Toward Breaking the Curse of Dimensionality: An FPTAS for Stochastic Dynamic Programs with Multidimensional Action and Scalar State

Nir Halman (HUJI)

Based upon joint work with Giacomo Nannicini (IBM)

Other cited coauthors:

C.L. Li, M. Mostagir, D. Klabjan, J. Orlin and D. Simchi-Levi

Optimization and Discrete Geometry: Theory & Practice, Tel Aviv U. April 24th, 2018

Talk Outline

- Apx. algs. and FPTASs
- Stochastic DPs
- The 3 curses of dimensionality
- Assumptions and hardness results
- Constructive results

Relative Approximation

- OPT – a minimization problem
- $OPT(x)$ – value of problem on instance x
- A K-apx. alg. $A(x)$ (K>1) returns $A(x) \le K\ OPT(x)$
- A PTAS (Polynomial Time Approximation Scheme) is a ($1+\varepsilon$)-apx. that runs in poly($|x|$) time, e.g., $O(|x|^{3/\varepsilon})$
- An FPTAS (Fully Poly. Time Approximation Scheme) is a PTAS that runs in poly($|x|, 1/\varepsilon$) time, e.g., $O((|x|/\varepsilon)^2)$

FPTASs are considered as the "Holy Grail" among apx. algs. because one can get close to $opt(x)$ as much as one wants in poly. time

A Very Brief History of FPTASs

- **Knapsack Problem**
  - Horowitz and Sahni [1974]
  - Ibarra and Kim [1975]
- **General Frameworks**
  - Korte and Schrader [1981]
  - Woeginger [2000]
  - H., Klabjan, Li, Orlin and Simchi-Levi [2014]
- **Multi-criteria**
  - Orlin [1982]
  - Safer and Orlin [1995]
  - Papadimitriou and Yannakakis [2000]
- **Google Scholar -- ~7000 listings for FPTASs**

Talk Outline

- √ Apx. algs. and FPTASs
- Stochastic DPs
- The 3 curses of dimensionality
- Assumptions and hardness results
- Constructive results

Inventory management under uncertainty

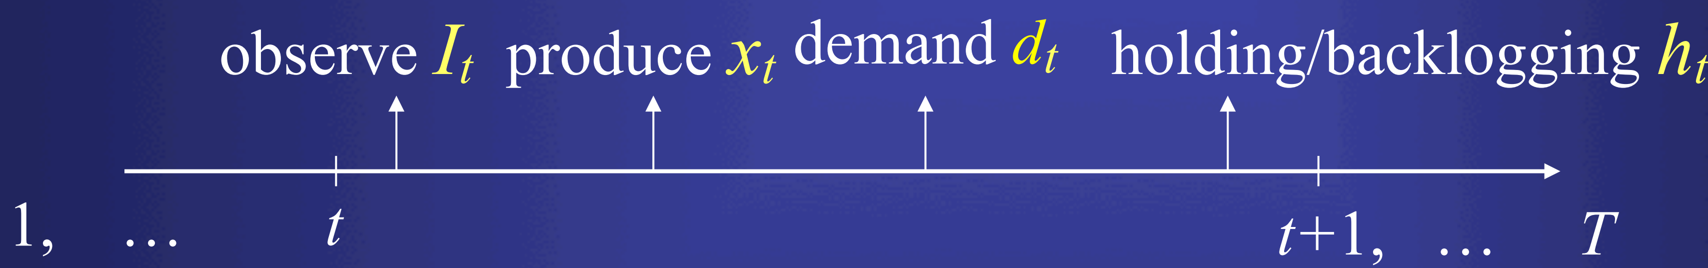

Goal:

- Minimize total costs given by: production, holding and backlogging

Stochastic DPs

Notation: T- number of periods, $I \in S_t$ - system state, $x \in A_t(I_t)$ = action, D_t random variable

