

Mapping Distributions through Hybrid Dynamical Systems and its Application to Kalman Filtering[★]

Nathan J. Kong^a, J. Joe Payne^a, George Council^a, Aaron M. Johnson^a

^a*Department of Mechanical Engineering, Carnegie Mellon University, Pittsburgh, Pennsylvania*

Abstract

Many state estimation and control algorithms require knowledge of how probability distributions propagate through dynamical systems. However, despite hybrid dynamical systems becoming increasingly important in many fields, there has been little work on how to map probability distributions through hybrid transitions. Here, we derive a propagation law that employs the saltation matrix (a first-order update to the sensitivity equation) to formally compute how a distribution's second moment is mapped through an isolated transition in a hybrid dynamical system. This saltation matrix update for the second moment of a distribution is compared to both the true distribution and a naive method which utilizes the differential of the reset map. Using this covariance propagation law, we propose the Salted Kalman Filter (SKF), a natural extension of the Kalman Filter and Extended Kalman Filter to hybrid dynamical systems. Away from hybrid events, the SKF is a standard Kalman filter. When a hybrid event occurs, the saltation matrix plays an analogous role as that of the system dynamics, subsequently inducing a discrete modification to both the prediction and update steps. Simulation results from the SKF show a reduced mean squared error in state estimation compared to using the differential of the reset map, especially immediately after a hybrid transition event.

Key words: Hybrid Systems; State Estimation; Kalman Filters; Nonlinear Systems.

1 Introduction

From legged robots to manipulator systems, many important contemporary control problems revolve around systems that make and break contact with their environments. These contact events are often represented as a discrete change to the system dynamics which introduces complexity for state estimation and control, as classic methods assume smoothness [6, 7, 16, 31]. These

“hybrid systems” [2, 15, 26] are systems with both continuous states (such as the position and velocity of a robot's center of mass and joints) and discrete states (such as whether or not a limb is in contact with the ground). Lacking out-of-the-box solutions, state estimation for these systems is a frontier with novel difficulties [8, 29], including how to deal with nonlinear dynamics on the continuous phases [5], discrete jumps in the continuous state [3], and real time computation [33].

[★] This material is based upon work supported by the U.S. Army Research Office under grant #W911NF-19-1-0080 and the National Science Foundation under grants #IIS-1704256 and #ECCS-1924723. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Office, National Science Foundation, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein. The material in this paper was not presented at any conference. Corresponding author N. J. Kong.

Email addresses: njkong@andrew.cmu.edu (Nathan J. Kong), jjpayne@andrew.cmu.edu (J. Joe Payne), gcouncil@andrew.cmu.edu (George Council), amj1@cmu.edu (Aaron M. Johnson).

In this work we take a step towards creating a Kalman-like filter compatible with hybrid systems while also avoiding the apparently-combinatorial effects of considering multiple modes simultaneously. To do this, we apply the saltation matrix (a standard tool from non-smooth analysis [25]) to propagate state uncertainty distributions through hybrid state transitions. The saltation matrix provides a first order approximation of the effects of a hybrid domain change based on the dynamics in the individual modes, the reset functions, and the location of the reset. It might be assumed that the propagation of uncertainty through hybrid transitions could be approximated by simply examining the first order approximation of the reset map itself, i.e. the Jacobian of the reset map. However, this naive ap-

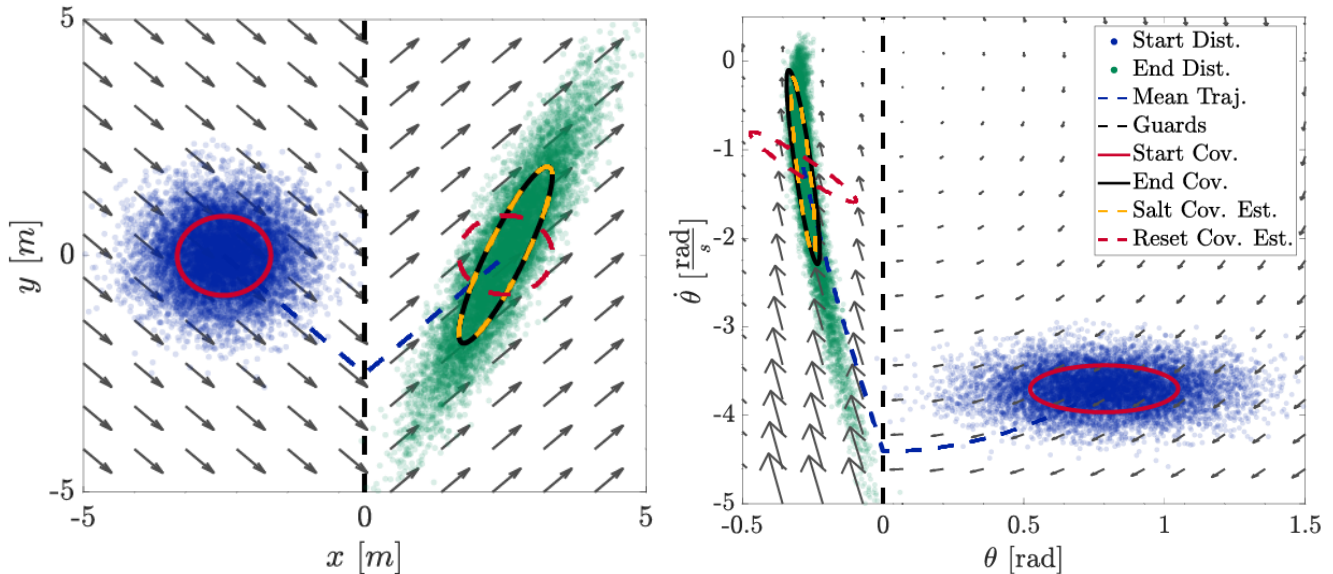


Fig. 1. Flowing an initial distribution (blue dots) with covariance (red solid line) along a nominal trajectory (blue dashed line) through hybrid dynamical systems with dynamics (gray arrows) and a single guard (black dashed line). The final distribution (green dots) is overlaid with the actual covariance (black line). Estimated covariance using the Jacobian of the reset map (red dashed line) is compared against our proposed estimate using the saltation matrix (gold dashed line). Left: constant flow vector fields. Right: pendulum hitting a torsional spring damper.

proximation does not take into account the differing dynamics in the distinct modes. The inaccuracy of the naive approach can be seen in Fig. 1, where each system has an identity reset and thus the Jacobian would be an identity matrix, but this does not accurately capture the effects of the hybrid transition on the distributions. For example, [7, 16] assumed that the hybrid transition does not affect the second moment of the distribution; i.e the reset map is identity and therefore the second moment is propagated with the identity matrix, with additional book keeping steps. As such, attempting to use the Jacobian of the reset map, while a “natural” idea, is ultimately incorrect in general.

The remainder of this paper is organized in the following manner. Section 2 provides a brief review of the hybrid system estimation literature. Section 3 defines the problem that we seek to solve in this work. Section 4 proves that using the saltation matrix is the correct first-order approach for propagating covariance matrices through state driven hybrid transitions. Section 5 introduces the “Salted Kalman Filter,” which is a Kalman Filter augmented with the capability to propagate the estimated first and second moments through hybrid transitions. Section 6 explains the experiments used to validate the performance of the saltation matrix on covariance propagation and in a Kalman filter. Section 7 compares results from using the saltation matrix to results from the more naive approach of using the Jacobian of the reset map. Finally, Section 8 provides a discussion of the work presented and potential future work.

2 Related Work

There has been a variety of work on the topic of state estimation for systems with differing dynamics and discrete modes, however current approaches either do not consider systems with state-driven mode transitions (e.g. the “switched system” case has received considerable attention) [4, 8, 9, 13, 19, 29], such as contact systems, or are computationally expensive and difficult to run in an online filtering setting [23, 33].

Our work seeks to understand how distributions are propagated through hybrid systems by applying knowledge from non-smooth systems literature [1, 18, 20] in order to make simplifying assumptions which retain sufficient information for the purposes of online state estimation.

2.1 Hybrid System Estimators

One approach to filtering on hybrid systems with linear dynamics is to use a filter bank where a filter is assigned to each discrete mode and the output of the filter with the lowest residual is used as the current state estimate [4]. Another style of filter bank method mixes the outputs of individual filters by utilizing a probability weight calculated based on measurement residuals and *a posteriori* estimate likelihoods such as the interacting multiple model (IMM) [8]. These filtering methods have been extended to hybrid systems with nonlinear dynamics [5] and hybrid systems with non-identity reset maps during hybrid transitions [3]. However, these filtering

bank strategies consider hybrid systems with transitions that do not depend on continuous state and therefore do not account for the effect that the continuous state dependent transitions have on the distribution. This is an issue because the first 2 moments of the distribution are not guaranteed to be captured after a transition. Particle filtering approaches such as the manifold particle filter (MPF) seek to represent uncertainty distributions directly with a variety of sample points rather than by representing belief as a parametric (e.g. Gaussian) distribution [23]. The MPF in particular is deliberate in the way it samples points when on a constrained surface in order to ensure proper coverage of possible states. One of the major drawbacks of the MPF and related methods is that they are computationally expensive due to combinatorial complexity, so it is difficult to utilize them for high degree of freedom systems in a real-time setting.

Some optimization based methods which seek to circumvent this issue of combinatorial complexity simultaneously select the continuous and discrete states over all timesteps to minimize the error associated with the measurements and the dynamics [14, 33]. The resulting optimization problem requires a much higher computational load compared to filtering methods and as such may be limited to offline estimation settings.

Online state estimation methods have been created for complex systems with continuous states and discrete modes, such as the case for legged robots making and breaking contact with the ground [7, 16]. In these settings, an extended Kalman filter is used to estimate the continuous states and the discrete mode is directly measured through contact sensors. The primary focus of these works is on the continuous phases rather than the discrete mode transitions due to the presence of direct mode sensing. Therefore, these estimators do not directly work for general hybrid systems, because there might not be a sensor to determine the hybrid event and there might be discontinuous jumps in the state that would need to be accounted for.

2.2 Non-smooth systems and the saltation matrix

This work makes extensive use of the saltation matrix [1, 10, 18, 25], which is a discontinuous update to the variational equation solution [22] and is a key part of linearizing hybrid dynamics around a chosen trajectory. They have previously been used to analyze stability of periodic solutions [1], trajectory sensitivity [18], and infinitesimal contraction [10]. In this work we utilize the saltation matrix to update the covariance of an uncertainty distribution.

3 Problem Formulation

The specific problem we seek to address in this work is the estimation of continuous states of a hybrid dynamical

system given:

- (1) A model of the dynamics in each mode.
- (2) A model of how the state resets between modes.
- (3) The location of the hybrid guards.
- (4) Measurements of the system's continuous state.

We are specifically not considering:

- (1) The probability of the discrete state
- (2) While our formal derivation of the second moment of a distribution holds for domains of differing dimension, estimation is restricted to domains of the same dimension
- (3) Hybrid systems with intersecting guards [28, § 3-4] e.g. for a walking system, intersecting guards occur when multiple feet impact simultaneously.

As many of these terms have multiple possible mathematical meanings, in this section we provide the essential definitions used in this work.

