorithm 2 can be used to solve this problem approximately to within a positive additive-error bound of ε (i.e., Theorems 1 and 2 are applicable).

In §7.1, we describe INV-LS. Then in §7.2, we show that the DP formulation corresponding to this problem satisfies our assumptions of §5.

7.1. Problem INV-LS

The objective is to minimize the expected sum of costs for a single-product inventory system over a finite horizon. Next, we describe the sequence of events in period $t \in [T]$:

1. The order that is due in period t arrives.
 2. A new order (if any) is placed at a purchase cost of c per unit. This order will arrive after a constant lead time of L periods.
 3. Demand is realized. Any demand that cannot be immediately satisfied is lost. If there is inventory left over at the end of the period, a unit holding cost of h is incurred. Any lost sales result in a penalty cost of p per unit.
- For ease of exposition, we assume that the discount factor is 1. We assume that demands across periods are independent. We also assume that *all inventory levels, order quantities, cost parameters (c , h , and p), and all possible demand realizations are nonnegative integers.*

We now describe the formulation using our DP (1). Let q_t denote the order quantity in period t . Let y_t denote the inventory level after the order that is due in period t is delivered. The ordering cost incurred in period t is cq_t . Let $v_{0,t} = y_t, v_{1,t} = y_t - q_t, \dots, v_{L-1,t} = y_t - (L-1)q_t$. Then the state vector in period t is $\mathbf{v}_t = (v_{0,t}, v_{1,t}, \dots, v_{L-1,t})$. Zipkin (2008b) transforms the state and action space as follows: Let $\mathbf{x}_t = (x_{0,t}, x_{1,t}, \dots, x_{L-1,t})$, where $x_{k,t} = \sum_{i=k}^{L-1} v_{i,t} \geq 0, k \in [L-1]$. Let w_t denote the random demand in period t for the single product. The possible realizations of w_t and the corresponding rational probabilities are explicitly specified in the same way as described in §5. Let $a_t = q_t$. It is easy to see that since the cost parameters are stationary, it is strictly suboptimal to order more than w_t in any period, where $w_t = \max_{i \in [L-1]} x_{i,t}$. Thus, we let the feasible set of a_t be $\mathcal{A}_t = \{a_t \in \mathbb{Z}^+ : a_t \leq w_t\}$. The transformation function \mathbb{T} from period t to $t+1$ is given by

$$\mathbf{x}_{t+1} = (x_{0,t+1}, x_{1,t+1}, \dots, x_{L-1,t+1})$$

$$x_{k,t+1} = x_{k,t} - a_t + w_t, \quad k \in [L-1]$$

For our purpose in this paper, we use the same DP formulation as that in the proof of Theorem 4 in Zipkin

(2008b). Specifically, for period t , we use two additional integer decision variables s_t and r_t that denote the selling decision amount and the remaining inventory after sale, respectively. Note that the decisions s_t and r_t are contingent upon the realization of demand w_t .

Assume that the last period in which ordering is allowed is T , i.e., $\mathcal{A}_{T+1} = \mathcal{A}_{T+2} = \dots = \mathcal{A}_{T+L} = \{0\}$. Also, assume that holding and penalty costs continue to accrue until period $T+L$. By convention, let $f_{T+L+1}(\cdot) = 0$. The expected cost $f_t(\mathbf{x}_t)$, incurred from period t until the end of the horizon $T+L$, is given by the following DP:

$$f_t(\mathbf{x}_t) = \min_{a_t \in \mathcal{A}_t} -ca_t + \sum_{w_t} u_t(\mathbf{x}_t, a_t, w_t) \tag{10}$$

where

$$u_t(\mathbf{x}_t, a_t, w_t) = \min_{\substack{s_t, r_t \in \mathbb{Z} \\ 0 \leq s_t, r_t \leq x_{0t} - x_{1t}}} hr_t + p(w_t - s_t) + f_{t+1}(x_{1t} + r_t, x_{2t}, \dots, x_{L-1,t}, 0) - a_t \mathbf{1}.$$

Using $s_t + r_t = x_{0t} - x_{1t}$ to remove s_t we obtain

$$u_t(\mathbf{x}_t, a_t, w_t) = \min_{\substack{r_t \in \mathbb{Z}, 0 \leq r_t \leq x_{0t} - x_{1t} \\ 0 \leq x_{0t} - x_{1t} - r_t}} hr_t + p(w_t + r_t + x_{1t} - x_{0t}) + f_{t+1}(x_{1t} + r_t, x_{2t}, \dots, x_{L-1,t}, 0) - a_t \mathbf{1}.$$

Letting $x_{+t} = x_{1t} + r_t$ we have

$$u_t(\mathbf{x}_t, a_t, w_t) = \min_{\substack{x_{+t} \in \mathbb{Z}, x_{+t} \geq x_{1t} \\ 0 \leq x_{0t} - x_{+t}}} h(x_{+t} - x_{1t}) + p(w_t + x_{+t} - x_{0t}) + f_{t+1}(x_{+t}, x_{2t}, \dots, x_{L-1,t}, 0) - a_t \mathbf{1}.$$

Observe that in the above equation, the feasible set of x_{+t} is dependent on the demand w_t . However, recall that in our DP (1), the feasible set of actions is independent of the demand. Thus, to move toward the form (1), we need some additional notation.