System dynamics (transition): $I_{t+1} = f_t(I_t, x_t, D_t) = I_t + x_t - D_t$

Single period cost function: $g_t(I_t, x_t, D_t)$

Objective:

$$z^*(I_1) = \min_{x_1 \in A_1(I_1), \ldots, x_T \in A_T(I_T)} \sum E_{D_t}\, g_t(I_t, x_t, D_t) + g_{T+1}(I_{T+1})$$

Theorem [Bellman 57]: optimal cost is $z_1(I_1)$ in recursion:

$$z_t(I) = \min_{x \in A_t(I)} E_{D_t}[g_t(I, x, D_t) + z_{t+1}(f_t(I, x, D_t))], \quad \text{(DP)}$$

where $z_t(I)$ = total opt. cost in periods $t, \ldots, T+1$

Remark: $g_t(I_t, x_t, D_t)$ may be given as oracle (a.k.a black boxes)

Types of uncertainties/DPs

$$z_t(I_t) = \min_{x_t \in A_t(I_t)} \mathrm{E}_{D_t} \{g_t(I_t, x_t, D_t) + z_{t+1}(f_t(I_t, x_t, D_t))\}$$

Explicitly-stochastic: support and PDF of D_t are given explicitly as pairs $(d_{t,i}, \mathrm{Prob}(D_t = d_{t,i}))$

Implicitly-stochastic: supports of D_t are given explicitly but CDF of D_t are given as oracles

Data-driven: when supports of D_t are given explicitly and samples to true distributions of D_t are available

Convex DP: $g_t(\cdot,\cdot,d)$ are "convex" and $f_t(\cdot,\cdot,d)$ are "linear"

Non-decreasing DP: $g_t(I,x,d)$, $f_t(I,x,d)$ are non-decreasing in I and monotone in x

Non-increasing DP: similarly

Frameworks for Monotone/Convex DP

- Discrete explicit stochastic DP → FPTAS [HKLOS14]
- Discrete implicit stochastic DP → FPTAS [H15]
- Data-driven DP → (δ, ε)-apx. scheme [H15]
- κ-Lipschitz continuous DP → (Σ, Π)-apx. scheme [HN17]

(δ, ε)-apx. scheme: observes poly($1/\varepsilon$,log($1/\delta$)) data points, returns a feasible sol. that with prob. > $1-\delta$ is a $(1+\varepsilon)$-apx.

(Σ,Π)-apx. scheme ($\Pi=1+\varepsilon$): returns in poly($1/\varepsilon$,log($1/\Sigma$)) time a feasible sol. with cost $\leq \Sigma+\Pi OPT$

In all of these frameworks the state and action spaces are scalars, as well as the random variables

Applications

Applied to derive (first) FPTAS/approximation schemes for:

- logistics and supply chain management (lotsizing…)
- knapsack (improves upon Discrete nonlinear [Ho95], continuous nonlinear [CERSW05], stochastic [DGV04])
- machine/project scheduling
- revenue management
- economics and financial engineering
- approximate counting – integer knapsack, *m*-tuples and 2-contingency tables
- more are emerging…

Inventory management - revisited

Problem archetype:

- Consider the problem of managing a resource of non-discrete nature at a central storing facility, over a finite time horizon T
- The state I_t of the system is the amount of the resource at the central facility
- At any stage t there is an unknown arrival of resource, e.g., cargo ships, described by a random variable
- There are m locations that consume resource, but do not have storing capacity. Their demand is given by continuous r.v.s
- At any stage t the resource can be moved over a capacitated flow network, potentially with losses over the arcs
- Unused resource at the central storing facility carries over to the next time period (but is lost at the m locations)

Goal:

- Minimize total costs given by: transportation, storage and shortage at each location

