Curse-of-Complexity Attenuation in the Curse-of-Dimensionality-Free Method for HJB PDEs

William M. McEneaney, Ameet Deshpande and Stephane Gaubert

***Abstract*— Recently, a curse-of-dimensionality-free method was developed for solution of Hamilton-Jacobi-Bellman partial differential equations (HJB PDEs) for nonlinear control problems, using semiconvex duality and max-plus analysis. The curse-of-dimensionality-free method may be applied to HJB PDEs where the Hamiltonian is given as (or well-approximated by) a pointwise maximum of quadratic forms. Such HJB PDEs also arise in certain switched linear systems. The method constructs the correct solution of an HJB PDE from a max-plus linear combination of quadratics. The method completely avoids the curse-of-dimensionality, and is subject to cubic computational growth as a function of space dimension. However, it is subject to a curse-of-complexity. In particular, the number of quadratics in the approximation grows exponentially with the number of iterations. Efficacy of such a method depends on the pruning of quadratics to keep the complexity growth at a reasonable level. Here we apply a pruning algorithm based on semidefinite programming. Computational speeds are exceptional, with an example HJB PDE in six-dimensional Euclidean space solved to the indicated quality in approximately 30 minutes on a typical desktop machine.**

I. Introduction

Dynamic programming is an extremely robust tool for solving nonlinear optimal control problems. In the case of deterministic optimal control, or in the case of deterministic games where one player's feedback is prespecified, the dynamic programming equation reduces to a Hamilton-Jacobi-Bellman (HJB) PDE. The difficulty is that one must solve the HJB PDE.

Various approaches have been taken to solving the HJB PDE. The most common methods are grid-based methods (c.f., [6], [7]). Although highly refined at this point, these methods still suffer from the curse-of-dimensionality, as the number of grid points and computations grow exponentially with the space dimension. However, in recent years, entirely new classes of numerical methods for HJB PDEs have emerged (c.f., [1], [2], [12], [15], [8]). These methods exploit the max-plus linearity of the associated semigroup.

In the previous work of the first author [12], [14], a new method based on above semigroup linearity was proposed for certain nonlinear HJB PDEs, and this method was free from the curse-of-dimensionality. In fact, the computational growth in state-space dimension is on the order of n^3. However, there is exponential computational growth in a certain measure of complexity of the Hamiltonian. Under this measure, the minimal complexity Hamiltonian is the linear/quadratic Hamiltonian – corresponding to solution by a Riccati equation. If the Hamiltonian is given as a pointwise maximum of M linear/quadratic Hamiltonians, then one could say the complexity of the Hamiltonian is M. Such PDEs can also arise in switched linear systems.

The algorithm transforms the original problem to its max-plus dual form, where the dual of the value function is expressed as a max-plus sum, i.e., a pointwise maximum, of certain quadratic functions. An infinite time-horizon problem is considered, and as such, the value function is approximated by iterating a finite-horizon semigroup until a large enough propagation horizon is reached. With the curse-of-dimensionality-free method, this finite-horizon semigroup is approximated by a semigroup whose semiconvex dual is represented as a finite number of quadratic forms. The dual of the approximate value at each iteration is stored as a set of quadratic functions. Acting on this dual with the above dual semigroup leads to a new approximation, where the number of quadratics grows by a fixed factor (the number of quadratics representing the approximate finite-horizon dual semigroup) at each iteration. This is the curse-of-complexity. To attenuate this computational growth, we develop a pruning method based on semidefinite programming (SDP).

II. Problem statement and assumptions

The HJB PDEs we consider arise in infinite-horizon nonlinear optimal control problems, and their Hamiltonians are given as (or well-approximated by) pointwise maxima of linear-quadratic functions. Note that pointwise maxima of quadratic forms can approximate, arbitrarily closely, any semiconvex function. More specifically, we consider

$$0 = -\widetilde{H}(x, \nabla V) = -\max_{m\in\{1,2,\ldots,M\}} \{H^m(x, \nabla V)\} \quad (1)$$

$$V(0) = 0 \quad (2)$$

(i.e., with boundary condition $V = 0$ at the origin) where each of the constituent Hamiltonians has the form

$$H^m(x,p) = \tfrac{1}{2}x^T D^m x + \tfrac{1}{2}p^T \Sigma^m p + (A^m x)^T p + (l_1^m)^T x + (l_2^m)^T p + \alpha^m, \quad (3)$$

where D^m, Σ^m are $n \times n$ symmetric matrices, $l_1^m, l_2^m \in I\!R^n$ and $\alpha^m \in I\!R$.

Research partially supported by NSF grant DMS-0307229, AFOSR grant FA9550-06-1-0238 and joint RFBR-CNRS grant 05-01-02807.

William McEneaney is with the Dept. of Mech. and Aero. Eng., University of California San Diego, San Diego, CA 92093-0411, USA. wmcenean@ucsd.edu

Ameet deshpande is with the Dept. of Mech. and Aero. Eng., University of California San Diego, San Diego, CA 92093-0411, USA. adeshpan@ucsd.edu

Stephane Gaubert is with INRIA, Domaine de Voluceau 78153 LeChesnay Cédex, France. Stephane.Gaubert@inria.fr

Hamiltonian $\widetilde{H}$ is associated with an optimal control problem for switched linear systems. Let $\mathcal{M} \doteq \{1, 2, \dots M\}$. The corresponding value function is

$$\begin{aligned}\widetilde{V}(x) &= \sup_{w\in\mathcal{W}} \sup_{\mu\in\mathcal{D}_\infty} \sup_{T<\infty} \widetilde{J}(x, T; w, \mu) \\ &= \sup_{w\in\mathcal{W}} \sup_{\mu\in\mathcal{D}_\infty} \sup_{T<\infty} \int_0^T L^{\mu_t}(\xi_t) - \tfrac{\gamma^2}{2}|w_t|^2\, dt \quad (4)\end{aligned}$$

where

$$\begin{aligned}&L^{\mu_t}(x) = \tfrac{1}{2}x^T D^{\mu_t} x + (l_1^{\mu_t})^T x + \alpha^{\mu_t},\\ &\mathcal{D}_\infty = \{\mu : [0,\infty) \to \mathcal{M} : \text{ measurable}\},\\ &\mathcal{W} \doteq L_2^{loc}([0,\infty); \mathbb{R}^k),\end{aligned}$$

and the state dynamics are given by

$$\dot{\xi} = A^{\mu_t}\xi + l_2^{\mu_t} + \sigma^{\mu_t} w_t, \quad \xi_0 = x \qquad (5)$$

where σ^m and γ are such that $\Sigma^m = \frac{1}{\gamma^2}\sigma^m(\sigma^m)^T$. Here μ_t is a switching control which appears in addition to the control w.

