Scalable Calculation of Reach Sets and Tubes for Nonlinear Systems with Terminal Integrators*

A Mixed Implicit Explicit Formulation

Ian M. Mitchell
Department of Computer Science
University of British Columbia
2366 Main Mall, Vancouver, BC, Canada
mitchell@cs.ubc.ca

ABSTRACT

The solution of a particular Hamilton-Jacobi (HJ) partial differential equation (PDE) provides an implicit representation of reach sets and tubes for continuous systems with nonlinear dynamics and can treat inputs in either worst-case or best-case fashion; however, it can rarely be determined analytically and its numerical approximation typically requires computational resources that grow exponentially with the state space dimension. In this paper we describe a new formulation—also based on HJ PDEs—for reach sets and tubes of systems where some states are *terminal integrators*: states whose evolution can be written as an integration over time of the other states. The key contribution of this new mixed implicit explicit (MIE) scheme is that its computational cost is linear in the number of terminal integrators, although still exponential in the dimension of the rest of the state space. Application of the new scheme to four examples of varying dimension provides empirical evidence of its considerable improvement in computational speed.

Categories and Subject Descriptors

J.2 [**Physical Sciences & Engineering**]: Engineering; I.6.4 [**Simulation & Modeling**]: Model Validation & Analysis

General Terms

Verification, Algorithms

Keywords

nonlinear systems, continuous reachability, Hamilton-Jacobi equations, optimal control

1. INTRODUCTION

Reachability has proved a powerful tool in the verification of discrete systems, but it is still impractical for most hybrid and continuous systems because there is as yet no scalable method of computing reachability for nonlinear continuous systems in high dimensions. It was shown in previous work [8, 10] that an implicit representation of the backwards reach tube of nonlinear systems with adversarial inputs is the viscosity solution of a particular Hamilton-Jacobi (HJ) partial differential equation (PDE). While this formulation permits the representation of nonconvex sets and can treat inputs in either a best-case (for control) or worst-case (for robustness) manner, its practical implementation so far requires computational resources exponential in the dimension of the state space.

The contribution of this paper is a new formulation based on HJ PDEs that leads to a significant reduction in computational complexity for systems with dynamics of a certain form. We divide the state space into two sets of states:

- *Coupled* states, which have the same weak restrictions on their dynamics and are treated in the same manner as in the traditional implicit formulation, but consequently have computational cost that scales exponentially with dimension.
- *Terminal integrator* states, whose dynamics can only depend on the coupled states but which can be treated in a more efficient manner that requires computational resources linear in dimension. The key restriction on terminal integrator states is that no other states may depend on their value(s).

Terminal integrator states often appear in dynamic models of mechanical systems, where position states are simply integrators for velocity states. We call the new scheme a *mixed implicit explicit* (MIE) formulation because the boundary of the reach set or tube is represented implicitly in the coupled dimensions but explicitly (in the form of intervals) in the terminal integrator dimensions.

The remainder of the paper is organized as follows. Section 2 discusses the continuous reachability problem that we seek to solve, the traditional HJ PDE based implicit formulation upon which we build, and related work. Section 3 provides the new formulation for the case of a single terminal integrator state, proves that it can be used to determine the backwards maximal reach set, and outlines how it can be modified to determine other reach sets and tubes. Section 4 demonstrates the application of the new scheme on two examples. Section 5 discusses a generalization of the terminal integrator's dynamics which has less theoretical support, but which is demonstrated on a third example. Section 6 extends the simpler dynamics to the vector case, and demonstrates it on a fourth example. Finally, section 7 compares the implicit and MIE formulations.

2. CONTINUOUS REACHABILITY

This paper concerns itself with the calculation of backwards reach sets and tubes for deterministic nonlinear continuous systems with

*Research supported by a Discovery Grant from the National Science and Engineering Research Council of Canada

HSCC'11, April 12–14, 2011, Chicago, Illinois, USA.

dynamics given by the ordinary differential equation (ODE)

$$\dot{z} = f(z, u) \tag{1}$$

for state variables $z \in \mathbb{R}^{d_z}$, input $u \in \mathcal{U}$ where $\mathcal{U}$ is compact and convex, and nonlinear dynamics $f : \mathbb{R}^{d_z} \times \mathcal{U} \to \mathbb{R}^{d_z}$ which are bounded and Lipschitz continuous in z for fixed u. Backwards reachability seeks to determine the set of states which are initial conditions to trajectories of (1) that intersect a specified target set $\mathcal{R}_0$. The reach set is the set of initial states which give rise to trajectories whose endpoints at a specified time lie within the target set, while the reach tube is the set of initial states which give rise to trajectories which arrive at or pass through the target set during a specified time interval; typically the interval is from time zero to a specified time.

Two methods of treating inputs to the system dynamics are outlined in [9]: Maximal reachability uses the inputs to make the reach set or tube as large as possible, while minimal reachability uses the inputs to make the reach set or tube as small as possible. For example, we can formalize the backward maximal reach set as

$$\mathcal{R}(\mathcal{R}_0, t) = \{z_s \mid \exists u(\cdot), \exists z_f \in \mathcal{R}_0, z(t) = z_f\}, \tag{2}$$

and the backward minimal reach tube as

$$\mathcal{R}(\mathcal{R}_0, t) = \{z_s \mid \forall u(\cdot), \exists z_f \in \mathcal{R}_0, \exists \sigma \in [t, 0], z(\sigma) = z_f\}, \tag{3}$$

where $t < 0$ and trajectory $z(\cdot)$ solves (1) with initial conditions $z(t) = z_s$ and input signal $u(\cdot)$. It is also possible to formulate versions with adversarial inputs, in which some inputs seek to maximize the size of the reach set or tube, while others seek to minimize it. Formalization of these versions is complicated by the need to consider the knowledge available to the competing players when choosing their inputs; consequently, we do not further pursue the adversarial case here.

2.1 Related Work

Computation of reachable sets for deterministic nonlinear continuous systems remains a challenge despite several decades of research. The methods discussed in this paper are based on the Hamilton-Jacobi(-Bellman)(-Isaacs) equation [8, 7, 10], as explained in the next section. These schemes, as well as the very closely related schemes arising from viability theory, such as [3, 13], share several important characteristics. On the positive side, they are designed to deal directly with nonlinear dynamics and are able to automatically determine optimal inputs in both the best-case and the worst-case senses. On the negative side, their implementation typically requires computational resources exponential in the number of state space dimensions.

There have been a number of attempts to reduce this computational complexity. For systems in which the dynamics decouple, the high dimensional PDE can be broken naturally into multiple, lower-dimensional PDEs. Complete decoupling of the dynamics occurs rarely, so a projection-based scheme whereby coupling terms could be treated as disturbances was proposed in [12] (see section 6 for a discussion of how that approach can be applied to the class of dynamics studied in this paper), and a time-based decoupling where the system dynamics could be separated into fast and slow components was proposed in [6]. The exponential cost of HJ PDE based schemes arises from the requirement in traditional PDE solvers to grid the state space, so [4] proposes to reduce that cost by using a neural network representation of the solution instead.

There is also a vast body of research on other schemes for approximating reachability and/or solving related verification problems in continuous and hybrid systems; so vast that we cannot survey it here. Previous proceedings of the *Hybrid Systems: Computation and Control* workshop provide many suitable entries into the literature for the interested reader.

2.2 The Traditional Implicit Formulation

A method for determining an implicit representation of the backward reach tube of nonlinear dynamic systems with adversarial inputs was proposed in [8]. An implicit representation of the reach set or tube takes the form of a function $\psi : \mathbb{R} \times \mathbb{R}^{d_z} \to \mathbb{R}$ whose zero sublevel set is the reach set or tube:

$$\mathcal{R}(\mathcal{R}_0, t) = \{z \mid \psi(t, z) \leq 0\}. \tag{4}$$

In [7] it was shown for a minimal reach tube (3) that ψ is the solution of the terminal value HJ PDE

$$D_t\psi(t, z) + H\left(t, z, D_z\psi(t, z)\right) = 0 \tag{5}$$

with Hamiltonian

$$H_{\text{tube}}(t, z, p) = \min\left[0, \max_{u \in \mathcal{U}} \left(p \cdot f(z, u)\right)\right], \tag{6}$$

and terminal conditions

$$\psi(0, z) = \psi_0(z) \tag{7}$$

such that $\mathcal{R}_0 = \{z \mid \psi_0(z) \leq 0\}$.