The 3 Curses of Dimensionality

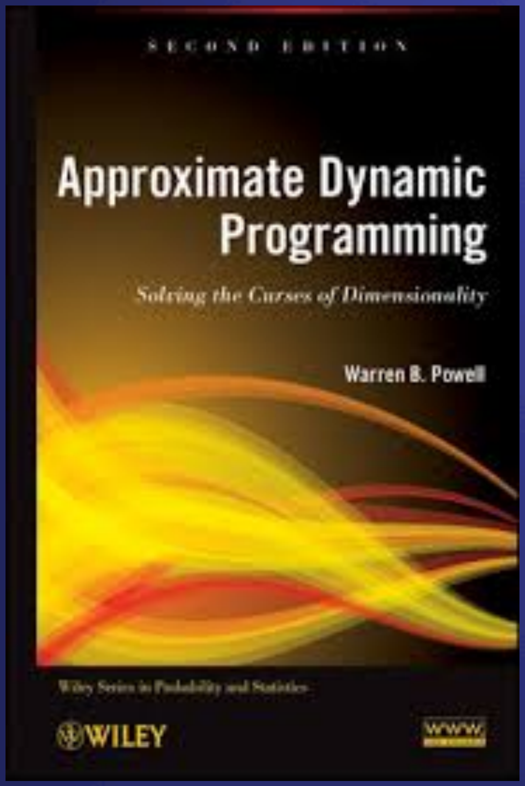

1. The state space. If $S_t=(S_{t1},\ldots,S_{tI})$ and each S_{ti} can take on L values then have L^I different states
2. The outcome space. If $D_t=(D_{t1},\ldots,D_{tJ})$ and each D_{tj} can take on M values then have M^J different outcomes
3. The action space. If $X_t=(X_{t1},\ldots,X_{tK})$ and each X_{tk} can take on N values then have N^K different actions