To motivate the assumptions for this rather general problem class, we consider $\widetilde{H}$ as being constructed so as to resemble some given nonlinear control problem which has a (finite) solution. That is, we think of $\widetilde{H}$ as being chosen to resemble some other Hamiltonian, which may correspond to the originating object of interest. In particular, we suppose that problem

$$0 = -\widetilde{\widetilde{H}}(x, \nabla V), \qquad V(0) = 0 \qquad (6)$$

has finite value, and that we are choosing $\widetilde{H}$ to approximate $\widetilde{\widetilde{H}}$. Let $Q_K = \{\phi : \mathbb{R}^n \to \mathbb{R} \,|\, \phi \text{ is semiconvex, and } 0 \le \phi(x) \le (K/2)|x|^2 \ \forall x \in \mathbb{R}^n\}$. We may take Q_K as the domain of the semigroup. We make the following block of assumptions.

Assume there exists unique viscosity solution, $\widetilde{\widetilde{V}}$, to (6) in Q_K for some $K \in (0,\infty)$.
Assume that $\widetilde{H}(x,p) = \max_{m\in\mathcal{M}} H^m(x,p) \le \widetilde{\widetilde{H}}(x,p)$ for all $x, p \in \mathbb{R}^n$.
Assume $H^1(x,p)$ has coefficients satisfying the following: $l_1^1 = l_2^1 = 0$; $\alpha^1 = 0$; there exists $c_A \in (0,\infty)$ such that $x'A^1x \le -c_A|x|^2\ \forall x \in \mathbb{R}^n$; D^1 is positive definite, symmetric; and $\gamma^2/|\sigma^1|^2 > c_D/c_A^2$, where c_D is such that $x'D^1x \le c_D|x|^2$ $\forall x \in \mathbb{R}^n$. $\qquad (A.m)$
Assume that system (5) is controllable in the sense that given $x, y \in \mathbb{R}^n$ and $T > 0$, there exist processes $w \in \mathcal{W}$ and μ measurable with range in $\mathcal{M}$, such that $\xi_T = y$ when $\xi_0 = x$ and one applies controls w, μ

Note that the last of these assumptions, the controllability assumption, is satisfied if there exists at least one $m \in \mathcal{M}$ such that $\sigma^m(\sigma^m)^T$ (which is $n \times n$) has n positive eigenvalues.

Assume there exist $c_1, c_2 < \infty$ such that for any ε–optimal pair, $\mu^\varepsilon, w^\varepsilon$ for the $\widetilde{H}$ problem, one has

$$\|w^\varepsilon\|^2_{L_2[0,T]} \le c_1 + c_2|x|^2 \qquad (A.w)$$

for all $\varepsilon \in (0,1]$, all $T < \infty$ and all $x \in \mathbb{R}^n$.

Note that the behavior specified in $(A.w)$ is proved in the purely quadratic case (c.f., [12]) under reasonable assumptions on the constituent-Hamiltonian matrices, but in this more general case, we assume it instead. Lastly, we make the following assumption.

Assume there exist $\underline{T}, c_3 \in (0,\infty)$ such that for all $x \in \mathbb{R}^n$, all $\varepsilon \in (0,1]$, and all $\mu^\varepsilon, w^\varepsilon$ which are ε–optimal for $\widetilde{V}$ (i.e., such that $\widetilde{J}(x, \mu^\varepsilon, w^\varepsilon) \ge \widetilde{V}(x) - \varepsilon$), one has

$$\int_0^T L^{\mu_t^\varepsilon}(\xi_t^\varepsilon)\, dt \ge c_3 \int_0^T |\xi_t^\varepsilon|^2\, dt \qquad \forall\, T \ge \underline{T} \qquad (A.\xi)$$

where $\dot{\xi}_t^\varepsilon = A^{\mu_t^\varepsilon}\xi_t^\varepsilon + l_2^{\mu_t^\varepsilon} + \sigma^{\mu_t^\varepsilon} w_t^\varepsilon$, $\xi_0^\varepsilon = x$.

Note that these last two assumptions might be difficult to verify. Easily verifiable assumptions appear in [12], [14], but these generate a significantly smaller class of systems than those for which these methods apply.

Now, define the operator

$$\widetilde{S}_T[\phi] = \sup_{w\in\mathcal{W}} \sup_{\mu\in\mathcal{D}_T} \int_0^T L^{\mu_t}(\xi_t) - \tfrac{\gamma^2}{2}|w_t|^2\, dt + \phi(\xi_T)$$

where $\mathcal{D}_T = \{\mu : [0,T) \to \mathcal{M} : \text{ measurable}\}$. Under the above assumptions, a viscosity solution, $\widetilde{V}$ of (1),(2) exists, satisfies $0 \le \widetilde{V} \le \widetilde{\widetilde{V}}$ and is given by $\widetilde{V} = \lim_{T\to\infty} \widetilde{S}_T[V_0]$ for any $V_0 \in Q_K$ such that $0 \le V_0 \le \widetilde{V}$, [13], [14].

In the max-plus algebra, addition and multiplication are defined as $a \oplus b = \max\{a, b\}$ and $a \otimes b = a + b$, respectively. It is well known that $\widetilde{S}_T$ forms a max-plus linear semigroup.

