Stat-duality based method for rapid solution of high-dimensional first-order HJ PDE problems with low-dimensional nonlinearities

AFOSR D&C Review, 12 Sept. 2023

W.M. McEneaney, UC San Diego

Collaborators: P. Dower, U. Melbourne; Y. Zheng, UCSD

Topic

- Current effort includes:
  - Complex-valued diffusion processes and representations of Schrödinger equation solutions under Coulomb potentials.
  - Isolation of complexity, and rapid solution of certain first-order HJ PDE problems (and associated control problems).
  - Applications in astrodynamics.

Complexity

- Classical computational methods for dynamic programming were grid based.
- This induced the dreaded curse-of-dimensionality that limited the development of realistically-useful computational tools from the time of the "discovery" of dynamic programming in the 1950s to recent times.
- The complexity of computation is bounded below by a function of the complexity of the representation of the solution.
- In general, function complexity has no rigorous relation to domain-space dimension.
  - In the LQ/LQG case with state in $\boldsymbol{R}^n$, the value function may at any time by represented by $\frac{(n+1)(n+2)}{2}$ real numbers.
  - A scalar brownian path on $[0, 1]$ cannot be represented by a finite number of real numbers.
- The curse-of-complexity is a more appropriate description of the effect.
- This effort addresses problems of high-dimensionality with relatively low complexity.
- Will consider the particular case where the problem is primarily linear-quadratic, but has a "low-dimensional" nonlinearity.

Problem Definition and Goal

- Problem definition:

$$\begin{aligned}
\dot{\xi}_s &= A\,\xi_s + L^0 f(M^0\xi_s) + \sigma\, u_s, \qquad \xi_t = x \in \boldsymbol{R}^n, \\
\tilde{J}_t(x,u) &\doteq \int_t^T \ell(M^0\xi_s) + \tfrac{1}{2}\xi_s'\, C\,\xi_s + \tfrac{1}{2}\,|u_s|^2\, ds, \\
\tilde{W}_t(x) &\doteq \operatorname*{stat}_{u\in\mathcal{U}[t,T]} \tilde{J}_t(x,u).
\end{aligned}$$

- $f : \boldsymbol{R}^k \to \boldsymbol{R}^\ell$ and $\ell : \boldsymbol{R}^k \to \boldsymbol{R}$ are both nonlinear.
- $M^0 \in \boldsymbol{R}^{k\times n}$ and $(L^0)' \in \boldsymbol{R}^{\ell\times n}$ are projections onto lower-dimensional spaces.
- Although the state lives in $\boldsymbol{R}^n$, the nonlinearities may operate on a significantly lower-dimensional space.
- We will exploit this.

Section 1:
Background on Staticization

Staticization

- Staticization is the search for stationary (static) points of functionals.
- Let $\bar{y} \in \mathcal{G}_\mathcal{Y}$ where $\mathcal{G}_\mathcal{Y}$ is an open subset of a Hilbert space. We say

$$\bar{y} \in \operatorname*{argstat}_{y \in \mathcal{G}_\mathcal{Y}} F(y) \quad \text{if} \quad \limsup_{y \to \bar{y}, y \in \mathcal{G}_\mathcal{Y}} \frac{|F(y) - F(\bar{y})|}{|y - \bar{y}|} = 0,$$

- If f is differentiable and $\mathcal{G}_\mathcal{Y}$ is open, then $\operatorname{argstat}_{y \in \mathcal{G}_\mathcal{Y}} F(y) = \{y \in \mathcal{G}_\mathcal{Y} \,|\, F_y(y) = 0\}$.
- Define set-valued $\overline{\mathrm{stat}}$ by

$$\overline{\mathrm{stat}}_{y \in \mathcal{G}_\mathcal{Y}} F(y) \doteq \left\{ F(\bar{y}) \,\middle|\, \bar{y} \in \operatorname*{argstat}_{y \in \mathcal{G}_\mathcal{Y}} \{F(y)\} \right\} \text{ if } \operatorname{argstat}\{F(y) \,|\, y \in \mathcal{G}_\mathcal{Y}\} \neq \emptyset.$$

- If there exists a s.t. $\overline{\mathrm{stat}}_{y \in \mathcal{G}_\mathcal{Y}} F(y) = \{a\}$, then $\mathrm{stat}_{y \in \mathcal{G}_\mathcal{Y}} F(y) \doteq a$.
- Staticization subsumes minimization, maximization and saddle-point searches for C^1 functionals.

Staticization-Based Representation for the Gravitational Potential

- Classic gravitational potential energy expression for bodies at x and origin with masses m and m_0:

$$-V(x) = \frac{G m_0 m}{|x|}.$$

- Inverse norm is difficult.

