49th IEEE Conference on Decision and Control
December 15-17, 2010
Hilton Atlanta Hotel, Atlanta, GA, USA

Improved Patchy Solution to the Hamilton-Jacobi-Bellman Equations

Thomas Hunt
Department of Mathematics
University of California - Davis
Davis, CA 95616
hunttw@math.ucdavis.edu

Arthur J. Krener
Department of Applied Mathematics
Naval Postgraduate School
Monterey, CA 93943-5216
ajkrener@nps.edu

***Abstract*— We test a new patch type for the patchy approximate solution to the Hamilton-Jacobi-Bellman equations, and we see an improvement in the worst relative error on a nonlinear test problem.**

I. Hamilton-Jacobi-Bellman PDEs

Consider the infinte horizon optimal control problem of minimizing the integral

$$\int_0^\infty l(x,u)dt \tag{1}$$

of a Lagrangian $l(x,u)$ subject to the controlled dynamics

$$\begin{aligned} \dot{x} &= f(x,u) \\ x(0) &= x^0 \end{aligned} \tag{2}$$

where f and l are both smooth, and l is strictly convex in $u \in \mathbb{R}^m$ for all $x \in \mathbb{R}^n$.

Suppose the dynamics and Lagrangian have Taylor expansions

$$\begin{aligned} \dot{x} &= Fx + Gu + \sum_{k=2}^\infty f^{[k]}(x,u) \\ l(x,u) &= \frac{1}{2}\big(x^TQx + u^TRu\big) + \sum_{k=3}^\infty l^{[k]}(x,u) \end{aligned} \tag{3}$$

where $^{[d]}$ denotes degree d terms in the power series. The optimal control problem is said to be *nice* if $Q \geq 0$, $R > 0$, (F,G) is stabilizable and $(Q^{1/2},F)$ is controllable.

A special case of the optimal control problem of (1),(2) is the linear-quadratic regulator (LQR) where one minimizes a quadratic cost

$$\int_0^\infty \frac{1}{2}\big(x^TQx + u^TRu\big)dt$$

subject to the linear dynamics

$$\dot{x} = Fx + Gu.$$

If the LQR problem is nice, then there is a unique nonnegative definite solution to the algebraic Riccati equaion

$$0 = F^TP + PF + Q - PGR^{-1}G^TP \tag{4}$$

that gives the optimal cost

$$\pi(x^0) = \frac{1}{2}\big(x^0\big)^TPx^0 = \min_u \frac{1}{2}\int_0^\infty \big(x^TQx + u^TRu\big)dt.$$

The optimal control can be expressed in feedback form

$$u(t) = \kappa(x(t)) = Kx(t)$$

where

$$K = -R^{-1}G^TP$$

and the closed loop dynamics

$$\dot{x} = \big(F + GK\big)x$$

is exponentially stable.

If the nonlinear optimal control problem has a smooth optimal cost $\pi(x)$ and a smooth optimal control control $u = \kappa(x)$ around $x = 0$, then it is widely known that they satisfy the Hamilton-Jacobi-Bellman (HJB) PDEs

$$\begin{aligned} 0 &= \min_u\Big(\frac{\partial\pi}{\partial x}(x)f(x,u) + l(x,u)\Big) \\ \kappa(x) &= \arg\min_u \Big(\frac{\partial\pi}{\partial x}(x)\frac{\partial f}{\partial u}(x,u) + \frac{\partial l}{\partial u}(x,u)\Big) \end{aligned} \tag{5}$$

If the LQR part of the problem is nice then the optimal cost is smooth in some neighborhood of the origin. Under the assumption that $\frac{\partial\pi}{\partial x}(x)f(x,u) + l(x,u)$ is strictly convex in u locally around $x = 0$, $u = 0$, then the HJB PDEs (5) can be rewritten as

$$\begin{aligned} 0 &= \frac{\partial\pi}{\partial x}(x)f(x,\kappa(x)) + l(x,\kappa(x)) \\ 0 &= \frac{\partial\pi}{\partial x}(x)\frac{\partial f}{\partial u}(x,\kappa(x)) + \frac{\partial l}{\partial u}(x,\kappa(x)). \end{aligned} \tag{6}$$

