Splitting for Hamilton-Jacobi Equations Arising from Optimal Control and Differential Games

Alex Tong Lin, Yat Tin Chow, Stanley Osher

({atlin, sjo}@math.ucla.edu, yattinc@ucr.edu)

UCLA

Introductions

- In this work, we apply a new splitting method based on the Primal Dual Hybrid Gradient algorithm (a.k.a. Chambolle-Pock) to nonlinear optimal control (OC) and differential games (DG) problems, based on using the direct collocation method, but with a Hamiltonian twist.
- This allow us to compute solutions at specified points directly, i.e. without the use of grids in space. And it also gives us the ability to create trajectories directly.
- Thus we are able to lift the curse of dimensionality a bit, and therefore compute solutions in much higher dimensions than before. And in our numerical experiments, we actually observe that our computations scale polynomially in time. Furthermore, this new algorithm is embarrassingly parallelizable.

Hamilton-Jacobi, Optimal Control, and Differential Games

- The goal of optimal control is to find a control policy that will drive a system while optimizing a criterion. This is mathematically defined as the ODE,

$$\begin{cases} \dot{\mathbf{x}}(s) = \mathbf{f}(\mathbf{x}(s), \boldsymbol{\alpha}(s), s), \quad t < s < T \\ \mathbf{x}(t) = x \end{cases} \tag{1}$$

where x is the initial point, T is the terminal time, and $\mathbf{u}$ is the control. And the functional we want to optimize is

$$J_{x,t}[\mathbf{u}] := g(\mathbf{x}(T)) + \int_t^T L(\mathbf{x}(s), \mathbf{u}(s), s)\, ds.$$

where J is a cost, which we minimize. Then we define the *value function*

$$\varphi(x,t) = \min_{\alpha(\cdot) \in \mathcal{A}} J_{x,t}[\mathbf{u}].$$

Under mild conditions on $\mathbf{f}$, g, and L, this value function satisfies the *terminal-valued* Hamilton-Jacobi PDE (HJ PDE)

$$\begin{cases} \partial_t \varphi(x,t) + H(x, \nabla_x \varphi(x,t), t) &= 0, \qquad (x,t) \in \mathbb{R}^n \times (0,T) \\ \varphi(x,T) &= g(x). \end{cases}$$

where $H(x,p,t) = \min_{a \in A} \{\langle \mathbf{f}(x,a,t), p\rangle + L(x,a,t)\}$.

- A (two-person, zero-sum) differential game now has competing controls:

$$\begin{cases} \dot{\mathbf{x}}(s) = \mathbf{f}(\mathbf{x}(s), \mathbf{u}(s), \boldsymbol{\beta}(s), s), \quad t < s < T \\ \mathbf{x}(t) = x \end{cases}$$

and the performance function,

$$J_{x,t}[\mathbf{u}, \boldsymbol{\beta}] := g(\mathbf{x}(T)) + \int_t^T L(\mathbf{x}(s), \mathbf{u}(s), \boldsymbol{\beta}(s), s)\, ds.$$

and the value function,

$$\varphi(x,t) = \inf_{\Psi[\cdot] \in \mathcal{B}(t)} \sup_{\mathbf{u}(\cdot) \in A(t)} J_{x,t}[\mathbf{u}, \Psi[\mathbf{u}]]$$

and the HJE PDE (which satisfies the min max = max min condition),

$$\begin{cases} \partial_t \varphi(x,t) + \max_{a \in A} \min_{b \in B} \{\langle \mathbf{f}(x,a,b,t), \nabla_x \varphi(x,t)\rangle + L(x,a,b,t)\} = 0 \\ \varphi(x,T) = g(x) \end{cases}$$

Primal-Dual Optimization Splitting

The PDHG algorithm [5, 3], which also goes by the name Chambolle-Pock [1], attempts to solve problems of the form

$$\min_{x \in X} \quad f(Ax) + g(x)$$

PDHG takes the Lagrangian dual formulation and seeks to find a saddle point:

$$\min_{x \in X} \max_{y \in Y} \langle Ax, y\rangle + g(x) - f^*(y)$$

where $f^*(y) = \sup_{x \in X} \{\langle x, y\rangle - f(y)\}$ is the *convex conjugate* of f. PDHG is also an alternating minimization technique that uses proximal operators. It iterates:

$$\begin{cases} y^{k+1} &= (I + \sigma \partial f^*)^{-1}(y^k + \sigma A \bar{x}^k) \\ x^{k+1} &= (I + \tau \partial g)^{-1}(x^k - \sigma A^* y^{k+1}) \\ \bar{x}^{k+1} &= x^{k+1} + \theta(x^{k+1} - x^k). \end{cases}$$

where $\sigma, \tau > 0$ satisfy $\sigma\tau \|A\|^2 < 1$, and $\theta \in [0,1]$. In practice, we put $\theta = 1$.