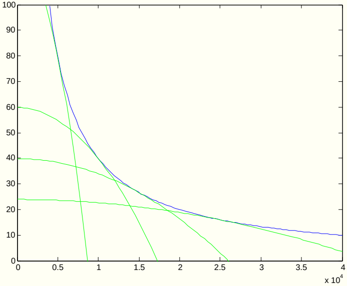

- Additive inverse of potential as optimized quadratic (with $\widehat{G} \doteq (3/2)^{3/2} G$).

$$-V(x) = \frac{\widehat{G} m_0 m}{|x|} = \widehat{G} m_0 m \sup_{\alpha\in[0,\infty)} \left\{ \alpha - \frac{\alpha^3 |x|^2}{2} \right\}.$$

- Argument is convex cubic on $[0,\infty)$; replace sup with stat:

$$-V(x) = \widehat{G} m_0 m \operatorname*{stat}_{\alpha\in[0,\infty)} \left\{ \alpha - \frac{\alpha^3 |x|^2}{2} \right\}.$$

Staticization-based extension of Coulomb potential to $\mathbb{C}^3$

- Although min and max are valid only for real-valued functionals, staticization is valid for complex-valued functions.
- For use in a certain diffusion representation, the Coulomb potential must be extended to $\mathbb{C}^3$, and nonetheless has stat representation:

$$-V(x) = \frac{\mu_c}{\sqrt{x^T x}} =\hat{\mu} \operatorname*{stat}_{\alpha \in \mathcal{H}^+} \left[\alpha - \frac{\alpha^3 (x^T x)}{2} \right],$$

where $\mathcal{H}^+ \doteq \{ \alpha = re^{i\theta} \mid r > 0,\, \theta \in (-\pi/2, \pi/2] \}$.

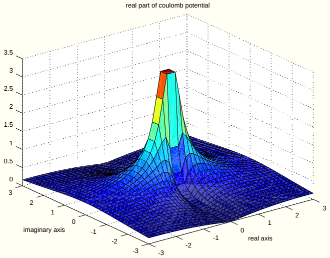

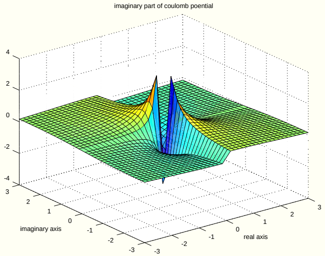

Part 2:
Application to Our Problem

Problem Definition and Goal

Recall our problem:

$$\dot{\xi}_s = A\,\xi_s + L^0 f(M^0\xi_s) + \sigma\, u_s, \qquad \xi_t = x \in \boldsymbol{R}^n,$$

$$\tilde{J}_t(x,u) \doteq \int_t^T \ell(M^0\xi_s) + \tfrac{1}{2}\xi_s'\,C\,\xi_s + \tfrac{1}{2}\,|u_s|^2\,ds,$$

$$\tilde{W}_t(x) \doteq \operatorname*{stat}_{u\in\mathcal{U}[t,T]} \tilde{J}_t(x,u).$$

- Note: stat subsumes minimization, maximization, minimax solution.
- $f : \boldsymbol{R}^k \to \boldsymbol{R}^\ell$ and $\ell : \boldsymbol{R}^k \to \boldsymbol{R}$ are both nonlinear.
- $M^0 \in \boldsymbol{R}^{k\times n}$ and $(L^0)' \in \boldsymbol{R}^{\ell\times n}$ are projections onto lower-dimensional spaces.
- Although the state lives in $\boldsymbol{R}^n$, the nonlinearities operate on a lower-dimensional subspace.
- We will exploit this.

Flowchart of the (Apparently) Required Mathematics

- Quite a large number of major steps are required to obtain this particular form.
- Approach requires Hamilton-Jacobi methods, stat-quad duality technique and a variety of tools from functional analysis.

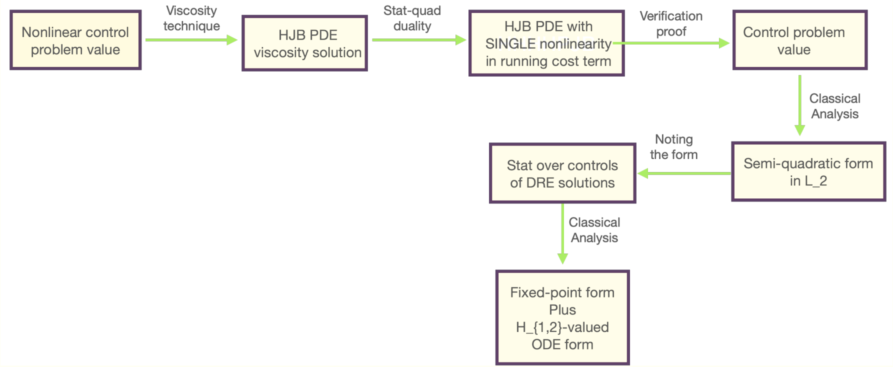

Step 1.