This result was extended to reach tubes for systems with two adversarial inputs (and therefore maximal reach tubes as well) in [10]. It is straightforward to adapt these PDEs to compute reach sets with various treatments of the inputs. For example, an implicit representation of the maximal reach set (2) is given by the solution of the HJ PDE (5) with Hamiltonian

$$H_{\text{set}}(t, z, p) = \min_{u \in \mathcal{U}} \left(p \cdot f(z, u)\right).$$

and terminal conditions (7).

For most systems of interest, (5) cannot be solved analytically. Typical approximation schemes involve creating a grid over the state space $\mathbb{R}^{d_z}$, and hence are only practical for $d_z \lesssim 4$.

3. THE MIXED IMPLICIT EXPLICIT FORMULATION

In this section we outline a new formulation for computing reachability that involves solving PDEs of dimension lower than (5). We derive the results for the backward maximal reach set (2), although the modifications to compute the other sets and tubes are straightforward.

Consider a system where the state can be decomposed as $z = (y, x)$ and the dynamics (1) as

$$\dot{y} = f(y, u) \tag{8}$$

$$\dot{x} = b(y) \tag{9}$$

where $y \in \mathbb{R}^{d_{\text{cp}}}$ are the *coupled* state variables, $x \in \mathbb{R}$ is the scalar *terminal integrator* state variable, and $u \in \mathcal{U}$ is the input as above. We adopt this particular choice of variables because the terminal integrator state(s) often represent positions and are hence typically plotted horizontally. The functions f and b are assumed to be Lipschitz continuous and bounded.

Instead of the fully implicit representation of the reach set (4), we will adopt the MIE representation

$$\mathcal{R}(\mathcal{R}_0, t) = \{(y, x) \mid \underline{\phi}(t, y) \leq x \leq \overline{\phi}(t, y)\} \tag{10}$$

with target set

$$\mathcal{R}_0 = \{(y, x) \mid \underline{\phi}_0(y) \leq x \leq \overline{\phi}_0(y)\}. \tag{11}$$

Consider first the upper boundary of the reach set. Let $\overline{x}(t, y)$ be the upper boundary of $\mathcal{R}(\mathcal{R}_0, t)$ at time t and coupled state y; in other words $\overline{x}(t, y) = \overline{\phi}(t, y)$. Formally by (9) and the chain rule

$$b(y) = \frac{d}{dt}\overline{x}(t, y) = \frac{d}{dt}\overline{\phi}(t, y) = D_t\overline{\phi}(t, y) + D_y\overline{\phi}(t, y) \cdot f(y, u). \quad (12)$$

Rearranging, we arrive at a Hamilton-Jacobi equation

$$D_t\overline{\phi}(t, y) + D_y\overline{\phi}(t, y) \cdot f(y, u) - b(y) = 0,$$

whose form is identical to those that arise in finite horizon optimal control problems.

The derivation above is entirely formal, but the finite horizon optimal control interpretation provides rigourous support for the resulting formulation. The terminal integrator's dynamics in (9) can be rewritten for $t \leq 0$ as

$$x(0, y(0)) = x(t, y(t)) + \int_t^0 b(y(s))ds \quad (13)$$

which is simply rearranged into the form of a finite horizon cost function

$$x(t, y(t)) = \int_t^0 -b(y(s))ds + x(0, y(0)) \quad (14)$$

with running cost $-b(y)$ and terminal cost $x(0, y)$ for trajectories $y(\cdot)$ generated by (8). Setting the terminal cost to be $x(0, y) = \overline{\phi}_0(y)$, the value function representing the maximum cost to go from each coupled state y at time $t \leq 0$ is the viscosity solution to the terminal value Hamilton-Jacobi-Bellman equation

$$D_t\overline{\phi}(t, y) + H\left(t, y, D_y\overline{\phi}(t, y)\right) = 0 \quad (15)$$

with Hamiltonian

$$\overline{H}_{\text{set}}(t, y, p) = \max_{u \in \mathcal{U}} \left(p \cdot f(y, u) - b(y)\right), \quad (16)$$

and terminal condition

$$\overline{\phi}(0, y) = \overline{\phi}_0(y). \quad (17)$$

Proposition 1. *For a coupled state $\hat{y}$ in the reach set, let $\overline{x}(t, \hat{y})$ be the upper boundary of the interval which is the reach set's projection at time $t < 0$ onto the terminal integrator dimension at that $\hat{y}$*

$$\overline{x}(t, \hat{y}) = \max\{x \mid (\hat{y}, x) \in \mathcal{R}(\mathcal{R}_0, t)\}$$

The solution $\overline{\phi}(t, y)$ of (15)–(17) provides $\overline{x}$

$$\overline{x}(t, \hat{y}) = \overline{\phi}(t, \hat{y}). \quad (18)$$

Note that $\overline{x}(\cdot, \cdot)$ is not assumed to be a trajectory of the system.

Proof. For a system with dynamics (8), initial condition $y(\tau_0) = \hat{y}$ and cost function

$$\gamma(\tau_0, \hat{y}, u(\cdot)) = \int_{\tau_0}^{\tau_f} \beta(y(\sigma), u(\sigma))d\sigma + \rho(y(\tau_f))$$

we can define the value function

$$\nu(\tau_0, \hat{y}) = \sup_{u(\cdot)} \gamma(\tau_0, \hat{y}, u(\cdot)).$$

The value function is the viscosity solution of the terminal value Hamilton-Jacobi-Bellman PDE (see, for example, [5, 2])

$$D_\tau\nu(\tau, y) + H(\tau, y, \nu(\tau, y)) = 0 \quad (19)$$

with Hamiltonian

$$H(t, y, p) = \max_{u \in \mathcal{U}} \left(p \cdot f(y, u) + \beta(y, u)\right) \quad (20)$$

and terminal condition

$$\nu(\tau_f, y) = \rho(y). \quad (21)$$

Set

$$\begin{aligned} \beta(y, u) &= -b(y), & \rho(y) &= \overline{\phi}_0(y), \\ \tau_0 &= t < 0, & \tau_f &= 0, \end{aligned} \quad (22)$$

so that (19)–(21) are the same as (15)–(17) and $\nu(t, y) = \overline{\phi}(t, y)$.

For any $\epsilon > 0$, we can pick a $\hat{u}(\cdot)$ such that

$$\gamma(t, \hat{y}, \hat{u}(\cdot)) \leq \nu(t, \hat{y}) < \gamma(t, \hat{y}, \hat{u}(\cdot)) + \epsilon \quad (23)$$

and let $y(\cdot)$ be the solution of (8) arising from $\hat{u}(\cdot)$ with initial condition $y(t) = \hat{y}$. Consider a trajectory $x^{(1)}(\cdot)$ solving (9) along $y(\cdot)$ with initial conditions $x^{(1)}(t)$ (we drop the dependence of the terminal integrator's trajectory $x(\cdot)$ on the coupled states' trajectory $y(\cdot)$ because only one coupled state trajectory is considered in the remainder of the proof). Let $x^{(1)}(0) = \overline{x}(0, y(0)) = \overline{\phi}_0(y(0))$. By (14) and (22),

$$\begin{aligned} x^{(1)}(t) &= \int_t^0 -b(y(\sigma))d\sigma + x^{(1)}(0), \\ &= \int_{\tau_0}^{\tau_f} \beta(y(\sigma), u(\sigma))d\sigma + \rho(y(\tau_f)), \\ &= \gamma(t, \hat{y}, u(\cdot)), \end{aligned}$$

and consequently by (23)

$$x^{(1)}(t) \leq \nu(t, \hat{y}) < x^{(1)}(t) + \epsilon.$$

Letting $\epsilon \to 0$, we find $x^{(1)}(t) = \nu(t, \hat{y}) = \overline{\phi}(t, \hat{y})$.

Now consider a second trajectory $x^{(2)}(\cdot)$ solving (9) along $y(\cdot)$ with initial conditions $x^{(2)}(t) > x^{(1)}(t) = \overline{\phi}(t, \hat{y})$. By (13)

$$\begin{aligned} x^{(2)}(0) &= x^{(2)}(t) + \int_t^0 b(y(s))ds \\ &> x^{(1)}(t) + \int_t^0 b(y(s))ds = x^{(1)}(0) = \overline{\phi}(0, y(0)), \end{aligned} \quad (24)$$

which implies $(x^{(2)}(0), y(0)) \notin \mathcal{R}_0$, which implies that $(x^{(2)}(t), \hat{y}) \notin \mathcal{R}(\mathcal{R}_0, t)$.