Collocation on OC and DG, with a Hamiltonian Twist

- We start by discretizing the value function. This is the procedure followed in [2]: We have the value function equals

$$\varphi(x,t) = \min_{\mathbf{x}(\cdot), \mathbf{u}(\cdot)} \left\{ g(\mathbf{x}(0)) + \int_0^t L(\mathbf{x}(s), \mathbf{u}(s), s)\, ds \right\}$$

where $\mathbf{x}(\cdot)$ and $\mathbf{u}(s)$ satisfy the ODE (1),

$$\min_{\{x_j\}, \{u_j\}} \left\{ g(x_0) + \delta \sum_{j=1}^N L(x_j, u_j, s_j) \mid \{x_j - x_{j-1} = \delta f(x_j, u_j, s_j)\}_{j=1}^N \right\}$$

As usual in constrained optimization problems, we compute the Lagrangian function (i.e. Lagrange multipliers) to get:

$$g(x_0) + \delta \sum_{j=1}^N L(x_j, u_j, s_j) + \sum_{j=1}^N \langle p_j, x_j - x_{j-1} - \delta f(x_j, u_j, s_j)\rangle + \langle p_N, x - x_N\rangle$$

Note that the constraint $x_N = x$ is trivially unneeded in the Lagrangian function. Then we minimize over $\{x_j\}_{j=0}^N$, while maximizing over $\{p_j\}_{j=1}^N$, and by moving the minimiziation with respect to $\{u_j\}_{j=1}^N$ we get,

$$\max_{\{p_j\}} \min_{\{x_k\}} \left\{ g(x_0) + \sum_{j=1}^N \langle p_j, x_j - x_{j-1}\rangle + \langle p_N, x - x_N\rangle - \delta \sum_{j=1}^N H(x_j, p_j, s_j) \right\}$$

- There is tremendous advantage in having a Hamiltonian. This is because if we want to instead perform optimization of the value function directly, we will be solving for the controls and this requires a constrained optimization technique.

  The miraculous advantage of having a Hamiltonian for optimization purposes is it *encodes* information from both the running cost function L, as well as the dynamics $\dot{\mathbf{x}}(s) = \mathbf{f}(\mathbf{x}(s), \mathbf{u}(s), s)$. And now we are free to perform *unconstrained optimization*. Plus, we lower the dimension of the numerical optimization by analytically minimizing over u, and conjuring the Hamiltonian.

Numerical Results

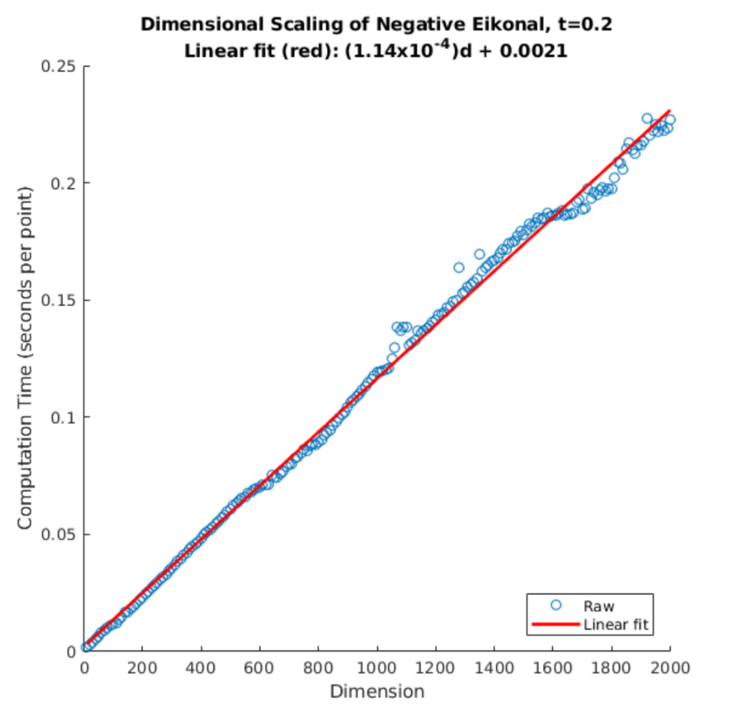

Linear scaling of computation time versus dimension.