For a possible realization $w_{[t]}$ of w_t , let x_{+t}^i denote the corresponding x_{+t} , $i \in \{1, 2, \dots, n_t\}$. Let the feasible set of x_{+t}^i be

$$\mathcal{C}^i(\mathbf{x}_t) = \{x_{+t} \in \mathbb{Z} : x_{+t} \in [x_{1t}, x_{0t} - x_{+t} - w_{[t]}]\}.$$

Define $\tilde{\mathbf{a}}_t = (\mathcal{C}^1, \mathcal{C}^2, \dots, \mathcal{C}^{n_t})$. Then the feasible set of $\tilde{\mathbf{a}}_t$ is

$$\tilde{\mathcal{A}}_t(\mathbf{x}_t) = \mathcal{A}_t \times \mathcal{C}^1(\mathbf{x}_t) \times \mathcal{C}^2(\mathbf{x}_t) \times \dots \times \mathcal{C}^{n_t}(\mathbf{x}_t).$$

Let

$$G_t(\mathbf{x}_t, \tilde{\mathbf{a}}_t) = -ca_t + \sum_{i=1}^{n_t} p_{i,t} h(x_{+t}^i - x_{1t}) + p(w_{[t]} + x_{+t}^i - x_{0t}),$$

where $p_{i,t}$ is the probability that the random variable w_t realizes as $w_{[t]}$, $i \in \{1, 2, \dots, n_t\}$, as specified in §5.

Now (10) can be rewritten as

$$f_t(\mathbf{x}_t) = \min_{\tilde{\mathbf{a}}_t \in \tilde{\mathcal{A}}_t(\mathbf{x}_t)} G_t(\mathbf{x}_t, \tilde{\mathbf{a}}_t) + \sum_{w_t} f_{t+1}(\mathbb{T}(\mathbf{x}_t, \tilde{\mathbf{a}}_t, w_t)) \tag{11}$$

where $\mathbb{T}(\mathbf{x}_t, \tilde{\mathbf{a}}_t, w_t) = (x_{+t}^1, x_{+t}^2, \dots, x_{+t}^{n_t}, 0) - a_t \mathbf{1}$ if $w_t = w_{[t]}$ for $i \in \{1, 2, \dots, n_t\}$.

7.2. Application of the Approximation Algorithm to INV-LS

To establish the applicability of our developments in §§4 and 5, we show that the results of Theorems 1 and 2, which characterize the running time and accuracy of Algorithm 2, hold for INV-LS. To this end, Theorem 3 below verifies that INV-LS satisfies Assumptions 1, 2, and 4, which were used to prove Theorems 1 and 2. Although Assumption 3 is not directly satisfied, we will discuss a simple modification in Algorithm 2 to make it applicable to INV-LS and ensure that the results of Theorems 1 and 2 hold.

THEOREM 3. *The DP formulation (11) of INV-LS satisfies Assumptions 1, 2, and 4.*

PROOF. First, we will show that Assumption 1 is satisfied. That is, for all t , $f_t(\mathbf{x}_t)$ and $G_t(\mathbf{x}_t, \tilde{\mathbf{a}}_t)$ are L^b -convex and the set $\tilde{\mathcal{A}}_t = \{(\mathbf{x}_t, \tilde{\mathbf{a}}_t) : \tilde{\mathbf{a}}_t \in \tilde{\mathcal{A}}_t(\mathbf{x}_t)\}$ is L^b -convex. Note that Zipkin (2008b) has already shown that when demand, states, and actions are all integer-valued, the corresponding cost-to-go function f_t is L^b -convex; see (Zipkin 2008b, p. 940). Also, since $G_t(\mathbf{x}_t, \tilde{\mathbf{a}}_t)$ is linear, it is also L^b -convex. From Definition 9, it is easy to see that $\tilde{\mathcal{A}}_t$ is L^b -convex. We conclude that INV-LS satisfies Assumption 1.

Second, consider Assumption 2; i.e., the transformation \mathbb{T} satisfies the following condition: if $f(\mathbf{x}_t)$ is L^b -convex, then $f(\mathbb{T}(\mathbf{x}_t, \tilde{\mathbf{a}}_t, w_t))$ is also L^b -convex in $(\mathbf{x}_t, \tilde{\mathbf{a}}_t)$, for every integral w_t . For $i \in \{1, 2, \dots, n_t\}$, if $w_t = w_{[t]}$, we have that $\mathbb{T}(\mathbf{x}_t, \tilde{\mathbf{a}}_t, w_t) = (x_{+t}^i, x_{2t}, \dots, x_{L-1,t}, 0) - a_t \mathbf{1}$. Therefore, it is sufficient to prove the following claim. The proof of this claim is essentially due to Zipkin (2008b), with a few additional clarifications.

CLAIM 2. *If $f(\mathbf{x}_t)$ is L^b -convex, then for every $i \in \{1, 2, \dots, n_t\}$, $f((x_{+t}^i, x_{2t}, \dots, x_{L-1,t}, 0) - a_t \mathbf{1})$ is L^b -convex in $(\mathbf{x}_t, \tilde{\mathbf{a}}_t)$.*

PROOF OF CLAIM 2. Since f is L^b -convex, by Definition 8, we know that the function $\psi(x_{+t}^i, x_{2t}, x_{3t}, \dots, x_{L-1,t}, w_t, u) = f((x_{+t}^i, x_{2t}, x_{3t}, \dots, x_{L-1,t}, 0) - a_t \mathbf{1} + u)$ (u is a dummy variable) is submodular in $(x_{+t}^i, x_{2t}, x_{3t}, \dots, x_{L-1,t}, w_t)$. Thus, $f((x_{+t}^i, x_{2t}, x_{3t}, \dots, x_{L-1,t}, 0) - a_t \mathbf{1})$ is submodular in $(\mathbf{x}_t, \tilde{\mathbf{a}}_t, w_t)$. Consequently, by Lemma 1 of Zipkin (2008b), we know that $f((x_{+t}^i, x_{2t}, \dots, x_{L-1,t}, 0) - a_t \mathbf{1})$ is L^b -convex in $(\mathbf{x}_t, \tilde{\mathbf{a}}_t, w_t)$. By restricting $u = 0$, from item 5 of Theorem 6.19 of Murota (2007), we conclude that $f((x_{+t}^i, x_{2t}, \dots, x_{L-1,t}, 0) - a_t \mathbf{1})$ is L^b -convex in $(\mathbf{x}_t, \tilde{\mathbf{a}}_t)$. \square