III. Curse-of-dimensionality-free algorithm

The key steps in the curse-of-dimensionality-free algorithm developed in [14] are given below. Since we are interested in understanding how the curse-of-complexity arises in this algorithm, we shall sidestep the theoretical foundations which are well covered in [14], [12], and focus on the algorithmic flow.

A. Approximate propagation

Define the consituent-Hamiltonian semigroup operators as

$$S_\tau^m[\phi] = \sup_{w\in\mathcal{W}} \int_0^\tau L^m(\xi_t) - \tfrac{\gamma^2}{2}|w_t|^2\, dt + \phi(\xi_\tau).$$

Importantly, propagation of a quadratic ϕ by an S_τ^m operator can be reduced to solution of a differential Riccati equation. Define the time-indexed operators

$$\bar{S}_\tau[\phi](x) = \max_{m\in\mathcal{M}} S_\tau^m[\phi](x) = \bigoplus_{m\in\mathcal{M}} S_\tau^m[\phi](x).$$

Fix any $T < \infty$. Under the above assumptions, we have (c.f., [13])

$$\lim_{N\to\infty} \left\{\overline{S}_{T/N}\right\}^N [\phi] = \widetilde{S}_T[\phi]$$

where the superscipt N represents repeated application of the operator, N times.

B. Duals of the semigroup operators

This algorithm uses the concept of *semiconvex dual* (c.f., [12]). For a function ϕ which is uniformly semiconvex with constant c, the semiconvex dual, a, is given by

$$a(z) = -\max_{x\in\mathbb{R}^N} \psi(x,z) - \phi(x), \tag{7}$$

and the dual relationship is given by

$$\phi(x) = \max_{z\in\mathbb{R}^n} [\psi(x,z) + a(z)] \tag{8}$$

where $\psi(x,z) = (-c/2)|x-z|^2$.

We may obtain the duals of the S_τ^m operators, $\mathcal{B}_\tau^m$, and these are also max-plus linear semigroup operators. In particular, they are max-plus integral operators with kernels

$$B_\tau^m(x,z) = -\max_{y\in\mathbb{R}^N} \{\psi(y,x) - S_\tau^m[\psi(\cdot,z)](y)\}.$$

Importantly, note that as the $S_\tau^m[\psi(\cdot,z)](y)$ are quadratic functions, the B_τ^m are quadratic functions. Each of these is obtained only once, at the outset of the algorithm.

C. Dual space propagation and the curse-of-complexity

Once the kernels B_τ^m of the dual semigroup are obtained, one can begin the iteration. One may begin with an initial quadratic (in the dual space), say

$$\overline{a}^0(z) = \widehat{a}^0(z) = (1/2)(z-\widehat{z})^T\widehat{Q}(z-\widehat{z}) + \widehat{r}.$$

Given approximation, $\overline{a}^k$ at step k, one obtains the next iterate from

$$\begin{aligned}\overline{a}^{k+1}(z) &= \bigoplus_{m\in\mathcal{M}} \int_{\mathbb{R}^n}^{\oplus} B_\tau^m(z,y)\otimes\overline{a}^k(y)\,dy \\ &= \max_{m\in\mathcal{M}}\max_{y\in\mathbb{R}^n}[B_\tau^m(z,y) + \overline{a}^k(y)].\end{aligned}$$

If $\overline{a}^k$ has the form $\overline{a}^k(z) = \bigoplus_{\{m_i\}_{i=1}^k\in\mathcal{M}^k} \widehat{a}^k_{\{m_i\}_{i=1}^k}(z)$ where each $\widehat{a}_{\{m_i\}_{i=1}^k}$ is a quadratic form, then $\overline{a}^{k+1}$ takes the form

$$\overline{a}^{k+1}(z) = \bigoplus_{\{m_i\}_{i=1}^{k+1}\in\mathcal{M}^{k+1}} \widehat{a}^{k+1}_{\{m_i\}_{i=1}^{k+1}}(z)$$

where the $\widehat{a}^{k+1}_{\{m_i\}_{i=1}^{k+1}}$ are also quadratic. Consequently, the computations reduce to obtaining the coefficients of these quadratics at each step, and these computations are analytic (modulo matrix inverses). This is the reason that the computational growth in the space dimension is only cubic. However, note that the number of quadratics comprising the $\overline{a}^k$ grows by a factor of M at each iteration — hence the curse-of-complexity. It has been noted that quite typically, most of the quadratics do not contribute to the value (as they never achieve the maximum at any point, z), and may be pruned wihtout consequence. However, over-pruning, saving only a limited number of quadratics at each iteration, has proven to be an excellent computational technique. We now proceed to discuss the use of semidefinite programming as a means of pruning the constituent quadratics, the $\widehat{a}^k_{\{m_i\}_{i=1}^{k+1}}$.

Lastly, note that once one has propagated sufficiently far (say $k = K$ steps), the value function approximation is recovered from $\overline{a}^K$ via (8).

IV. Pruning Algorithms

In the above curse-of-dimensionality-free algorithm, at step k, $\overline{a}^k$ is represented as a max-plus sum of quadratics. Let us index the elements of this sum by integers $i \in \mathcal{I}_k$ (rather than by the sequences $\{m_i\}_{i=1}^k$). That is, we have

$$\overline{a}^k(z) = \bigoplus_{i\in\mathcal{I}^k} \widehat{a}_i^k(z)$$

where we let each $\widehat{a}_i^k$ be given in the form

$$\widehat{a}_i^k(z) = \widehat{a}_i(z) = z^T A_i z + 2b_i^T z + c_i$$

where we delete the superscript k for simplicity of notation here and in the sequel.