Conversely, consider that trajectory $x^{(2)}(\cdot)$ starts from $x^{(2)}(t) < x^{(1)}(t)$. Reversing the inequality in (24), it can be shown that $(x^{(2)}(t), \hat{y}) \in \mathcal{R}(\mathcal{R}_0, t)$.

Since states $(x^{(2)}(t), \hat{y})$ are in the reach set if and only if $x^{(2)}(t) \leq x^{(1)}(t)$, we conclude that $\overline{x}(t, \hat{y}) = x^{(1)}(t) = \overline{\phi}(t, \hat{y})$. □

Following the derivation in [10], the reach tube can be computed by adjusting the Hamiltonian so that the computed set only grows; in this case solve (15) with Hamiltonian

$$\overline{H}_{\text{tube}}(t, y, p) = \max\left[0, \max_{u \in \mathcal{U}} \left(p \cdot f(y, u) - b(y)\right)\right].$$

The derivation for the lower boundary function $\underline{\phi}(t, y)$ of the reach set is similar except that using the input to maximize the size of the reach set requires using the input to minimize the height of the lower boundary, and in the case of a reach tube the lower boundary should only decrease as time goes backwards. Therefore, we need to use Hamiltonians

$$\begin{aligned} \underline{H}_{\text{set}}(t, y, p) &= \min_{u \in \mathcal{U}} \left(p \cdot f(y, u) - b(y)\right) \\ \underline{H}_{\text{tube}}(t, y, p) &= \min\left[0, \min_{u \in \mathcal{U}} \left(p \cdot f(y, u) - b(y)\right)\right] \end{aligned}$$

in (15) for the reach set and tube respectively, and terminal condition $\underline{\phi}(0, y) = \underline{\phi}_0(y)$ instead of (17).

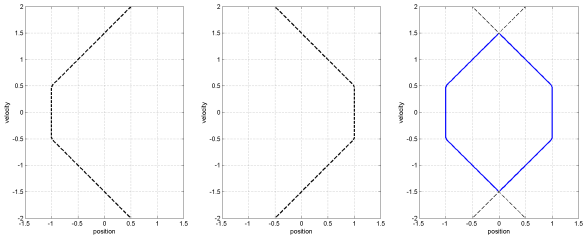

Figure 1: Terminal conditions for the double integrator. The axes are position x and velocity y. Left: $\underline{\phi}_0(y)$. Middle: $\overline{\phi}_0(y)$. Right: $\mathcal{R}_0$ is the region *outside* the solid curve.

Following [10], we can similarly compute minimal backward reachability or even backward reachability under adversarial inputs by appropriate adjustment of the sense of the optimization over inputs in the Hamiltonian.

4. SCALAR EXAMPLES

We demonstrate the MIE scheme on some examples in which the terminal integrator state is a scalar.

4.1 Computational Setting

In the examples that follow, the solution of the HJ PDEs are approximated in MATLAB with the Toolbox of Level Set Methods [11] (ToolboxLS). The numerical schemes used by ToolboxLS require that solutions remain continuous (although not necessarily differentiable). Under the assumptions placed on the dynamics, viscosity solution theory ensures that the solutions of (5) and (15) will remain continuous provided that their terminal condition functions are continuous; consequently, we restrict the target sets in the examples to ensure that $\psi_0(y,x)$, $\overline{\phi}_0(y)$ and $\underline{\phi}_0(y)$ are continuous.

While the software requires continuity, viscosity solution theory allows for more general semicontinuous terminal conditions for finite horizon optimal control problems [2, section V.5.2]. We plan to adapt numerical schemes from conservation laws to permit approximation of solutions with these more general terminal conditions.

All computations were done with ToolboxLS version 1.1 in MATLAB version 7.11 under 64-bit Windows 7 on a Lenovo x200 tablet with 4GB RAM and an Intel Core2 Duo L9400 CPU at 1.86 GHz. MATLAB code for all examples is available at the author's web site.

4.2 The Double Integrator

To demonstrate the technique, consider the simplest possible dynamics with input and a terminal integrator: the traditional double integrator

$$\dot{y} = f(y,u) = u, \qquad \dot{x} = b(y) = y \tag{25}$$

with $y \in \mathbb{R}$ and $u \in \mathcal{U} = [-u_{\max}, +u_{\max}]$ for some constant $u_{\max} \geq 0$.

The usual double integrator target set is the complement of a rectangle in position x cross velocity y space, which can be translated as finding the set of states such that the system will not exceed specified upper and lower bounds on position and velocity, given upper and lower bounds on acceleration u. To represent the complement, we adjust the interpretation of $\underline{\phi}$ and $\overline{\phi}$ to

$$\mathcal{R}(\mathcal{R}_0, t) = \{(y,x) \mid x \leq \underline{\phi}(t,y) \text{ or } x \geq \overline{\phi}(t,y)\},$$

and similarly for $\mathcal{R}_0$. However, using the complement of a rectangle as $\mathcal{R}_0$ requires piecewise continuous $\overline{\phi}_0(y)$ and $\underline{\phi}_0(y)$, since the maximum and minimum position are discontinuous functions of

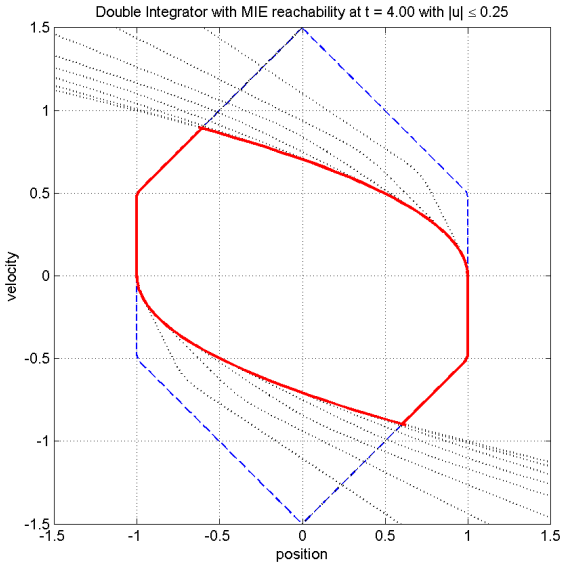

Figure 2: The reach tube for the double integrator computed with the MIE formulation. The reach tube is everything outside the solid curve. The dashed curve shows the target set, while the dotted curves show the evolution of $\overline{\phi}$ and $\underline{\phi}$.

velocity at the upper and lower bounds on velocity. Because ToolboxLS requires continuous functions, we instead use the hexagonal shape shown in figure 1 as the target set. The initial functions in figure 1 are given by

$$\begin{aligned} \overline{\phi}_0(y) &= \min\left[+1, y + \frac{3}{2}, -y + \frac{3}{2}\right], \\ \underline{\phi}_0(y) &= \max\left[-1, y - \frac{3}{2}, -y - \frac{3}{2}\right]. \end{aligned} \tag{26}$$

Since these functions are the minimum or maximum of three linear functions, they are continuous.

For this reach tube, we choose Hamiltonians to minimize the lower boundary $\underline{\phi}$ subject to the restriction that it can only increase, and maximize the upper boundary $\overline{\phi}$ subject to the restriction that it can only decrease. Plugging (25) into (15) with appropriate Hamiltonians yields the PDEs

$$\begin{aligned} 0 &= D_t\overline{\phi}(t,y) + \min\left[0, \max_{u\in\mathcal{U}}\left(D_y\overline{\phi}(t,y)\cdot u - y\right)\right] \\ &= D_t\overline{\phi}(t,y) + \min\left[0, \left(+u_{\max}|D_y\overline{\phi}(t,y)| - y\right)\right] \end{aligned}$$

and

$$\begin{aligned} 0 &= D_t\underline{\phi}(t,y) + \max\left[0, \min_{u\in\mathcal{U}}\left(D_y\underline{\phi}(t,y)\cdot u - y\right)\right] \\ &= D_t\underline{\phi}(t,y) + \max\left[0, \left(-u_{\max}|D_y\underline{\phi}(t,y)| - y\right)\right]. \end{aligned}$$

Using ToolboxLS we approximate the solutions $\overline{\phi}$ and $\underline{\phi}$ of these PDEs. The results are shown in figure 2, where the final reach tube is everything outside the solid lines.