Al'brekht [1] has shown that the HJB PDEs for nice optimal control problems can be solved by Taylor series methods locally around the origin. Assume the dynamics and Lagrangian have power series expansions (3), and the optimal cost and optimal control have the expansions

$$\begin{aligned} \pi(x) &= \frac{1}{2}x^TPx + \pi^{[3]}(x) + \pi^{[4]}(x) + \cdots \\ \kappa(x) &= Kx + \kappa^{[2]}(x) + \kappa^{[3]}(x) + \cdots \end{aligned} \tag{7}$$

If we plug the expansions of (7) into (6) and collect the lowest order terms, then we obtain

$$\begin{aligned} 0 &= x^T\big(F^TP + PF + Q - K^TRK\big)x \\ 0 &= x^T\big(PG + K^TR\big). \end{aligned}$$

We can solve the second equation for K in terms of P, and substitute into the first equation to arrive at the familiar algebraic Riccati equation of (4). K can then be computed from the second equation once P is known.

After determining P and K, we can extract the next lowest terms from (6) to arrive at the new equations

$$\begin{aligned} 0 &= \frac{\partial \pi^{[3]}}{\partial x}(x)\big(F+GK\big)x + x^T P f^{[2]}(x,Kx) \\ &\quad + l^{[3]}(x,Kx) \\ 0 &= \frac{\partial \pi^{[3]}}{\partial x}(x)G + x^T P \frac{\partial f^{[2]}}{\partial x}(x,Kx) \\ &\quad + \frac{\partial l^{[3]}}{\partial x}(x,Kx) + \big(\kappa^{[2]}(x)\big)^T R. \end{aligned} \tag{8}$$

Note that the unknowns $\pi^{[3]}(x)$ and $\kappa^{[2]}(x)$ appear linearly, the equations are triangular, and R is invertible, so we can solve the second equation for $\kappa^{[2]}(x)$ if we can solve the first equation for $\pi^{[3]}(x)$.

To determine the solvability of the first equation for $\pi^{[3]}(x)$ in (8), consider the linear operator

$$\pi^{[3]}(x) \mapsto \frac{\partial \pi^{[3]}}{\partial x}(x)\big(F+GK\big)x \tag{9}$$

that maps cubic polynomials to cubic polynomials. Its eigenvalues are of the form $\lambda_i + \lambda_j + \lambda_k$ where λ_i, λ_j, λ_k are eigenvalues of $F+GK$, thus it is invertible since the eigenvalues of $F+GK$ lie in the open left half of the complex plane.

The higher degree terms of the series expansions of $\pi(x)$ and $\kappa(x)$ can be found in a similar fashion by plugging in the expansions of $\pi(x)$ and $\kappa(x)$ up to degree $d+1$ and d into (6) and setting the terms of degree $d+1$ and d equal to zero.

This process can be carried out to compute polynomial approximations of arbitrary degree to the optimal cost and optimal control if f and l are sufficiently smooth. However, in practice, the computation is limited by memory and computational speed. There are $n+d-1$ choose d monomials of degree d in n variables, so the problem size exhibits superpolynomial growth in the degree d for fixed state space dimension n.

A limitation of the power series approach to solving the HJB PDEs is its local in nature, meaning the series solution is close to the true solution in a neighborhood of the origin of the state space. Increasing the degree of the approximating polynomials may only increase the accuracy in a neighborhood without enlarging the domain of validity. Furthermore, the HJB PDEs need not have globally smooth solutions. The underlying optimal control problem may have conjugate or focal points, and it is for this reason that the theory of viscosity solutions was developed [5], [4].

