Active Estimation for Switching Linear Dynamic Systems

Lars Blackmore and Senthooran Rajamanoharan and Brian C. Williams

***Abstract*— Switching Linear Dynamic Systems are convenient models for systems that exhibit both continuous dynamics and discrete mode changes. Estimating the hybrid discrete-continuous state of these systems is important for control and fault detection. Existing solutions for hybrid estimation approximate the belief state by maintaining a subset of the possible discrete mode sequences. This approximation can cause the estimator to lose track of the true mode sequence when the effects of discrete mode changes are subtle.**

In this paper we present a method for *active* hybrid estimation, where control inputs can be designed to discriminate between possible mode sequences. By probing the system for the purposes of estimation, such a sequence of control inputs can greatly reduce the probability of losing the true mode sequence compared to a nominal control sequence. Furthermore, by using a constrained finite horizon optimization formulation, we are able to guarantee that a given control task is achieved, while optimally detecting the hybrid state.

I. Introduction

Stochastic hybrid discrete-continuous models have been used to represent a large number of physical and biological systems, from Mars rovers to dancing bees[1][2] [3][4]. In these models, the system dynamics depend on which discrete mode the system is in, and discrete mode transitions occur stochastically. Typically the continuous and discrete state is only partial observable, which means that estimation of the hybrid system state is a challenging problem. However since tasks such as robot fault detection and pilot intent recognition can be posed as hybrid state estimation problems, such problems are of great interest.

Exact state estimation in such systems is, in general, intractable[5]. A number of tractable algorithms have been proposed that approximate the true belief state[6][7][8]. One common approach is to store a finite subset of the possible discrete mode sequences[9][10]. However, by approximating the true belief state it is possible to lose track of the true mode sequence, at which point the estimator diverges. Previous work has highlighted this problem and suggested a number of solutions[2][11].

These approaches are ‘passive’ in the sense that they attempt to do the best possible with the observations that are made available during nominal operation. In many cases, however, it is possible to to obtain a great deal more information about the state of a hybrid system by issuing appropriate control inputs. For example, in the case of detecting a fault in a drive motor, a change in the motor dynamics will not be apparent in the observations unless some effort is requested from that motor.

In this paper we introduce an *active* hybrid estimation approach that generates control inputs to minimize the probability of the estimator losing the true mode sequence. This approach applies to Switching Linear Dynamic Systems; here the system is described by a discrete-time stochastic linear dynamic model whose parameters depend on the discrete mode. The system switches at random between modes; the discrete mode is governed by a Markov process. Switching Linear Dynamic Systems are an important class of hybrid discrete-continuous systems that have been used in a number of applications[3][4].

Previous work designed control inputs to discriminate between a finite set of linear dynamic models[12][13]. This approach minimized a tractable upper bound on the probability of model selection error, while constraining the control inputs and expected system state to ensure that a control task was achieved. In this paper, we extend the approach to Switching Linear Dynamic Systems. The key insight is that for a given mode trajectory, the system dynamics, although time-varying, are fully known. By extending the bound derived in [13] to time-varying systems, we can create a tractable upper bound on the probability of the true mode trajectory being pruned. In a similar manner to [13], we use a constrained finite horizon control design approach to ensure that a given control task is achieved, conditioned on nominal system operation.

A finite horizon control approach such as this suffers from the fact that the number of possible mode sequences is exponential in the number of discrete modes and in the length of the design horizon. In practice this means that an active hybrid estimation approach can only consider a subset of the possible mode sequences. We therefore introduce an efficient pruning method that enables sequences that are *a priori* unlikely to contribute to the probability of losing the true mode sequence to be discarded from the control design problem. The result is a tractable optimization problem that can be solved using Sequential Quadratic Programming[14], for example.

We demonstrate the new active hybrid estimation approach using a satellite fault detection scenario and show that the new approach significantly reduces an upper bound on the probability of losing the true mode sequence.