While there are many similar definitions for a hybrid dynamical system, e.g. [2, 15, 26], in this work we define a C^r hybrid dynamical system, closely following [20, Def. 2]:

Definition 1. A C^r hybrid dynamical system, for continuity class $r \in \mathbb{N}_{>0} \cup \{\infty, \omega\}$, is a tuple $\mathcal{H} := (\mathcal{J}, \Gamma, \mathcal{D}, \mathcal{F}, \mathcal{G}, \mathcal{R})$ where the constituent parts are defined as:

- (1) $\mathcal{J} := \{I, J, \dots, K\} \subset \mathbb{N}$ is the finite set of discrete **modes**.
- (2) $\Gamma \subset \mathcal{J} \times \mathcal{J}$ is the set of discrete **transitions** forming a directed graph structure over \mathcal{J} .
- (3) $\mathcal{D} := \prod_{I \in \mathcal{J}} D_I$ is the collection of **domains** where D_I is a C^r manifold with corners [21, 24].
- (4) $\mathcal{F} := \prod_{I \in \mathcal{J}} F_I$ is a collection of C^r time-varying **vector fields**, $F_I : \mathbb{R} \times D_I \rightarrow \mathcal{T}D_I$.
- (5) $\mathcal{G} := \prod_{(I,J) \in \Gamma} G_{(I,J)}(t)$ is the collection of **guards**, where $G_{(I,J)}(t) \subset D_I$ for each $(I, J) \in \Gamma$ is defined as a sublevel set of a C^r function, i.e. $G_{(I,J)}(t) = \{x \in D_I | g_{(I,J)}(t, x) \leq 0\}$.
- (6) $\mathcal{R} : \mathbb{R} \times \mathcal{G} \rightarrow \mathcal{D}$ is a C^r map called the **reset** that restricts as $R_{(I,J)} := \mathcal{R}|_{G_{(I,J)}(t)} : G_{(I,J)}(t) \rightarrow D_J$ for each $(I, J) \in \Gamma$.

An execution of a hybrid system [20, Def. 4] starts with initializing a state in some hybrid domain \mathcal{D}_I , where I is a discrete mode in \mathcal{J} . The dynamics on I , F_I , are followed until the trajectory reaches a guard $G_{(I,J)}$, where (I, J) is a discrete transition in Γ . This triggers the hybrid transition from mode I to mode J and the reset map $R_{(I,J)}$ is applied to the state to initialize the new state in hybrid domain \mathcal{D}_J . The execution is defined over a hybrid time domain [20, Def. 3],

Definition 2. A **hybrid time domain** is a disjoint union of intervals $\prod_{i=1}^N T_i$ such that:

- (1) $N \in \mathbb{N} \cup \{\infty\}$.
- (2) $T_I := [\underline{t}_i, \bar{t}_i] \subset \mathbb{R}$ is a closed interval from the start of an interval to the end of an interval for all $i < N$ and if $N < \infty$ then $T_N \subset \mathbb{R}$ is also a closed interval.
- (3) $T_i \cap T_{i+1}$ is nonempty and consists of a single element, $\bar{t}_i = \underline{t}_{i+1}$, for all $i < N$.

A classic result [17, Thm. 1, § 15.2] for a smooth system $\dot{x} = f(x)$ is that we can linearize around a trajectory $\phi^t(x)$ using the so-called *variational equation*

$$\frac{d}{dt} D_x \phi^t(x_0) = Df_x(\phi^t(x_0)) D_x \phi^t(x_0) \quad (1)$$

For the type of hybrid systems we consider, an analogous equation exists, but additional care must be taken to treat hybrid events consistently. As shown in [1, 18], if for some time τ the solution $\phi^\tau(x_0)$ intersects a single surface of discontinuity, the variational equation must be updated discontinuously with the saltation matrix $\Xi_{(I,J)}(\bar{t}_i, x(\bar{t}_i))$, which is defined at time \bar{t}_i at the end of hybrid time interval i with state $x(\bar{t}_i) \in G_{(I,J)}$.

Definition 3 ([10, Prop. 2]). *The **saltation matrix**,*

$$\Xi := D_x R + \frac{(F_J - D_x R \cdot F_I - D_t R) D_x g}{D_t g + D_x g \cdot F_I} \quad (2)$$

where

$$\begin{aligned} \Xi &:= \Xi_{(I,J)}(\bar{t}_i, x(\bar{t}_i)), & F_I &:= F_I(\bar{t}_i, x(\bar{t}_i)) \\ D_x R &:= D_x R_{(I,J)}(\bar{t}_i, x(\bar{t}_i)), & D_t R &:= D_t R_{(I,J)}(\bar{t}_i, x(\bar{t}_i)) \\ D_x g &:= D_x g_{(I,J)}(\bar{t}_i, x(\bar{t}_i)), & D_t g &:= D_t g_{(I,J)}(\bar{t}_i, x(\bar{t}_i)) \\ F_J &:= F_J(\bar{t}_{i+1}, R_{(I,J)}(x(\bar{t}_i))) \end{aligned}$$

is the first order approximation of variations at hybrid transitions from mode I to J and maps perturbations to first order from pre-transition $\delta x(\bar{t}_i)$ to post-transition $\delta x(\bar{t}_{i+1})$ during the i th transition in the following way,

$$\delta x(\bar{t}_{i+1}) = \Xi_{(I,J)}(\bar{t}_i, x(\bar{t}_i)) \delta x(\bar{t}_i) + \text{h.o.t.} \quad (3)$$

where h.o.t. represents higher order terms, i.e. $o(\|\delta x\|)$.

To ensure that the saltation matrix is well defined for all transitions and to avoid degenerate cases, such as Zeno behavior, we accept the assumptions (Assumption 1) from [10] to limit the class of hybrid dynamic systems under consideration.

Assumption 1. *Consider a hybrid dynamical system, $\mathcal{H} = (\mathcal{J}, \Gamma, \mathcal{D}, \mathcal{F}, \mathcal{G}, \mathcal{R})$, with the following **local** properties,*

- (1) Each vector field F_I is Lipschitz continuous and continuously differentiable for all $I \in \mathcal{J}$.
- (2) Discrete transitions are isolated, $\mathcal{R}(t, \mathcal{G}(t)) \cap \mathcal{G}(t) = \emptyset$ for all $t \in [0, \infty)$.
- (3) Guards and resets are differentiable.

- (4) The vector field is transverse to the guards, $\frac{d}{dt} g_{(I,J)}(t, x(t)) = D_t g_{(I,J)}(t, x) + D_x g_{(I,J)}(t, x) \cdot F_I(t, x) < 0$, for all $I, J \in \mathcal{J}, t \geq 0$ and $x \in G_{(I,J)}(t)$.
- (5) No trajectory undergoes an infinite number of resets in finite time (there are no Zeno events).
- (6) Hybrid system flows are continuous.

For full technical detail refer to [10]. Note that Assumption 1.4 restricts the definition of the guard from Definition 1 to be both a sublevel set and only exist when the vector field is transverse to it at the boundary. Equivalently, we can write each guard set $G_{(I,J)}$ as the following, where $g := g_{(I,J)}$, and $x(t)$ is a trajectory in D_I

$$G_{(I,J)} := \left\{ x \in D_I \mid g(t, x) \leq 0, \frac{d}{dt} g(t, x(t)) < 0 \right\} \quad (4)$$

This assumption also ensures the denominator in (2) does not approach zero.

In this paper, we start by considering linear systems before extending to the full nonlinear case. As such, through a slight abuse of notation,

Definition 4. *A C^r **linear hybrid dynamical system** is a C^r hybrid dynamical system (Definition 1) with the following restrictions:*

- The flow follows linear dynamics, $\dot{x} = F_I(t, x) = A_I(t)x + B_I(t)u(t)$, for some input function $t \mapsto u(t) \in \mathbb{R}^{m_I}$, where $A_I \in \mathbb{R}^{n_I \times n_I}$, $B_I \in \mathbb{R}^{n_I \times m_I}$, $n_I = \dim D_I$, and $m_I = \dim u$.
- Each guard set is given by a sublevel set of a linear function $g_{(I,J)}(x, t) = g_{ij}(t) \cdot x$, where $g_{ij}(t) \in \mathbb{R}^{n_I}$, so that $G_{(I,J)}(t) = \{x \in D_I \mid g(t) \cdot x \leq 0\}$.
- The reset is a linear function, mapping, $x_J = R_{(I,J)}(t, x) = R_{(I,J)}(t) \cdot x_I$ for $x_I \in G_{(I,J)}$, $x_J \in D_J$, and $R_{(I,J)} \in \mathbb{R}^{n_J \times n_I}$.
- All functions of t are C^r -smooth in t .

Note that where $D_x R_{(I,J)}$ is used in the text, it refers to the Jacobian of the map $R_{(I,J)}(\cdot)$, which in the linear case is simply the matrix $R_{(I,J)}$.

The goal of this work is to understand how distributions are mapped through hybrid transitions and to ultimately apply this knowledge to estimation algorithms that operate in discrete time,

Definition 5. *The discrete time $\mathcal{K} := \{t_k\}_{k \in \mathbb{N}}$ is a (possibly infinite) set of times with discrete index $k \in \mathbb{N}$ defined for a fixed time step Δ such that the corresponding time t_k is the product $k\Delta$.*

4 Mapping Covariance through Hybrid Transitions

In this section, we show how the saltation matrix can be used to approximately map the second moment of a dis-

tribution through hybrid transitions on hybrid dynamical systems (Def. 1) and prove that this idea is exact for linear hybrid systems (Def. 4). Note that the proposed method does not predict the final distribution but rather just the first two moments. We introduce this method in Sec. 4.1 and compare it to a naive approach using the Jacobian of the reset map in Sec. 4.2.

4.1 Proposed approach to mapping covariances through hybrid transitions

Without loss of generality, consider the unforced case where $\forall t \ u(t) = 0$. The linearization about \tilde{x} of the dynamics $F_I(t, \tilde{x}, 0)$ are captured by the dynamic matrix:

$$\hat{A}_I(t, \tilde{x}) := D_x F_I(t, \tilde{x}, 0) \quad (5)$$

Assume a discrete update law of the nonlinear dynamics $F_I(t, x(t), 0)$ with timestep Δ is given to be $f_{I,\Delta}(x(t))$ and the discretization of $\hat{A}_I(t, \tilde{x})$ to be $\hat{A}_{I,\Delta}$.¹ Note that time t and input $u(t)$ are dropped for brevity.