For comparison purposes, the same double integrator problem can be solved using the traditional implicit formulation from section 2.2. The single HJ PDE in this case is

$$\begin{aligned} 0 &= D_t\psi(t,y,x) + \min\left[0, \min_{u\in\mathcal{U}}\left(D_y\psi(t,y,x)\cdot u + D_x\psi(t,y,x)\cdot y\right)\right] \\ &= D_t\psi(t,y,x) + \min\left[0, \left(-u_{\max}|D_y\psi(t,y,x)| + D_x\psi(t,y,x)\cdot y\right)\right]. \end{aligned}$$

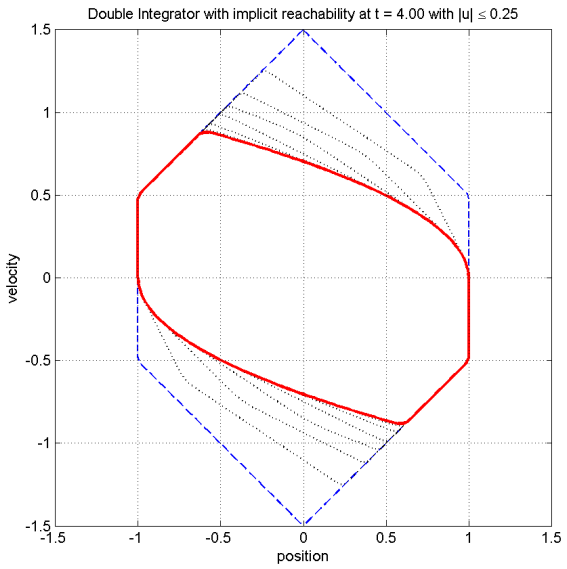

Figure 3: The reach tube for the double integrator computed with the implicit formulation. The reach tube is everything outside the solid curve. The dashed curve shows the target set, while the dotted curves show the evolution of the zero level set of ψ.

The target set is still the complement of a hexagon, so the terminal conditions $\psi_0(y, x)$ are constructed by taking the complement of the intersection of six half-spaces (with implicit surface functions, these operations can be accomplished through some linear algebra and maximum operations). The results are shown in figure 3.

The MIE formulation requires solving two PDEs over a one dimensional grid, while the implicit formulation requires solving a single PDE over a two dimensional grid. Consequently the latter requires at least an order of magnitude more computational effort. In this case the MIE formulation took less than 0.4 seconds on a grid of size 101 (although most of that time is probably consumed by generating the plots), while the implicit formulation took less than 4 seconds on a grid of size 101^2.

4.3 The Rotating Double Integrator

We modify the double integrator so that $y \in \mathbb{R}^2$ and the optimal input is no longer constant along trajectories. The dynamics are

$$\begin{bmatrix}\dot{y}_1\\ \dot{y}_2\end{bmatrix} = \dot{y} = f(y,u) = \begin{bmatrix}-y_2\\ +y_1\end{bmatrix} + \mu(\|y\|_2)\begin{bmatrix}u_1\\ u_2\end{bmatrix}, \qquad \dot{x} = b(y) = \|y\|_2 \tag{27}$$

where $u \in \mathcal{U} = \{u \in \mathbb{R}^2 \mid \|u\|_2 \le u_{\max}\}$ for some constant $u_{\max} \ge 0$ and $\mu : \mathbb{R} \to \mathbb{R}$.

If $\mu(\alpha) \equiv 1$, the system behaves radially like the first quadrant of a traditional double integrator. For this example, we choose

$$\mu(\alpha) = 2\sin(4\pi\alpha)$$

so that the optimal vector input not only changes direction and sign, but has variable effect depending on the current state.

For this example we compute only the lower boundary of a backwards minimal reach tube

$$\mathcal{R}(\mathcal{R}_0, t) = \{(y, x) \mid x \ge \underline{\phi}(t, y)\},$$

where $\underline{\phi}_0$, the lower boundary of $\mathcal{R}_0$, is given by

$$\underline{\phi}_0(y) = \min\left[1, 5(1.2 - y_1), 5(1.2 + y_1), 5(1.2 - y_2), 5(1.2 + y_2)\right];$$

and is shown on the left of figure 4. Note that the "position" variable x is the vertical axis (unlike the previous example, where it was shown on the horizontal axis). We use this truncated pyramid shape for the target set boundary to ensure that $\underline{\phi}_0$ is continuous; however, note that the slopes of the sides are much steeper than those in figure 1.

Plugging in the dynamics (27), the Hamiltonian for (15) is

$$\begin{aligned} H(t,y,p) &= \min\left[0, \max_{u\in\mathcal{U}}\left(p\cdot f(y,u) - b(y)\right)\right] \\ &= \min\left[0, (-y_2p_1 + y_1p_2 + u_{\max}|\mu(\|y\|_2)|\|p\|_2 - \|y\|_2)\right] \end{aligned}$$

The results are shown in the middle of figure 4, where the reach tube is everything above the solid surface.

For comparison, the same reach tube computed with the standard implicit scheme uses PDE (5) with Hamiltonian

$$H(t,z,p) = \min\left[0, \left(-y_2p_{y_1} + y_1p_{y_2} + u_{\max}|\mu(\|y\|_2)|\|p\|_2 + \|y\|_2p_x\right)\right]$$

and target set constructed by taking the complement of the intersection of six half-spaces. Results are shown on the right of figure 4.

The MIE formulation requires a single PDE on a two dimensional grid and took 18 seconds on a grid of size 121^2. The implicit formulation requires a single PDE on a three dimensional grid and took around 3000 seconds on a grid of size $121^2 \times 61$.

5. GENERALIZING THE TERMINAL INTEGRATOR

Consider a generalization of the terminal integrator's dynamics

$$\dot{x} = a(y)x + b(y,v) \tag{28}$$

where $v \in \mathcal{V}$ is an input with $\mathcal{V}$ compact convex. Following the same formal derivation as in (12) for $\overline{\phi}$ for a maximal reach set, for example, we arrive at a terminal value PDE of the form

$$D_t\overline{\phi}(t,y) + H\left(t, y, \overline{\phi}(t,y), D_y\overline{\phi}(t,y)\right) = 0 \tag{29}$$

with Hamiltonian

$$\overline{H}_{\text{set}}(t,y,q,p) = \max_{u\in\mathcal{U}}\max_{v\in\mathcal{V}}\left(p\cdot f(y,u) - a(y)q - b(y,v)\right), \tag{30}$$

and terminal conditions (17). The major difference in these equations compared to (15) and (16) is that the Hamiltonian now depends on the function value $\overline{\phi}$.

Hamilton-Jacobi PDEs of the form (29) and (30) arise in optimal control theory for discounted finite horizon problems [2, section III.3]. If the linear term is a positive constant $a(y) \equiv a > 0$, then a unique viscosity solution to (29) and (30) exists. Unfortunately, the cost function corresponding to (14) is in this case

$$\int_t^0 -b(y(s),v)e^{a(s-t)}ds + e^{-at}x(0,y(0)) \ne x(t,y(t)),$$

so there is no simple extension of Proposition 1 to make a rigourous argument. Furthermore, the case of $a < 0$ is more likely to occur in practice, since $a > 0$ gives rise to unstable growth in state variable x; however, $a < 0$ in (30) breaks a key monotonicity property assumed of Hamiltonians in viscosity solution theory. The theoretical basis for the MIE formulation with terminal integrator dynamics of the form (28) therefore remains an area of future research.