II. Problem Statement

A. Switching Linear Dynamic Systems

In this section we define a Switching Linear Dynamic System (SLDS) and described how approximate state estimation can be carried out for such systems. We consider the following discrete-time stochastic system:

$$\begin{aligned}\mathbf{x}_{t+1} &= A(m_t)\mathbf{x}_t + B(m_t)\mathbf{u}_t + \omega_t \\ \mathbf{y}_{t+1} &= C(m_t)\mathbf{x}_{t+1} + D(m_t)\mathbf{u}_t + \nu_t\end{aligned} \tag{1}$$

where m_t is a Markov chain that evolves according to a transition matrix T such that:

$$p(m_{t+1} = j | m_t = i) = T_{ij}. \quad (2)$$

The variables ω_t and ν_t are zero-mean Gaussian white noise processes. This system is a SLDS; the continuous dynamics depend on the discrete mode m_t, which switches stochastically. There are D discrete modes, such that $m_t \in \{1, \ldots, D\}$. We define the problem of hybrid estimation in an SLDS as that of estimating the probability distribution $p(\mathbf{x}_t, m_t | \mathbf{y}_{1:t})$ over the hybrid discrete-continuous state, conditioned on the sequence of all observations $\mathbf{y}_{1:t}$. This probability can be written as a sum over all possible mode sequences that end in the mode m_t:

$$p(\mathbf{x}_t, m_t | \mathbf{y}_{1:t}) = \sum_{m_{1:t-1}} p(\mathbf{x}_t, m_{1:t} | \mathbf{y}_{1:t}). \quad (3)$$

Each summand can be further expanded as a product of the posterior probability of the discrete mode sequence $m_{1:t}$ and the posterior distribution over the continuous state, conditioned on this mode sequence:

$$p(\mathbf{x}_t, m_{1:t} | \mathbf{y}_{1:t}) = p(m_{1:t} | \mathbf{y}_{1:t}) p(\mathbf{x}_t | m_{1:t}, \mathbf{y}_{1:t}) \quad (4)$$

For a given mode sequence, the system dynamics are fully known, although time-varying. This means that the probability distribution $p(\mathbf{x}_t | m_{1:t}, \mathbf{y}_{1:t})$ can be calculated exactly using the Kalman filter recursion[15]. The probability of a given mode sequence $p(m_{1:t} | \mathbf{y}_{1:t})$ can also be calculated using the residuals in the Kalman filter equations. In principle, therefore, it is possible to calculate the distribution over the hybrid state $p(\mathbf{x}_t, m_t | \mathbf{y}_{1:t})$ exactly, yielding a sum-of-Gaussians expression. In practice, however, this *exact* hybrid state estimation is infeasible since the number of mode sequences $m_{1:t}$ grows exponentially with time and with the number of possible modes.

B. Approximate Hybrid Estimation

A large number of approximate methods have been proposed that make the problem tractable by approximating the probability $p(\mathbf{x}_t, m_t | \mathbf{y}_{1:t})$ [4][16][17]. One common approach is to discard mode sequences that have a low posterior probability $p(m_{1:t} | \mathbf{y}_{1:t})$. Such *pruning* approaches typically ensure that a fixed number of mode trajectories are tracked. In this paper, we assume that a pruning approach is used so that k individual mode sequences are tracked. While pruning is in fact usually carried out at every time step, for the purposes of finite-horizon control design, we assume that pruning is carried out at the end of the control horizon. The reasons for, and implications of this assumption are discussed in Section VIII.

It is possible for the true mode sequence to be discarded in this pruning process. If this occurs, the hybrid estimator typically diverges and the approximated state distribution no longer resembles the true distribution. Fig. 1 shows the pruning process for a time horizon of h time steps and $k = 4$ tracked mode sequences.