II. The patchy method

Krener and Navasca have proposed a patchy method [3], which is an extension of Al'brekht's method [1], the Cauchy-Kovalevskaya technique [6], the method of Tsitsiklis [7] and the patchy method of Ancona and Bresnan [2], which we review here.

To simplify notation, we assume that the control is one dimensional ($m=1$), and the dynamics and Lagrangian take on the form

$$\begin{aligned} f(x,u) &= f(x) + g(x)u \\ l(x,u) &= q(x) + \frac{1}{2} r(x) u^2. \end{aligned} \tag{10}$$

Suppose we have computed $\pi^0(x)$ and $\kappa^0(x)$, the approximate optimal cost and optimal control in a neighborhood of the origin. We find a sublevel set $\{x \mid \pi^0(x) \le c\}$ on which the error of $\pi^0(x)$ is acceptable, where we use the relative HJB error

$$\mathrm{err}^{\mathrm{HJB}}_{\mathrm{rel}} \equiv \frac{\frac{\partial \pi^i}{\partial x}(x) f(x,\kappa^i(x)) + l(x,\kappa^i(x))}{\pi^0(x)} \tag{11}$$

to measure the error of the approximate optimal cost and optimal control. The HJB error measures how much the approximate optimal cost and optimal control fail to satisfy the HJB PDEs. We choose x^1 on the level set $\pi^0(x) = c$, and seek to extend the solution to the HJB PDEs in a *patch* containing x_1. To do so, we approximate each partial derivative of the optimal cost and optimal control up to some prescribed order at x_1, and determine which surfaces bound the new patch.

Let $\pi^1(x)$ and $\kappa^1(x)$ denote the new optimal cost and optimal control approximating polynomials that are centered at x^1. The coefficients of these polynomials are computed in order from lowest to highest degree. Under the assumption of (10), the HJB PDEs becomes

$$\begin{aligned} 0 &= \frac{\partial \pi}{\partial x_\sigma}(x)\big(f_\sigma + g_\sigma \kappa(x)\big) \\ &\quad + q(x) + r(x) + \frac{1}{2}\big(\kappa(x)\big) \\ 0 &= \frac{\partial \pi}{\partial x_\sigma}(x) g_\sigma(x) + r(x)\kappa(x). \end{aligned} \tag{12}$$

where we invoke the summation convention that a σ subscript indicates summation over all values of the index unless explicitly indicated otherwise.

First, we determine the *characteristic* index, which is the index k that maximizes the quantity

$$|f_k(x^1) + g_k(x^1)\kappa^0(x^1)|.$$

For notational convenience, we assume that $k=n$. We compute the constant term of π^1 and its *noncharacteristic* first order partial derivatives by inheritance, meaning we set

$$\begin{aligned} \pi^1(x^1) &= \pi^0(x^1) \\ \frac{\partial \pi^1}{\partial x_i}(x^1) &= \frac{\partial \pi^0}{\partial x_i}(x^1) \end{aligned}$$

for $1 \le i \le n-1$. We can solve the second HJB equation (12) for $\kappa(x^1)$, substitute it into the first, and arrive at a quadratic equation in the *characteristic* first order partial derivative of π^1 at x^1

$$0 = a\Big(\frac{\partial \pi^1}{\partial x_n}(x^1)\Big)^2 + b\Big(\frac{\partial \pi^1}{\partial x_n}(x^1)\Big) + c$$

where the coefficients a, b, and c depend on the problem data f, g, q, and r as well as the noncharacteristic first order partial derivatives $\frac{\partial \pi^1}{\partial x_i}(x^1)$, $1 \leq i \leq n-1$. Assuming this quadratic equation has real roots, we set $\frac{\partial \pi^1}{\partial x_1}(x^1)$ to be the root closest to $\frac{\partial \pi^0}{\partial x_1}(x^1)$, and then solve the second equation of (12) for $\kappa^1(x^1)$.