Finally, consider the closed-form expression of G_t :

$$G_t(\mathbf{x}_t, \tilde{\mathbf{a}}_t) = -ca_t + \sum_{i=1}^{n_t} p_{i,t} h(x_{+t}^i - x_{1t}) + p(w_{[t]} + x_{+t}^i - x_{0t}).$$

Since the parameters are either nonnegative integers or rational numbers, we conclude that $G_t(\mathbf{x}_t, \tilde{\mathbf{a}}_t)$ is a rational number. Assumption 4 is satisfied after proper scaling of the DP recursion (11). \square

In Assumption 3 of §5, we have assumed, for ease of exposition, that (i) the dimensionality of $\mathbf{a}_{t|}$ is fixed and (ii) $\mathbf{x}_{t|}$ belongs to some hypercube $L_{t|} \times \mathcal{U}_{t|}^{k_{t|}}$. However, for INV-LS, these assumptions are too restrictive. Specifically, (i) observe that the dimensionality of $\tilde{\mathbf{a}}_{t|}$ in DP (11) depends on the number of different realizations of $w_{t|}$ which is part of the input. Since this dependence is only a linear dependence, it does not affect the running time of our approximation algorithm. (ii) The feasible set of $\mathbf{x}_{t|}$ in (11) is not a rectangular set. We now discuss simple modifications to Algorithm 2 so that it can be applied to INV-LS.

Recall that w_t^* denotes the maximum possible demand over all periods. We assume that the starting vector $\mathbf{v}_1 = (v_{10}, v_{11}, \dots, v_{1, L-1}) \in \mathbb{Z}^L$ satisfies

$$\begin{aligned} 0 & v_{10} & 2^{\lceil \log(w_t^*) \rceil} \\ 0 & v_{11} & 2^{\lceil \log(w_t^*) \rceil} \\ & \vdots & \\ 0 & v_{1, L-1} & 2^{\lceil \log(w_t^*) \rceil}. \end{aligned}$$

Recall that the order quantity $q_{t|} = \lceil a_{t|} / w_t^* \rceil$. This, together with the above assumption on \mathbf{v}_1 , implies that the untransformed state vector $\mathbf{v}_{t|} = (y_{t|}, h_{t+1-L}, \dots, h_{t-1}) \in \mathbb{Z}^L$ satisfies the following:

$$\begin{aligned} 0 & y_{t|} & t 2^{\lceil \log(w_t^*) \rceil} \\ 0 & q_{t+1-L} & 2^{\lceil \log(w_t^*) \rceil} \\ & \vdots & \\ 0 & q_{t-1} & 2^{\lceil \log(w_t^*) \rceil}. \end{aligned}$$

Denote the set of all possible $\mathbf{v}_{t|}$ that satisfy the above inequalities by $V_{t|}$.

In terms of the transformed state $\mathbf{x}_{t|} = (x_{0t|}, x_{1t|}, \dots, x_{L-1, t|}) \in \mathbb{Z}^L$, these constraints become

$$\begin{aligned} 0 & x_{0t|} - x_{1t|} & t 2^{\lceil \log(w_t^*) \rceil} \\ 0 & x_{1t|} - x_{2t|} & 2^{\lceil \log(w_t^*) \rceil} \\ & \vdots & \\ 0 & x_{L-1, t|} & 2^{\lceil \log(w_t^*) \rceil}. \end{aligned}$$

Denote the set of all possible $\mathbf{x}_{t|}$ that satisfy the above inequalities by $X_{t|}$. Note that there is a linear correspondence between $X_{t|}$ and $V_{t|}$ represented by

$$X_{t|} = \tau_{t|} \cdot V_{t|}$$

and

$$V_{t|} = \tau_{t|}^{-1} \cdot X_{t|}$$

where $\tau_{t|}$ is the following $L \times L$ matrix:

$$\begin{pmatrix} 1 & 1 & \dots & // & 1 & 1 \\ 0 & 1 & \dots & // & 1 & 1 \\ 0 & 0 & \dots & \cdot & 1 & 1 \\ 0 & \dots & // & 0 & 1 & 1 \\ 0 & 0 & \dots & // & 0 & 1 \end{pmatrix}.$$

The main obstacle in applying Algorithm 2 to INV-LS is the following: since we are working with the transformed state $\mathbf{x}_{t|}$, its domain $X_{t|}$ is no longer a rectangular set. Therefore, we cannot construct the ε -coarse cube of the cost-to-go functions directly using the definition in §4.2. This difficulty is resolved as follows. We construct the ε -coarse cube on the untransformed state (we can do this since the corresponding domain is a rectangular set) and then use the transformation $\tau_{t|}$ to transform the ε -coarse cube.