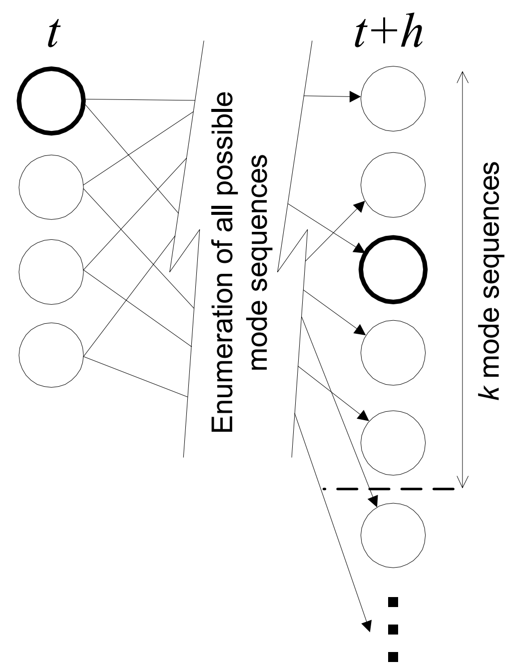

Fig. 1. Pruning approach to approximate hybrid estimation. At time t, the estimator is tracking $k = 4$ distinct mode sequences. We assume that at time $t + h$, the posterior probabilities of all possible mode sequences are calculated. The top k sequences are retained, while the remaining sequences are pruned. The true mode sequence is shown in bold; in this case it is not pruned.

Before the start of the horizon, at time t, the observations $\mathbf{y}_{t:t+h}$ are unknown. However the probability of pruning the true mode sequence $m^*_{1:t}$ can be expressed by marginalizing over all possible observations:

$$p(prune) = \int_{\mathbf{y}_{t:t+h}} p(A | \mathbf{y}_{1:t+h}) p(\mathbf{y}_{t:t+h} | \mathbf{y}_{1:t}) d\mathbf{y}_{t+1:t+h}, \quad (5)$$

where A is the event that the posterior probability of the true mode sequence $m^*_{1:t}$ is not in the top k posteriors:

$$A \Longleftrightarrow p(m^*_{1:t} | \mathbf{y}_{1:t+h}) < p(m^{(i)}_{1:t} | \mathbf{y}_{1:t+h}) \quad (6)$$

for k or more i.

The posterior probability $p(m^{(i)}_{1:t} | \mathbf{y}_{1:t+h})$ can be calculated for a given mode sequence $m^{(i)}_{1:t}$ and a given observation sequence $\mathbf{y}_{t:t+h}$, since the past observations $\mathbf{y}_{1:t-1}$ are known. The probability of a given observation sequence $p(\mathbf{y}_{t+1:t+h} | \mathbf{y}_{1:t-1})$ can be calculated as follows:

$$p(\mathbf{y}_{t+1:t+h} | \mathbf{y}_{1:t-1}) = \sum_i p(\mathbf{y}_{t+1:t+h} | m^{(i)}_{1:t+h}, \mathbf{y}_{1:t}) p(m^{(i)}_{1:t+h} | \mathbf{y}_{1:t}). \quad (7)$$

In this equation, $p(m^{(i)}_{1:t+h} | \mathbf{y}_{1:t})$ is the prior probability of the mode sequence $p(m^{(i)}_{1:t+h})$. Calculation of this value is straightforward, as described in Section IV. Conditioned on a mode sequence, the distribution $p(\mathbf{y}_{t:t+h} | m^{(i)}_{1:t+h}, \mathbf{y}_{1:t-1})$ over the observation sequences can be calculated using the Kalman filter update equations. This is described in detail in Section III. The key point is that the terms in the integral

(5) can be calculated in closed form; however the integral itself cannot be evaluated. Hence the probability of pruning the true mode sequence can not be used as a criterion for optimization. In a similar spirit to [12] and [13], in this paper we therefore derive a tractable upper bound on the probability of pruning the true mode sequence.

III. Bounding the Probability of Pruning

In [13] we introduced a new upper bound on the probability of error when selecting between an arbitrary number of stochastic linear dynamic systems. In this section we extend this work to create a bound on the probability of pruning the true mode sequence in hybrid state estimation for SLDS. The key insight behind this extension is that, enumerating each possible mode sequence over a finite horizon, the dynamics for the system conditioned on the mode sequence are fully known. Hence each mode sequence can be considered a *hypothesis* in the sense that only one mode sequence is the true one, and the approximate hybrid estimation algorithm described in Section II chooses the most *a posteriori* likely k hypotheses based on the available observations. The ideas for selection between known linear models developed in [13] therefore extend naturally to active hybrid estimation for SLDS.