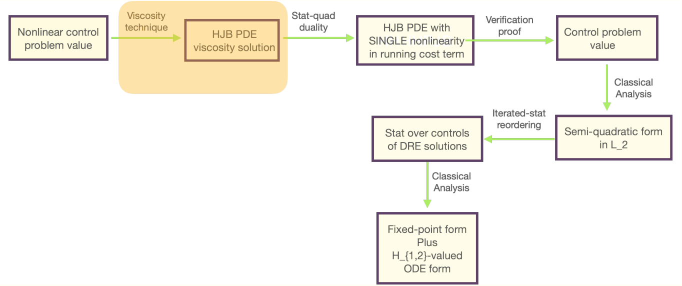

Nonlinear control problem value
Viscosity technique
HJB PDE viscosity solution
Stat-quad duality
HJB PDE with SINGLE nonlinearity in running cost term
Verification proof
Control problem value
Classical Analysis
Semi-quadratic form in L_2
Iterated-stat reordering
Stat over controls of DRE solutions
Classical Analysis
Fixed-point form Plus H_{1,2}-valued ODE form

Problem Definition and the HJ PDE

Recall problem:
$$\dot{\xi}_s = A\,\xi_s + L^0 f(M^0\xi_s) + \sigma\,u_s, \qquad \xi_t = x \in \boldsymbol{R}^n,$$
$$\tilde{J}(t,x,u) \doteq \int_t^T \ell(M^0\xi_s) + \tfrac{1}{2}\xi_s'\,C\,\xi_s + \tfrac{1}{2}\,|u_s|^2\,ds,$$
$$\tilde{W}(t,x) \doteq \operatorname*{stat}_{u\in\mathcal{U}[t,T]} \tilde{J}(t,x,u).$$

- Use staticization and viscosity-solution theories $\Longrightarrow$ $\tilde{W}$ is the viscosity solution of the associated HJ PDE.
- This implies existence of such a solution.

$$\begin{cases} 0 = -U_s + \tilde{H}(x, \nabla_x U), & (s,x) \in (t,T) \times \boldsymbol{R}^n, \\ U(T,x) = 0, & x \in \mathbb{R}^n, \end{cases}$$

$$\tilde{H}(x,p) \doteq -\Big[\tfrac{1}{2}\,x'\,C\,x + p'\,A\,x - \tfrac{1}{2}\,p'\,\sigma\,\sigma'\,p + \underbrace{\ell(M^0x) + [(L^0)'p,]'\,f(M^0x)}_{\doteq \tilde{N}(M^0x,(L^0)'p)}\Big]$$

- All the nonlinearities reside in $\tilde{N}$.

Stat-Quad Duality Application

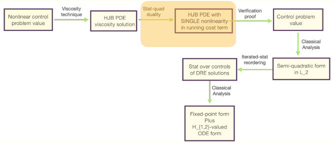

- Take the stat-quad dual of $\tilde{N}$:

$$\tilde{N}(M^0 x, (L^0)' p) = \operatorname*{stat}_{(a,b)\in\mathbb{R}^{k+\ell}} \left\{ \tilde{\Theta}(a,b) + Q^1(x,p,a,b) \right\} \qquad \forall\, (x,p) \in \boldsymbol{R}^{2n},$$

$$\tilde{\Theta}(a,b) = \operatorname*{stat}_{(x,p)\in\boldsymbol{R}^{2n}} \left\{ \tilde{N}(M^0 x, (L^0)' p) - Q^1(x,p,a,b) \right\} \qquad \forall\, (a,b) \in \mathbb{R}^{k+\ell},$$

where $Q^1(x,p,a,b) \doteq -\frac{c_1}{2}|M^0 x - a|^2 - \frac{c_2}{2}|(L^0)' p - b|^2$.

- Note that $\tilde{\Theta}$ is a function on $k+\ell$ dimensional space – not on $\boldsymbol{R}^n$.

Stat-Quad Duality Application

- Employing the stat-quad dual yields

$$
\begin{aligned}
0 = -\Big\{ & U_s + \tfrac{1}{2}x'Cx + (\nabla_x U)'Ax - \tfrac{1}{2}(\nabla_x U)'\tilde{\Gamma}\nabla_x U \\
& + \operatorname*{stat}_{(a,b)\in\mathbb{R}^{k+\ell}} \left[\tilde{\Theta}(a,b) - \tfrac{c_1}{2}|M^0x - a|^2 - \tfrac{c_2}{2}|b|^2 + c_2(\nabla_x U)'L^0 b \right] \Big\}
\end{aligned}
$$

$$
U(T,x) = 0 \quad x \in \boldsymbol{R}^n,
$$

where $\tilde{\Gamma} \doteq \sigma\sigma' + c_2 L^0 (L^0)'$.