Formally, for all $t \in \{1, 2, \dots, T\}$, let $\tilde{F}_{t|}(\mathbf{v}_{t|}) = F_{t|}(\tau_{t|}^{-1} \mathbf{x}_{t|})$, where $F_{t|}$ is the integer version of $f_{t|}$ in (11). We construct the ε -coarse cube of $\tilde{F}_{t|}$ $V_{t|} \rightarrow \mathbb{R}^{+L}$ as follows. Let $E_{1t|}, E_{2t|}, \dots, E_{L2^{L-1}t|}$ denote the set of points on the edges of $V_{t|}$. Let the restriction of $\tilde{F}_{t|}$ to $E_{1t|}, E_{2t|}, \dots, E_{L2^{L-1}t|}$ be $\tilde{F}_{t|, E_{1t|}}, \tilde{F}_{t|, E_{2t|}}, \dots, \tilde{F}_{t|, E_{L2^{L-1}t|}}$, respectively. Since the function $F_{t|}$ is L^2 -convex and the transformation $\tau_{t|}$ is linear, it is easy to see that the single-dimensional functions $\tilde{F}_{t|, E_{1t|}}, \tilde{F}_{t|, E_{2t|}}, \dots, \tilde{F}_{t|, E_{L2^{L-1}t|}}$ are convex.

Next, we construct the canonical uniform-interval ε -approximation sets of the functions $\tilde{F}_{t|, E_{1t|}}, \tilde{F}_{t|, E_{2t|}}, \dots, \tilde{F}_{t|, E_{L2^{L-1}t|}}$. Let $l_{1t|}, l_{2t|}, \dots, l_{L2^{L-1}t|}$ be the lengths of the intervals in the canonical uniform-interval ε -approximation sets of the functions $\tilde{F}_{t|, E_{1t|}}, \tilde{F}_{t|, E_{2t|}}, \dots, \tilde{F}_{t|, E_{L2^{L-1}t|}}$, respectively. Let $l_t = \min\{l_{1t|}, l_{2t|}, \dots, l_{L2^{L-1}t|}\}$. Let $S_{1t|} = \{0, l_t, 2l_t, \dots, \lfloor 2^{\lceil \log(w_t^*) \rceil} \rfloor\}$ and $S_{2t|} = \{0, l_t, 2l_t, \dots, \lfloor 2^{\lceil \log(w_t^*) \rceil} \rfloor\}$. For a given ε , we refer to the set $D_{t|} = S_{1t|} \times S_{2t|}^{L-1}$ as the ε -coarse cube of $\tilde{F}_{t|}$. Then $\tau_{t|} D_{t|}$ is defined to be the ε -coarse cube for $F_{t|}(\mathbf{x}_{t|})$. With this definition of the ε -coarse cube, it is not hard to check that our approximation results in §4.2 continue to hold.

We can now apply Algorithm 2 to INV-LS. The following result notes the additive-error guarantee and the running time of Algorithm 2 for INV-LS.

COROLLARY 1. *Given $\varepsilon > 0$, under Assumptions 1 through 4, Algorithm 2 solves INV-LS to within an additive error of ε . The corresponding running time is pseudo-polynomial in the binary size of the input and polynomial in $1/\varepsilon$. Furthermore, the corresponding ε -approximate policy can also be obtained (see §6.3).*

8. Computational Experiments

In this section, we report computational results on instances of the INV-LS problem to demonstrate the practical value of Algorithm 2. Section 8.1 describes the generation of the instances and §8.2 reports the results.

8.1. The Test Bed

Our test bed consists of instances of Problem INV-LS with lead time $L = 2$. The other parameters are set as follows:

1. Without loss of generality, the ordering cost in each period is set to 0 (Janakiraman and Muckstadt 2004).

2. We consider two sets of holding and penalty costs. In the first set, the unit holding cost h is \$1 per period and the unit penalty cost p for lost sales is \$1. In the second set, these costs are \$1 and \$4, respectively.

3. The demands across the periods are independent and identically distributed. We use two distributions: a uniform distribution with support $[0, 1,024]$ and a symmetric triangular distribution with the same support.

For the four combinations of the problem parameters (namely, the two choices of holding/penalty costs and the two choices of the demand distribution), we generate 10 instances for each combination by varying the horizon length $T \in \{2, 4, \dots, 20\}$ for a total of $4 \times 10 = 40$ instances. Note that for a horizon length of T , we let the costs accrue until period $T + L = T + 2$.

To generate informed choices of the additive-error guarantee ϵ that is required by Algorithm 2, we first obtain a lower bound on the optimal expected cost starting from any initial state \mathbf{v}_1 ; i.e., $f_1(\mathbf{v}_1) \geq LB$, for all \mathbf{v}_1 . This lower bound, referred to here as LB , is calculated as follows:

Let

$$l = \min_{q \in \{0, 1, \dots, 1,024\}} \left[h \sum_{d=0}^{q-1} P(\text{Demand} = d)(q - d) + p \sum_{d=q}^{1,024} P(\text{Demand} = d)(d - q) \right]$$

where $P(\text{Demand} = d) = 1/1,025$ for the uniform distribution, and $P(\text{Demand} = d) = \min\{d, 1,024 - d\} / (512 \times 513)$ for the symmetric triangular distribution. Notice that l is a lower bound on the expected cost achievable in any period t . Therefore,

$$LB = (T + L) \times l$$

is a lower bound on the optimal total expected cost over the entire horizon, over all feasible policies.

Using the lower bound LB , we consider three choices of the additive-error guarantee ϵ as follows: $\epsilon = \epsilon' \times LB$, where $\epsilon' \in \{0.01, 0.02, 0.05\}$. Note that an additive guarantee of $\epsilon = \epsilon' \times LB$ implies a relative-error guarantee of ϵ' . Recall from the proof of Theorem 1 that the time complexity of Algorithm 2 increases with the value of the ratio s/ϵ' , where $\epsilon' = \epsilon / (2(T + L))$. In the computational results presented in the next section, we also report an estimated value of this ratio to understand its actual effect on the computational time of Algorithm 2.