In Bayes-optimal selection between several hypotheses, the most *a posteriori* likely hypothesis is selected. An error is defined as this selection rule choosing a hypothesis which is not the true hypothesis. In this case, [13] showed that for Gaussian observation distributions, such that $p(\mathbf{y}|H_i) = \mathcal{N}(\mu_i, \Sigma_i)$ for all i, an upper bound on the probability of error is given by:

$$P(error) \leq \sum_i \sum_{j>i} P(H_i)^{\frac{1}{2}} P(H_j)^{\frac{1}{2}} e^{-k(i,j)} \qquad (8)$$

where:

$$\begin{aligned} k(i,j) \quad &= \quad \frac{1}{4}(\mu_j - \mu_i)^T \left[\Sigma_i + \Sigma_j\right]^{-1} (\mu_j - \mu_i) \\ &\quad + \frac{1}{2} ln \frac{\left|\frac{\Sigma_i + \Sigma_j}{2}\right|}{\sqrt{|\Sigma_i||\Sigma_j|}}. \end{aligned} \qquad (9)$$

This upper bound can be evaluated in closed form, while the true expression for the probability of error cannot, in general. In [12] we showed how control inputs can be designed to minimize this bound, as shown in Fig. 2. Considering each mode sequence $m^{(i)}_{t:t+h}$ as a hypothesis H_i, the expression in (8) gives an upper bound on the probability that the true mode sequence is not selected as the *most likely* hypothesis. This is greater than or equal to the probability that the true mode sequence is not selected among the k most likely hypotheses, which is in turn the probability of pruning defined in (5). Hence (8) gives an upper bound on the probability of the true hypothesis being pruned, as required.

The bound in (8) applies for a general vector of observations $\mathbf{y}$ with a multivariate Gaussian distribution. In order for this bound to be tractable for optimization, however, we must be able to calculate the mean and covariance of this distribution. In [12] we showed that for selection between

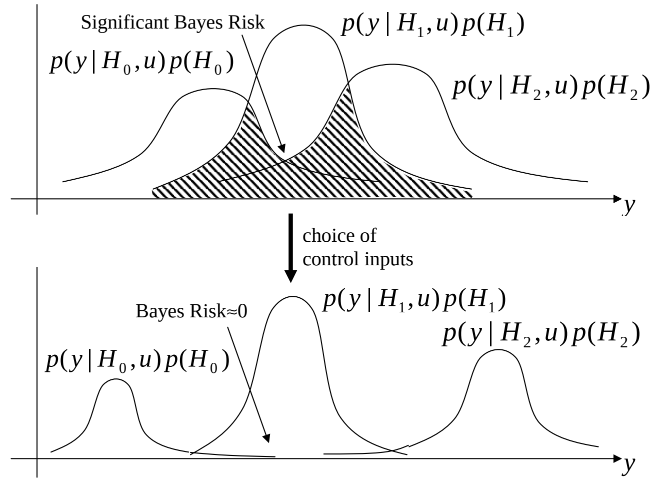

Fig. 2. Designing control inputs to minimize the probability of error (Bayes Risk) when selecting between multiple hypotheses. We employ a similar approach for active hybrid estimation by considering each mode sequence as a hypothesis.

multiple linear models, this mean and covariance for a finite horizon of observations can be calculated efficiently. Extending this work, we now show that for Switching Linear Dynamic Systems, the finite horizon observation distribution can also be calculated efficiently.

We define:

$$\begin{aligned} \mathbf{y} &= \left[\begin{array}{cccc} \mathbf{y}^T_{t+1} & \mathbf{y}^T_{t+2} & \dots & \mathbf{y}^T_{t+h} \end{array} \right]^T \\ \mathbf{u} &= \left[\begin{array}{cccc} \mathbf{u}^T_{t} & \mathbf{u}^T_{t+1} & \dots & \mathbf{u}^T_{t+h-1} \end{array} \right]^T \end{aligned} \qquad (10)$$