- Minor detail: The purple term appears because of the $(L^0)'p$ term in Q^1. (Recall $Q^1(x,p,a,b) = -\frac{c_1}{2}|M^0x - a|^2 - \frac{c_2}{2}|(L^0)'p - b|^2$
- a, b are staticizing-controllers.
- All of the nonlinearities are now confined to a control-cost term in the running cost.

Stat-Quad Duality Application

- Let $\tilde{\Gamma} = \tilde{B}'\tilde{B}$. We obtain

$$-\tfrac{1}{2}(\nabla_x U)'\tilde{\Gamma}\nabla_x U = \operatorname*{stat}_{v\in R^n}\{(\nabla_x U)'\tilde{B}v + \tfrac{1}{2}|v|^2\}$$

- This yields

$$0 = -\Big\{ U_s + \tfrac{1}{2}x'Cx + (\nabla_x U)'Ax + \operatorname*{stat}_{v\in R^n}\big[(\nabla_x U)'\tilde{B}v + \tfrac{1}{2}|v|^2\big] + \operatorname*{stat}_{(a,b)\in\mathbb{R}^{k+\ell}}\big[\tilde{\Theta}(a,b) - \tfrac{c_1}{2}|M^0x - a|^2 - \tfrac{c_2}{2}|b|^2 + c_2(\nabla_x U)'L^0 b\big]\Big\}$$

- a, b and v are staticizing controllers.
- Above is a separated staticization case, (analogous to Isaacs cond.); hence:

$$0 = -\Big\{ U_s + \tfrac{1}{2}x'Cx + (\nabla_x U)'Ax + \operatorname*{stat}_{(a,b,v)\in R^{k+\ell+n}}\big[(\nabla_x U)'\tilde{B}v + \tfrac{1}{2}|v|^2 + \tilde{\Theta}(a,b) - \tfrac{c_1}{2}|M^0x - a|^2 - \tfrac{c_2}{2}|b|^2 + c_2(\nabla_x U)'L^0 b\big]\Big\}$$

Step 3.

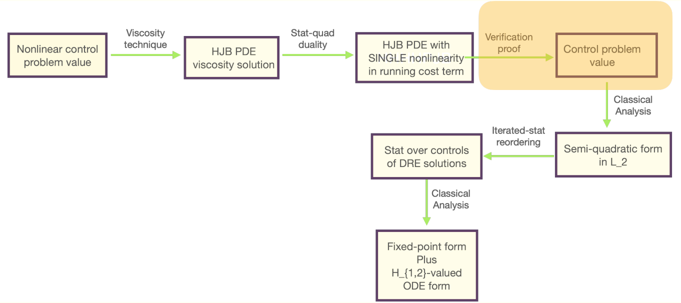

Nonlinear control problem value
Viscosity technique
HJB PDE viscosity solution
Stat-quad duality
HJB PDE with SINGLE nonlinearity in running cost term
Verification proof
Control problem value
Classical Analysis
Semi-quadratic form in L_2
Iterated-stat reordering
Stat over controls of DRE solutions
Classical Analysis
Fixed-point form Plus H_{1,2}-valued ODE form

Stat-Control Verification Result

Recall: $$0 = -\Big\{ U_s + \tfrac{1}{2}x'Cx + (\nabla_x U)'Ax + \operatorname*{stat}_{(a,b,v)\in \mathbb{R}^{k+\ell+n}} \big[(\nabla_x U)'\tilde{B}v + \tfrac{1}{2}|v|^2 + \tilde{\Theta}(a,b) - \tfrac{c_1}{2}|M^0x - a|^2 - \tfrac{c_2}{2}|b|^2 + c_2(\nabla_x U)'L^0 b \big] \Big\}.$$

- Extend viscosity-solution verification methods to the staticization case.
- Under sufficient smoothness (C^1!), the HJ PDE problem solution is the value of the following control problem:

$$\dot{\zeta}_s = A\zeta_s + \tilde{B}\nu_s + c_2 L^0 \beta_s, \quad \zeta_t = x \in \mathbb{R}^n,$$

$$\breve{J}(t,x,\nu,\alpha,\beta) \doteq \int_t^T \tfrac{1}{2}\zeta_s' C\zeta_s + \tfrac{1}{2}|\nu_s|^2 + \tilde{\Theta}(\alpha_s,\beta_s) - \tfrac{c_1}{2}|M^0\zeta_s - \alpha_s|^2 - \tfrac{c_2}{2}|\beta_s|^2 \, ds$$

$$\breve{W}(t,x) \doteq \operatorname*{stat}_{(\nu,\tilde{\alpha},\tilde{\beta})\in L_2(t,T;\mathbb{R}^{n+k+\ell})} \breve{J}(t,x,\nu,\alpha,\beta).$$

- Aside: This implies uniqueness of the HJ PDE problem solution.