8.2. Results

Algorithms 1 and 2 were implemented in C under Linux.¹⁵ The computations were carried out on a computer with an Intel i7-870 processor and 4GB of memory. For each of the 40 instances in our test bed, Tables 1 through 4 report (i) the CPU time taken by Algorithm 1, (ii) the CPU times required by Algorithm 2 and the solution quality for the three choices of the additive error, and (iii) the ratio s/ϵ' . To speed up the computation, we use a state-space reduction technique due to Zipkin (2008a) in both Algorithms 1 and 2.

We now explain the entries of each of these tables. The first column indicates the horizon length T . The second column shows the CPU time (in seconds) for Algorithm 1. In the third, fourth, and fifth columns, we report the time, cost performance, and the ratio s/ϵ' of Algorithm 2 when the guaranteed upper bound ϵ on the additive error equals $0.01 \times LB$. The third column shows the CPU time (in seconds). The fourth column shows the maximum (over all initial states) actual percentage error, i.e.,

$$\max_{\mathbf{v}_1} \frac{f_1(\mathbf{v}_1) - f_1(\mathbf{v}_1)}{f_1(\mathbf{v}_1)} \times 100 \%$$

where the maximum is taken over all possible initial states \mathbf{v}_1 in Zipkin's reduced state space. Note that the calculation of the maximum actual percentage error uses the result of Algorithm 1. The fifth column shows the ratio s/ϵ' . The remaining columns of each table are organized in a similar fashion.

Tables 1 through 4 demonstrate that Algorithm 2 is indeed of practical value. For example, in Table 1, when $T = 20$ and $\epsilon = 0.05 \times LB$, Algorithm 2 runs in three minutes, whereas Algorithm 1 takes approximately three hours. Note that as the guaranteed additive error increases from $0.01 \times LB$ to $0.05 \times LB$, the ϵ -coarse cubes used by Algorithm 2 become coarser, resulting in a significant reduction of computing time.

The effect of s/ϵ'

- (Changing ϵ) We refer the readers to Tables 1 through 4. For a fixed value of T , as the imposed additive error ϵ reduces (and therefore $\epsilon' = \epsilon / (2(T + L))$ reduces), the ratio s/ϵ' increases and therefore the time taken by Algorithm 2 increases. However, even for an additive error of $\epsilon = 0.01 \times LB$, this time is significantly below that of Algorithm 1.

- (Changing T) We refer the readers to Tables 1 through 4. For a fixed value of the relative error ϵ' (i.e., additive error of $\epsilon' \times LB$), as T increases, we find that the ratio s/ϵ' increases slightly and then remains stable. Correspondingly, we find that the ratio of the running time of Algorithm 2 over that of Algorithm 1 also remains stable.

- (Changing the unit penalty cost p): Table 5 shows the effect of increasing the unit penalty cost p which, in turn, has the overall effect of increasing the ratio s/ϵ' . In this table, we see that the ratio of the time taken by Algorithm 2 relative to that taken by Algorithm 1 indeed increases. In particular, there is a significant jump in this ratio when p changes from 16 to 32. This sudden change is because Algorithm 2 shifts to a finer discretization (i.e., the lengths of the coarse cubes reduce).

- We have also investigated other situations where the ratio s/ϵ' changes. Overall, we observe that an increase in this ratio is usually accompanied by an increase in the ratio of the time taken by Algorithm 2 to that of Algorithm 1. However, for reasonable parameter ranges, we find that Algorithm 2 continues to exhibit significant computational time savings over Algorithm 1, even at a relative error guarantee of 1%.

Table 1. Experimental results for $p]=h]=\lambda$ and uniform distribution.

T]	Algorithm 1 time	Performance of Algorithm 2								
		$\varepsilon]=0.01 \times LB]$			$\varepsilon]=0.02 \times LB]$			$\varepsilon]=0.05 \times LB]$		
		Time	% error	$s/\varepsilon'/'$	Time	% error	$s/\varepsilon'/'$	Time	% error	$s/\varepsilon'/'$
2	1,499.8	97.7	0.0343	3.0545	27.0	0.0900	1.5261	20.6	0.2077	0.6104
4	2,472.9	160.0	0.0257	3.7815	43.6	0.0825	1.8862	36.6	0.1487	0.7544
6	3,677.1	220.9	0.0235	3.9457	59.7	0.0784	1.9670	53.6	0.1268	0.7868
8	4,727.1	280.8	0.0221	3.9685	76.3	0.0761	1.9782	70.8	0.1143	0.7913
10	5,253.1	347.9	0.0215	3.9706	94.4	0.0745	1.9792	88.0	0.1064	0.7917
12	6,687.3	409.2	0.0210	3.9707	111.2	0.0735	1.9793	105.1	0.1009	0.7917
14	7,611.8	467.9	0.0207	3.9707	127.2	0.0728	1.9793	120.6	0.0967	0.7917
16	8,729.9	536.6	0.0204	3.9707	145.0	0.0722	1.9793	138.2	0.0935	0.7917
18	9,839.1	589.4	0.0202	3.9707	160.0	0.0717	1.9793	153.7	0.0909	0.7917
20	10,885.4	651.6	0.0200	3.9707	176.6	0.0713	1.9793	170.5	0.0888	0.7917

Table 2. Experimental results for $p]=h]=\lambda$ and triangular distribution.