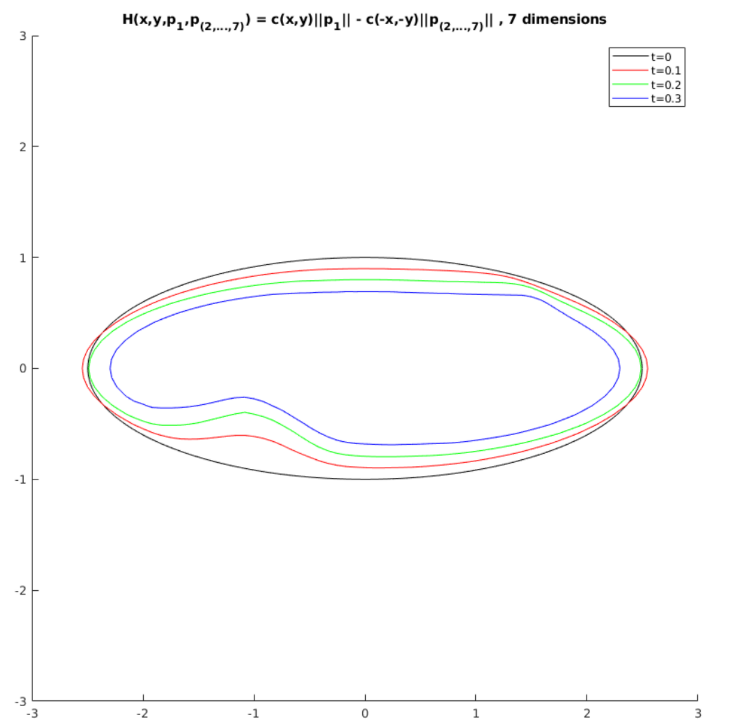

Seven-dimensional differential games problem

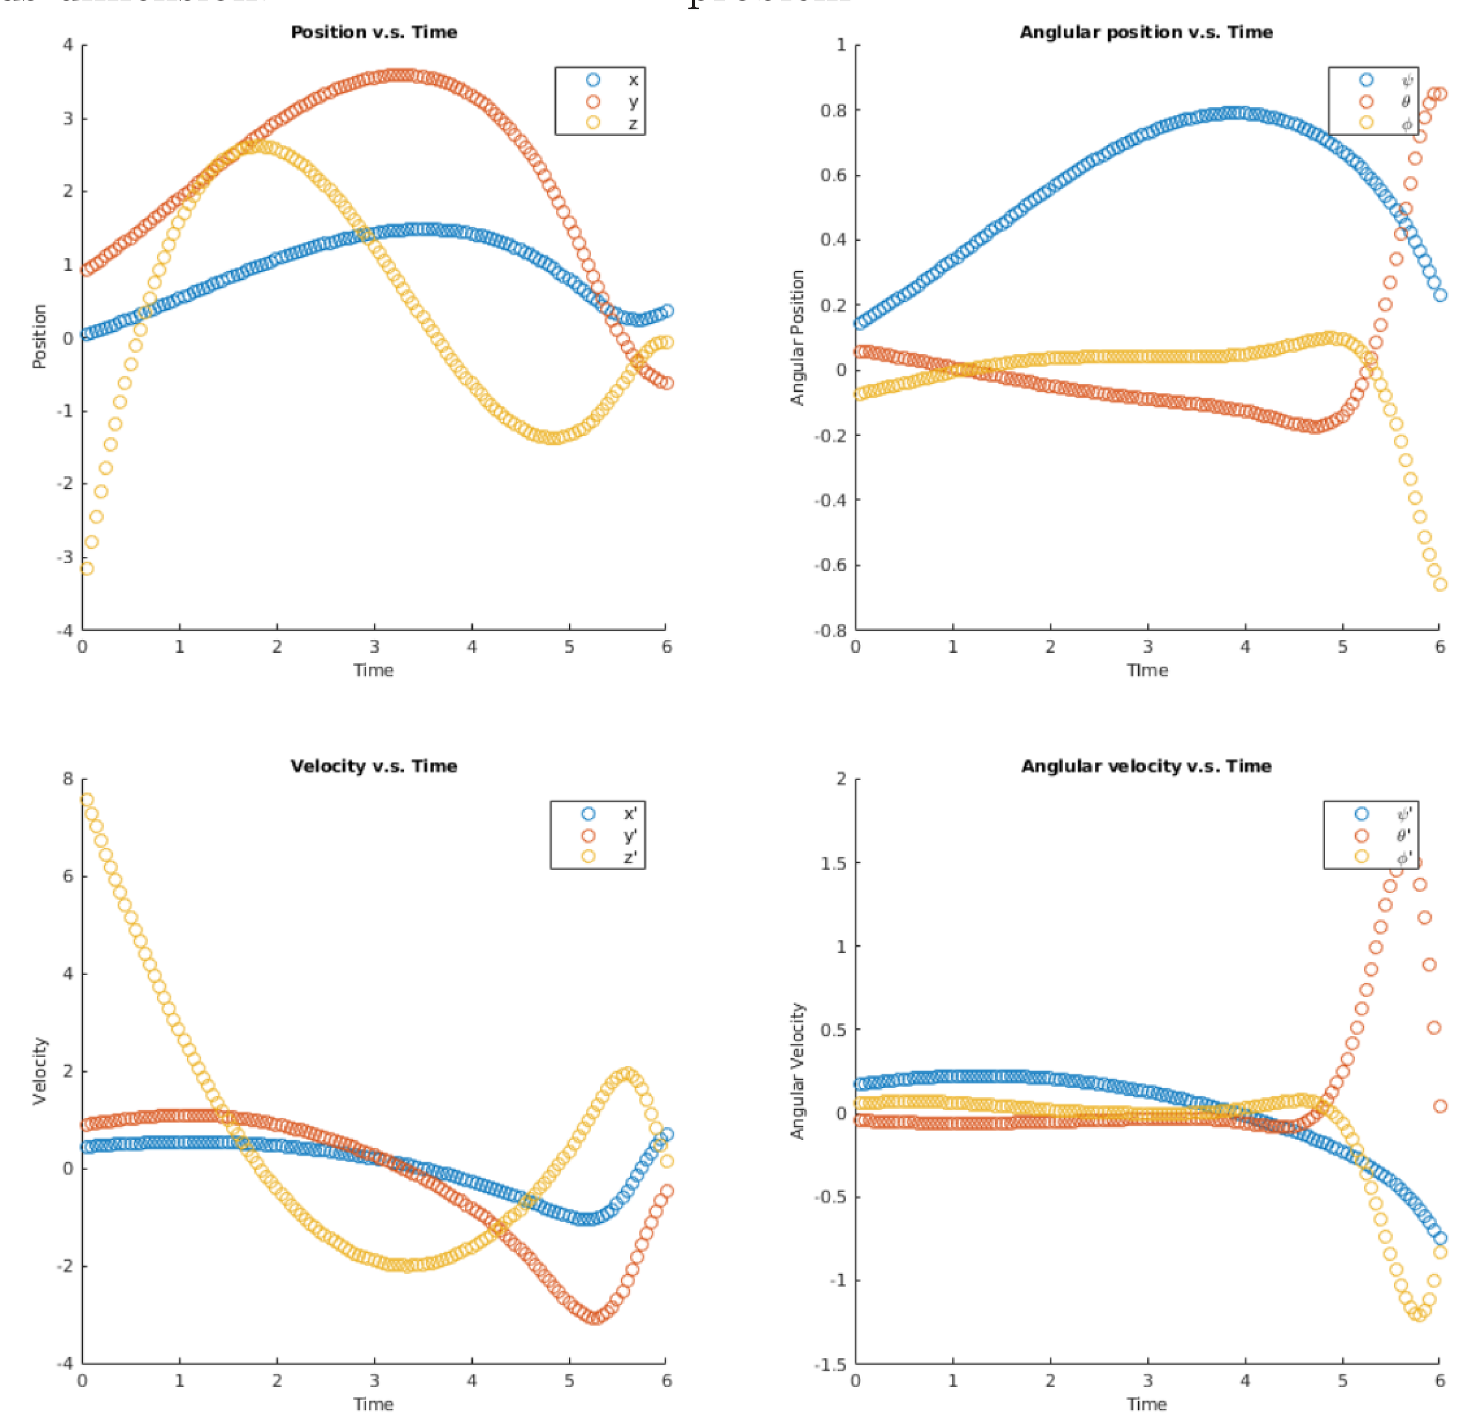

Twelve-dimensional quadcopter problem

References

[1] A. Chambolle and T. Pock. A first-order primal-dual algorithm for convex problems with applications to imaging. *Journal of Mathematical Imaging and Vision*, 40(1):120–145, May 2011.

[2] Y.-T. Chow, J. Darbon, S. Osher, and W. Yin. Algorithm for overcoming the curse of dimensionality for state-dependent hamilton-jacobi equations. UCLA CAM 17-16, Apr. 2017.

[3] E. Esser, X. Zhang, and T. F. Chan. A general framework for a class of first order primal-dual algorithms for convex optimization in imaging science. *SIAM Journal on Imaging Sciences*, 3(4):1015–1046, 2010.

[4] A. T. Lin, Y. T. Chow, and S. Osher. A splitting method to compute solutions to possibly nonconvex state-and-time-dependent hamilton-jacobi equations arising from optimal control and differential games. *To appear in Communications in Mathematical Sciences*, 2018.

[5] M. Zhu and T. Chan. An efficient primal-dual hybrid gradient algorithm for total variation image restoration. UCLA CAM, 2008.