Step 4. Classical Functional Analysis

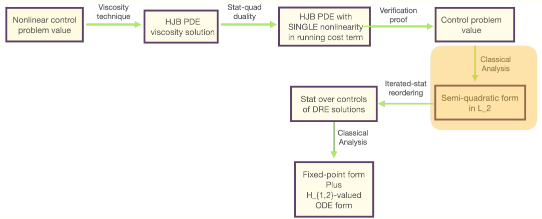

- The above problem takes the form

$$\breve{J}(t,x,\nu,\alpha,\beta) = f_1(\alpha,\beta;t,x) + \langle f_2(\alpha,\beta;t,x),\nu\rangle_{L_2} + \tfrac{1}{2}\langle \nu, \bar{B}_3(t)\nu\rangle_{L_2},$$

$$\breve{W}(t,x) \doteq \operatorname*{stat}_{(\nu,\tilde{\alpha},\tilde{\beta})\in L_2(t,T;\mathbb{R}^{n+k+\ell})} \breve{J}(t,x,\nu,\alpha,\beta).$$

Step 5.

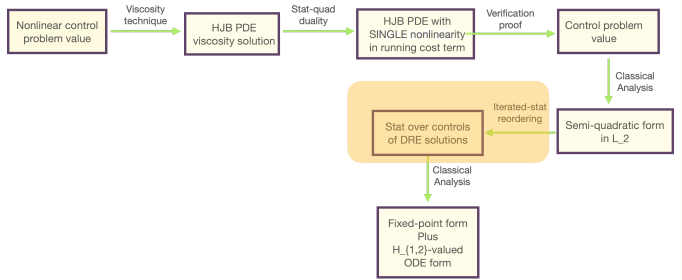

Nonlinear control problem value
Viscosity technique
HJB PDE viscosity solution
Stat-quad duality
HJB PDE with SINGLE nonlinearity in running cost term
Verification proof
Control problem value
Classical Analysis
Semi-quadratic form in L_2
Iterated-stat reordering
Stat over controls of DRE solutions
Classical Analysis
Fixed-point form Plus H_{1,2}-valued ODE form

Iterated-Stat Reordering

- One can generally reorder minimization [maximization] operations.
- The ability to reorder differentiation can be misleading with regard to stat, where $\frac{d}{dx}\frac{d}{dy}f(x,y) = \frac{d}{dy}\frac{d}{dx}f(x,y)$.
- Consider $f(x,y) = y(x^2 - 1)$.

$$\operatorname*{stat}_{x\in \boldsymbol{R}} \operatorname*{stat}_{y\in \boldsymbol{R}} y(x^2-1) = \operatorname*{stat}_{(x,y)\in \boldsymbol{R}^2} y(x^2-1) = 0,$$

$$\operatorname*{stat}_{y\in \boldsymbol{R}} \operatorname*{stat}_{x\in \boldsymbol{R}} y(x^2-1) \text{ does not exist.}$$

- Cases where stat may be reordered are where the function is quadratic in at least one argument, or where it is Morse in at least one argument, both with additional conditions.

Employing Iterated Staticization Results

- We use iterated-staticization theory.
- Convert $\mathrm{stat}_{(\nu,\tilde{\alpha},\tilde{\beta})\in L_2(t,T;\mathbb{R}^{n+k+\ell})}$ into $\mathrm{stat}_{(\tilde{\alpha},\tilde{\beta})\in L_2(t,T;\mathbb{R}^{k+\ell})}\,\mathrm{stat}_{\nu\in L_2(t,T;\mathbb{R}^n)}$.

$$\breve{J}(t,x,\nu,\alpha,\beta) = f_1(\alpha,\beta;t,x) + \langle f_2(\alpha,\beta;t,x),\nu\rangle_{L_2} + \tfrac{1}{2}\langle \nu, \bar{B}_3(t)\nu\rangle_{L_2},$$

$$\breve{W}(t,x) \doteq \underset{(\tilde{\alpha},\tilde{\beta})\in L_2(t,T;\mathbb{R}^{k+\ell})}{\mathrm{stat}} W^{\tilde{\alpha},\tilde{\beta}}(t,x),$$

where
$$W^{\tilde{\alpha},\tilde{\beta}}(t,x) \doteq \underset{\nu\in L_2(t,T;R^n)}{\mathrm{stat}} \breve{J}(t,x,\nu,\alpha,\beta)$$
$$= \underset{\nu\in L_2(t,T;R^n)}{\mathrm{stat}} \left\{ f_1(\alpha,\beta;t,x) + \langle f_2(\alpha,\beta;t,x),\nu\rangle_{L_2} + \tfrac{1}{2}\langle \nu, \bar{B}_3(t)\nu\rangle_{L_2}\right\},$$

- Note the LQ form of that last stat argument.