Recall that we are reducing computational cost by pruning quadratics ($\widehat{a}_i$) which do not contribute to the solution approximation (not achieving the maximum at any $z \in \mathbb{R}^n$). Consequently, we want to determine whether the p^{th} quadratic contributes to the pointwise maximum. In other words, we need to determine whether there is a region where it is greater than all other quadratics. Fix $p \in \mathcal{I}^k$. We need to check the feasibility of the following inequalities. That is, we want find z such that

$$z^T A^p z + 2z^T b^p + c^p \geq z^T A^i z + 2z^T b^i + c^i \quad \forall i \neq p. \tag{9}$$

Alternatively, we consider the problem:

$$\begin{aligned}&\text{Minimize } G(z,\nu) \doteq \nu \quad \text{subject to} \\ &z^T(A_i - A_p)z + 2z^T(b_i - b_p) + (c_i - c_p) \leq \nu \quad \forall i \neq p.\end{aligned} \tag{10}$$

Then, the minimum value of ν, $\bar{\nu}$, is the minimum amount by which the p^{th} quadratic needs to be raised in order to contribute to the max-plus sum. If $\bar{\nu} > 0$, then p^{th} quadratic does not contribute to the max-plus sum, and hence it can be pruned without consequence. If $\bar{\nu} < 0$, it implies that quadratic contributes to the max-plus sum, and moreover it is $|\bar{\nu}|$ above the pointwise maximum of all other quadratics at some point. In this case, $|\bar{\nu}|$ can serve as some measure of contribution of the p^{th} quadratic to the value function – useful for over-pruning.

A. Pairwise pruning

Before undertaking the pruning using semidefinite programming, pairwise pruning is used. This is a simple, fast and effective technique, which checks between all pairs of quadratic basis functions, and prunes those which are completely dominated by another. Let $A = A_{i_1} - A_{i_2}$, $b = b_{i_1} - b_{i_2}$, $c = c_{i_1} - c_{i_2}$, and define $Q(z) = z^T Az + 2b^T z + c$. Then Q is nonnegative everywhere if and only if the homogeneous quadratic form is nonnegative everywhere. Further, $Q(t^{-1}z) \geq 0$ for all $z \in \mathbb{R}^n$ and all $t \neq 0$ is equivalent to $\widehat{Q}(z,t) \doteq z^T Az + 2tb^T z + ct^2 \geq 0$ for all $z \in \mathbb{R}^n$ and all $t \in \mathbb{R}$. This latter statement is true if and only if

$$\begin{bmatrix} c & b^T \\ b & A \end{bmatrix} \succeq 0, \tag{11}$$

i.e., nonnegative definite. The i_2 quadratic can be pruned if (11) is true. This pairwise pruning reduces the computational

effort of the semidefinite pruning by getting rid of obviously dominated quadratics.

B. Shor's semidefinite relaxation based SDP

The problem of evaluating an individual quadratic $\widehat{a}_p(z)$ for pruning, (10), can be rephrased as below. Let $q_i(z) = \widehat{a}_i(z) - \widehat{a}_p(z)$ for all $i \neq p$. Then, $\widehat{a}_p$ can be pruned if

$$\min_{z,\nu} \{\nu : \ \nu - q_i(z) \geq 0 \ \ \forall i\} \geq 0. \tag{12}$$

Introducing the slack variables, u_i, (and letting $u = \{u_i\}$) we see that $\widehat{a}_p$ can be pruned if

$$\min_{z,\nu,u} \left\{\nu : \nu - q_i(z) - {u_i}^2 = 0 \ \ \forall i\right\} \geq 0. \tag{13}$$

Let $f_i(z,\nu,u) = \nu - q_i(z) - {u_i}^2$. Since the f_i are zero only on the constraint set, we can add their λ-weighted combination to the objective function without changing its value, where $\lambda \in I\!R^{\#\mathcal{I}^k - 1}$. That is,

$$\begin{aligned}
&\min_{z,\nu,u} \left\{\nu : \nu - q_i(z) - {u_i}^2 = 0 \ \ \forall i\right\} \\
&= \min_{z,\nu,u} \left\{ \nu - \sum_i \lambda_i f_i(z,\nu,u) : \ \ f_i(z,\nu,u) = 0 \right\} \\
&\geq \min_{z,\nu,u} \left\{ \nu - \sum_i \lambda_i f_i(z,\nu,u) \right\} \doteq \zeta(\lambda).
\end{aligned}$$

Note that in the last form, the domain for minimization is no longer the constraint set, but the whole space, i.e., $z \in I\!R^n, \nu \in I\!R, u \in I\!R^{\#\mathcal{I}^k - 1}$. Also, we see that $\zeta(\lambda)$ is a lower bound for the minimum in (13). We wish to maximize lower bound $\zeta(\lambda)$ by varying λ. Consequently, the dual problem is

$$\max_{\lambda \in I\!R^{\#\mathcal{I}^k - 1}} \zeta(\lambda). \tag{14}$$

Next, we narrow the search for the optimal λ by considering the following arguments. If $\sum \lambda_i \neq 1$, then $\forall z, u$

$$\zeta(\lambda) \leq \min_{\nu} \left(1 - \sum \lambda_i\right) \nu + \sum \lambda_i \left(q_i(z) + u_i^2\right) = -\infty.$$

Hence, for a finite lower bound, it is fruitful to restrict λ such that $\sum \lambda_i = 1$. This makes the objective independent of ν. Now, if $\lambda_j < 0$, for some j, then for all z and all $\{u_i : i \neq j\}$,

$$\zeta(\lambda) \leq \min_{u_j} \left(\sum_i \lambda_i q_i(z) + \sum_{i \neq j} \lambda_i u_i^2 + \lambda_j u_j^2 \right) = -\infty$$

Hence, we may further restrict the domain to $\lambda_i \geq 0$, $\forall i$. Thus, λ lies within the simplex $\mathcal{S}$, $\sum \lambda_i = 1$ and $\lambda_i \geq 0 \ \forall i$. Note, also, that this makes the objective function independent of u_i, since the above minimum with respect to u_i is always achieved at $u_i = 0$. Consequently,

$$\zeta(\lambda) = \min_z \sum_i \lambda_i q_i(z) \qquad \forall \lambda \in \mathcal{S}. \tag{15}$$