Define a random variable $X(t_k)$ which is drawn from a distribution at time t_k with mean $\mu(t_k) := x(t_k)$

$$X(t_k) := \mu(t_k) + \delta x(t_k) \quad (6)$$

where perturbations, $\delta x(t_k)$, are in the tangent space of D_I and are sampled from a distribution with zero mean and covariance $\Sigma(t_k)$,

$$\Sigma(t_k) = \mathbb{E}[\delta x(t_k) \delta x(t_k)^T] \quad (7)$$

The dynamic update [32, Eqns. 2.14–2.15] for the first two moments in the case where the hybrid system does not switch modes is given by,

$$\mu(t_{k+1}) = f_{I,\Delta}(\mu(t_k)) \quad (8)$$

$$\Sigma(t_{k+1}) = \hat{A}_{I,\Delta} \Sigma(t_k) \hat{A}_{I,\Delta}^T + \text{h.o.t.} \quad (9)$$

where again *h.o.t.* represents higher order terms, i.e. $o(t_k)$. In the linear hybrid systems case (Def. 4), there are no higher order terms and the dynamic update is exactly,

$$\mu(t_{k+1}) = A_{I,\Delta} \mu(t_k) \quad (10)$$

$$\Sigma(t_{k+1}) = A_{I,\Delta} \Sigma(t_k) A_{I,\Delta}^T \quad (11)$$

Next, consider the update when the hybrid system does transition modes. Assume for simplicity that only a single transition occurs during a timestep, i.e. $t_k \in T_i$ and $t_{k+1} \in T_{i+1}$ are in consecutive intervals of the hybrid

¹ Since there are many ways of discretizing a system, we are merely assuming that one has been chosen.

time domain (though the derivation may be easily extended to the multiple isolated transition case). The time periods from $[t_k, \bar{t}_i]$ and $[\underline{t}_{i+1}, t_{k+1}]$ each follow the smooth case and the update to the mean and covariance are the same as (8)–(9) with the appropriate hybrid modes and time steps. To derive the updates to the mean and covariance during the transition $[\bar{t}_i, \underline{t}_{i+1}]$, the post-impact mean state $\mu(\underline{t}_{i+1})$ is mapped from the pre-impact mean state $\mu(\bar{t}_i)$ through the reset map

$$\mu(\underline{t}_{i+1}) = R_{(I,J)}(\bar{t}_i, \mu(\bar{t}_i)) \quad (12)$$

and the perturbations are mapped using the saltation matrix as in (3),

$$\begin{aligned} \mu(\underline{t}_{i+1}) + \delta x(\underline{t}_{i+1}) &= R_{(I,J)}(\bar{t}_i, \mu(\bar{t}_i)) \\ &+ \Xi_{(I,J)}(\bar{t}_i, \mu(\bar{t}_i)) \delta x(\bar{t}_i) \\ &+ \text{h.o.t.} \end{aligned} \quad (13)$$

The state just after the transition, $X(\underline{t}_{i+1}) := \mu(\underline{t}_{i+1}) + \delta x(\underline{t}_{i+1})$, is a multivariate distribution with mean $\mu(\underline{t}_{i+1})$ and covariance $\Sigma(\underline{t}_{i+1})$ with an update analogous to (10)–(11), which we summarize in Proposition 1.

Proposition 1. *When the higher order terms are zero, the dynamic update to the mean and covariance of a hybrid system at the time of a reset is, defining $\mu^* := \mu(\bar{t}_i)$,*

$$\mu(\underline{t}_{i+1}) = R_{(I,J)}(\bar{t}_i, \mu^*) \quad (14)$$

$$\Sigma(\underline{t}_{i+1}) = \Xi_{(I,J)}(\bar{t}_i, \mu^*) \Sigma(\bar{t}_i, \mu^*) \Xi_{(I,J)}^T(\bar{t}_i, \mu^*) \quad (15)$$

See Appendix A for the proof. Note that this update holds as an approximation when higher order terms are small.

Now that the dynamic update for the first and second moments at the time of transition are known, the full update for a timestep Δ during a hybrid transition can be constructed. At a high level, the smooth update starting at t_k , (10)–(11), is applied up until transition at time \bar{t}_i , then the transition update using the saltation matrix, (14)–(15), is applied, and then the smooth update resumes in the new domain at time \bar{t}_i until t_{k+1} . Specifically, we divide the interval Δ into two regions, $\Delta_1 = \bar{t}_i - t_k$, and $\Delta_2 = t_{k+1} - \bar{t}_i$. The update to the mean and the covariance during the transition case is then,

$$\mu(t_{k+1}) = f_{J,\Delta_2}(R_{(I,J)}(\bar{t}_i, f_{I,\Delta_1}(\mu(t_k)))) \quad (16)$$

$$\begin{aligned} \Sigma(t_{k+1}) &= \hat{A}_{J,\Delta_2} \Xi_{(I,J)} \hat{A}_{I,\Delta_1} \Sigma(t_k) \hat{A}_{I,\Delta_1}^T \Xi_{(I,J)}^T \hat{A}_{J,\Delta_2}^T \\ &+ \text{h.o.t.} \end{aligned} \quad (17)$$

and is exactly in the linear hybrid systems case (Def. 4),

$$\mu(t_{k+1}) = A_{J,\Delta_2} R_{(I,J)} A_{I,\Delta_1} \mu(t_k) \quad (18)$$

$$\Sigma(t_{k+1}) = A_{J,\Delta_2} \Xi_{(I,J)} A_{I,\Delta_1} \Sigma(t_k) A_{I,\Delta_1}^T \Xi_{(I,J)}^T A_{J,\Delta_2}^T \quad (19)$$

The saltation matrix is evaluated at $\Xi_{(I,J)}(\bar{t}_i, A_{I,\Delta_1}\mu(t_k))$. In the case of multiple hybrid transitions within a time period, the update equations can similarly be built by successively composing (10)–(11) and (14)–(15).

4.2 Naive approach to mapping covariances through hybrid transitions

One naive approach to updating mean and covariance through a hybrid transition is to simply use the Jacobian of the reset function. This idea is particularly interesting because the update to the mean is just to apply to reset map, so intuitively the Jacobian of the reset map might properly map deviations from the mean. This is because if we consider only the effect of the reset map and apply a first order approximation (Taylor expansion) we get,

$$\begin{aligned} X(\underline{t}_{i+1}) &= R_{(I,J)}(\bar{t}_i, X(\bar{t}_i)) = R_{(I,J)}(\bar{t}_i, \mu(\bar{t}_i) + \delta x(\bar{t}_i)) \\ &= R_{(I,J)}(\bar{t}_i, \mu(\bar{t}_i)) + D_x R_{(I,J)}(\bar{t}_i, \mu(\bar{t}_i)) \delta x(\bar{t}_i) \\ &\quad + \text{h.o.t.} \end{aligned} \quad (20)$$

We would then expect the propagation of covariances to follow as,

$$\Sigma(\underline{t}_{i+1}) = D_x R_{(I,J)} \Sigma(\bar{t}_i) (D_x R_{(I,J)})^T + \text{h.o.t.} \quad (21)$$

However, this approach is not correct because it only considers variations about the reset map and not in the state of both domains and in the guard surface whereas the saltation matrix encompasses all variations, i.e. an expansion needs to account for the flow away from the guards, as well as the reset map. The basic reason for this requirement is that even though in a hybrid dynamical system we may have $\mu(\bar{t}_i) \in G_{(I,J)}$, we do not automatically have that $\mu(\bar{t}_i) + \delta x(\bar{t}_i) \in G_{(I,J)}$. For example, if δx is in the opposite direction of the flow then the state needs to continue flowing along the prior dynamics for some time before applying the reset function. The saltation matrix captures this phenomenon, and the approximation becomes

$$\begin{aligned} X(\underline{t}_{i+1}) &= R_{(I,J)}(\bar{t}_i, \mu(\bar{t}_i)) + \Xi_{(I,J)}(\bar{t}_i, \mu(\bar{t}_i)) \delta x(\bar{t}_i) \\ &\quad + \text{h.o.t.} \end{aligned} \quad (22)$$

Note that the difference between the Jacobian of the reset map and the saltation matrix is a rank one update as shown in (2) – the right hand side consists of $D_x R$ plus an outer product. Therefore, using the Jacobian of the reset map captures a portion of the distribution mapping but not all of it.

To illustrate the difference between this naive approach and the proposed, we compare using the Jacobian of the reset map instead of the saltation matrix in all experiments.

5 Kalman filtering for hybrid systems

In this section, we present the Salted Kalman Filter (SKF) by applying the prior result on the mapping of second moments to Kalman filters which enables their use on hybrid dynamical systems. To simplify expressions, we abuse notation and use $a(k) := a(t_k)$ for any relevant function a . Again, without loss of generality, we assume the autonomous case $u(k) = 0 \forall k$. To start, the stochastic dynamics considered for the standard Kalman filter [32, Eqn. 1.1] on domain I of a linear hybrid dynamical system are given by

$$x(k+1) := A_{I,\Delta} x(k) + \omega_{I,\Delta}(k) \quad (23)$$

where $\omega_{I,\Delta}$ is the process noise that is sampled from a zero mean Gaussian distribution with covariance $W_{I,\Delta}$. The stochastic measurement equation [32, Eqn. 1.2] is defined to be

$$y(k) := C_I x(k) + v_I(k) \quad (24)$$

where C_I is the measurement matrix, and v_I is the measurement noise that is sampled from a zero mean Gaussian distribution with covariance V_I . The *a priori* update equations are,

$$\hat{x}(k+1|k) = A_{I,\Delta} \hat{x}(k) \quad (25)$$

$$\hat{\Sigma}(k+1|k) = A_{I,\Delta} \hat{\Sigma}(k) A_{I,\Delta}^T + W_{I,\Delta} \quad (26)$$

and the *a posteriori* update is:

$$K_{k+1} = \hat{\Sigma}(k+1|k) C_I^T \left[C_I \hat{\Sigma}(k+1|k) C_I^T + V_I \right]^{-1} \quad (27)$$

$$\begin{aligned} \hat{x}(k+1|k+1) &= \hat{x}(k+1|k) \\ &\quad + K_{k+1} [y(k+1) - C_I \hat{x}(k+1|k)] \end{aligned} \quad (28)$$

$$\hat{\Sigma}(k+1|k+1) = \hat{\Sigma}(k+1|k) - K_{k+1} C_I \hat{\Sigma}(k+1|k) \quad (29)$$

where K_{k+1} is the Kalman gain [32, Eqns. 1.9–1.13]. These merely amount to the conventional Kalman filter equations in discrete time.

While the standard Kalman filter is adequate when a trajectory is confined to a single domain, we must also account for hybrid events. In this setting, we assume that the true time of impact to the guard \bar{t}_i is unknown to the filter and is estimated by determining when a hybrid transition occurs for the mean. In this filter, we allow both the *a priori* and *a posteriori* update to trigger a hybrid transition. Therefore, both updates are modified such that the mean and covariance are properly transformed during the hybrid transition. Recall, the conditions for hybrid transition are that the state needs to be within the guard set and the vector field evaluated at that state needs to be transverse to the guard (4). Since

these conditions define the guard locations, we can use them as testable statements in discrete time to assess if a hybrid event has occurred.

In this section we first show these changes for a Kalman filter on a linear hybrid dynamical systems (Sec. 5.1–5.2), then the same changes are similarly applied for the Extended Kalman filter on general hybrid dynamical systems (Sec. 5.3).

5.1 Hybrid transition during a priori update

For the *a priori* update, the state is propagated from the previous estimate for a single timestep Δ . If the guard and transversality conditions (4) are not met during the propagation, no hybrid transition is considered and the standard Kalman filter *a priori* update is used (25)–(26). If the conditions are met, then the forward simulation is stopped, a hybrid transition is assumed, and the time of impact $\bar{t}_i = t_k + \Delta_1$ is estimated to be the stopping time. Because we assume that a finite number of isolated transitions occur, this process can be repeated until the entire timestep is simulated. Without loss of generality, in this section we only consider the case where a single transition occurs, but appending additional transitions can be computed in a similar fashion.