Step 6.

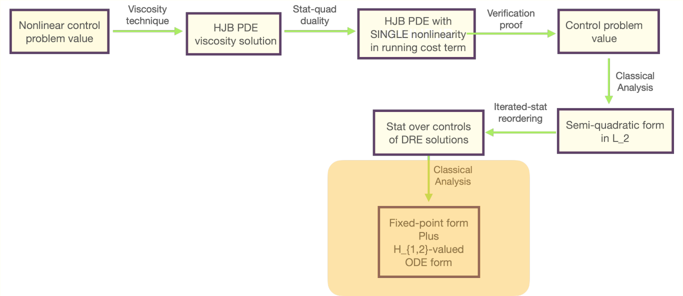
Nonlinear control problem value
Viscosity technique
HJB PDE viscosity solution
Stat-quad duality
HJB PDE with SINGLE nonlinearity in running cost term
Verification proof
Control problem value
Classical Analysis
Semi-quadratic form in L_2
Iterated-stat reordering
Stat over controls of DRE solutions
Classical Analysis
Fixed-point form Plus H_{1,2}-valued ODE form

DRE Extraction

- We have

$$\breve{W}(t,x) \doteq \operatorname*{stat}_{(\tilde{\alpha},\tilde{\beta})\in L_2(t,T;\mathbb{R}^{k+\ell})} W^{\tilde{\alpha},\tilde{\beta}}(t,x),$$

$$W^{\tilde{\alpha},\tilde{\beta}}(t,x) = \operatorname*{stat}_{\nu\in L_2(t,T;R^n)} \{f_1(\alpha,\beta;t,x) + \langle f_2(\alpha,\beta;t,x),\nu\rangle_{L_2} + \tfrac{1}{2}\langle \nu, \bar{B}_3(t)\nu\rangle_{L_2}\},$$

- We can solve the inner LQ problem!
- The DRE-based solution may be written as

$$W^{\tilde{\alpha},\tilde{\beta}}(t,x) = \tfrac{1}{2}x' P_t x + x' q_t^{\tilde{\alpha},\tilde{\beta}} + r_t^{\tilde{\alpha},\tilde{\beta}},$$

$$\dot{P}_t = -S_{c_1} - A' P_t - P_t A + P_t \tilde{\Gamma} P_t, \quad P_T = 0_{n\times n}, \quad \text{(control-independent!)},$$

$$\dot{q}_t^{\mu} = [P_t[\tilde{\Gamma} - A']q_t^{\tilde{\alpha},\tilde{\beta}} - \left[(M^0)', L^0[P_t]\right]\hat{C}\mu_t, \quad q_T^{\tilde{\alpha},\tilde{\beta}} = 0_{n\times 1}$$

$$\dot{r}_t^{\mu} = (q_t^{\tilde{\alpha},\tilde{\beta}})'\tilde{\Gamma} q_t^{\tilde{\alpha},\tilde{\beta}} + \mu_t'\hat{C}\mu_t - 2\left[c_2(q_t^{\tilde{\alpha},\tilde{\beta}})' L^0 \mathcal{I}^2 \mu_t + \tilde{\Theta}(\mu_t)\right], \quad r_T^{\tilde{\alpha},\tilde{\beta}} = 0,$$

- N.B.: Here, we've used μ_t in place of $(\tilde{\alpha}_t', \tilde{\beta}_t')'$ for readability.

DRE Extraction

- We have

$$\breve{W}(t,x) \doteq \operatorname*{stat}_{(\tilde\alpha,\tilde\beta)\in L_2(t,T;\mathbb{R}^{k+\ell})} W^{\tilde\alpha,\tilde\beta}(t,x),$$

$$W^{\tilde\alpha,\tilde\beta}(t,x) = \tfrac{1}{2}x' P_t x + x' q_t^{\tilde\alpha,\tilde\beta} + r_t^{\tilde\alpha,\tilde\beta}$$

$$\dot P_t = -S_{c_1} - A' P_t - P_t A + P_t \tilde\Gamma P_t, \quad P_T = 0_{n\times n}, \quad \text{(control-independent!)},$$

$$\dot q_t^{\mu} = [P_t[\tilde\Gamma - A']q_t^{\tilde\alpha,\tilde\beta} - \left[(M^0)', L^0[P_t]\right]\hat C\mu_t, \quad q_T^{\tilde\alpha,\tilde\beta} = 0_{n\times 1}$$

$$\dot r_t^{\mu} = (q_t^{\tilde\alpha,\tilde\beta})'\tilde\Gamma q_t^{\tilde\alpha,\tilde\beta} + \mu_t'\hat C\mu_t - 2\left[c_2(q_t^{\tilde\alpha,\tilde\beta})' L^0\mathcal{I}^2\mu_t + \tilde\Theta(\mu_t)\right], \quad r_T^{\tilde\alpha,\tilde\beta} = 0,$$