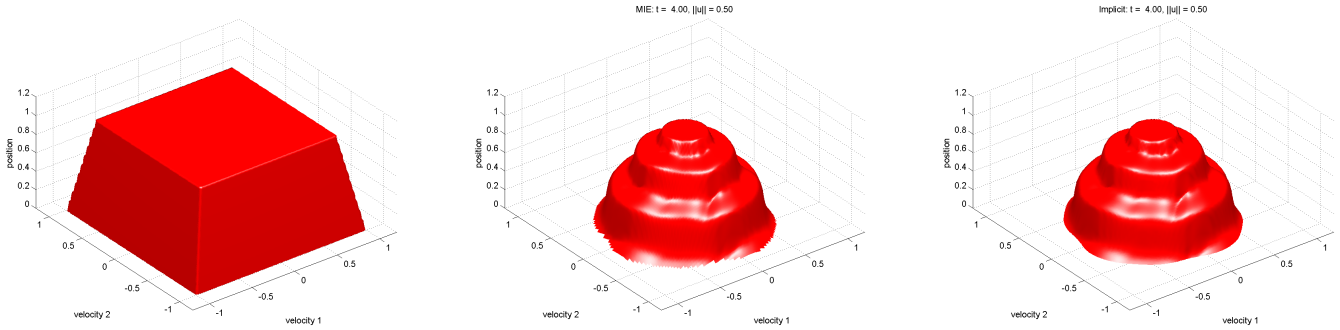

Figure 4: The rotating double integrator example. Note that the "position" variable x is the vertical axis, and in every case the relevant set is *above* the surface shown. Left: The lower boundary $\underline{\phi}_0$ of the target set. Middle: The lower boundary $\underline{\phi}$ of the minimal backwards reach tube computed with the MIE formulation. The ragged bottom edge is a visualization artifact arising from the truncation of the surface for $x < 0$. Right: The lower boundary of the minimal backwards reach tube computed with the implicit formulation.

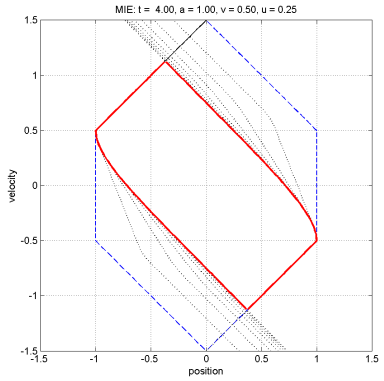

Figure 5: The reach tube for the modified-beyond-recognition double integrator (a double integrator with terminal integrator dynamics of the form (28)) computed with the MIE formulation. The reach tube is everything outside the solid curve. The dashed curve shows the target set, while the dotted curves show the evolution of $\overline{\phi}$ and $\underline{\phi}$.

5.1 The Modified-Beyond-Recognition Double Integrator

Despite the theoretical issues outlined above, it is trivial to adapt the numerical schemes in ToolboxLS to handle (29) and (30). To demonstrate, we again modify the double integrator example, this time by adding a linear term and an input to the position variable's dynamics

$$\dot{y} = f(y,u) = u, \qquad \dot{x} = ax + b(y,v) = ax + y + v \tag{31}$$

with inputs $u \in \mathcal{U} = [-u_{\max}, +u_{\max}]$ and $v \in \mathcal{V} = [-v_{\max}, +v_{\max}]$ for some constants $u_{\max} \geq 0$ and $v_{\max} \geq 0$. We seek to approximate the minimal reach tube for the same target set as in section 4.2, so the Hamiltonian for $\overline{\phi}$ will be

$$\begin{aligned} H(t,y,q,p) &= \min\left[0, \max_{u\in\mathcal{U}} \max_{v\in\mathcal{V}} (p \cdot f(y,u) - aq - b(y,v))\right] \\ &= \min\left[0, (+u_{\max}|p| - aq - y + v_{\max})\right] \end{aligned}$$

with PDE (29) and terminal conditions (26). This terminal value HJ PDE is known to have a viscosity solution for $a > 0$. The results computed for $a = +1$, $u_{\max} = 0.25$ and $v_{\max} = 0.5$ are shown in figure 5, and took less than one second to compute on a grid of size 151. The results using a fully implicit scheme are similar, and so are not shown, although they take about twenty seconds to complete on a grid of size 151^2. While the implicit formulation is much more computationally intensive, it is fully supported by viscosity solution theory for any bounded and Lipschitz continuous $a(y)$ on any bounded domain of x.

6. MULTIPLE TERMINAL INTEGRATORS

We extend to systems with independent terminal integrator states $x \in \mathbb{R}^{d_{\text{ti}}}$ for $d_{\text{ti}} > 1$ whose dynamics are

$$\dot{x}_i = b_i(y) \qquad \text{for } i = 1, 2, \ldots, d_{\text{ti}}.$$

There are three straightforward adaptations of the procedures discussed above to systems of this form. First, one can solve an HJ PDE of the form (5) (with appropriate Hamiltonians and terminal conditions) in the full state space $z \in \mathbb{R}^{d_z}$ to get an implicit representation of the full dimensional reach set or tube. Second, because the terminal integrator variables are fully decoupled from one another, one can use the projection ideas from [12] (without any need to resort to artificial disturbance inputs) and solve d_{ti} separate HJ PDEs of the form (5) in the state spaces $(y, x_i) \in \mathbb{R}^{d_{\text{cp}}+1}$ for $i = 1, 2, \ldots, d_{\text{ti}}$ to get d_{ti} separate implicit representations of the projections of the reach set or tube. Finally, one can solve $2d_{\text{ti}}$ separate HJ PDEs of the form (15) in the state space $y \in \mathbb{R}^{d_{\text{cp}}}$ to get d_{ti} separate MIE representations of the projections of the reach set or tube.

While the decoupled implicit and MIE schemes are obviously appealing because of the considerably lower dimension in which the PDEs need to be solved, there is a challenge: the Hamiltonians (such as (16)) of the resulting vector HJ PDE are coupled through their choice of input u. For maximal reach sets and tubes choosing the input independently for each component of the solution is sound if potentially pessimistic—it may be that to achieve the maximal reach set or tube for x_i one must choose a $\hat{u}_i$ which produces a reach set or tube for x_j that is much smaller than the maximal $\hat{u}_j$. However, this independent choice is unlikely to be too pessimistic, because in general there will be states $(\hat{x}_i, \hat{x}_j)$ such that $\hat{x}_i$ is in the interior of its reachable interval but $\hat{x}_j$ is at one of the boundaries of its reachable interval, and thus $\hat{u}_j$ is the appropriate maximal choice of input.

The situation for minimal reach sets and tubes is more complicated. Simply choosing the inputs independently risks introducing leaky corners into the reach set or tube: There may be states $(\hat{x}_i, \hat{x}_j)$

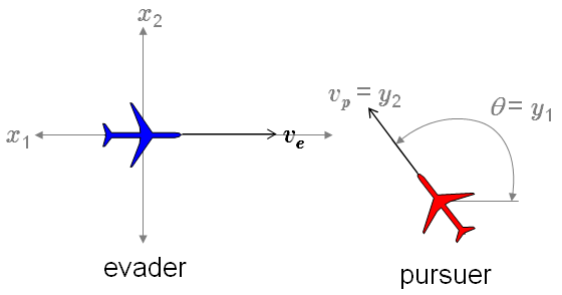

Figure 6: The relative coordinate system for the pursuit of an oblivious vehicle problem. The (oblivious) evader is fixed at the origin facing right and cannot modify its speed or heading. The pursuer can modify its speed and heading.

such that $\hat{u}_i$ must be chosen to avoid the target set in dimension i, $\hat{u}_j$ must be chosen to avoid the target set in dimension j, and $\hat{u}_i \neq \hat{u}_j$. In the fully implicit formulation with a single Hamiltonian such as (6), $D_z\psi(t,z)$ (the costate of the corresponding optimal control problem) provides a definitive choice of u. In the MIE formulation, only $D_y\overline{\phi}_i(t,y)$ and $D_y\underline{\phi}_i(t,y)$ are available for $i = 1, 2, \ldots, d_y$. A similar issue exists for the decoupled implicit formulation. Fortunately, there is again a sound if potentially pessimistic solution: Use any single $u \in \mathcal{U}$; for example,

$$\arg\min_{u\in\mathcal{U}} \left(\sum_{i=1}^{d_x} D_y\overline{\phi}_i(t,y) - \sum_{i=1}^{d_x} D_y\underline{\phi}_i(t,y) \right) \cdot f(y,u). \tag{32}$$

Because the choice may be suboptimal in some or all dimensions, the resulting set may not be the true minimal reach set or tube; however, it will not have leaky corners.