Discrete Convexity

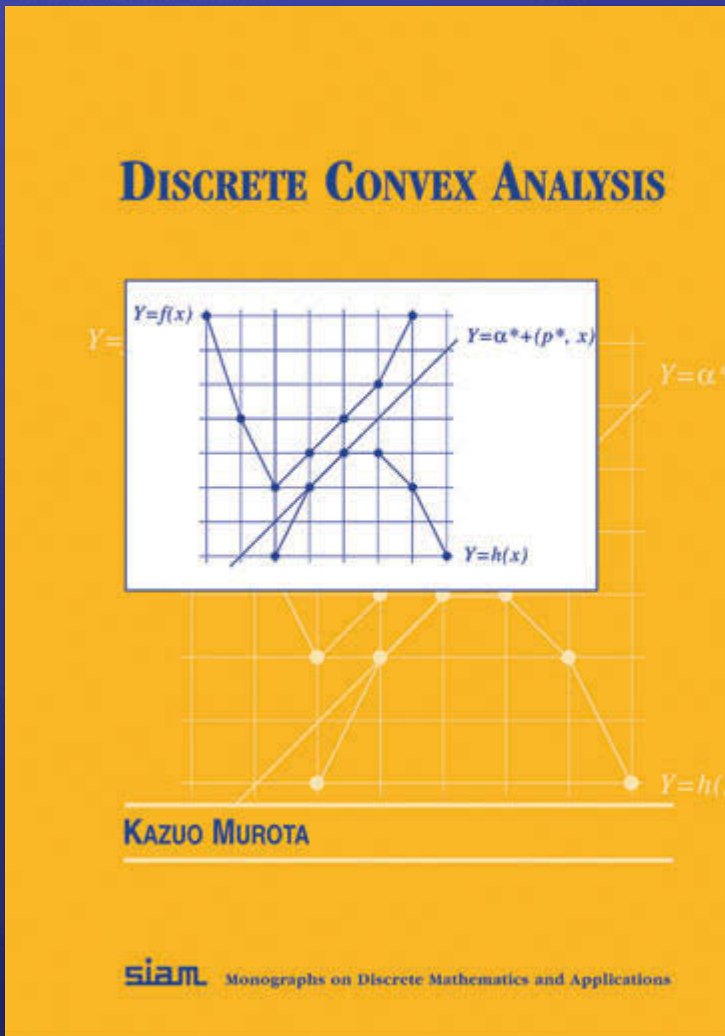

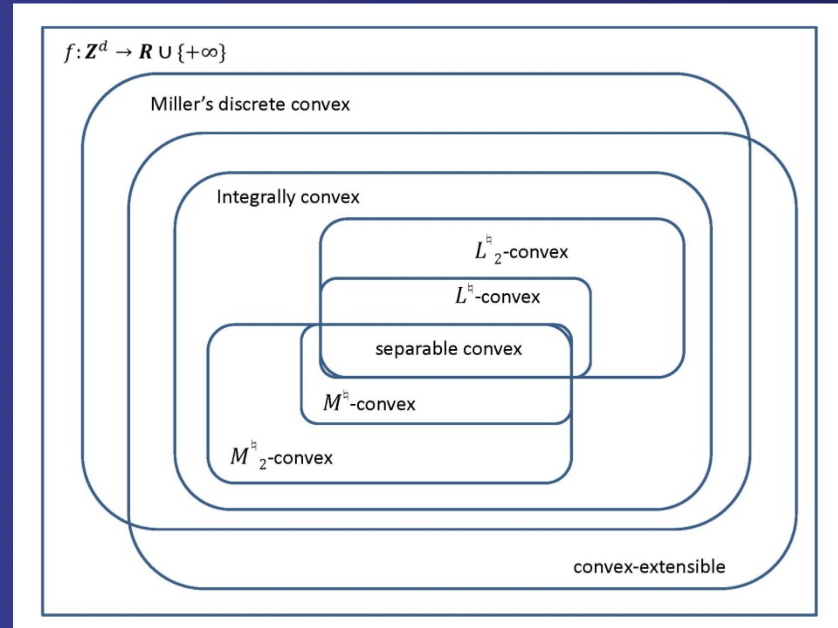

Table 1.2. *Operations for discrete convex sets and functions (f: function, S: set; ◯: Yes [cf. Theorem, Prop.], ×: No).*

	Miller's discrete convex	convex extensible	integrally convex	separable convex
$f_1 + f_2$	×	◯	×	◯
$S_1 \cap S_2$	×	◯	×	◯
$f+$ sep-conv	×	◯	◯ [3.24]	◯
$S \cap [a, b]$	◯	◯	◯	◯
$f+$ affine	×	◯	◯ [3.25]	◯
$f_1 \Box_{\mathbf{Z}} f_2$	×	×	×	◯
$S_1 + S_2$	×	×	×	◯
$f^\bullet$	×	◯	×	◯
dom f	◯	◯	◯ [3.28]	◯
arg min f	◯	◯	◯ [3.28]	◯

	$\mathrm{M}^\natural_2$-convex	$\mathrm{L}^\natural_2$-convex	$\mathrm{M}^\natural$-convex	$\mathrm{L}^\natural$-convex
$f_1 + f_2$	×	×	× ($\mathrm{M}^\natural_2$-conv)	◯ [7.11]
$S_1 \cap S_2$	×	×	× ($\mathrm{M}^\natural_2$-conv)	◯ [5.7]
$f+$ sep-conv	◯	×	◯ [6.15]	◯ [7.11]
$S \cap [a, b]$	◯	×	◯	◯
$f+$ affine	◯	◯	◯ [6.15]	◯ [7.11]
$f_1 \Box_{\mathbf{Z}} f_2$	×	×	◯ [6.15]	× ($\mathrm{L}^\natural_2$-conv)
$S_1 + S_2$	×	×	◯ [4.23]	× ($\mathrm{L}^\natural_2$-conv)
$f^\bullet$	× ($\mathrm{L}^\natural_2$-conv)	× ($\mathrm{M}^\natural_2$-conv)	× ($\mathrm{L}^\natural$-conv)	× ($\mathrm{M}^\natural$-conv)
dom f	◯ [8.29]	◯ [8.39]	◯ [6.7]	◯ [7.8]
arg min f	◯ [8.30]	◯ [8.40]	◯ [6.29]	◯ [7.16]

Tackling the Curses of Dimensionality – An overview

Thm. Can deal with multi-dim separable convex functions

Thm. Miller's discrete convex functions cannot be approximated even in R^2, monotone and (Σ, Π)-apx

Thm. Discrete convex extensible functions cannot be approximated even in R^2, monotone and (Σ, Π)-apx

Still Open. Remaining 5 classes of discrete convex functions

Thm. Can deal with multi-dim outcome spaces if appear in transition functions as affine combinations (via convolutions)

Thm. Can deal with multi-dim action spaces if single-period cost functions are linear (via parametric LP)

Assumptions (1)

Domains

The state spaces $\mathscr{S}_t$ for $t = 1, \ldots, T+1$ are intervals on the real line.