T]	Algorithm 1 time	Performance of Algorithm 2								
		$\varepsilon]=0.01 \times LB]$			$\varepsilon]=0.02 \times LB]$			$\varepsilon]=0.05 \times LB]$		
		Time	% error	$s/\varepsilon'/'$	Time	% error	$s/\varepsilon'/'$	Time	% error	$s/\varepsilon'/'$
2	1,727.4	213.1	0.0409	4.6590	104.3	0.0534	2.3292	28.6	0.2140	0.9315
4	2,442.3	493.7	0.0268	5.8735	170.3	0.0478	2.9331	46.6	0.1787	1.1702
6	3,461.5	765.6	0.0202	6.0364	234.5	0.0447	3.0137	63.9	0.1764	1.2019
8	4,449.0	1,031.5	0.0180	6.0472	299.8	0.0425	3.0190	81.6	0.1741	1.2039
10	5,857.5	1,305.5	0.0166	6.0475	369.4	0.0409	3.0192	100.7	0.1732	1.2040
12	6,888.0	1,594.4	0.0156	6.0476	430.5	0.0397	3.0192	117.3	0.1725	1.2040
14	7,631.0	1,851.5	0.0148	6.0476	495.9	0.0388	3.0192	135.0	0.1720	1.2040
16	8,540.7	2,120.5	0.0143	6.0475	569.5	0.0381	3.0192	154.9	0.1716	1.2040
18	9,330.6	2,486.5	0.0138	6.0475	639.9	0.0376	3.0192	172.6	0.1713	1.2040
20	10,098.4	2,703.1	0.0134	6.0475	691.5	0.0371	3.0192	190.8	0.1710	1.2040

Table 3. Experimental results for $h]=\lambda, p]=\lambda$, and uniform distribution.

T]	Algorithm 1 time	Performance of Algorithm 2								
		$\varepsilon]=0.01 \times LB]$			$\varepsilon]=0.02 \times LB]$			$\varepsilon]=0.05 \times LB]$		
		Time	% error	$s/\varepsilon'/'$	Time	% error	$s/\varepsilon'/'$	Time	% error	$s/\varepsilon'/'$
2	1,538.2	95.0	0.0225	1.9470	26.0	0.0891	0.9714	25.7	0.1240	0.3885
4	1,846.2	157.5	0.0206	2.1333	43.0	0.0815	1.0632	42.7	0.1033	0.4252
6	2,779.5	218.5	0.0193	2.1816	59.5	0.0773	1.0868	59.2	0.0937	0.4347
8	3,241.0	276.5	0.0188	2.1860	75.1	0.0751	1.0890	75.0	0.0879	0.4356
10	3,375.8	337.2	0.0184	2.1863	91.6	0.0734	1.0891	91.3	0.0841	0.4356
12	3,547.7	397.6	0.0181	2.1863	107.9	0.0723	1.0891	107.6	0.0814	0.4357
14	4,746.9	458.1	0.0179	2.1863	124.3	0.0715	1.0892	124.1	0.0794	0.4356
16	7,869.5	525.1	0.0177	2.1863	140.6	0.0708	1.0891	140.5	0.0779	0.4357
18	8,126.6	579.2	0.0176	2.1863	159.0	0.0703	1.0892	158.8	0.0766	0.4356
20	8,816.7	647.2	0.0174	2.1863	175.5	0.0698	1.0891	175.3	0.0756	0.4356

9. Future Research Direction

The following question is open: Does an FPTAS exist for our problem under Assumptions 1 through 4?

Related to this question, in Halman et al. (2011), the authors prove that if f_{ij} is two-dimensional and convex in the sense of Miller (1971), then a generalization of the K -approximation technique is unlikely to succeed, in a way such that the amount of information stored is still polyno-

mial in the size of the binary input and $1/\varepsilon$. The following two comments are relevant. First, since Miller’s notion of discrete convexity is more general than L^1 -convexity, we do not know whether the above result holds for two-dimensional L^1 -convex functions. Our efforts to answer this question using the counting argument in Halman et al. (2011) were unsuccessful. Second, even if we can show that the specific technique of a K -approximation is unlikely to succeed, we cannot rule out the possibility of deriv-

Table 4. Experimental results for $h]=\lambda$, $p]=\beta$, and triangular distribution.

T]	Algorithm 1 time	Performance of Algorithm 2								
		$\varepsilon]=\theta.01 \times LB]$			$\varepsilon]=\theta.02 \times LB]$			$\varepsilon]=\theta.05 \times LB]$		
		Time	% error	s/ε'	Time	% error	s/ε'	Time	% error	s/ε'
2	1,705.1	105.2	0.0532	2.3881	103.1	0.0532	1.1940	28.5	0.2100	0.4776
4	2,726.1	169.6	0.0476	2.7774	166.8	0.0476	1.3887	46.2	0.1905	0.5538
6	3,658.1	233.7	0.0443	2.8247	233.9	0.0443	1.4124	64.0	0.1810	0.5631
8	4,437.0	299.2	0.0424	2.8274	299.3	0.0424	1.4137	81.7	0.1744	0.5636
10	5,297.1	364.9	0.0413	2.8275	360.2	0.0413	1.4138	99.8	0.1698	0.5636
12	6,742.7	430.7	0.0404	2.8275	429.4	0.0404	1.4137	117.7	0.1665	0.5636
14	7,840.2	495.9	0.0398	2.8275	495.1	0.0398	1.4138	137.3	0.1641	0.5636
16	9,049.4	561.2	0.0394	2.8275	553.5	0.0394	1.4138	154.0	0.1622	0.5636
18	9,994.8	628.4	0.0390	2.8275	618.1	0.0390	1.4138	172.3	0.1607	0.5636
20	11,186.7	694.7	0.0387	2.8275	690.8	0.0387	1.4138	189.2	0.1594	0.5636