If a transition occurs, the stochastic dynamics (23) are defined to be,

$$x(k+1) := A_{J,\Delta_2} (R_{(I,J)} [A_{I,\Delta_1} x(k) + \omega_{I,\Delta_1}(k)] + \omega_{R_{(I,J)}}(k)) + \omega_{J,\Delta_2}(k) \quad (30)$$

where $\omega_{R_{(I,J)}}$ is the reset process noise, sampled from a zero mean Gaussian distribution with covariance $W_{R_{(I,J)}}$, ω_{I,Δ_1} is the process noise in domain I with timestep Δ_1 , and ω_{J,Δ_2} is the process noise in domain J with timestep Δ_2 . The dynamic update at transition (14)–(15) augmented with the reset process noise is,

$$x(\underline{t}_{i+1}) = R_{(I,J)} x(\bar{t}_i) \quad (31)$$

$$\Sigma(\underline{t}_{i+1}) = \Xi_{(I,J)} \Sigma(\bar{t}_i) \Xi_{(I,J)}^T + W_{R_{(I,J)}} \quad (32)$$

where the saltation matrix is evaluated at $\Xi_{(I,J)} = \Xi_{(I,J)}(\bar{t}_i, x(\bar{t}_i))$. Combined with the continuous *a priori* updates before and after transition, (25)–(26), the *a priori* update over a full timestep is,

$$\hat{x}(k+1|k) = A_{J,\Delta_2} R_{(I,J)} A_{I,\Delta_1} \hat{x}(k) \quad (33)$$

$$\hat{\Sigma}(k+1|k) = A_{J,\Delta_2} [\Xi_{(I,J)} (A_{I,\Delta_1} \Sigma(k) A_{I,\Delta_1}^T + W_{I,\Delta_1}) \Xi_{(I,J)}^T + W_{R_{(I,J)}}] A_{J,\Delta_2}^T + W_{J,\Delta_2} \quad (34)$$

where the saltation matrix is evaluated at $\Xi_{(I,J)} = \Xi_{(I,J)}(\bar{t}_i, A_{I,\Delta_1} \hat{x}(k))$.

5.2 Hybrid transition during a posteriori update

Next, we consider the case where the measurement update pulls the mean estimate into a guard set (4), i.e. if $\exists J$ s.t. $\hat{x}(k+1|k+1) \in G_{(I,J)}$. In that case, the *a posteriori* update is modified by applying the reset to the mean and the saltation update to the covariance,

$$\tilde{x}(k+1|k) = R_{(I,J)} \hat{x}(k+1|k) \quad (35)$$

$$\tilde{\Sigma}(k+1|k) = \Xi_{(I,J)} \hat{\Sigma}(k+1|k) \Xi_{(I,J)}^T + W_{R_{(I,J)}} \quad (36)$$

where the saltation matrix is evaluated at $\Xi_{(I,J)} = \Xi_{(I,J)}(\bar{t}_i, \hat{x}(k+1|k))$. These $\tilde{x}(k+1|k)$ and $\tilde{\Sigma}(k+1|k)$ are the updated *a posteriori* mean and covariance in the new hybrid domain, J . Note that this update is identical to (31)–(32).

5.3 Extended Kalman Filter

Similar to the Kalman filter, the standard Extended Kalman Filter (EKF) [32, Eqn. 2.1] can be directly applied for nonlinear hybrid systems when no transition occurs. The nonlinear stochastic dynamics are given by

$$x(k+1) = f_{I,\Delta}(x(k), u(k), \omega(k)) \quad (37)$$

$$\hat{A}_{I,\Delta} = D_x f_{I,\Delta}(x(k), u(k), \omega(k)) \quad (38)$$

$$\hat{W}_{I,\Delta} = D_\omega f_{I,\Delta}(x(k), u(k), \omega(k)) \quad (39)$$

$$y(k) = h_I(x(k)) + v_I(k) \quad (40)$$

$$\hat{C}_I = D_x h_I(x(k)) \quad (41)$$

where $f_{I,\Delta}$ is the discrete nonlinear update for the continuous dynamics F_I , $\hat{A}_{I,\Delta}$ is the linear approximation of the dynamic matrix, $\hat{B}_{I,\Delta}$ is the linear approximation of the process noise input matrix, h_I is the measurement function and \hat{C}_I is the linear approximation of the measurement matrix.

When there is a hybrid transition during the *a priori* update, the dynamic updates for the nonlinear transition case (16)–(17) are substituted in the same manner as the linear case into (33)–(32). When there is a hybrid transition during the *a posteriori* update, the mean update equation (35) is applied with the full nonlinear reset map, while the covariance update (36) is the same for both the linear and nonlinear hybrid systems because the saltation matrix is already a linearization. With these updates, the nonlinear extension to the Salted Kalman Filter follows naturally.

5.4 Summary and pseudocode

The Salted Kalman Filter (SKF) as presented above is summarized in Algorithm 1. Note that the only difference from the standard Kalman Filter algorithm is

Algorithm 1 Salted Kalman Filter (SKF)

```

1: input  $(t_k, x_k, \Sigma_k, m_k, y_{k+1})$ 
2:  $\hat{t} \leftarrow t_k, \hat{x} \leftarrow x_k, \hat{\Sigma} \leftarrow \Sigma_k, I \leftarrow m_k$ 
3: while  $(\hat{t} < t_k + \Delta)$  do
4:    $(\hat{t}^+, \hat{x}) \leftarrow$  integrate  $F_I(\hat{t}, \hat{x})$ 
      until  $(\hat{t}^+ = t_k + \Delta)$  or  $(\exists J$  s.t.  $\hat{x} \in G_{(I,J)})$ 
5:    $\Delta_1 \leftarrow \hat{t}^+ - \hat{t}, \hat{t} \leftarrow \hat{t}^+$ 
6:    $\Sigma_k \leftarrow A_{I,\Delta_1} \hat{\Sigma} A_{I,\Delta_1}^T + W_{I,\Delta_1}$  ▷ (26)
7:   if  $\exists J$  s.t.  $\hat{x} \in G_{(I,J)}$  then
8:      $\hat{x} \leftarrow R_{(I,J)}(\hat{t}, \hat{x})$  ▷ (31)
9:      $\hat{\Sigma} \leftarrow \Xi_{(I,J)} \hat{\Sigma} \Xi_{(I,J)}^T + W_{R_{(I,J)}}$  ▷ (32)
10:     $I \leftarrow J$ 
11:   end if
12: end while
13:  $K \leftarrow \hat{\Sigma} C_I^T [C_I \hat{\Sigma} C_I^T + V_I]^{-1}$  ▷ (27)
14:  $\hat{x} \leftarrow \hat{x} + K [y_{k+1} - C_I \hat{x}]$  ▷ (28)
15:  $\hat{\Sigma} \leftarrow \hat{\Sigma} - K C_I \hat{\Sigma}$  ▷ (29)
16: if  $\exists J$  s.t.  $\hat{x} \in G_{(I,J)}$  then
17:    $\hat{x} \leftarrow R_{(I,J)}(\hat{t}, \hat{x})$  ▷ (35)
18:    $\hat{\Sigma} \leftarrow \Xi_{(I,J)} \hat{\Sigma} \Xi_{(I,J)}^T + W_{R_{(I,J)}}$  ▷ (36)
19:    $I \leftarrow J$ 
20: end if
21:  $t_{k+1} \leftarrow \hat{t}, x_{k+1} \leftarrow \hat{x}, \Sigma_{k+1} \leftarrow \hat{\Sigma}, m_{k+1} \leftarrow I$ 
22: return  $(t_{k+1}, x_{k+1}, \Sigma_{k+1}, m_{k+1})$ 

```

applying the proposed moment updates when the estimated state satisfies the guard condition (lines 7–11 and 16–20). The SKF is in many ways similar to the EKF because the saltation matrix is a linearization about the hybrid transition – if the transition is linear or the prediction is close to the actual then the filter performs well. This property holds for the nonlinear Extended SKF as well, and in general this filter suffers from the same pitfalls as the EKF. Furthermore, like the EKF this linearization means that the optimal belief may not remain Gaussian, and thus that the filter may fail to have the optimally properties we obtain in the linear case.

For the measurement update, if a hybrid transition is triggered, the approach presented here simply transforms the already updated estimates. However, a more accurate approach might include breaking up the measurement update into sub-updates over each domain. In this work, we assume the updates are small enough such that this isn't an issue, but as the measurement update magnitude increases, this may be worth investigating. While the filter is not optimal, like the EKF, we expect that it will perform well when the covariances and timesteps are relatively small so that the local linearizations hold.

6 Experiments

We present two sets of experiments: the mapping of the second moments of uncertainty distributions (Sec. 6.1)

and Kalman filtering of sensor feedback (Sec. 6.2). In each experiment, the proper mapping of second moments of distributions through hybrid transitions using the saltation matrix (Sec. 4.1) is compared against a naive mapping using the Jacobian of the reset map (Sec. 4.2). Experiments are performed in simulation to ensure consistency and accurate model knowledge. The first set of experiments compares the resulting distributions to the ground truth distribution using divergence metrics. The second set of experiments compares the utility of the saltation matrix as a dynamic update in a Kalman filtering setting by comparing the mean squared error of the two propagation methods in a series of Monte Carlo tests.

6.1 Covariance Propagation Experiments

This experiment seeks to validate the claim that the saltation matrix correctly maps the second moment of distributions that experience hybrid transitions by analyzing the distributions of particles simulated through hybrid transitions. Distributions are created by randomly sampling 1000 particles from a known mean and covariance. The particles are then simulated forwards through the hybrid dynamical system using Matlab's ODE45. The starting mean and covariance of the population are calculated then propagated for each timestep using the update rules shown in Section 4 both with our proposed solution (using the saltation matrix Ξ) and with the naive solution (using the Jacobian of the reset map $D_x R$). The final covariance of the population and the estimated covariances are compared using Kullback–Leibler (KL) divergence [30] – where the output is a natural unit of information (nat) therefore, the smaller the KL divergence, the more closely the distributions match.

$$KL(\Sigma_0 \parallel \Sigma_1) = \frac{1}{2} \left(\Sigma_1^{-1} \Sigma_0 - \dim(\Sigma_0) + \ln \frac{|\Sigma_1|}{|\Sigma_0|} \right) \quad (42)$$

Each hybrid system example is tested with two different initial covariances – a higher and lower value – to capture any nonlinearities.

6.2 Kalman Filter Experiments

To test the Salted Kalman Filter (SKF) as defined in Sec. 5, we ran a Monte Carlo simulation showing improvement compared to a Kalman filter that uses the Jacobian of the reset to propagate uncertainty. We call this comparison algorithm the Jacobian of the Reset Kalman Filter (JRKF), and it follows the same code as Algorithm 1 but with the saltation matrix Ξ replaced by the Jacobian of the reset map $D_x R$. Tests were run with a range of measurement noise, process noise, and differing time steps. The starting covariance, starting mean, reset covariance, chosen measurements, and simulation time were held constant between trials.

The effectiveness of the filter for each trial is evaluated by calculating the mean squared error (MSE) along a simulated trajectory,

$$\text{MSE} = \frac{1}{K} \sum_{k=1}^K ((x(t_k) - \hat{x}(t_k))^T (x(t_k) - \hat{x}(t_k))) \quad (43)$$

where K is the number of time steps, $\hat{x}(t_k)$ is the state estimate at time t_k , and $x(t_k)$ is the true state at time t_k . For each measurement noise, process noise, and time step combination, the filter is run on 1000 randomly sampled starting conditions with randomly sampled process noise and randomly sampled measurements. The same random trials are then passed to each filter for comparison. Each set of trials with the SKF and JRKF are compared using the *sign test* [12]. The null hypothesis is that the median difference between the pairs is zero,

$$H_0 : M_1 - M_2 = 0 \quad (44)$$

The sign test is chosen because the data are not normally distributed, which rules out the paired t-test, and are not necessarily symmetric, which rules out the Wilcoxon Signed Rank test.