The action spaces $\mathscr{A}_t(I_t) := \{\vec{x}_t : A_t\vec{x}_t \geq \vec{b}_t + \vec{\delta}_{b_t} I_t, \vec{x}_t \geq 0\} \subset \mathbb{R}^p$ are p-dimensional polyhedral sets expressed as the feasible set of a parametric LP with p variables, where the right-hand side vector is an affine function of the parameter I_t, and where $A_t \in \mathbb{Q}^{m \times p}$, and $\vec{b}_t, \vec{\delta}_{b_t} \in \mathbb{Q}^m$ for $t = 1, \ldots, T$.

Intuitive explanation:

- We require that the state space is a (single-dimensional) interval. This is necessary because of hardness results for two-dimensional problems (discussed later)
- We require that the optimal action can be computed as the solution of a r.h.s.-parametric LP, where the r.h.s. parameter is the (scalar) state. This is consistent with our inventory problem

Assumptions (2)

Description of random events (simplified)

For every $t = 1, \ldots, T+1$, there is a vector of ℓ r.v.s $\vec{D}_t \in \mathbb{R}^\ell$. Each one of the random variables $\vec{D}_{t,i}$ satisfies one of the following

(i) $\vec{D}_{t,i}$ is a truncated continuous random variable with compact support $\vec{D}_t = [\vec{D}_{t,i}^{\min}, \vec{D}_{t,i}^{\max}] \subset \mathbb{R}$, and its CDF is Lipschitz continuous.

(ii) $\vec{D}_{t,i}$ is a discrete random variable with finite support $\vec{D}_{t,i} \subset [\vec{D}_{t,i}^{\min}, \vec{D}_{t,i}^{\max}] \subset \mathbb{R}$.

Furthermore, $\Pr(\vec{D}_{t,i} = \vec{D}_{t,i}^{\min}) > 0$ and $\Pr(\vec{D}_t = \vec{D}_{t,i}^{\max}) > 0$, and the information about the random events is given via value oracles to the CDF. All $\vec{D}_{t,i}$ are independent.

Intuitive explanation:

- We need r.v.s with bounded support and positive probability mass at the endpoints so that we can approximately compute expectations in finite time
- We require independence due to hardness results in the non-independent case [HKLOS14]

Assumptions (3)

Structure of the functions (simplified)

The terminal cost function $g_{T+1} : \mathscr{S}_{T+1} \to \mathbb{R}^+$ is strictly positive piecewise linear convex.

The transition function $f_t(I_t, \vec{x}_t, \vec{D}_t) : \mathscr{S}_t \otimes \mathscr{A}_t \times \mathscr{D}_t \to \mathscr{S}_{t+1}$ is affine.

The function $g_t : \mathscr{S}_t \otimes \mathscr{A}_t \times \mathscr{D}_t \to \mathbb{R}^+$ can be expressed as

$$g_t(I_t, \vec{x}_t, \vec{D}_t) = g_t^I(I_t, \vec{x}_t) + g_t^D(f_t^g(I_t, \vec{x}_t, \vec{D}_t)),$$

where g_t^I, g_t^D are piecewise linear convex, and f_t^g is affine.

All piecewise linear convex functions are described as the pointwise maximum of a finite number of hyperplanes.

Intuitive explanation:

- We need, in general, nonnegative Lipschitz-continuous convex functions (may be black boxes) to compute an approximation
- Nonnegativity cannot be dropped due to hardness result [HNO18]
- The affine transition function preserves convexity

Hardness of approximation in 2D

Remark. The assumption that the state space is an interval on the real line is very restrictive. Can we get rid of it?