6.1 Pursuit of an Oblivious Vehicle

To demonstrate these approaches for a system with multiple terminal integrators, consider a pursuit problem played in a planar workspace. A pursuer vehicle wishes to collide with an evader vehicle. The evader is oblivious to the pursuer in the sense that it travels at constant linear speed v_e and at constant heading, which we choose without loss of generality to be zero. The pursuer travels at speed v_p and heading θ, and may modify its speed through linear acceleration input $a_p \in A_p$ and its heading through angular velocity input $\omega_p \in \Omega_p$. Because collision only depends on relative spatial position, we fix the evader at the origin and use the spatial variables x_1 and x_2 to represent the relative position of the pursuer. See figure 6 for a diagram of the problem. The dynamics can be written as

$$\frac{d}{dt}\begin{bmatrix}\theta\\ v_p\\ x_1\\ x_2\end{bmatrix} = \begin{bmatrix}\omega_p\\ a_p\\ -v_e + v_p\cos\theta\\ v_p\sin\theta\end{bmatrix}$$

which we can decompose into coupled dynamics

$$\dot{y} = \frac{d}{dt}\begin{bmatrix}y_1\\ y_2\end{bmatrix} = \begin{bmatrix}\theta\\ v_p\end{bmatrix} = \begin{bmatrix}\omega_p\\ a_p\end{bmatrix} = f(y,u)$$

and terminal integrators

$$\dot{x} = \frac{d}{dt}\begin{bmatrix}x_1\\ x_2\end{bmatrix} = \begin{bmatrix}-v_e + y_2\cos y_1\\ y_2\sin y_1\end{bmatrix} = b(y).$$

Note that the coupled dynamics do not depend on the coupled variables and are linear in the inputs, while the terminal integrators are nonlinear in y_1.

Traditionally the target set for collision problems would be circular in the position variables and independent of other variables; however, in order to treat the position variables as terminal integrators we must decouple their target set components. As a consequence, we use an interval as the target set for each position variable. We also constrain the pursuer's speed such that $0 < v_p^{\min} \leq v_p \leq v_p^{\max}$. Speed constraints of this sort would be appropriate, for example, if the vehicles are fixed-wing aircraft. The resulting target set is a square in $x_1 \times x_2$ space for all headings $\theta = y_1$ and all valid speeds $v_p = y_2$.

While the constraint on v_p certainly affects the target set, it goes further than that. For the backward reach tube in this problem we are only interested in states which can reach the target set within the specified time without violating the constraints on v_p. The construction of the target set ensures that the state within the target set which the trajectory reaches satisfies the constraints, but we need additional effort to ensure that all other states along the trajectory also satisfy the constraints. In the implicit formulation state constraints are enforced by applying a constraint on the solution of the HJ PDE $\psi(t,z) \geq \zeta_{\text{implicit}}(z)$ for all t, where ζ_{implicit} is an implicit surface function for the state constraints; equivalently, $\psi(t,z)$ becomes the solution of a variational inequality. For state constraints on the coupled variables (such as the constraint on $v_p = y_2$ in this problem), the MIE formulation can also apply a constraint of the form $\overline{\phi}(t,y) \leq \zeta^{\text{upper}}_{\text{indicator}}(y)$ (and a corresponding lower bound on $\underline{\phi}$); however, $\zeta^{\text{upper}}_{\text{indicator}}$ is now a discontinuous function with value $+\infty$ for states which satisfy the constraint and value 0 for those which do not.

As mentioned earlier, our current implementation of MIE requires that $\overline{\phi}_i$ and $\underline{\phi}_i$ are continuous with respect to y. That is trivially true with respect to $y_1 = \theta$ (the target set is constant with respect to this variable), but not true with respect to $y_2 = v_p$ because of the constraints on v_p. In order to maintain continuity, the size of the target square is shrunk gradually to zero as $y_2 = v_p$ approaches $v_p^{\min}$ or $v_p^{\max}$. In order to maintain consistency across all formulations, this same smoothed target set is used in all cases. With regard to the constraint functions, no smoothing is necessary for ζ_{implicit} in the implicit formulations because it is chosen as a signed distance function and is hence continuous already, and no smoothing is applied to $\zeta_{\text{indicator}}$ in the MIE formulation beyond the effect of grid discretization.

We approximate the reach set for parameters $v_e = 1$, $A_p = [-0.2, +0.2]$, $\Omega_p = [-0.2, +0.2]$, $v_p^{\min} = 1.0$, $v_p^{\max} = 3.0$ and target set $x \in [-1, +1]^2$ for valid v_p. We compute the backward reach tube out to $t = 2.0$ in each of the three formulations: full dimensional implicit (one PDE in four dimensions), decoupled implicit (two PDEs in three dimensions) and decoupled MIE (four PDEs in two dimensions). For the decoupled calculations, inputs are chosen independently in each PDE, and thus the computed tube is conservative. The results are somewhat difficult to visualize because the reach tube lies in a four dimensional space.

Figure 7 shows the projections of the reach tube into the (x_1, v_p, θ) and (x_2, v_p, θ) subspaces. These projections are computed directly for the two decoupled formulations, and for the full dimensional implicit case we simply take a minimum over the missing dimension of the implicit surface function ψ.

Of course, projections almost always lose some information. Figure 8 shows a sample slice of the reach tube in the (x_1, x_2, θ) subspace for $v_p = 2.0$. This slice is a subset of the data computed by the full dimensional implicit formulation. For the decoupled formulations, this slice can be recreated by backprojecting the lower dimensional reach tubes into full dimensional prisms and then in-