Thus, dual problem (14) can be reposed as

$$\begin{aligned}
\max_{\lambda} \zeta(\lambda) \ &= \max_{\lambda \in \mathcal{S}} \zeta(\lambda) = \max_{\lambda \in \mathcal{S}} \min_z \sum_i \lambda_i q_i(z) \\
&= \max_{\lambda \in \mathcal{S}, \zeta \in I\!R} \left\{ \zeta : \zeta \leq \sum_i \lambda_i q_i(z) \ \forall z \right\}. \qquad (16)
\end{aligned}$$

Suppose that quadratics $q_i(z)$ are specified by parameters $\bar{A}_i$, $\bar{b}_i$ and $\bar{c}_i$. Then, using linear superposition and result (11), the condition $\zeta \leq \sum_i \lambda_i q_i(z) \ \forall z$ from (16) can be posed as following linear matrix inequality:

$$\max_{\lambda \in \mathcal{S}, \zeta} \left\{ \zeta : \left[\begin{array}{cc} \sum_i \lambda_i \bar{c}_i - \zeta & \sum_i \lambda_i \bar{b}_i^T \\ \sum_i \lambda_i \bar{b}_i & \sum_i \lambda_i \bar{A}_i \end{array} \right] \succeq 0 \right\}. \tag{17}$$

Note that if the maximal ζ value, $\bar{\zeta}$ satisfies $\bar{\zeta} \geq 0$, then the minimal ν value in (12), $\bar{\nu}$, is greater than $\bar{\zeta}$, and thus positive. This indicates that the quadratic, $\widehat{a}_p$, needs to be raised by $\bar{\nu}$ units before it can contribute to the max-plus sum. Hence it can be pruned. If $\bar{\zeta} < 0$, then its "prunability" is not conclusive, as $\bar{\zeta}$ is a lower bound for $\bar{\nu}$, and the gap is not guaranteed to be zero. Nevertheless, it does give us a working indication of the importance of the quadratic, since $\bar{\zeta} < 0$ indicates that the p^{th} quadratic has to come down by at least $|\bar{\zeta}|$ units, before it is dominated by the convex hull of the remaining quadratics.

An additional way to develop intuition for result (17) is as follows. The above test evaluates the p^{th} quadratic $\widehat{a}_p$, which can be pruned if $\zeta = 0$ satisfies the inequality in (16). Lets now substitute $q_i(z) = \widehat{a}_i(z) - \widehat{a}_p(z)$ in (16). We see that $\widehat{a}_p$ can be pruned if for some $\lambda \in \mathcal{S}$

$$0 \leq \sum_i^{K-1} \lambda_i q_i(z) = \sum_{i \neq p} \lambda_i \widehat{a}_i(z) - \widehat{a}_p(z) \qquad \forall\, z \in I\!R^n,$$

or equivalently,

$$\left[\begin{array}{cc} c_p & b_p^T \\ b_p & A_p \end{array} \right] \preceq \sum_{i \neq p} \lambda_i \left[\begin{array}{cc} c_i & b_i^T \\ b_i & A_i \end{array} \right]. \tag{18}$$

Thus, if the convex hull of remaining quadratics intersects the semidefinite cone of quadratics greater than $\widehat{a}_p$, then $\widehat{a}_p$ can be pruned.

In our context, the substantial benefits of Shor's relaxation are that it accommodates both concave and convex quadratics, calculates an importance/quality metric and gives a clean analytical framework for analysis of pruning. One efficient way to implement the above scheme is to prune all but the quadratics lying on the periphery of the convex hull, and then to evaluate each vertex quadratic one by one by the above semidefinite program (17). A seeming drawback of this algorithm is that it gives the optimum in λ space, and not z, making it difficult to assign locally weighted importance. That is, we might choose to assign more importance to the quadratics which are vital near origin than those far from it. The next method helps us do precisely that.

C. Schur complement-based SDP

In this method, we again aim to solve (10) without resorting to dual relaxation schemes. It is based on following well-known Schur's complement lemma, which can be found in [4]

Lemma 4.1: Let

$$E = \begin{pmatrix} B & C^T \\ C & D \end{pmatrix}$$

be a symmetric matrix with $k \times k$ block B, and $l \times l$ block D. Assume that B is positive definite. Then E is positive semi-definite if and only if $D - CB^{-1}C^T$ is positive semi-definite.

Now starting form (10), if we define $\bar{A}_i = A_i - A_p$, $\bar{b}_i = b_i - b_p$ and $\bar{c}_i = c_i - c_p$, and assume that $\bar{A}_i \succ 0$, then (10) can be rephrased as:

$$\begin{aligned} &\text{Minimize } G(z,\nu) \doteq \nu \quad \text{subject to} \\ &z^T \bar{A}_i z + 2z^T \bar{b}_i + \bar{c}_i \leq \nu \ \ \forall i \neq p. \end{aligned} \tag{19}$$

By Lemma 4.1, this is equivalent to the following linear matrix inequality (LMI).

$$\begin{bmatrix} -(2z^T\bar{b}_i + \bar{c}_i) + \nu & z^T \\ z & \bar{A}_i^{-1} \end{bmatrix} \succeq 0 \ \ \forall i \neq p$$

Note that if $\bar{A}_i \succ 0$, $\forall i \neq p$, then (19) is equivalent to the LMI given by

$$\min_{z,\nu} \left\{ \nu : \begin{bmatrix} -2z^T\bar{b}_i - \bar{c}_i + \nu & z^T \\ z & \bar{A}_i^{-1} \end{bmatrix} \succeq 0 \ \ \forall i \neq p \right\}. \tag{20}$$

If the minimal ν value, $\bar{\nu}$ satisfies $\bar{\nu} \geq 0$, then the quadratic $\widehat{a}_p$ can be pruned without consequence. If it is negative, then $|\bar{\nu}|$ serves as a metric of importance of $\widehat{a}_p(z)$, as it says that at some point z, $\widehat{a}_p$ is at least $|\bar{\nu}|$ higher than all other quadratics. $\bar{\nu}$ and the z at which it occurs can be combined to create a new importance metric catered to region of interest. For example, $|\bar{\nu}|/(1+|z|^2)$ ensures tighter error bounds near origin than far from it.