Thm. Let $A, U \in \mathbb{Z}^+$, and let $\phi : [1,U]^2 \rightarrow [A, A+U]$ be a nondecreasing convex function described as a pointwise maximum of hyperplanes. Then, any approximation of ϕ that attains relative error less than $(A+1)/A$, or absolute error less than 1, requires $\Omega(\sqrt{U})$ space, regardless of the scheme used to represent the function

Idea of the proof.

- Construct a family of $\Omega(2^{\sqrt{U}})$ convex functions, all of which take the value A at different parts of the integer lattice in the domain, and at least A+1 everywhere else.
- Any approximation must distinguish between functions of this family

Hardness of apx. of multidimensional DP

Consider the following multidimensional generalization of problem (DP):

$$z^*(\vec{I}_1) = \min_{\pi_1,\ldots,\pi_T} \mathbb{E}\left[\sum_{t=1}^{T} g_t(\vec{I}_t, \pi_t(\vec{I}_t), \vec{D}_t) + g_{T+1}(\vec{I}_{T+1})\right], \qquad \text{(MD-DP)}$$

$$\text{subject to: } \vec{I}_{t+1} = \vec{f}_t(\vec{I}_t, \pi_t(\vec{I}_t), \vec{D}_t), \qquad \forall t = 1,\ldots,T,$$

where $\vec{I}_t \in \mathbb{R}^d$ and $\vec{f}_t$ is a vector-valued function.

Cor. There does not exist a PTAS for continuous DPs of the form (MD-DP) with cost functions discribed by a value oracle, even if $d = 2$, $T = 0$, and g_1 is nonincreasing, piecewise linear convex and bounded from below by a given positive number

An FPTAS

Algorithm APXSCHEME(ε):

$K \leftarrow \sqrt[2T]{1+\varepsilon}$, $\hat{z}_{T+1} \leftarrow g_{T+1}$
for $t := T$ **downto** 1 **do**
  $(W_{F_g}, \tilde{F}_g) \leftarrow$ COMPRESSCONVOLUTION$(\vec{D}_t, \vec{\sigma}^D, K)$.
  $(W_{F_z}, \tilde{F}_z) \leftarrow$ COMPRESSCONVOLUTION$(\vec{D}_t, \vec{\theta}^D, K)$.
  For fixed I_t, define $(W_G, \tilde{G}_t^D) \leftarrow$ COMPRESSEXPVAL$(g_t^D, (1, \sigma^I I_t, 1), \tilde{F}_g)$
  /* $\tilde{G}_t^D(\cdot)$ is an approximation of $\mathbb{E}[g_t^D(\cdot + \sigma^I I_t + \vec{\sigma}^D \cdot \vec{D}_t)]$ */
  For fixed I_t, define $(W_Z, \tilde{Z}_{t+1}) \leftarrow$ COMPRESSEXPVAL$(\hat{z}_{t+1}, (1, \theta^I I_t, 1), \tilde{F}_z)$
  /* $\tilde{Z}_{t+1}(\cdot)$ is an approximation of $\mathbb{E}[z_{t+1}(\cdot + \theta^I I_t + \vec{\theta}^D \cdot \vec{D}_t)]$ */
  $(W_z, \hat{z}_t) \leftarrow$ SCALEDCOMPRESSCONV$(\bar{z}_t, [\mathscr{S}_t^{\min}, \mathscr{S}_t^{\max}], K)$
return $\hat{z}_1$

Thm. Given stochastic DP satisfying Assumptions 1, 2, 3 APXSCHEME(ε) computes a (1+ε)-approximation of the optimal value function z_1, and runs in time polynomial in the binary input size and 1/ε

5 Building blocks

1. K-approximation functions and the Calculus of Approximation
2. K-approximation sets: a framework to compute succinct, efficient representation of (structured) functions
3. An algorithm to efficiently compute approximate convolutions of discrete/continuous r.v.s
4. An algorithm to efficiently compute approximate expected values over r.v.s described by value oracles to their CDF
5. A proof that all these computations can be done with bounded-size numbers, avoiding the exponential growth resulting from recursive LP solutions