The next unknown coefficients in the power series expansions of the optimal cost and optimal control centered at x^1 are $\frac{\partial^2 \pi^1}{\partial x_i \partial x_j}(x^1)$ and $\frac{\partial \kappa^1}{\partial x_i}(x^1)$ for $1 \leq i \leq j \leq n$. We compute the noncharacteristic second order partial derivatives of the optimal cost by inheritance, that is

$$\frac{\partial^2 \pi^1}{\partial x_i \partial x_j}(x^1) = \frac{\partial^2 \pi^0}{\partial x_i \partial x_j}(x^1)$$

for $1 \leq i \leq j \leq n-1$. We still need to determine the remaining n characteristic second order partial derivatives of π^1 at x^1, as well as the n first order partial derivatives of κ^1 at x^1. To get formulas for these unknowns, we take the partial derivative of both equations in (12) with respect to x_i and get $2n$ equations.

$$\begin{aligned}
0 = &\frac{\partial^2 \pi^1}{\partial x_i \partial x_\sigma}(x^1)\big(f_\sigma(x^1) + g_\sigma(x^1)\kappa^1(x^1)\big) \\
&+ \frac{\partial \pi^1}{\partial x_\sigma}(x^1)\left(\frac{\partial f_\sigma}{\partial x_i}(x^1) + \frac{\partial g_\sigma}{\partial x_i}(x^1)\kappa^1(x^1)\right) \\
&+ \frac{\partial q}{\partial x_i}(x^1) + \frac{1}{2}\frac{\partial r}{\partial x_i}(x^1)\big(\kappa^1(x^1)\big)^2 \qquad (13) \\
0 = &\frac{\partial^2 \pi^1}{\partial x_i \partial x_\sigma}(x^1) g_\sigma(x^1) + \frac{\partial \pi^1}{\partial x_\sigma}(x^1)\frac{\partial g_\sigma}{\partial x_i}(x^1) \\
&+ \frac{\partial r}{\partial x_i}(x^1)\kappa^1(x^1) + r(x^1)\frac{\partial \kappa^1}{\partial x_i}(x^1)
\end{aligned}$$

The unknowns $\frac{\partial \kappa^1}{\partial x_i}(x^1)$ fall out of the first equation in (13) because of (12). Furthermore, the first n equations decouple and can be solved for the unknowns $\frac{\partial^2 \pi^1}{\partial x_i \partial x_n}(x^1)$ one by one by iterating over i from 1 to n. We can compute $\frac{\partial \kappa^1}{\partial x_i}(x^1)$ for $1 \leq i \leq n$ from the second set of n equations in (13) once all the second order partial derivatives of π^1 have been computed at x^1.

We next find the third order partials of π^1 at x^1. We compute the noncharacteristic partial derivatives of π^1 by inheritance, so we set

$$\frac{\partial^3 \pi^1}{\partial x_i \partial x_j \partial x_k}(x^1) = \frac{\partial^3 \pi^0}{\partial x_i \partial x_j \partial x_k}(x^1)$$

for $1 \leq i \leq j \leq k \leq n-1$. We derive equations for the remaining characteristic third order partial derivatives of π^1 by taking the partial derivative of the first equation in (13) with respect to x_j. This yields $(n+1)n/2$ equations in the unknowns $\frac{\partial^3 \pi^1}{\partial x_i \partial x_j \partial x_n}(x^1)$, $1 \leq i \leq j \leq n$, which can be solved one by one in lexicographical order. Once we have computed all the third order partial derivatives of π^1 at x^1, we can compute all the second order partial derivatives of κ^1 at x^1 by taking the derivative of the second equation in (13) with respect to x_j, which yields $n(n+1)/2$ equations for the $n(n+1)/2$ unknowns $\frac{\partial^2 \kappa^1}{\partial x_i \partial x_j}(x^1)$, where $1 \leq i \leq j \leq n$.