- The problem is (finally!!!) reduced to:

$$\breve{W}(t,x) = \tfrac{1}{2}x' P_t x + \operatorname*{stat}_{(\tilde\alpha,\tilde\beta)\in L_2(t,T;\mathbb{R}^{k+\ell})} \{x' q_t^{\tilde\alpha,\tilde\beta} + r_t^{\tilde\alpha,\tilde\beta}\}.$$

Pointwise Solution

- Recall that we only need to solve

$$\tilde{W}(t,x) \doteq \operatorname*{stat}_{(\tilde{\alpha},\tilde{\beta})\in L_2(t,T;\mathbb{R}^{k+\ell})} \{x' q_t^{\tilde{\alpha},\tilde{\beta}} + r_t^{\tilde{\alpha},\tilde{\beta}}\}.$$

- Space variable, x, is only a parameter.
- Pointwise in x, one need only solve integral equation

$$\hat{C}(\tilde{\alpha}_s,\tilde{\beta}_s) - \nabla\tilde{\Theta}(\tilde{\alpha},\tilde{\beta}) = c_2(L^0\mathcal{I}^2)' q_s^{\tilde{\alpha},\tilde{\beta}} + \hat{C}(\bar{\Phi}_{T,s}\bar{D}_s)'x \tag{1}$$
$$+ \int_t^s \hat{C}(\bar{\Phi}_{\tau,s}\bar{D}_s)'(-\tilde{\Gamma} q_\tau^{\tilde{\alpha},\tilde{\beta}} + c_2 L^0\mathcal{I}^2(\tilde{\alpha}_\tau,\tilde{\beta}_\tau))\,d\tau, \ \text{ a.e. } s\in(t,T),$$

where $$q_s^{\tilde{\alpha},\tilde{\beta}} = \int_s^T \bar{\Phi}_{s,\tau}\bar{D}_\tau\hat{C}(\tilde{\alpha}_\tau,\tilde{\beta}_\tau)\,d\tau \qquad \forall\, s\in[t,T]. \tag{2}$$

- The curse of dimensionality is removed, and replaced by the curse of complexity (possibly as effective dimensionality).
- Equations (1)–(2) are an indexed set of integral equations – not a PDE.

Pointwise Solution Method Details

- For each x, we solve (1)–(2), i.e.,

$$\hat{C}(\tilde{\alpha}_s, \tilde{\beta}_s) - \nabla\tilde{\Theta}(\tilde{\alpha}, \tilde{\beta}) = c_2(L^0\mathcal{I}^2)' q_s^{\tilde{\alpha},\tilde{\beta}} + \hat{C}(\bar{\Phi}_{T,s}\bar{D}_s)' x$$
$$+ \int_t^s \hat{C}(\bar{\Phi}_{\tau,s}\bar{D}_s)'(-\tilde{\Gamma} q_\tau^{\tilde{\alpha},\tilde{\beta}} + c_2 L^0\mathcal{I}^2(\tilde{\alpha}_\tau, \tilde{\beta}_\tau))\, d\tau, \ \text{ a.e. } s \in (t, T),$$

where $\quad q_s^{\tilde{\alpha},\tilde{\beta}} = \int_s^T \bar{\Phi}_{s,\tau}\bar{D}_\tau \hat{C}(\tilde{\alpha}_\tau, \tilde{\beta}_\tau)\, d\tau \qquad \forall s \in [t, T].$

- Our solution method:
  - Equivalent to a two-point boundary value problem.
  - Due to low-dimensionality of the reduced-complexity problem, there was no need to curse-of-dimensionality-free methods for examples so far considered.
  - This is a contraction for $(\tilde{\alpha}, \tilde{\beta})$ if T is sufficiently small.
  - That solution is efficiently propagated to larger T by a function-valued ODE (over fixed longer interval with $s, T \in (0, \bar{T})$.
  - Notably, we have guaranteed convergence to the unique correct solution.