1a. *K*-approximations of Functions

- Assume that functions are non-negative, i.e., $\varphi(x) \geq 0\ \forall x \in [0, U]$
- We say that $\varphi^*(\cdot)$ is a *K*-approximation of $\varphi(\cdot)$ if

$$\varphi(x) \leq \varphi^*(x) \leq K\varphi(x), \quad \forall x \in [0, U]$$

- We denote it as $\varphi^* \cong_K \varphi$

1b. Calculus of K-apx. Functions

Let $\alpha,\beta \geq 0$, $\varphi_i^* \cong_{K_i} \varphi_i$, and $K_1 \geq K_2$

Operation name	Operation and apx. ratio
summation	$\alpha+\beta\varphi_1^*+\varphi_2^* \cong_{K_1} \alpha+\beta\varphi_1+\varphi_2$
minimization	$\min\{\varphi_1^*, \varphi_2^*\} \cong_{K_1} \min\{\varphi_1, \varphi_2\}$
composition	$\varphi_1^*(\psi) \cong_{K_1} \varphi_1(\psi)$
approximation	$\varphi_2^* \cong_{K_1 K_2} \varphi_1$ when $\varphi_2 = \varphi_1^*$

$$\boxed{z_t(I) = \min_{x \in A_t(I)} \{g_t(x) + z_{t+1}(I+x)\}}$$

Corollary: (mimization of summation of composition)

Let g_t^*, z^*_{t+1} be K_1,K_2-apx. functions of g_t, z_{t+1} $(K_1 \geq K_2)$ then $z_t^*(I)=\min_{x\in A_t(I)}\{g_t^*(x)+z^*_{t+1}(I+x))\}$ is a K_1-apx of $z_t(I)$

2a. *K*-approximation Sets

Definition: Let φ:[0,...,U]$\rightarrow Z^+$ be monotone/convex. A K-apx. set of φ is $W \subset D$ with argminφ, 0, $U \in W$ and ratio between values of φ on each two consecutive points in W is at most K=1+ ε

Construction: φ^* is the apx. of φ induced by W:

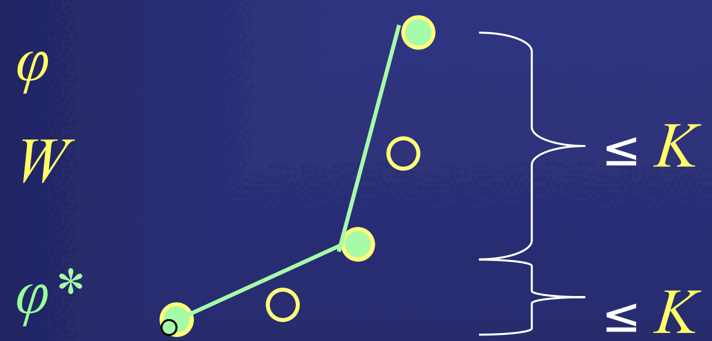

Thm: K-apx. sets of φ of size poly in the (binary) input size and 1/ε, and constructed in such time

Thm: φ^* - the apx. of φ induced by a K-apx. set of φ is a K-apx. function of φ and can be constructed in time poly. in the (binary) input size and 1/ε

2b. *K*-approximation Sets

Powerful framework:

- Introduced in [HMOS09] for functions over discrete domains
- Extended to functions over continuous domains in [HN16]

What we need to know:

- We have algorithms to construct K-approximation sets for convex or monotone nonnegative univariate Lipschitz-continuous functions ϕ in polytime
- Work under the oracle model for ϕ, i.e., do not require access to derivatives, only function values
- Since a K-approximation set for function ϕ over domain $[A,B]$ has size $O(\log_K(B-A))$, interpolation between the points in the approximation set is efficient