Table 5. Experimental results for $h]=\lambda$, $T]=\beta$, $\varepsilon]=\theta.01 \times LB$, and uniform distribution.

p]	Algorithm 1 time	Performance of Algorithm 2				
		Time	% error	s/ε'	s]	Lengths of the coarse cubes for periods 3, 2, 1, 0
2	1,384.73	97.26	0.032402	2.2656	3.87	32 16 16 16
4	1,338.25	98.03	0.022459	1.9470	3.99	16 16 16 16
8	1,249.25	97.78	0.025333	3.5083	7.99	16 16 16 16
16	1,293.11	97.73	0.032359	6.6302	15.99	16 16 16 16
32	1,259.59	364.81	0.011637	12.8760	31.99	8 8 8 8
64	1,250.81	364.83	0.018686	25.3639	63.99	8 8 8 8

ing an FPTAS for our problem via a completely different technique.

For our inventory application INV-LS, if the lead time is either 0 or 1, the corresponding DP is single-dimensional. Thus, one can obtain an FPTAS in each of these two cases; see §§7.2 and 7.3 of Halman et al. (2009, 2011). When the lead time is greater than 1, the existence of an FPTAS for INV-LS is an open question.

We briefly discuss other possible research directions. Given $\varepsilon, \delta] >]0$ and oracle access to a submodular function $f:] \{0,]^n \rightarrow]0, \mathbb{K}]_{\mathbb{Z}}$, Raskhodnikova and Yaroslavtsev (2013) have shown that there exists a randomized algorithm that can output the value of $f]$ correctly (with probability $1 - \delta$) on at least $(1 - \varepsilon)2^n]$ points of the domain. The running time of the algorithm is polynomial in n , $(2k)^{O(2k \log(k/\varepsilon))}]$ and $\log(1/\delta)$. It would be interesting to investigate the applicability of their results to approximate a dynamic program. With the objective of robust optimization, Iancu et al. (2013) have characterized convexity and supermodularity based conditions on the value function of a dynamic program to identify simple optimal policies. It would also be interesting to investigate the applicability of discrete convexity in this setting.

Endnotes

1. The details about how the input is specified are in §5.
2. As we will see later in §3, the notion of discrete convexity in $\mathbb{Z}^n]$ we use (L^2 -convexity) coincides with the notion of discrete convexity used in Halman et al. (2008, 2009) when $n]=\lambda$.

3. $K]$ is an upper bound on the ratio of the values of the approximate function to those of the original function; for an approximation scheme, $K]=\lambda + \varepsilon$.

4. Note that we are using the word “extension” in a loose manner, in the sense that a continuous extension of a discrete function $f]$ does not necessarily coincide with $f]$ on all of its domain points.

5. The fact that unconstrained minimization preserves L^2 -convexity is established in Theorem 6.19, item (6), of Murota (2007). However, the result extends to the case of constrained minimization as long as the constraint set is L^2 -convex. This lemma and the proof idea were suggested to the authors by Professor Kazuo Murota.

6. We are thankful to Professor Kazuo Murota for suggesting this proof.

7. Since Murota (2007) is in Japanese, we also provide the English translations of Theorems 7.1(1) and 7.3(1):

“Theorem 7.1(1):]The discretization (7.1)] of an L^2 -convex function in continuous variables is an L^2 -convex function in discrete variables.

Theorem 7.3(1):]The elongation (7.2)] of an L^2 -convex function in discrete variables is an L^2 -convex convex function in continuous variables. (Elongation means the convex extension.)”

These English translations are due to Professor Kazuo Murota.

8. To be precise, we should attach the sub/superscript $\varepsilon]$ to \check{f} and $H]$ in Definition 16. We avoid using this sub/superscript to keep the notation simple.

9. Assumption 3 is only for notational simplicity. This assumption can be replaced with the weaker requirement that $\mathbf{x}_{i]} \in] \prod_{i=1}^{k_1} L_{1, \check{f}}] \cup]_{1, \check{f}}]^{k_2} L_{2, \check{f}}] \cup]_{2, \check{f}}]^{k_3} L_{3, \check{f}}] \cup]_{3, \check{f}}]$.

10. Note that for period $T]$ $f_{T]}$ is constructed using $f_{T]}$ directly; for period $T]-\lambda$ (and other periods), we use $\check{f}_{T-1]}$ to construct $\check{f}_{T-1]}$.

11. In our case, by Assumption 1 the constraint set is assumed to be L^3 -convex.
12. To clarify, \check{F}_i is not necessarily integer-valued. This is due to the linear interpolation step used in taking the Lovász Extension of the function obtained by restricting F_i to the ε -coarse cube. This fact will be considered when calculating the time complexity of our algorithm.
13. From Halman et al. (2009), given a function $f: L, \mathcal{U} \rightarrow \mathbb{R}^+$ and a subset S of L, \mathcal{U} , we refer to the continuous function obtained by making f linear between successive points in S as the piecewise-linear extension of f induced by S .
14. The actual procedure in Halman et al. (2009) involves additional steps; we only capture the basic idea here.
15. The source code, as well as instructions for carrying out the experiments, can be downloaded from <http://www.utdallas.edu/~wei.chen/dp.tar.bz2>.