Pointwise Solution Method Details (continued)

- The scalar-α case ($L^0 = 0,\ M^0 \in \boldsymbol{R}^{1\times n}$) yields fixed-point problem

$$\tilde{\alpha}_s^T - \frac{1}{c_1}\nabla\tilde{\Theta}(\tilde{\alpha}_s^T) = M^0\bar{\Phi}'_{T,s}x - c_1 M^0\Big[\int_0^s D'_{\sigma,s}(M^0)'\tilde{\alpha}_\sigma^T\,d\sigma + \int_s^T D_{s,\sigma}(M^0)'\tilde{\alpha}_\sigma^T\,d\sigma\Big] \qquad \forall\, s \in [t, T].$$

- The $\tilde{\alpha}_\cdot^T$ ODE is

$$\frac{d\tilde{\alpha}_s^T}{dT} = \big[1 - \frac{1}{c_1}\Theta''(\tilde{\alpha}_s^T)\big]^{-1}\Big\{M^0\bar{\Phi}'_{T,s}\bar{A}'x - c_1 M^0\bar{D}_{T,s}(M^0)'\tilde{\alpha}_T^T - c_1\int_0^T M^0\bar{D}_{\sigma,s}(M^0)'\frac{d\tilde{\alpha}_\sigma^T}{dT}\,d\sigma\Big\}.$$

- ODE solved on fixed longer interval with $s, T \in (0, \bar{T})$.

Pointwise Solution Method Details (continued)

- Propagated $\tilde{\alpha}$ solutions, scalar running-cost nonlinearity case:

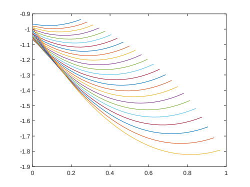

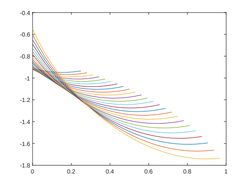

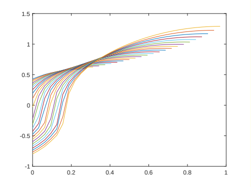

- A k-dimensional nonlinearity in an n-dimensional systemgenerates a $2k$-dimensional nonlinear control problem that may be solved with curse-of-dimensionality methods.

Stat Expansions: Graphs for Heuristics

- One has discrete-stat over finite sets of functions.
- Discrete stat-quad dual converges to the stat-quad dual as the density of the functions approaches the continuum.
- Nasty example below.

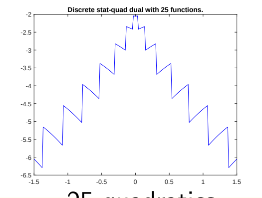

25 quadratics

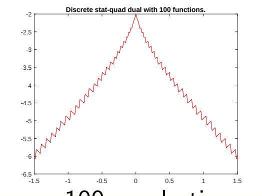

100 quadratics

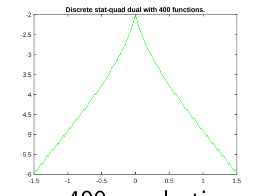

400 quadratics

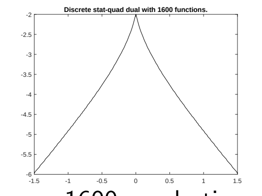

1600 quadratics

Example

- Space dimension: $n = 5$.
- One-dimensional nonlinearity in running cost; no nonlinearity in dynamics.
- $\tilde{\Theta}(\alpha, \beta) = \tilde{\Theta}(\alpha) = -k\sqrt{\varepsilon + \alpha^2}$. Yields a nonconvex running cost.
- Solution on any plane in 45 secs on a 5 y.o. surface pro ($\simeq$1500 grid points).
- Requires additional 150 secs to obtain backsubstitution errors in the original HJ PDE there.

Example

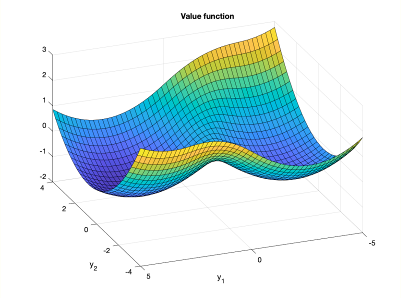

Figure: E1: Value function

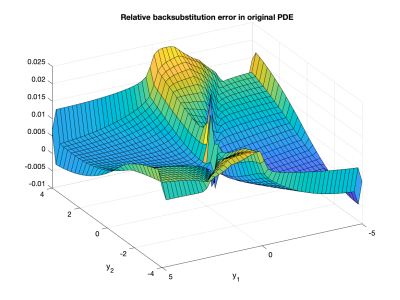

Figure: E2: Relative backsubs errors

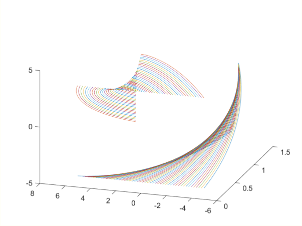

Figure: E3: Projections of state trajectories

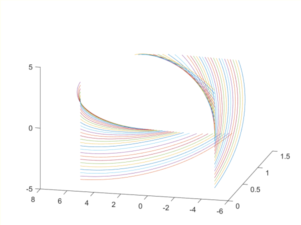

Figure: E4: Projections of state trajectories

Thank you.