We compute the fourth order partial derivatives of π^1 in the same way we computed its third order partial derivatives at x^1. We compute the noncharacteristic partial derivatives $\frac{\partial^4 \pi^1}{\partial x_i \partial x_j \partial x_k \partial x_\sigma}(x^1)$, $1 \leq i \leq j \leq k \leq l \leq n-1$ by inheritance, and then derive two sets of triangular equations for the remaining fourth order characteristic partial derivatives of π^1 and all the third order partial derivatives of κ^1 at x^1. We first solve for the remaining fourth order characteristic partial derivatives of π^1 at x^1 one by one in lexicographical order, and then solve for the third order partials of κ^1 at x^1.

After we have computed polynomial approximations to the optimal cost and optimal control along the level surface $\pi^0(x) = c$, we subdivide part of the state space into a set of nonoverlapping patches where the optimal control and optimal cost are approximated at every point in the patch with a single pair of optimal cost and optimal control approximating polynomials. All patches, except the patch where the Al'brekht solution is valid, are defined by four intersecting boundaries when $n = 2$.

It is an open question on how to shape the patches to reduce the error of the patchy solution. We consider two types of patches, one new and one old, that are differentiated by how their boundaries are constructed. The preliminary numerical evidence suggests that the new patch type leads to lower error in the patchy optimal cost. For the sake of definiteness, assume that we are constructing a patch that is adjacent to the existing Al'brekht patch (the patch that contains $x = 0$), and that $\pi^1(x)$ and $\kappa^1(x)$ are the optimal cost and optimal control approximating polynomials associated with the patch that is under construction. For a patch of each type, two of its boundaries are defined by the level curves $\pi^0(x) = c$ and $\pi^1(x) = d$, where we choose d so that the error inside the patch is acceptable. We choose the remaining two *lateral* boundaries to be line segments that are anchored to the level curve $\pi^0(x) = c$ at two *patch boundary points* on the curve midway between the patch point x^1 and its two neighboring patch points. The two lateral boundaries of the old patch type are line segments that coincide with the the two lines that emanate from $x = 0$ and pass through the two patch boundary points. The new patch type improves on the first by orienting the lateral boundaries in the direction of the optimal trajectory with dynamics $\dot{x} = f(x, \kappa^0(x))$. The error may be further reduced by altering the lateral boundaries so that they are invariant manifolds that coincide with optimal state trajectories. This requires further investigation.

III. TEST PROBLEM

As a test for the patchy method, we consider the optimal control problem of minimizing

$$\pi(x^0) = \inf_u \frac{1}{2}\int_0^\infty \left(\sin^2(x_1) + \left(x_2 - \frac{1}{3}x_1^3\right)^2 + u^2\right) dt$$

subject to the dynamics

$$\begin{aligned}\dot{x}_1 &= \left(x_2 - \frac{1}{3}x_1^3\right)\sec(x_1)\\ \dot{x}_2 &= \left(x_1^2 x_2 - \frac{1}{3}x_1^5\right)\sec(x_1) + u.\end{aligned}$$

This problem was derived from the linear-quadratic regulator problem of minimizing the quadratic cost

$$\pi(y^0) = \inf_u \frac{1}{2}\int_0^\infty (y^T y + u^2)dt$$

subject to the linear dynamics

$$\begin{aligned}\dot{y}_1 &= y_2\\ \dot{y}_2 &= u.\end{aligned}$$

The linear quadratic optimal control problem is transformed to the nonlinear optimal control problem by the change of coordinates $y_1 = \sin(x_1)$ and $y_2 = x_2 - x_1^3/3$.

We computed the patchy optimal cost and optimal control for test problem for both patch types. We used 73 patches of each type and the same levels to define the optimal cost patch boundaries.

It took approximately 3 seconds to compute the patchy solution for both patch types on a 1.33 GHz iBook G4 with 1 GB of memory running Matlab 7.1.0.21 (R14) Service Pack 3 for this test problem.

Figure 2 displays the true and patchy optimal cost for the new patch type.