Acknowledgments

The authors thank Nir Halman, Retsef Levi, and Chung-Lun Li for answering their questions regarding Halman et al. (2008, 2009, 2011) and Levi et al. (2008b). They are also indebted to Kazuo Murota for answering their numerous questions, suggesting the proofs of some lemmas, and sending them a copy of his book (Murota 2007). The authors thank the associate editor and three anonymous referees for their valuable suggestions, which have significantly improved the paper.

References

- Ausiello G, Crescenzi P, Gambosi G, Kann V, Marchetti-Spaccamela A, Protasi M (1999) *Complexity and Approximation: Combinatorial Optimization Problems and Their Approximability Properties* (Springer-Verlag, Berlin).
- Falk JE, Hoffman KR (1976) A successive underestimation method for concave minimization problems. *Math. Oper. Res.* 1(3):251–259.
- Favati P, Tardella F (1990) Convexity in nonlinear integer programming. *Ricerca Operativa* 53:3–44.
- Fujishige S, Murota K (2000) Notes on L/M-convex functions and the separation theorems. *Math. Programming* 88(1):129–146.
- Halman N, Nannicini G, Orlin JB (2013) A computationally efficient FPTAS for convex stochastic dynamic programs. *Eur. Sympos. Algorithms* (Springer-Verlag, Berlin), 577–588.
- Halman N, Klabjan D, Li C-L, Orlin JB, Simchi-Levi D (2008) Fully polynomial time approximation schemes for stochastic dynamic programs. *Proc. ACM-SIAM Sympos. Discrete Algorithms* (Society for Industrial and Applied Mathematics, Philadelphia), 700–709.
- Halman N, Klabjan D, Li C-L, Orlin JB, Simchi-Levi D (2011) Fully polynomial time approximation schemes for stochastic dynamic programs. Technical Report 3918, <http://optimization-online.org>.
- Halman N, Klabjan D, Mostagir M, Orlin JB, Simchi-Levi D (2009) A fully polynomial-time approximation scheme for single-item stochastic inventory control with discrete demand. *Math. Oper. Res.* 34(3):674–685.
- Iancu DA, Sharma M, Sviridenko M (2013) Supermodularity and affine policies in dynamic robust optimization. *Oper. Res.* 61(4):941–956.
- Janakiraman G, Muckstadt JA (2004) Inventory control in directed networks: A note on linear costs. *Oper. Res.* 52(3):491–495.
- Levi R, Janakiraman G, Nagarajan M (2008a) A 2-approximation algorithm for stochastic inventory control models with lost sales. *Math. Oper. Res.* 33(2):351–374.
- Levi R, Lodi A, Sviridenko M (2008b) Approximation algorithms for the capacitated multi-item lot-sizing problem via flow-cover inequalities. *Math. Oper. Res.* 33(2):461–474.
- Levi R, Pál M, Roundy R, Shmoys DB (2007) Approximation algorithms for stochastic inventory control models. *Math. Oper. Res.* 32(2):284–302.
- Levi R, Roundy R, Shmoys DB, Truong VA (2008c) Approximation algorithms for capacitated stochastic inventory control models. *Oper. Res.* 56(5):1184–1199.
- Lovász L (1983) Submodular functions and convexity. Bachem A, Grötschel M, Korte B, eds. *Mathematical Programming—The State of the Art* (Springer-Verlag, Berlin), 235–257.
- Miller BL (1971) On minimizing nonseparable functions defined on the integers with an inventory application. *SIAM J. Appl. Math.* 21(1):166–185.
- Morton TE (1969) Bounds on the solution of the lagged optimal inventory equation with no demand backlogging and proportional costs. *SIAM Rev.* 11(4):572–596.
- Murota K (2003) *Discrete Convex Analysis*, SIAM Monographs on Discrete Mathematics and Applications (Society for Industrial and Applied Mathematics, Philadelphia).
- Murota K (2007) *Primer of Discrete Convex Analysis—Discrete versus Continuous Optimization* (Kyoritsu Publishing Company, Tokyo).
- Murota K, Shioura A (2000) Extension of M -convexity and L -convexity to polyhedral convex functions. *Adv. Appl. Math.* 25(4):352–427.
- Powell WB (2007) *Approximate Dynamic Programming: Solving the Curses of Dimensionality*, Wiley Series in Probability and Statistics (Wiley-Interscience, Hoboken, NJ).
- Raskhodnikova S, Yaroslavtsev G (2013) Learning pseudo-Boolean k -DNF and submodular functions. *Proc. ACM-SIAM Sympos. Discrete Algorithms* (Society for Industrial and Applied Mathematics, Philadelphia), 1356–1368.
- Woeginger GJ (2000) When does a dynamic programming formulation guarantee the existence of a fully polynomial time approximation scheme (FPTAS)? *INFORMS J. Comput.* 12(1):57–74.
- Zipkin PH (2008a) Old and new methods for lost-sales inventory systems. *Oper. Res.* 56(5):1256–1263.
- Zipkin PH (2008b) On the structure of lost-sales inventory models. *Oper. Res.* 56(4):937–944.

Wei Chen is a Ph.D. candidate of operations management in the Naveen Jindal School of Management at the University of Texas at Dallas. His research interests include inventory management and procurement auctions.

Milind Dawande is Ashbel Smith Professor of Operations Management in the Naveen Jindal School of Management at the University of Texas at Dallas. His research interests are in optimization theory and its applications to problems in manufacturing and service operations management.

Ganesh Janakiraman is professor of operations management in the Naveen Jindal School of Management at the University of Texas at Dallas. His research interests include stochastic inventory theory and auctions for procurement.