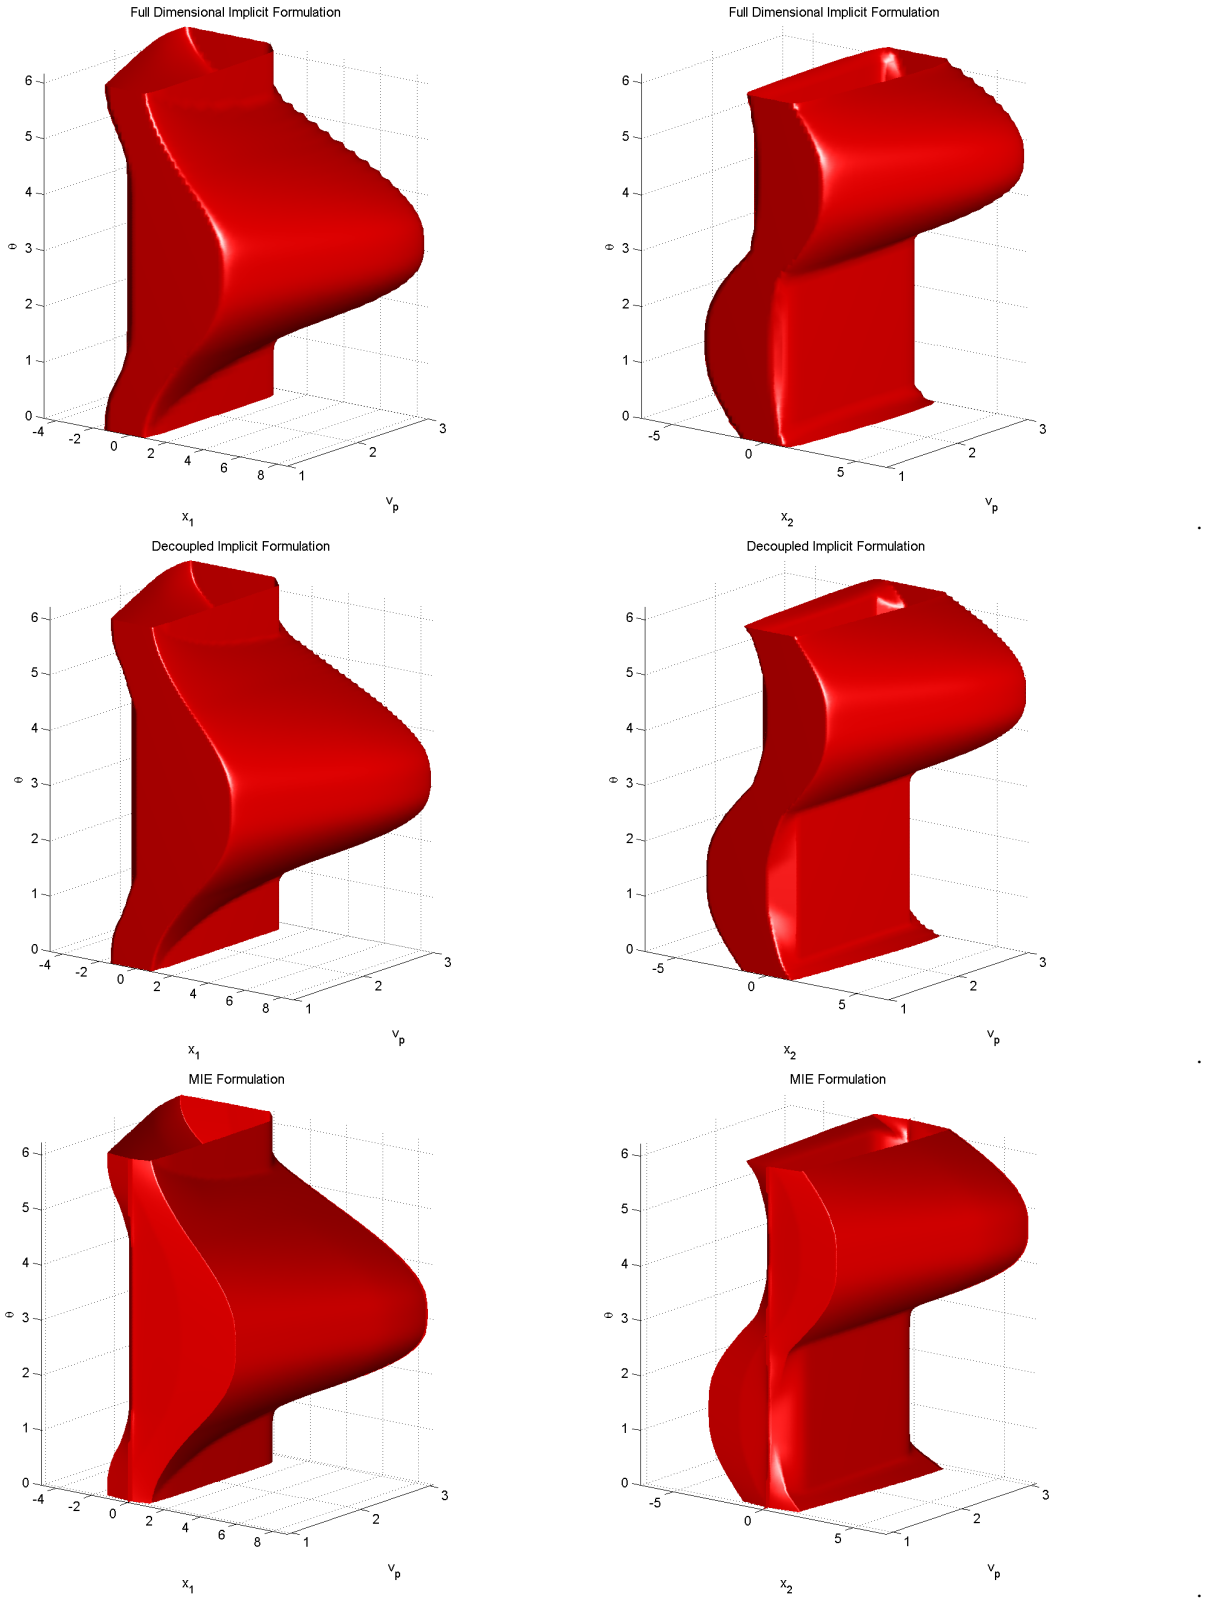

Figure 7: Projections of the reach tube for the pursuit of an oblivious vehicle problem, computed with the three formulations. Left column: (x_1, v_p, θ) projection. Right column: (x_2, v_p, θ) projection. Top row: full dimensional implicit formulation. Middle row: decoupled implicit formulation. Bottom row: decoupled MIE formulation.

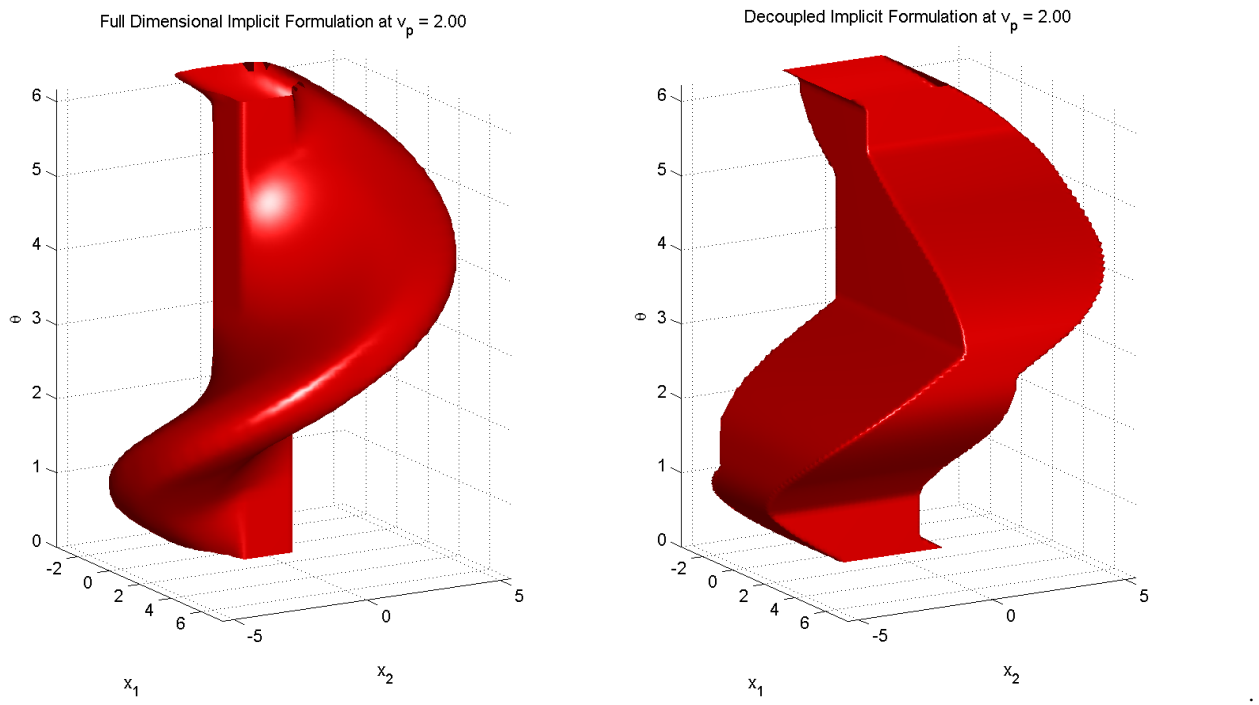

Figure 8: A slice of the reach tube in (x_1, x_2, θ) space for $v_p = 2.0$. Left: The reach tube as computed by the full dimensional implicit formulation. Right: The reach tube as computed by backprojection of the decoupled implicit formulation (the MIE result would be the same). Because the decoupled formulations (implicit or MIE) work in subsets of the state space, there is inevitably some overapproximation of the reach tube.

tersecting the prisms; the result from the decoupled implicit calculation was easier to backproject and is shown, although the result from the MIE calculation would be virtually identical. Note how the reach tube computed by projections is an overapproximation of the reach tube computed in the full dimensional state space.

While the decoupled formulations effectively result in an overapproximation of the reach tube, the trade-off is their speed. On a grid of size 65×100 the MIE formulation required 3.1 seconds (a nontrivial fraction of which is I/O) to solve four PDEs. On a grid of size $151 \times 65 \times 100$ the decoupled implicit formulation required 541 seconds to solve two PDEs. On a grid of size $151^2 \times 65 \times 100$ the full dimensional implicit formulation required too much memory, but extrapolation from a grid of half that size in each dimension provides a rough estimate of 30 *hours* to finish.

There are two obvious extensions of this problem: permit the evader to modify its speed and/or heading. By including v_e as a (coupled) state variable, it is straightforward to augment this model to permit the evader to modify its speed. The additional state dimension would make the full dimensional implicit formulation five dimensional and hence impractical to compute, but the MIE formulation would be only three dimensional and still quite feasible. It is possible to permit the evader to modify its heading without augmenting the state space—the pursuer heading state θ is replaced with a relative heading state, and the result is a version of the game of two identical vehicles (see, for example, [10] or [12]). Unfortunately, in this model the relative position variables are coupled and hence cannot be treated with the MIE formulation. The alternative is to augment the state space with the evader's heading, which has the same dimensional costs as adding v_e to the state space.