7 Results

In this section we present the results of the experiments detailed in the previous section on a series of hybrid systems.

7.1 Hybrid System Definitions and Covariance Propagation Results

7.1.1 Constant Flow

The simplest hybrid system we examine is the case where there are two hybrid modes that are linearly separated and which have constant, but distinct, dynamics in each mode. The dynamics in the hybrid modes are defined:

$$F_1 = [1, -1]^T, \quad F_2 = [1, 1]^T \quad (45)$$

The guard sets are defined at $x_1 = 0$, such that the domain of F_1 is the left half plane and the domain of F_2 is the right half plane (Fig. 1). The reset is an identity map.

For the propagation experiment, the starting mean was $x = [-2.5, 0]^T$, the total simulation time was 5 seconds, and the time steps were 0.01 seconds as shown in Fig. 1. The high covariance level was $0.1I$ and the low covariance level was $0.005I$. All samples from the system began in hybrid mode 1 and ended in hybrid mode 2.

This scenario is interesting because the covariance does not change in either hybrid mode as the flow is simply

translation in those regions – the only covariance changes are a result of the hybrid transition. The results, listed in Table 1 and apparent in Fig. 1, show that this change in covariance is captured well using the saltation matrix (as the KL-divergence is almost zero) but not well using the Jacobian of the reset map. Note that, as expected for a linear system, the KL-divergence does not significantly depend on the initial covariance and in this case the difference is practically zero.

7.1.2 Bouncing Ball

The bouncing ball is a hybrid dynamical system [15] which consists of 2 hybrid domains in the $[y, \dot{y}]^T$ plane, where the first domain is defined when the ball has negative velocity $\dot{y} < 0$ and the second domain is defined when the ball has non-negative velocity $\dot{y} \geq 0$. The guard sets are defined to be when the ball hits the ground $g_{1,2}(t, y, \dot{y}) := y$ and when the velocity changes sign $g_{2,1}(t, y, \dot{y}) := \dot{y}$. Note that this could equivalently be defined as a single domain with a self-reset, however to match our definition of a hybrid dynamical system (Definition 1) we treat it as a system with separate domains. The dynamics are standard ballistic dynamics in both domains

$$F_1 = F_2 = [\dot{y}, -a_g]^T \quad (46)$$

where a_g is the acceleration from gravity. The reset from 1 to 2 is defined by elastic impact

$$R_{1,2} = [y, -\alpha\dot{y}]^T \quad (47)$$

where α is the coefficient of restitution. The reset from 2 to 1 is an identity transformation. For the experiments, the gravitational acceleration is $a_g = 9.8$ and the coefficient of restitution is $\alpha = 0.75$. This system is an example of a system with linear dynamics, guards, and linear but non-identity resets.

For the propagation experiment, the starting mean was $x = [3, -2]^T$, the total simulation time was 1 second, and the time steps were 0.001 seconds as shown in Fig. 2. The high covariance level was $0.05I$ and the low covariance level was $0.001I$. Most of the samples from the system began in hybrid mode 1 and ended in hybrid mode 2, while some samples ended back in hybrid mode 1 (after a (2,1) transition). There is no effect on the mean or covariance through the (2,1) transition as $R_{(2,1)}$ is identity and $F_1 = F_2$ which means that $\Xi_{(2,1)} = D_x R + 0$ is also identity.

The results are listed in Table 1 and an example trial is shown in Fig. 2. Interestingly, even though the final distribution is no longer Gaussian, the second moment (covariance) is still tracked accurately through the impact using the saltation matrix (as the KL-divergence is small) but poorly with the Jacobian of the reset map.

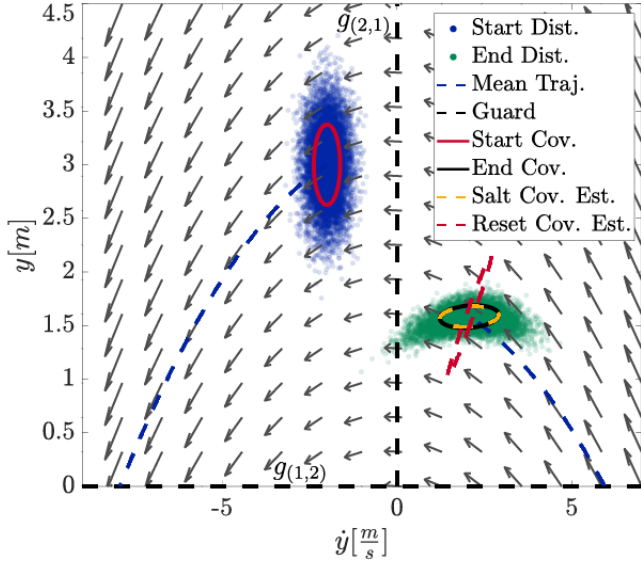


Fig. 2. Flowing an initial distribution (blue dots) with covariance (red solid line) along a nominal trajectory (blue dashed line) through the bouncing ball hybrid system’s dynamics (gray arrows) with guards (black dashed line and labeled). The final distribution (green dots) is overlaid with the actual covariance (black line). Estimated covariance using the Jacobian of the reset map (red dashed line) is compared against our proposed estimate using the saltation matrix (gold dashed line).

Again, as this is a linear system, the KL-divergence is independent of the initial covariance.

7.1.3 Pendulum hitting a spring damper

The pendulum hitting a spring damper, as shown in Fig. 3, is a hybrid system which consists of 2 hybrid domains over the $[\theta, \dot{\theta}]^T$ state space. The first domain is defined when the pendulum’s angular position is positive $\theta > 0$ and the second domain is defined when the angular position is non-positive $\theta \leq 0$, such that the guard functions are $g_{1,2}(t, \theta, \dot{\theta}) = g_{2,1}(t, \theta, \dot{\theta}) = \theta$. The dynamics are standard pendulum dynamics while in domain 1 and while in domain 2, the pendulum is in contact with a torsional spring and damper. The resulting dynamics are,

$$F_1 = \left[\dot{\theta}, -\frac{a_g}{l} \sin(\theta) \right]^T \quad (48)$$

$$F_2 = \left[\dot{\theta}, \frac{-(k\theta + c\dot{\theta} + ma_g l \sin(\theta))}{ml^2} \right]^T \quad (49)$$

where a_g is the acceleration from gravity, l is the length of the pendulum’s center of mass along the arm, and m is the mass of the pendulum. In the experiment, the constants were set to $a_g = 9.8$, $l = 1$, $m = 1$, $k = 10$, and $c = 10$. Both resets are identity transformations because there are no instantaneous changes in state during mode




Fig. 3. Pendulum hitting a spring damper hybrid system where the pendulum engages a spring damper at the $\theta = 0$.

transition. This system is nonlinear but with identity resets.

For the propagation experiment, the starting mean was $x = [\frac{\pi}{4}, -3.7]^T$, the total simulation time was 0.3 seconds, and the time steps were 0.001 seconds as shown in Fig. 1. The high covariance level was $0.05I$ and the low covariance level was $0.001I$. All samples from the system began in hybrid mode 1 and ended in hybrid mode 2. This example demonstrates the validity of the linear approximations when the higher order terms are small and as expected, using the saltation update decreases KL-divergence when compared against the Jacobian of the reset map as shown in Table 1.

7.1.4 Asymmetric Spring Loaded Inverted Pendulum (ASLIP)

The asymmetric spring loaded inverted pendulum (ASLIP) system consists of a spring leg, torsional spring hip, and a body with inertia in the sagittal plane as shown in Fig. 4. This system is similar to the one in [27] and a full derivation for the system dynamics can be found in Appendix B. In this system, the body configuration space is defined to be the position and orientation of the body $q_b := [x_b, y_b, \theta_b]^T \in \mathbb{R} \times \mathbb{R} \times \mathbb{S}^1$. The leg configuration space is defined to be the angle between the toe and the ground, the angle of the hip, and the length of the leg $q_l := [\theta_t, \theta_h, l]^T \in \mathbb{S}^1 \times \mathbb{S}^1 \times \mathbb{R}$, where impact location of the toe defines a pin joint for the body to pivot around. Once the location of the toe, $q_t = [x_t, y_t]^T \in \mathbb{R} \times \mathbb{R}$, is fixed to a ground location, either configuration can be used to define the full configuration space of the system. When the toe position is known, the transformation from the leg configuration to the body configuration $T_{lb} : (q_l, q_t) \mapsto q_b$. The inverse mapping can also be calculated which maps the body configuration to the leg configuration $T_{bl} : (q_b, q_t) \mapsto q_l$.

Hybrid mode 1 is defined to be when the toe is not in contact with the ground. The resulting domain \mathcal{D}_1 is chosen to be parameterized by the body’s configuration,

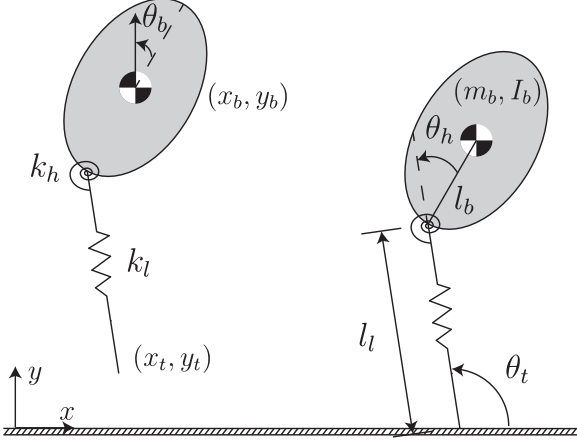


Fig. 4. Asymmetric Spring Loaded Inverted Pendulum (ASLIP) diagram showing the aerial phase hybrid mode on the left and the stance phase hybrid mode on the right and their corresponding configuration variables.

toe position, and body's velocity.

$$[x_b, y_b, \theta_b, x_t, y_t, \dot{x}_b, \dot{y}_b, \dot{\theta}_b]^T \in \mathcal{D}_1 \quad (50)$$

When the toe is in contact with the ground, the hybrid mode is 2. The domain \mathcal{D}_2 is chosen to be parameterized by the toe angle with the ground, hip angle, the leg extension, toe position, the time derivative of the toe angle, hip angle, and leg extension.

$$[\theta_t, \theta_h, l_l, x_t, y_t, \dot{\theta}_t, \dot{\theta}_h, \dot{l}_l]^T \in \mathcal{D}_2 \quad (51)$$

Note that the toe position is augmenting the state rather than being treated as an external parameter because variations in the toe placement affect the other configuration states. Because of this, the toe dynamics are constrained relative to the body in domain 1 and relative to the ground contact in domain 2. These dynamics F_1, F_2 are derived using a Lagrangian approach (see Appendix B). The system parameters and their experimental values are body mass $m_b = 1$, body inertia $I_b = 1$, leg spring constant $k_l = 1000$, hip spring constant $k_\theta = 400$, body length $l_b = 0.5$, acceleration due to gravity $a_g = 9.8$, resting leg length $l_{l0} = 1$, and resting angle of the hip spring $\theta_{h0} = -\frac{\pi}{8}$.