However, the assumption $\bar{A}_i \succ 0 \ \forall i \neq p$ is overly restrictive, as it only enables comparison with quadratics more convex than $\widehat{a}_p$, as we simply drop the inequalities for which this does not hold. This also ensures conservative pruning. To do a tighter pruning, it is advisable to rank the quadratics in the increasing order of lowest eigenvalues, before pruning.

If $\bar{A}_i \succeq 0$ instead of If $\bar{A}_i \succ 0$, then the above LMI can be modified to use a linear transformation of $\bar{A}_i$ using its nonzero singular values.

V. Computational complexity

Since our aim is to reduce the curse-of-complexity without letting go of our freedom from the curse-of-dimensionality, it is worthwhile to discuss the computational overhead involved in these pruning methods. They are polynomial in both dimensionality and the number of quadratic functions. In particular, we retain our freedom from the curse-of-dimensionality.

A generic semi-definite program $\mathcal{P}$ is given by

$$P_0 = \min_{\eta \in \mathbb{R}^N} \left\{ c^T\eta : A_0 + \sum_{j=1}^{N} \eta_j A_j \geq 0, ||\eta||_2 \leq R \right\}$$

where the A_j are symmetric matrices with $\widetilde{M}$ diagonal blocks of size $k_i \times k_i$, $i = 1, \ldots \widetilde{M}$. We say that η^ε is an ε–optimal solution if

$$\|\eta^\varepsilon\|_2 \leq R, \ A_0 + \sum_{j=1}^{N} \eta_j^\varepsilon A_j \geq -\epsilon I, \ c^T\eta^\varepsilon \leq P_0 + \epsilon.$$

In [4], the authors derive the computational complexity of obtaining such an η^ε.

$$\begin{aligned} \mathcal{C}(\mathcal{P},\epsilon) = \ &\mathcal{O}(1)\left(1 + \sum_{1}^{\widetilde{M}} k_i\right)^{1/2} \\ &\cdot N\left(N^2 + N\sum_{1}^{\widetilde{M}} k_i^2 + \sum_{1}^{\widetilde{M}} k_i^3\right)\mathcal{D}(\mathcal{P},\epsilon) \end{aligned} \tag{21}$$

where $\mathcal{D}(\mathcal{P},\epsilon)$ depends on the specific problem data (as indicated by the $\mathcal{P}$ in the argument as well as ϵ). Using this expression, we can obtain an upper bound for the arithmetic complexity of the pruning algorithms. Assuming the worst case scenario where no quadratic gets pruned, we find the complexity of testing one quadratic for pruning as follows.

For the Schur complement pruning of (20), we have $\eta = [z^T \ \nu]^T$. Hence $N = n + 1$. Further, the number of block LMIs is $I \doteq \#\mathcal{I}_k - 1$. Thus $\widetilde{M} = I$, and $k_1 = k_2 = \ldots = k_m = n + 1$. Substituting these into (21), we obtain the complexity as

$$\begin{aligned} \mathcal{C}(\mathcal{P},\epsilon) = \ &\mathcal{O}(1)\left(1 + I(n+1)\right)^{1/2}(n+1)^3 \\ &\cdot [1 + 2(n+1)(I-1)]\,\mathcal{D}(\mathcal{P},\epsilon), \end{aligned} \tag{22}$$

and we see that the complexity grows as $n^{4.5}$ and $I^{1.5}$.

For the Shor semidefinite relaxation approach, we apply a similar analysis, and find that the complexity grows as $n^{3.5}$ and $I^{3.5}$. Note that these computational cost growth rates as a function of space dimension, n, are poorer than the cubic growth of the basic algorithm. However, they are still tremendously helpful for real-world problems, and the curse-of-dimensionality is still very far off.

VI. Importance-based over-pruning

Both the Shor's relaxation and the Schur complement pruning schemes generate importance metrics for unpruned quadratics. This is very useful in controlling the complexity growth according to an error tolerance, region of interest, and/or computational limits. Currently, we set a bound $L(k)$ for the number of quadratics which will remain after over-pruning at the k^{th} step. The algorithm keeps only the $L(k)$ quadratics, $\widehat{a}_i^k$ with the highest values according to our measure of importance. However, there does not yet exist a theory which allows us to map the importance measure of a quadratic at step k to an error bound in the approximation

at the terminal step. Further, a bound on the errors induced by over-pruning of *multiple* quadratics is obviously a necessary additional step. Regardless, application of these methods has proven extremely fruitful, as can be seen in the example below.

VII. Six-dimensional, six-Hamiltonian example

We applied the curse-of-dimensionality-free method to the following problem over $\mathbb{R}^6$ with six constituent quadratic Hamiltonians. Note that the problem was tweaked to exhibit sufficiently complex and interesting behavior, such that there is interaction amongst dimensions, and each operator is important somewhere in the domain. Hence the following data yields a reasonably rich problem.

We shall specify the matrices in terms of the following building blocks for the dynamics:

$$A_a = \begin{bmatrix} -1 & .5 \\ .1 & -1 \end{bmatrix}, \qquad A_b = A_a,$$
$$A_c = A_a, \qquad A_d = \begin{bmatrix} -1 & .5 \\ .3 & -1 \end{bmatrix},$$
$$A_e = A_a, \qquad A_f = \begin{bmatrix} -1 & .5 \\ .1 & -1 \end{bmatrix},$$

$$\Sigma_a = 0.4 \times \begin{bmatrix} 0.27 & -.01 \\ -.01 & 0.27 \end{bmatrix}, \qquad \Sigma_b = 0.4\Sigma_a,$$
$$\Sigma_c = \Sigma_a, \qquad \Sigma_d = 0.4\Sigma_a,$$
$$\Sigma_e = \Sigma_a, \qquad \Sigma_f = 0.4\Sigma_a,$$