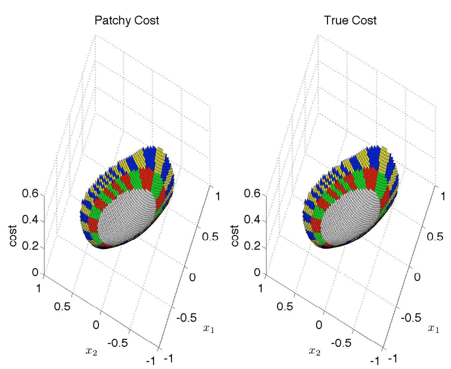

1: Patchy and true optimal cost, new patch type

Figures 3 and 4 display the relative error and the relative HJB error for the patchy optimal cost for both patch types. The relative error in each patch is measured in the typical way

$$\text{err}_{\text{rel}} \equiv \frac{\pi(x) - \pi^i(x)}{\pi(x)}$$

and the HJB error is defined in (11). The figures indicate that the worst relative error in the patchy optimal cost for the new patch type is significantly less than the error for the old patch type.

Note that regions of high relative HJB error correspond to regions of high relative error. This suggests that the HJB

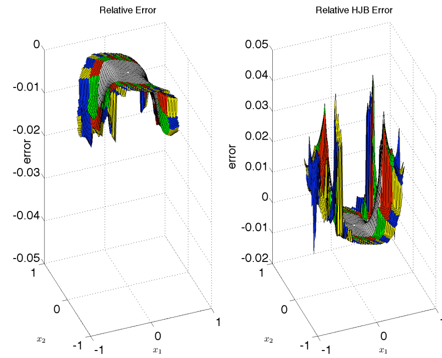

2: Relative error, new patch type

error, which can be computed in the realistic setting where the true optimal cost is unknown, serves as a good proxy for the relative error.

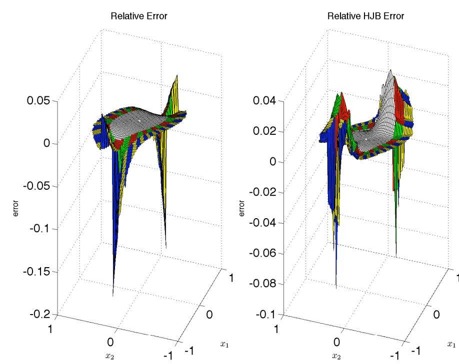

3: Relative error, old patch type

IV. Conclusion

We have applied the patchy method of Krener and Navasca [3] with a new patch type whose lateral boundaries are oriented in the direction of optimal state trajectories. We observed a significant reduction in the worst relative error by employing the new patch type. Given the improvement of the new patch type over the old, we propose using optimal state trajectories as lateral boundaries as a natural next step. This new patch type merits further study.

References

[1] E. G. Al'brekht. On the optimal stabilization of nonlinear systems. *Journal of Applied Mathematics and Mechanics*, 25(5):836 – 844, 1961.

[2] Fabio Ancona and Alberto Bressan. Nearly time optimal stabilizing patchy feedbacks. *Annales de l'Institut Henri Poincare (C) Non Linear Analysis*, 24(2):279 – 310, 2007.

[3] Carmeliza Navasca and Arthur J. Krener, Patchy Solutions of Hamilton Jacobi Bellman Partial Differential Equationst *Modeling, Estimation and Control: Festschrift in Honor of Giorgio Picci on the Occasion of his Sixty-Fifth Birthday*, Alessandro Chiuso, Augusto Ferrante, and Stefano Pinzoni, editors. Springer-Verlag, 2007.

[4] Michael G. Crandall and Pierre-Louis Lions. Viscosity solutions of hamilton-jacobi equations. *Transactions of the American Mathematical Society*, 277(1):1–42, 1983.

[5] Lawrence C. Evans. *Partial Differential Equations*. American Mathematical Society, 1998.

[6] Fritz John. *Partial Differential Equations*. Springer-Verlag, 4th edition, 1982.
[7] John N. Tsitsiklis. Efficient algorithms for globally optimal trajectories. *IEEE Transactions on Automatic Control*, 40:1528–1538, 1995.