The guard for mode 1 is defined to be when the toe touches the ground, $g_{(1,2)}(t, q, \dot{q}) = y_t$, and the guard for mode 2 is defined to be when the normal force of the toe reaches zero, i.e when the leg spring reaches the resting length, $g_{(2,1)}(t, q, \dot{q}) = l_l - l_{l0}$. The reset maps are defined to be the coordinate changes from the body

states to the leg states.

$$R_{1,2} = \begin{bmatrix} T_{bl}(q_b) \\ q_t \\ D_x T_{bl}(q_b, q_t) \dot{q}_b \end{bmatrix} \quad (52)$$

$$R_{2,1} = \begin{bmatrix} T_{lb}(q_l) \\ q_t \\ D_x T_{lb}(q_l, q_t) \dot{q}_l \end{bmatrix} \quad (53)$$

This hybrid system is especially useful to analyze because it includes both nonlinear dynamics and non-identity resets, wherein the switch from Cartesian to polar coordinates must be accounted for.

For the propagation experiment, the starting mean was $x = [0, 1.8, \frac{5\pi i}{12}, 0.0011, 0.3256, 0 - 10]^T$, the total simulation time was 0.55 seconds, and the time steps were 0.0001 seconds. The high covariance level was 0.001 and the low covariance level was 0.00005 for the non toe states. There was no additional noise injected into the toe states, because they are constrained to the body states. The system started in hybrid mode 1 and ended in hybrid mode 2. This scenario is the most complex out of all the considered hybrid systems because it includes both nonlinear dynamics and nonidentity resets. Nevertheless, the change in covariance is still captured well using the saltation matrix and is not captured well when using the Jacobian of the reset, Table 1. Similar to the pendulum example, as the initial covariance increases, the saltation update approximation gets worse due to the inaccuracies in the linearization.

7.1.5 Propagation experiment discussion

Overall, using the saltation update estimates the covariance very well for the linear cases as shown by the KL-divergence. For the non-linear cases, the saltation update is a good linear approximation. This is especially apparent in the linear cases where increasing the initial covariance did not have much affect on the KL-divergence while in the nonlinear cases the increase in initial covariance increased the resulting KL-divergence; this is unsurprising, as the linearization is local and will generally increasingly fail to predict the correct dynamics as the domain is enlarged. These results motivate the use of the saltation update in a Kalman filtering framework.

7.2 Kalman Filter Results

7.2.1 Constant Flow

The first Kalman filtering experiment uses the same constant flow system defined in Sec. 7.1.1 and shown in Fig. 1. The system was simulated for 5 seconds with 4

Table 1

KL-divergence comparison between the ground truth and estimated second moment of distributions that undergo hybrid transitions for the 4 hybrid system test cases using the saltation matrix $\Xi_{(I,J)}$ and the Jacobian of the reset map $D_x R_{(I,J)}$.

System	Nonlinear	Non-identity Reset	Initial Covariance	KL-div $D_x R_{(I,J)}$	KL-div $\Xi_{(I,J)}$
1 Constant Flow	No	No	Low	1.9193	1.8419×10^{-6}
			High	1.9219	2.2574×10^{-7}
2 Bouncing Ball	No	Yes	Low	32.2253	0.0035
			High	32.2253	0.0035
3 Pendulum	Yes	No	Low	29.2477	0.0011
			High	32.3330	0.2058
4 ASLIP	Yes	Yes	Low	26.8931	0.0011
			High	24.8121	0.0867

different time steps: $\Delta = 5, 1, 0.1, \text{ and } 0.05$ seconds. The process covariance levels ranged from $\|W_{I,\Delta}\| = 0.1\Delta^2$ to $0.0001\Delta^2$ and the measurement covariance was swept from $\|V_I\| = 1$ to 0.0001 in powers of 10, for a total of 4 process covariance levels and 5 measurement covariance levels. There was no noise added to reset because the reset map is an identity transformation. The initial covariance was set to $\Sigma(0) = 0.1I$. Lastly, the measurements for this system were chosen to be both states, i.e

$$h_I(x) = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} x = Cx \quad (54)$$

In total, 80 Monte Carlo simulations were tested with 1000 trials each.

Both filters perform well for the given noise levels and timesteps because they both have relatively low mean squared error for each Monte Carlo simulation. An example Monte Carlo simulation trial is shown in Fig. 5 where both filters seem to perform similarly, but the difference is seen when comparing the absolute mean error, and upper and lower quartiles – where the Kalman filter using the saltation update has a lower mean error. When comparing the mean squared error for each batch of Monte Carlo simulations, the saltation updated Kalman filter performed better than or as well as the Jacobian of the reset map. The saltation updated filter improved performance with statistical significance ($p < 0.05$) for 68 of the 80 combinations. In the 12 remaining cases they failed to reject the null hypothesis (44) and are statistically indistinguishable. For these cases, the time steps are high, the measurement noises are low, and the process noises are high. In these cases the filter depends mostly on the sensors and therefore the dynamic update becomes less important.

7.2.2 ASLIP

The Kalman filtering experiments were also run on the ASLIP system, defined in Sec. 7.1.4. For these tests, we simulated the dynamics for 2.5 seconds which resulted in 4 jumps (8 hybrid transitions). Timesteps were chosen to be $\Delta = 0.03, 0.01, 0.005, 0.001$ seconds. The process noise covariance levels were $\|W_{I,\Delta}\| = 0.01\Delta^2, 0.001\Delta^2, \text{ and } 0.0001\Delta^2$ (which were applied as wrenches), and the measurement noise covariance levels were $\|V_I\| = 0.005, 0.001, \text{ and } 0.0001$. The initial covariance was set to be $0.1I$, where the covariance in the toe position was set to match the constraint between the body configuration and the toe – this is necessary because the noise in the toe position is correlated to the noise in the body states. Again, reset noise is not applied because there isn't uncertainty in the coordinate transformation. For this system, only measurements of the body states are given. This is more realistic, because it is assumed that the hybrid mode is unknown. Therefore, in the aerial phase, hybrid mode 1, the measurement function is simply.

$$h_1(x) = \begin{bmatrix} q_b \\ \dot{q}_b \end{bmatrix} \quad (55)$$

However, in the stance phase, hybrid mode 2, the states are the leg states and the toes positions and cannot be directly compared against the body measurements. Therefore, the measurement function in hybrid mode 2 is the transformation from leg states to body states

$$h_2(x) = \begin{bmatrix} T_{lb}(q_l) \\ T_{lb}(q_l, q_t, \dot{q}_l) \end{bmatrix} \quad (56)$$

Both estimators have relatively low mean squared error for each Monte Carlo Simulation indicating that each filter performs well for these noise levels and timesteps. When comparing the two filters, the saltation updated Kalman filter outperformed the Jacobian of the reset

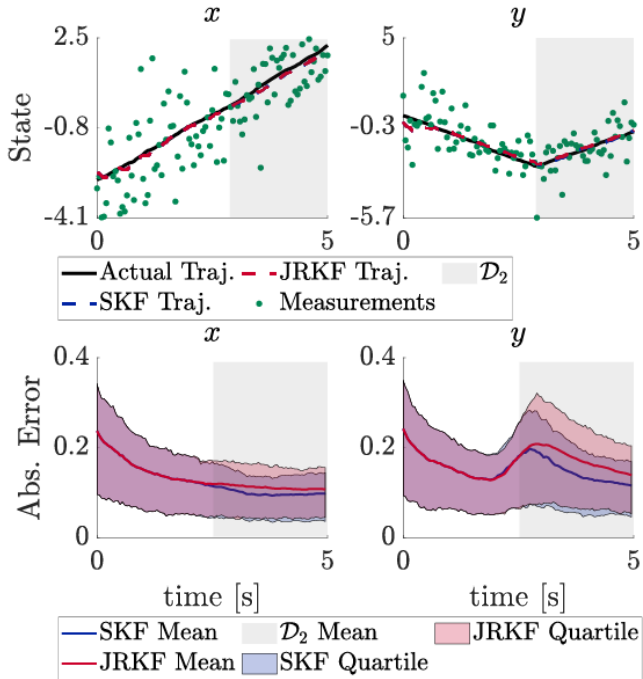


Fig. 5. Kalman filter results on the constant flow system comparing the SKF to the JRKF where the two methods differ only at the moment of transition. Note that the main differences are just after the transition because Kalman Filters are stable, these differences disappear as time goes on. Testing conditions for this example are timestep $\Delta = 0.05\text{s}$, process noise $\|W_{I,\Delta}\| = 0.1\Delta^2$, and measurement noise $\|V_I\| = 1$. Top: For a single trial, the SKF estimated trajectory (blue dashed), the JRKF estimated trajectory (red dashed), and the ground truth trajectory (black solid) are shown with the measurements (green dots) and highlighting (gray shaded) when the system is in \mathcal{D}_2 . Bottom: Absolute mean error is plotted for the SKF (blue) and JRKF (red) with 0.25 and 0.75 quantiles over 1000 trials, highlighting (gray shaded) the mean transition time to \mathcal{D}_2 .

map update for all 36 combination of $(\Delta, W_{I,\Delta}, V_I)$, with statistical significance ($p < 0.001$). We again expect the SKF to perform better than or equal to the JRKF because as the quality of measurements increase, the less we use the knowledge of the dynamics. An example run of one of the Monte Carlo simulations is shown in Fig. 6, where both filters again perform similarly for each state except for the vertical body position y_b . The deviation from the ground truth vertical body position y_b are more pronounced for the naive solution than our proposed solution.

This effect is magnified when comparing the absolute error mean and quartiles for the 1000 trials as shown in Fig. 6. It is clear that most of the improvement is in the body vertical position y_b . The difference between the saltation matrix and the Jacobian of the reset map on impact is in the column associated with the vertical height y_b . Therefore, the improvements in y_b are expected be-

cause the dynamics along this axis are accounted for. Similarly on lift off, the difference was in the column associated with the leg length. As the leg was close to vertical, most of the error was shifted into the vertical height of the body y_b .

7.2.3 Kalman Filter Experiment Discussion

The Kalman Filter experiments span from simple linear (even constant) dynamics to more complex nonlinear and nonidentity dynamics. The constant flow case is beneficial to analyze because the linearization is exact for the hybrid transition as well as the dynamics. Therefore, the initial conditions as well as noise parameters should not affect the fundamental quality of the estimator. The ASLIP example validates the extended Kalman filter variant through analyzing a nominal trajectory with several hybrid transitions with non-identity resets and highly nonlinear dynamics.

Both the naive solution, using the Jacobian of the reset map, and our proposed solution, using the saltation matrix, have low mean squared error for both examples. Since the naive solution and our proposed method have the same mean update and algorithm structure, the fact that they both perform well highlights the importance of having an accurate update for the mean as well as handling each transition case in the algorithm. However, for both examples using the saltation covariance update improves the mean squared error and is directly visible in the absolute error distribution plots, Figs. 5 and 6.

8 Conclusion

In this paper, we introduced the concept of utilizing saltation matrices for mapping the second moment of probability distributions through hybrid transitions. The mapping was validated through a variety of hybrid dynamical system cases and compared against a naive method of using the Jacobian of the reset map. We also apply this approach to create a new Kalman filtering algorithm which allows estimation on hybrid dynamical systems with state based transitions, including an extended Kalman filter variant which can handle nonlinear dynamics with non-identity reset maps. This Salted Kalman Filter was validated on both a linear and nonlinear system and then compared against the Jacobian of the reset map counterpart. The results show that using our proposed method is statistically better than or equivalent to the naive method in all tested cases.