and the following building blocks for the payoff functions:

$$D_a = \begin{bmatrix} 1.5 & .2 \\ .2 & 1.5 \end{bmatrix}, \qquad D_b = 1.4 \cdot \begin{bmatrix} 1.5 & .2 \\ .2 & 1.5 \end{bmatrix},$$
$$D_c = 1.4 \cdot \begin{bmatrix} 0.2 & 1.5 \\ 1.5 & 0.2 \end{bmatrix}, \qquad D_d = 1.2 \cdot \begin{bmatrix} 1.6 & 0 \\ 0 & 0 \end{bmatrix},$$
$$D_e = 1.1 \cdot \begin{bmatrix} 0.3 & 1.5 \\ 1.5 & 0.3 \end{bmatrix}, \qquad D_f = 1.3 \cdot \begin{bmatrix} 0 & 0 \\ 0 & 1.6 \end{bmatrix}.$$

We will use a parameter to adjust the interaction in the dynamics across the dimensions, and this will be $\gamma = -0.1$. Now we are ready to define each of the Hamiltonians. We need to specify parameters for the dynamics (A, Σ, l_2) and the payoff (D, l_1, α). For the example below, $l_1 = 0$ and $l_2 = 0$ for all the Hamiltonians. The remaining parameters are as follows.

For the first Hamiltonian, H^1, we let

$$A^1 = \begin{bmatrix} A_a & \gamma I & \gamma I \\ \gamma I & A_a & 0 \\ \gamma I & 0 & A_a \end{bmatrix}, \qquad \Sigma^1 = \begin{bmatrix} \Sigma_a & 0 & 0 \\ 0 & \Sigma_a & 0 \\ 0 & 0 & \Sigma_a \end{bmatrix}$$
$$D^1 = \begin{bmatrix} D_a & 0 & 0 \\ 0 & D_a & 0 \\ 0 & 0 & D_a \end{bmatrix}, \quad \alpha_1 = 0\,.$$

For the second Hamiltonian, H^2, we let

$$A^2 = \begin{bmatrix} A_b & \gamma I & \gamma I \\ \gamma I & A_b & 0 \\ \gamma I & 0 & A_b \end{bmatrix}, \qquad \Sigma^2 = \begin{bmatrix} \Sigma_b & 0 & 0 \\ 0 & \Sigma_b & 0 \\ 0 & 0 & \Sigma_b \end{bmatrix}$$
$$D^2 = \begin{bmatrix} D_b & 0 & 0 \\ 0 & D_b & 0 \\ 0 & 0 & D_b \end{bmatrix}, \quad \alpha_2 = -0.4\,.$$

For the third Hamiltonian, H^3, we let

$$A^3 = \begin{bmatrix} A_c & \gamma I & \gamma I \\ \gamma I & A_c & 0 \\ \gamma I & 0 & A_c \end{bmatrix}, \qquad \Sigma^3 = \begin{bmatrix} \Sigma_c & 0 & 0 \\ 0 & \Sigma_c & 0 \\ 0 & 0 & \Sigma_c \end{bmatrix}$$
$$D^3 = \begin{bmatrix} D_c & 0 & 0 \\ 0 & D_c & 0 \\ 0 & 0 & D_c \end{bmatrix}, \quad \alpha_3 = 0.$$

For the fourth Hamiltonian, H^4, we let

$$A^4 = \begin{bmatrix} A_d & \gamma I & \gamma I \\ \gamma I & A_d & 0 \\ \gamma I & 0 & A_d \end{bmatrix}, \qquad \Sigma^4 = \begin{bmatrix} \Sigma_d & 0 & 0 \\ 0 & \Sigma_d & 0 \\ 0 & 0 & \Sigma_d \end{bmatrix}$$
$$D^4 = \begin{bmatrix} D_d & 0 & 0 \\ 0 & D_d & 0 \\ 0 & 0 & D_d \end{bmatrix}, \quad \alpha_4 = -0.4\,.$$

For the fifth Hamiltonian, H^5, we let

$$A^5 = \begin{bmatrix} A_e & \gamma I & \gamma I \\ \gamma I & A_e & 0 \\ \gamma I & 0 & A_e \end{bmatrix}, \qquad \Sigma^5 = \begin{bmatrix} \Sigma_e & 0 & 0 \\ 0 & \Sigma_e & 0 \\ 0 & 0 & \Sigma_e \end{bmatrix}$$
$$D^5 = \begin{bmatrix} D_e & 0 & 0 \\ 0 & D_e & 0 \\ 0 & 0 & D_e \end{bmatrix}, \quad \alpha_5 = 0.$$

For the sixth Hamiltonian, H^6, we let

$$A^6 = \begin{bmatrix} A_f & \gamma I & \gamma I \\ \gamma I & A_f & 0 \\ \gamma I & 0 & A_f \end{bmatrix}, \qquad \Sigma^6 = \begin{bmatrix} \Sigma_f & 0 & 0 \\ 0 & \Sigma_f & 0 \\ 0 & 0 & \Sigma_f \end{bmatrix}$$
$$D^6 = \begin{bmatrix} D_f & 0 & 0 \\ 0 & D_f & 0 \\ 0 & 0 & D_f \end{bmatrix}, \quad \alpha_6 = -0.4\,.$$

For this example, we let the time-discretization step-size be $\tau = 0.2$, and propagation was carried out with the Shor's semidefinite relaxation based pruning. The overpruning threshold was set heuristically to $L(k) = 20 + 6k$. That is, a maximum of $L(k)$ quadratics, $\widehat{a}_i^k$, were retained at the k^{th} step. In this test, 25 iteration steps were carried out in 30 minutes, yielding a rather accurate solution in a compact domain in all six dimensions. This computation-time is for an Apple mac desktop, from roughly 2005. Slices of statistics for this value function along the 1-2 axes are shown in the accompanying figures. The backsubstitution error depends on the propagation as well as the time-discretization. The theoretical error bounds in [13] are of the form $\varepsilon(1 + |x|^2)$ (over the entire space) where $\varepsilon \downarrow 0$ as the number of propagation steps goes to infinity and time-discretization go to zero, with the required relative rates being discussed in