3. Approximation of the convolution

Thm. Suppose we are given ℓ truncated continuous r.v.s in $[A,B]$ with probability mass $\gamma > 0$ at the endpoints. Suppose each r.v. X_i admits a Lipschitz-continuous CDF. If $X_1, \ldots, X_\ell$ are independent, then for any $K = 1+\varepsilon$ we can compute a K-approximation set for the CDF of $\Sigma^{\ell}_{i=1} X_i$ in time poly. in ℓ, $1/\varepsilon$, $\log(B-A)$ and $\log 1/\gamma$

Main idea:

- Proceed by induction $\Sigma^{k}_{i=1} X_i$ for $k = 1,\ldots,\ell$.
- At each step k, construct a K-approximation set for the CDF of $\Sigma^{k}_{i=1} X_i$ using the one for $\Sigma^{k-1}_{i=1} X_i$, iterating over the points in the latter approximation set

$z_t(I_t) = \min_{x_t \in A_t(I_t)} g^k_t(I_t, x_t) + E_{D_t}\{g^D_t(\,f D \downarrow t \uparrow (I_t, x_t, D_t)) + z_{t+1}(f_t(I_t, x_t, D_t))\}$

Remark. Under our assumptions, r.v.s $\rightarrow D t \top$ only appear through affine

transformations. Since we now know how to handle $\Sigma^{\ell}_{i=1} w_i \rightarrow D t, i \top$,

4. Approximation of the expectation

$z_t(I_t) = \min_{x_t \in A_t(I_t)} g^I{}_t(I_t, x_t) + \mathrm{E}_{D_t}\{g^D{}_t(f D\downarrow t\uparrow (I_t, x_t, D_t)) + z_{t+1}(f_t(I_t, x_t, D_t))\}$

Issue:

- How do we compute expectations efficiently?
- Assume that D_t is a truncated continuous r.v. in $[A,B]$ with probability mass $\gamma > 0$ at the endpoints, and Lipschitz continuous CDF. (A similar approach works for discrete r.v.)

Our approach in a nutshell:

- Suppose we want to compute $\mathbb{E}_{D_t}(\varphi(x))$ with φ piecewise linear convex increasing.
- Compute $\tilde{F}$, a K-approximation of the CDF of D_t. This can be done in time (roughly) $O(\frac{1}{\varepsilon}\log\kappa(B-A)\log\frac{1}{\gamma})$.
- Decompose $\mathbb{E}_{D_t}(\varphi(x))$ as a finite sum involving $\tilde{F}$ for each piece of the piecewise linear function φ.

Is this it?

- Let us assume we can compute a piecewise linear convex approximation z_{t+1} of the value function at stage t +1. Can we compute it at stage t ?
- Define the following mathematical program:

$$
\left.
\begin{array}{rcl}
\bar{z}_t(I_t) := \min\limits_{x_t} \quad \mathbb{E}_{\vec{D}_t}[g_t(I_t, \vec{x}_t, \vec{D}_t) & + & \tilde{z}_{t+1}(f_t(I_t, \vec{x}_t, \vec{D}_t))] \\
A_t \vec{x}_t & \geq & \vec{b}_t + \vec{\delta}_{b_t} I_t \\
\vec{x}_t & \geq & 0,
\end{array}
\right\}
$$

where the two inequalities define the action space

- Using our techniques to apx. convolutions and expectations, we can compute a piecewise linear convex approximation of E_{D_t} [...], and the problem above can be cast as an LP.
- Thanks to the oracle model, we can let $z_t(I_t)$ be our convex function $\phi(\cdot)$, and compute a K-approximation set for it. This yields the desired z_t. But . . .

5. One more obstacle to overcome

Size of the numbers:

- The backward recursion uses z_{t+1} to build z_t.
- The construction requires solving multiple LPs, with input data generated through other LPs.
- The size of the numbers grows at most polynomially, and this is repeated $O(T)$ times, yielding numbers of size exponential in T in the worst case.

How to keep numbers "small":

- Instead of working with $\phi(x)$, work with $\alpha\phi(x/\beta)$ with large but polysize α,β.
- Show that in the transformed space, there is a fully integer approximation set.
- Transform back. Number size now depends (mostly) on α,β.

Thank you !