The proposed method, similar to the Extended Kalman Filter, suffers when model uncertainty is added to the hybrid dynamical system, when the local approximation is violated, or when the noise is non-Gaussian as stated in Sec. 3. Furthermore, as noted in the bouncing ball example (Sec. 7.1.2), even in the case of a linear hybrid dynamical system an uncertainty distribution that

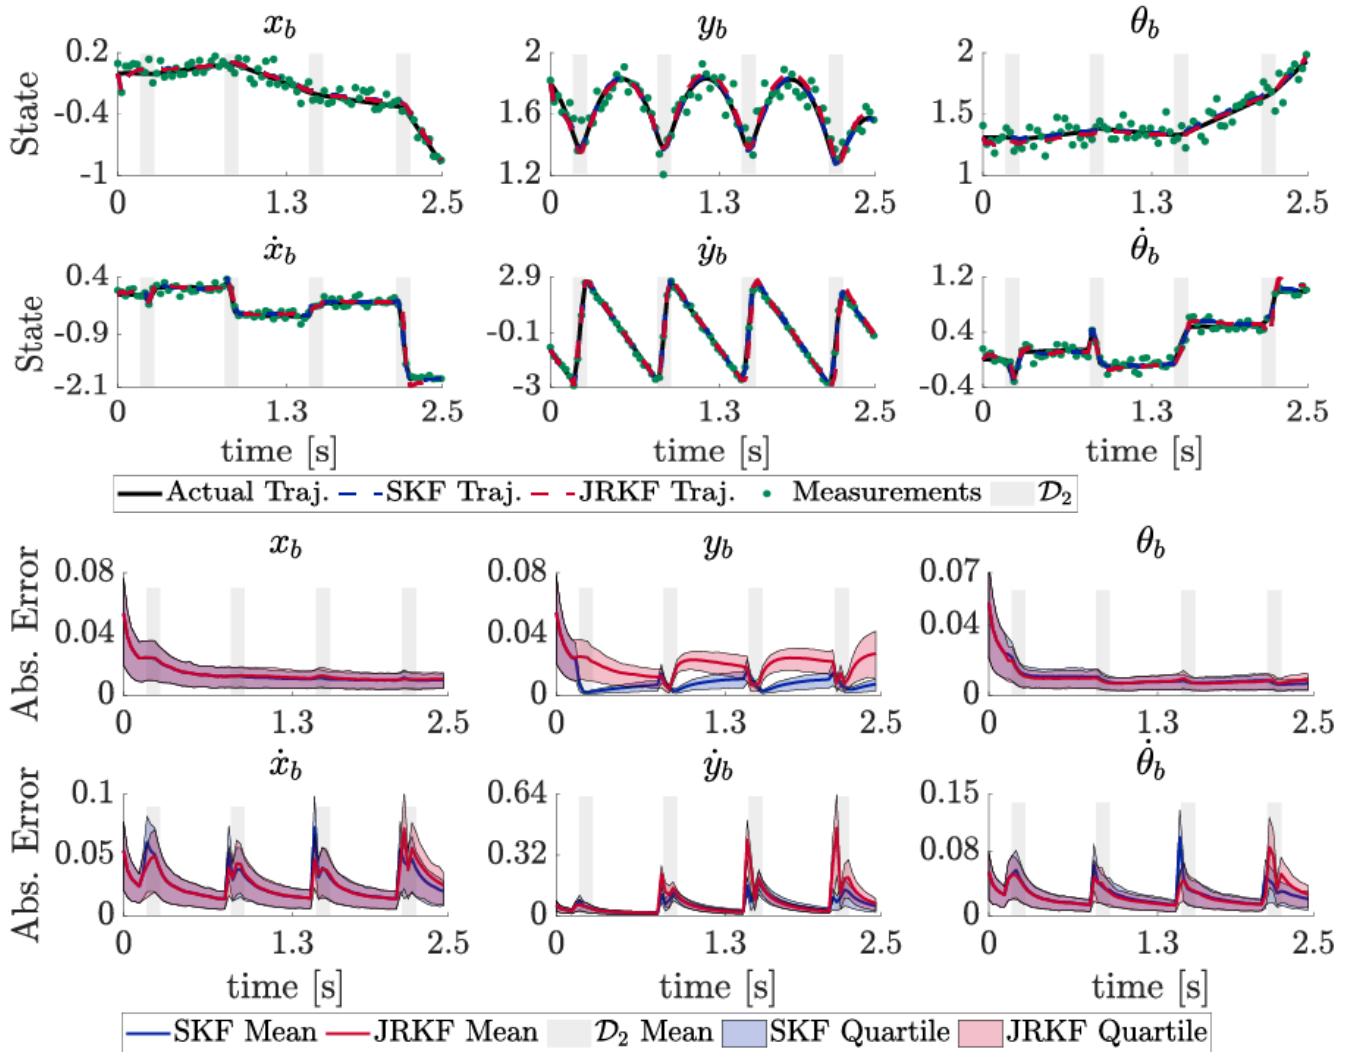


Fig. 6. ASLIP Kalman filter results comparing the SKF to the JRKF. Note that the main differences between the methods are at the transitions and also that the improvement is in one direction (here, mostly in the vertical body position y_b) because the saltation matrix is different from the Jacobian of the reset map by a rank 1 update. Testing conditions for this example are timestep $\Delta = 0.03\text{s}$, measurement noise $\|V_I\| = 0.005$, and process noise $\|W_{I,\Delta}\| = 0.01\Delta^2$. Top: For a single trial, the SKF estimated trajectory (blue dashed), the JRKF estimated trajectory (red dashed), and the ground truth trajectory (black solid) are shown with the measurements (green dots) and highlighting (gray shaded) when the system is in \mathcal{D}_2 . Bottom: Absolute mean error is plotted for the SKF (blue) and JRKF (red) with 0.25 and 0.75 quantiles over 1000 trials, highlighting (gray shaded) the mean transition times to \mathcal{D}_2 .

is initially Gaussian may not remain so (unlike a conventional linear system), which is assumed in the classical Kalman filter. The reset map and the saltation matrix still map the first two moments of the distribution correctly, but those two alone do not describe the entire distribution. In cases where the non-linearity, non-localness, or non-Gaussianness are significant, a particle filtering approach may be more appropriate.

Note that while using the saltation matrix captures the update for the covariance to first order, the saltation

matrix is model-dependent, and may require significant effort to obtain in practice in order to use (2) directly. However, as the saltation matrix is a linear map relating pre- and post-transition states, regression techniques may be able to approximate it with measured data without the need for complete (and differentiable) models of the hybrid system.

While this is a good start to developing an online hybrid state estimation system, there is still further work needed to improve the estimation. Our method does not

explicitly reason about the probability of a state or measurement being in a particular hybrid mode or guard, and an extension that reasons about this probability will be covered in future work. Additionally, we are also looking into the behaviors of distributions which pass through multiple intersections, in which case an extension based on the Bouligand derivative [11, 28] might be used to capture the propagation of uncertainty.

A Proof for Proposition 1

Proof: To find the mean of $X(\underline{t}_{i+1})$, we take the expectation of $X(\underline{t}_{i+1})$

$$\mu(\underline{t}_{i+1}) = \mathbb{E}[X(\underline{t}_{i+1})] = \mathbb{E}[x(\underline{t}_{i+1}) + \delta x(\underline{t}_{i+1})] \quad (\text{A.1})$$

$$= \mathbb{E}[x(\underline{t}_{i+1})] + \mathbb{E}[\delta x(\underline{t}_{i+1})] \quad (\text{A.2})$$

where the two terms are separable because the expectation is a linear operator, and the expectation post impact state is just its value, $\mathbb{E}[x(\underline{t}_{i+1})] = x(\underline{t}_{i+1})$. Substituting in $\delta x(\underline{t}_{i+1}) = \Xi_{(I,J)}(\bar{t}_i, x(\bar{t}_i))\delta x(\bar{t}_i)$.

$$\mu(\underline{t}_{i+1}) = x(\underline{t}_{i+1}) + \mathbb{E}[\Xi_{(I,J)}(\bar{t}_i, x(\bar{t}_i))\delta x(\bar{t}_i)] \quad (\text{A.3})$$

Again, because of the linear properties of the expectation we can pull out $\Xi_{(I,J)}(\bar{t}_i, x(\bar{t}_i))$ and because $\delta x(\bar{t}_i)$ is centered about zero, $\mathbb{E}[\delta x(\bar{t}_i)] = 0$

$$\mu(\underline{t}_{i+1}) = x(\underline{t}_{i+1}) + \Xi_{(I,J)}(\bar{t}_i, x(\bar{t}_i)) \cdot 0 = x(\underline{t}_{i+1}) \quad (\text{A.4})$$

By applying the reset map from (12), the mean is thus,

$$\mu(\underline{t}_{i+1}) = x(\underline{t}_{i+1}) = R_{(I,J)}x(\bar{t}_i) \quad (\text{A.5})$$

Using the definition of covariance,

$$\text{COV}[X] = \mathbb{E}[(X - \mathbb{E}[X])(X - \mathbb{E}[X])] \quad (\text{A.6})$$

the post-impact covariance $\Sigma(\underline{t}_{i+1})$ is

$$\begin{aligned} \Sigma(\underline{t}_{i+1}) &= \text{COV}[X(\underline{t}_{i+1})] \\ &= \text{COV}[x(\underline{t}_{i+1}) + \delta x(\underline{t}_{i+1})] \end{aligned} \quad (\text{A.7})$$

$$\begin{aligned} &= \mathbb{E}[(x(\underline{t}_{i+1}) + \delta x(\underline{t}_{i+1}) - \mu(\underline{t}_{i+1})) \\ &\quad (x(\underline{t}_{i+1}) + \delta x(\underline{t}_{i+1}) - \mu(\underline{t}_{i+1}))^T] \end{aligned} \quad (\text{A.8})$$

Since $\mu(\underline{t}_{i+1}) = x(\underline{t}_{i+1})$, this simplifies to

$$\Sigma(\underline{t}_{i+1}) = \mathbb{E}[\delta x(\underline{t}_{i+1})\delta x(\underline{t}_{i+1})^T] \quad (\text{A.9})$$

Using (3), $\delta x(\underline{t}_{i+1})$ can be substituted with $\Xi_{(I,J)} = \Xi_{(I,J)}(\bar{t}_i, x(\bar{t}_i))$ which results in

$$\Sigma(\underline{t}_{i+1}) = \mathbb{E}[(\Xi_{(I,J)}\delta x(\bar{t}_i))(\Xi_{(I,J)}\delta x(\bar{t}_i))^T] \quad (\text{A.10})$$

$$= \Xi_{(I,J)}\Sigma(\bar{t}_i)(\Xi_{(I,J)})^T \quad (\text{A.11})$$

■

B Derivation of the ASLIP system

The change of coordinate functions are,

$$T_{lb}(q_b, q_t) = \begin{bmatrix} l_l \cos(\theta_t) + l_b \cos(\theta_t + \theta_h) + x_t \\ l_l \sin(\theta_t) + l_b \sin(\theta_t + \theta_h) + y_t \\ \theta_t + \theta_h \end{bmatrix} \quad (\text{B.1})$$

$$T_{bl}(q_b, q_t) = \begin{bmatrix} \text{atan}\left(\frac{y_b - (l_b \sin(\theta_b) + y_t)}{x_b - (l_b \cos(\theta_b) + x_t)}\right) \\ \theta_b - \text{atan}\left(\frac{y_b - (l_b \sin(\theta_b) + y_t)}{x_b - (l_b \cos(\theta_b) + x_t)}\right) \\ \sqrt{(y_b - l_b \sin(\theta_b) - y_t)^2 + (x_b - l_b \cos(\theta_b) - x_t)^2} \end{bmatrix} \quad (\text{B.2})$$