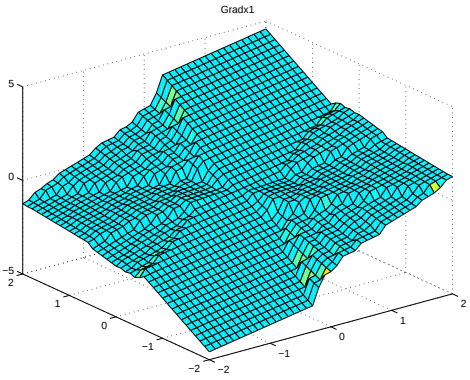

Fig. 1. x_1 partial on the x_1, x_2 plane

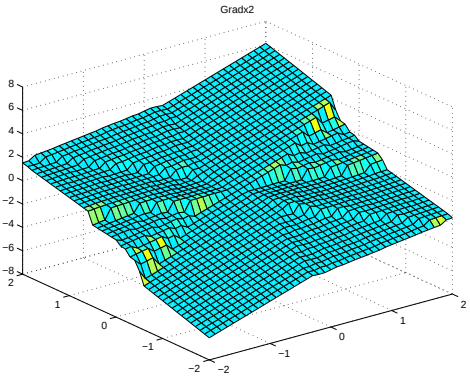

Fig. 2. x_2 partial on the x_1, x_2 plane

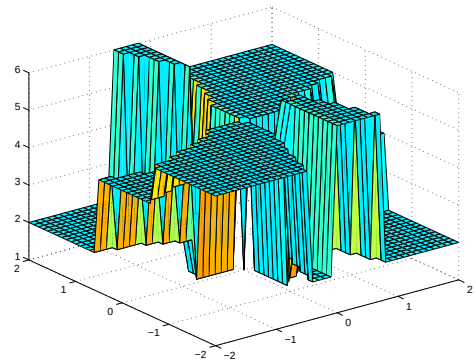

Fig. 3. Optimal switching policy on the x_1, x_2 plane

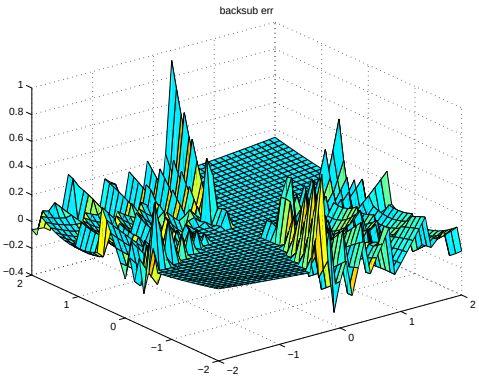

Fig. 4. Backsubstitution error on the x_1, x_2 plane

the reference. However, we found that, when looking only over a compact sub-domain, the component of the error due to the number of propagation steps went to zero after a finite number of steps, leaving only the error due to time-discretization. That is, the error reduction on a compact sub-domain saturated after a sufficient number of iterations were performed.

REFERENCES

[1] M. Akian, S. Gaubert, and A. Lakhoua, "The max-plus finite element method for optimal control problems: further approximation results", Proc. 44th IEEE Conf. on Dec. and Control (joint with ECC'05), Sevilla (2005).

[2] M. Akian, S. Gaubert and A. Lakhoua, "A max-plus finite element method for solving finite horizon determinsitic optimal control problems", Proc. 16th MTNS (2004).

[3] F.L. Baccelli, G. Cohen, G.J. Olsder and J.-P. Quadrat, *Synchronization and Linearity*, John Wiley (1992).

[4] A. Ben-Tal and A. Nemirovski, *Lectures on Modern Convex Optimization: analysis, algorithms and engineering applications*, MPS-SIAM Series on Optimization (2001).

[5] R.A. Cuninghame-Green, *Minimax Algebra*, Lecture Notes in Economics and Mathematical Systems 166, Springer (1979).

[6] P. Dupuis and A. Szpiro, *Second order numerical methods for first order Hamilton-Jacobi equations*, SIAM J. Num. Analysis, **40** (2002), pp. 113–118.

[7] M. Falcone and R. Ferretti, *Convergence analysis for a class of high-order semi-lagrangian advection schemes*, SIAM J. Num. Analysis, **35** (1998), pp. 909–940.

[8] W.H. Fleming and W.M. McEneaney, "A max-plus based algorithm for an HJB equation of nonlinear filtering", SIAM J. Control and Optim. **38** (2000), 683–710.

[9] V.N. Kolokoltsov and V.P. Maslov, *Idempotent Analysis and Its Applications*, Kluwer (1997).

[10] G.L. Litvinov, V.P. Maslov and G.B. Shpiz, "Idempotent Functional Analysis: An Algebraic Approach", Mathematical Notes, **69**, No. 5 (2001), 696–729.

[11] V.P. Maslov, "On a new principle of superposition for optimization problems", Russian Math. Surveys, **42** (1987) 43–54.

[12] W.M. McEneaney, *Max-Plus Methods in Nonlinear Control and Estimation*, Birkhauser (2006).

[13] W.M. McEneaney and J.L. Kluberg, "Convergence Rate for a Curse-of-Dimensionality-Free Method for HJB PDEs Represented as Maxima fo Quadratic Forms," Submitted to SIAM J. Control and Optim.

[14] W.M. McEneaney, "A Curse-of-Dimensionality-Free Numerical Method for Solution of Certain HJB PDEs," SIAM J. on Control and Optim., **46** (2007), 1239–1276.

[15] W.M. McEneaney, "Max–Plus Eigenvector Representations for Solution of Nonlinear H_∞ Problems: Basic Concepts", IEEE Trans. Auto. Control **48**, (2003) 1150–1163.

[16] R.T. Rockafellar and R.J. Wets, *Variational Analysis*, Springer–Verlag (1997).