7. COMPARING THE IMPLICIT AND MIE FORMULATIONS

For a system with $y \in \mathbb{R}^{d_{\text{cp}}}$ and $x \in \mathbb{R}^{d_{\text{ti}}}$, the traditional full dimensional implicit formulation described in section 2.2 requires solving a single PDE (5) over the entire state space. Approximating the solution of such a PDE is typically accomplished by gridding the entire state space and hence requires $O(n^{(d_{\text{cp}}+d_{\text{ti}})})$ memory for a grid of size n in each dimension, and computational time $O(n^{(d_{\text{cp}}+d_{\text{ti}})})$ or greater. As described in section 6, because the dynamics of the terminal integrator states are decoupled, the decoupled implicit formulation requires solving d_{ti} separate PDEs in $\mathbb{R}^{d_{\text{cp}}+1}$, so the cost would be $O(d_{\text{ti}} n^{d_{\text{cp}}+1})$. Finally, the MIE formulation for the same system requires solving two PDEs of the form (15) for each terminal integrator state, but these PDEs are solved over only the coupled state space $\mathbb{R}^{d_{\text{cp}}}$; thus, the cost is $O(2d_{\text{ti}} n^{d_{\text{cp}}})$.

Despite their higher computational cost, the implicit formulations (both full and decoupled) do have a few features which the MIE formulation lacks. First, in the implicit formulations the reach set or tube is defined by the zero level set of the solution. The implicit formulation can therefore represent sets with sharp features and some types of discontinuities while maintaining a continuous solution, which makes both theory and numerics easier. Such continuity benefits also accrue when applying constraints to the reach set or tube (such as the constraints on v_p in section 6.1), since those constraints can be represented by continuous implicit functions for the implicit formulations but may require discontinuous functions for the MIE formulation. Furthermore, since only the zero level set is of interest, narrowband or local level set schemes can be used to slightly reduce the memory and cost (although ToolboxLS does not implement such schemes). Finally, artificial boundary conditions

for the edges of the computational domain are easier to construct when only the zero level set is of interest, because errors caused by such boundary conditions are easily kept away from the zero level set through reinitialization or similar procedures. In the MIE formulation, the entire solution of the PDE is relevant, so none of these benefits apply.

Second, in the implicit formulations the gradient of the solution is the costate of the corresponding finite horizon optimal control problem (see [10] for details). Consequently, this costate can be used to choose optimal inputs along the boundary of the reach set or tube (or throughout the state space under appropriate circumstances). In the MIE formulation, only $D_y\overline{\phi}_i$ and $D_y\underline{\phi}_i$ are available; we are presently investigating how they might be used to deduce optimal inputs.

Given these benefits of implicit schemes, the decoupled implicit formulation might appear to be a "best of both worlds," since it requires only one more dimension than MIE. One dimension, however, is often quite significant computationally; for example, consider the reach tube in section 6.1 which took more than 150 times as long to approximate with the three dimensional decoupled implicit formulation as it did with the two dimensional MIE formulation. Furthermore, the decoupling of the terminal integrator states (whether in an implicit or MIE formulation) introduces two additional weaknesses when compared with the full dimensional implicit formulation. The first is the unresolved coupling through the inputs discussed in section 6. The second is the likely overapproximation of the reach set or tube caused by working in subspaces of the state space, as demonstrated in figure 8. Only the full dimensional implicit formulation can avoid these issues, and its dimension is often impractically large.

8. CONCLUSIONS AND FUTURE WORK

We have demonstrated the new mixed implicit explicit (MIE) formulation for the computation of reach sets and tubes using Hamilton-Jacobi type PDEs for a class of systems with terminal integrators. ODE models of mechanical systems often take this form, since position states are simple integrations of velocity states. While the traditional full dimensional implicit formulation requires computational resources exponential in the number of state space dimensions, the MIE formulation requires resources linear in the number of terminal integrators and exponential in the number of remaining state space dimensions. We have also discussed an intermediate, decoupled implicit formulation for these systems which almost matches the MIE formulation's asymptotic complexity, but which also shares some of its weaknesses.

In the examples, the solutions of the resulting PDEs were approximated using ToolboxLS, but other numerical methods could certainly be applied. In particular, ToolboxLS provides no guarantee on the sign of the error in the approximation. For rigourous computation of reach sets and tubes, schemes providing such guarantees would be desirable; for example, the numerical algorithms for approximating HJ PDEs developed in the viability community, such as those described in [1, 3].

We are currently investigating several aspects of the new formulation. Section 5 raised some theoretical questions regarding an extension of the terminal integrator's dynamics. As discussed in section 6, there are questions regarding the treatment of inputs to the coupled states in the case of vector terminal integrators. Section 7 raised a related issue regarding the inputs: The relationship between the gradients of the MIE solution and the costate of the implicit formulation, since the latter can be used to construct optimal feedback policies. As mentioned in section 4.1, we are planning to implement new schemes in ToolboxLS which will permit computation with piecewise continuous functions. Finally, the examples used in this paper were all toys; we are working on more realistic problems in five and higher dimensions, and seeking extensions of the MIE formulation which will reduce the computational costs even further.

9. REFERENCES

[1] J.-P. Aubin, A. Bayen, and P. Saint-Pierre. Dirichlet problems for some Hamilton-Jacobi equations with inequality constraints. *SIAM Journal of Control and Optimization*, 47(5):2348–2380, 2008.

[2] M. Bardi and I. Capuzzo-Dolcetta. *Optimal Control and Viscosity Solutions of Hamilton-Jacobi-Bellman equations*. Birkhäuser, Boston, 1997.

[3] P. Cardaliaguet, M. Quincampoix, and P. Saint-Pierre. Set-valued numerical analysis for optimal control and differential games. In M. Bardi, T. E. S. Raghavan, and T. Parthasarathy, editors, *Stochastic and Differential Games: Theory and Numerical Methods*, volume 4 of *Annals of International Society of Dynamic Games*, pages 177–247. Birkhäuser, 1999.

[4] B. Djeridane and J. Lygeros. Neural approximation of PDE solutions: An application to reachability computations. In *Proceedings of the IEEE Conference on Decision and Control*, pages 3034–3039, San Diego, CA, 2006.

[5] L. C. Evans. *Partial Differential Equations*. American Mathematical Society, Providence, Rhode Island, 1998.

[6] I. Kitsios and J. Lygeros. Final glide-back envelope computation for reusable launch vehicle using reachability. In *Proceedings of the IEEE Conference on Decision and Control*, pages 4059–4064, Seville, Spain, 2005.

[7] J. Lygeros. On reachability and minimum cost optimal control. *Automatica*, 40(6):917–927, 2004.

[8] J. Lygeros, C. Tomlin, and S. Sastry. Controllers for reachability specifications for hybrid systems. *Automatica*, 35(3):349–370, 1999.

[9] I. M. Mitchell. Comparing forward and backward reachability as tools for safety analysis. In A. Bemporad, A. Bicchi, and G. Buttazzo, editors, *Hybrid Systems: Computation and Control*, number 4416 in Lecture Notes in Computer Science, pages 428–443. Springer Verlag, 2007.

[10] I. M. Mitchell, A. M. Bayen, and C. J. Tomlin. A time-dependent Hamilton-Jacobi formulation of reachable sets for continuous dynamic games. *IEEE Transactions on Automatic Control*, 50(7):947–957, 2005.

[11] I. M. Mitchell and J. A. Templeton. A toolbox of Hamilton-Jacobi solvers for analysis of nondeterministic continuous and hybrid systems. In M. Morari and L. Thiele, editors, *Hybrid Systems: Computation and Control*, number 3414 in Lecture Notes in Computer Science, pages 480–494. Springer Verlag, 2005.

[12] I. M. Mitchell and C. J. Tomlin. Overapproximating reachable sets by Hamilton-Jacobi projections. *Journal of Scientific Computing*, 19(1–3):323–346, 2003.

[13] P. Saint-Pierre. Hybrid kernels and capture basins for impulse constrained systems. In C. J. Tomlin and M. R. Greenstreet, editors, *Hybrid Systems: Computation and Control*, number 2289 in Lecture Notes in Computer Science, pages 378–392. Springer Verlag, 2002.