The differential mappings are defined via chain rule

$$\begin{bmatrix} \dot{q}_b \\ \dot{q}_t \end{bmatrix} = D_{q_b, q_t} T_{lb}(q_b, q_t) \begin{bmatrix} \dot{q}_l \\ \dot{q}_t \end{bmatrix} \quad (\text{B.3})$$

$$\begin{bmatrix} \dot{q}_l \\ \dot{q}_t \end{bmatrix} = D_{q_l, q_t} T_{bl}(q_l, q_t) \begin{bmatrix} \dot{q}_b \\ \dot{q}_t \end{bmatrix} \quad (\text{B.4})$$

Since the toe is massless, the velocity of the toe is assumed to be zero when mapping velocities and is therefore removed from the differential mapping.

$$\dot{q}_l = D_{q_b} T_{lb}(q_b, q_t)\dot{q}_b, \quad \dot{q}_b = D_{q_l} T_{bl}(q_l, q_t)\dot{q}_t \quad (\text{B.5})$$

The dynamics for mode 1 are ballistic dynamics for the center of mass and because the toe is massless, both the hip and leg springs are kept at their resting locations θ_{h0} and l_{l0} respectively. Therefore, while in mode 1, the toe is kinematically constrained by the body configuration. Define $T_{bt} : q_b \mapsto q_t$ to be the transformation from the body configuration to the toe configuration

$$T_{bt}(q_b) = \begin{bmatrix} x_b - l_b \cos(\theta) - l_{l0} \cos(\theta_{h0} - \theta_b) \\ y_b - l_b \sin(\theta) + l_{l0} \sin(\theta_{h0} - \theta_b) \end{bmatrix} \quad (\text{B.6})$$

The velocity constraint is enforced through the differential mapping of T_{bt}

$$\dot{q}_t = DT_{bt}(q_b)\dot{q}_b = \begin{bmatrix} \dot{x}_b + \dot{\theta}_b(l_b \sin \theta - l_{l0} \sin(\theta_{h0} - \theta_b)) \\ \dot{y}_b - \dot{\theta}_b(l_b \cos \theta + l_{l0} \cos(\theta_{h0} - \theta_b)) \end{bmatrix} \quad (\text{B.7})$$

Therefore, the dynamics for mode 1 are

$$F_1 = \begin{bmatrix} \dot{x}_b \\ \dot{y}_b \\ \dot{\theta}_b \\ \dot{x}_b + \dot{\theta}_b(l_b \sin(\theta) - l_{l0} \sin(\theta_{h0} - \theta_b)) \\ \dot{y}_b - \dot{\theta}_b(l_b \cos(\theta) + l_{l0} \cos(\theta_{h0} - \theta_b)) \\ 0 \\ -a_g \\ 0 \end{bmatrix} \quad (\text{B.8})$$

The dynamics for mode 2 are derived using Lagrangian dynamics where the Lagrangian is defined to be the difference between the kinetic and potential energy.

$$\begin{aligned} \mathcal{L} = & \frac{1}{2}(m_b \dot{x}_b^2 + m_b \dot{y}_b^2 + I_b \dot{\theta}_b^2) \\ & - m_b a_g (l_t \sin(\theta_t) - l_b \sin(\theta_t + \theta_h)) \\ & - \frac{1}{2}(k(l_{l0} - l_t)^2 + k_h(\theta_{h0} - \theta_h)^2) \end{aligned} \quad (\text{B.9})$$

The body states are transformed to the leg states using T_{bl} and DT_{bl} . Also, in this example, the toe cannot penetrate the ground and a no slip condition is added. Therefore, the dynamics for the toe are calculated separately.

References

- [1] Mark A Aizerman and Felix R Gantmacher. Determination of stability by linear approximation of a periodic solution of a system of differential equations with discontinuous right-hand sides. *The Quarterly Journal of Mechanics and Applied Mathematics*, 11(4):385–398, 1958.
- [2] Allen Back, J. M. Guckenheimer, and Mark Myers. A dynamical simulation facility for hybrid systems. In *Hybrid Systems*, volume 736 of *Lecture Notes in Computer Science*, pages 255–267. Springer Berlin / Heidelberg, 1993.
- [3] Andrea Balluchi, Luca Benvenuti, Maria D Di Benedetto, and Alberto Sangiovanni-Vincentelli. The design of dynamical observers for hybrid systems: Theory and application to an automotive control problem. *Automatica*, 49(4):915–925, 2013.
- [4] Andrea Balluchi, Luca Benvenuti, Maria D Di Benedetto, and Alberto L Sangiovanni-Vincentelli. Design of observers for hybrid systems. In *International Workshop on Hybrid Systems: Computation and Control*, pages 76–89. Springer, 2002.
- [5] Nabil Barhoumi, Faouzi Msahli, Mohamed Djemaï, and Krishna Busawon. Observer design for some classes of uniformly observable nonlinear hybrid systems. *Nonlinear Analysis: Hybrid Systems*, 6(4):917–929, 2012.
- [6] G. Bledt, P. M. Wensing, S. Ingersoll, and S. Kim. Contact model fusion for event-based locomotion in unstructured terrains. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4399–4406, 2018.
- [7] Michael Bloesch, Marco Hutter, Mark A Hoepffinger, Stefan Leutenegger, Christian Gehring, C David Remy, and Roland Siegwart. State estimation for legged robots-consistent fusion of leg kinematics and IMU. In *Robotics: Science and Systems*, pages 17–24, 2012.
- [8] Henk AP Blom and Yaakov Bar-Shalom. The interacting multiple model algorithm for systems with markovian switching coefficients. *IEEE Transactions on Automatic Control*, 33(8):780–783, 1988.
- [9] Henk AP Blom and Edwin A Bloem. Particle filtering for stochastic hybrid systems. In *IEEE Conference on Decision and Control*, volume 3, pages 3221–3226, 2004.
- [10] Samuel A. Burden, Thomas Libby, and Samuel D. Coogan. On contraction analysis for hybrid systems, 2018. arXiv:1811.03956.
- [11] Samuel A Burden, S Shankar Sastry, Daniel E Koditschek, and Shai Revzen. Event-selected vector field discontinuities yield piecewise-differentiable flows. *SIAM Journal on Applied Dynamical Systems*, 15(2):1227–1267, 2016.
- [12] Wilfrid J Dixon and Alexander M Mood. The statistical sign test. *Journal of the American Statistical Association*, 41(236):557–566, 1946.
- [13] Wendy Y Eras-Herrera, Alexandre R Mesquita, and Bruno OS Teixeira. Equality-constrained state estimation for hybrid systems. *IET Control Theory & Applications*, 13(13):2018–2028, 2019.
- [14] Giancarlo Ferrari-Trecate, Domenico Mignone, and Manfred Morari. Moving horizon estimation for hybrid systems. *IEEE Transactions on Automatic Control*, 47(10):1663–1676, 2002.
- [15] Rafal Goebel, Ricardo G Sanfelice, and Andrew R Teel. Hybrid dynamical systems. *IEEE control systems magazine*, 29(2):28–93, 2009.
- [16] Ross Hartley, Maani Ghaffari, Ryan M Eustice, and Jessy W Grizzle. Contact-aided invariant extended kalman filtering for robot state estimation. *The International Journal of Robotics Research*, 39(4):402–430, 2020.
- [17] Morris W Hirsch, Stephen Smale, and Robert L Devaney. *Differential equations, dynamical systems, and an introduction to chaos*. Academic press, 2012.
- [18] Ian A Hiskens and MA Pai. Trajectory sensitivity analysis of hybrid systems. *IEEE Transactions on Circuits and Systems I: Fundamental Theory and Applications*, 47(2):204–220, 2000.
- [19] Inseok Hwang, Hamsa Balakrishnan, and Claire Tomlin. State estimation for hybrid systems: applications to aircraft tracking. *IET Proceedings-Control Theory and Applications*, 153(5):556–566, 2006.
- [20] Aaron M Johnson, Samuel A Burden, and Daniel E Koditschek. A hybrid systems model for simple manipulation and self-manipulation systems. *The International Journal of Robotics Research*, 35(11):1354–1392, 2016.
- [21] D. Joyce. On manifolds with corners. In *Advances in Geometric Analysis*, volume 21 of *Advanced Lectures in Mathematics*, pages 225–258. International Press of Boston, Inc., 2012.
- [22] Hassan K Khalil and Jessy W Grizzle. *Nonlinear systems*, volume 3. Prentice hall Upper Saddle River, NJ, 2002.
- [23] Michael C. Koval, Nancy S. Pollard, and Siddhartha S. Srinivasa. Pose estimation for planar contact manipulation with manifold particle filters. *The International Journal of Robotics Research*, 34(7):922–945, 2015.
- [24] J. M. Lee. *Introduction to smooth manifolds*. Springer-Verlag, New York, 2012.

- [25] Remco I Leine and Henk Nijmeijer. *Dynamics and bifurcations of non-smooth mechanical systems*, volume 18. Springer Science & Business Media, 2013.
- [26] J. Lygeros, K. H. Johansson, S. N. Simic, J. Zhang, and S. S. Sastry. Dynamical properties of hybrid automata. *IEEE Transactions on Automatic Control*, 48(1):2–17, 2003.
- [27] I. Poulakakis and J. W. Grizzle. The spring loaded inverted pendulum as the hybrid zero dynamics of an asymmetric hopper. *IEEE Transactions on Automatic Control*, 54(8):1779–1793, 2009.
- [28] Stefan Scholtes. *Introduction to piecewise differentiable equations*. Springer Science & Business Media, 2012.
- [29] Sarjoun Skaff, Alfred Rizzi, Howie Choset, and Pei-Chun Lin. A context-based state estimation technique for hybrid systems. In *IEEE International Conference on Robotics and Automation*, pages 3935–3940, April 2005.
- [30] Tim Van Erven and Peter Harremoos. Rényi divergence and kullback-leibler divergence. *IEEE Transactions on Information Theory*, 60(7):3797–3820, 2014.
- [31] Patrick Varin and Scott Kuindersma. A constrained kalman filter for rigid body systems with frictional contact. In *International Workshop on the Algorithmic Foundations of Robotics (WAFR)*, 2018.
- [32] Greg Welch and Gary Bishop. An introduction to the kalman filter. Technical Report 95-041, Department of Computer Science, University of North Carolina at Chapel Hill, 1995. Updated: July 24, 2006.
- [33] Jize Zhang, Andrew M. Pace, Samuel A. Burden, and Aleksandr Aravkin. Offline state estimation for hybrid systems via nonsmooth variable projection. *Automatica*, 115:108871, 2020.