Each observation vector $\mathbf{y}_{t+l}$ is a random variable, which we denote Y_{t+l}. Under the assumptions in Section II, Y_{t+l} is normally distributed given an initial hybrid state estimate, a sequence of inputs $\mathbf{u}$ and a mode sequence $m^{(i)}_{1:t+h}$. We now define $\mu^{(i)}_{t+l}$ and $\Sigma^{(i)}_{t+l}$ for time steps $l = 1, \dots, h$ and all mode sequences $m^{(i)}_{1:t+h}$ such that:

$$\begin{aligned} &p_{Y_{t+l}}(\mathbf{y}_{t+l}|m^{(i)}_{1:t+h}, \mathbf{y}_{1:t}, \mathbf{u}) \\ &\qquad = \mathcal{N}(\mu^{(i)}_{t+l}, \Sigma^{(i)}_{t+l}) \end{aligned} \qquad (11)$$

We also define $\mu^{(i)}$ and $\Sigma^{(i)}$ such that

$$\begin{aligned} &p_Y(\mathbf{y}|m^{(i)}_{1:t+h}, \mathbf{y}_{1:t}, \mathbf{u}) \\ &\qquad = \mathcal{N}(\mu^{(i)}, \Sigma^{(i)}) \end{aligned} \qquad (12)$$

We define $Y^{(i)}$ as the block vector of all observations over the time horizon, conditioned on mode sequence $m^{(i)}_{1:t+h}$. The distribution of $Y^{(i)}$ is given by:

$$\mu^{(i)} = \left[\mu^{(i)T}_{t+1} \dots \mu^{(i)T}_{t+k}\right]^T \qquad (13)$$

$$\begin{aligned} &\left[\Sigma^{(i)}\right]_{f,g} \\ &= E\Big[\big([Y]_f - [\mu^{(i)}]_f\big)\big([Y]_g - [\mu^{(i)}]_g\big)\Big| m^{(i)}_{1:t+h}, \mathbf{y}_{1:t+h}\Big], \end{aligned} \qquad (14)$$

where $[\cdot]_{f,g}$ denotes the (f,g)'th index into the matrix. The mean and covariance expressions can be calculated explicity in terms of the control inputs using repeated application of the Kalman Filter update equations. Given a distribution for the state at time t such that $p(\mathbf{x}_t|\mathbf{y}_{1:t}, m_{1:t}^{(i)}) = \mathcal{N}(\hat{x}^{(i)}, P^{(i)})$, the mean $\mu^{(i)}$ and covariance $\Sigma^{(i)}$ as defined in (13) and (14) can be calculated. The results are given here for a system where $\mathbf{y_t} \in \Re^n$.

Define:

$$f = n(p-1) + q, \tag{15}$$

where:

$$1 \le q \le n \quad 1 \le p \le k \quad p,q \in \mathbb{Z}. \tag{16}$$

Then the mean is given by:

$$\begin{aligned}
[\mu^{(i)}]_f &= [\mu_{t+p}^{(i)}]_q \\
&= C(m_{t+p})\Big(\prod_{l=1}^{p} A(m_{t+l})\Big)\hat{\mathbf{x}}_0 \\
&+ C(m_{t+p})\sum_{l=0}^{p-1}\Big(\prod_{v=l}^{p-1} A(m_{t+v})\Big)B(m_{t+l})u_{t+l} \\
&+ D(m_{t+p})u_{t+p-1}
\end{aligned} \tag{17}$$

Defining $g = n(r-1) + s$ in the same manner as (15), the expression for the covariance is:

$$\begin{aligned}
&\Big[\Sigma^{(i)}\Big]_{f,g} = R'(p,r) + \\
&C(m_{t+p})\Big(\prod_{v=1}^{p} A(m_{t+v})\Big)P^{(i)}\Big(A(m_{t+w})^T\prod_{w=1}^{r}\Big)C(m_{t+r})^T + \\
&\sum_{l=0}^{m-1} C(m_{t+p})\Big(\prod_{v=l+2}^{p} A(m_{t+v})\Big)Q\Big(A(m_{t+w})^T\prod_{w=l+2}^{r}\Big)C(m_{t+r})^T
\end{aligned} \tag{18}$$

where $m = min\{f,g\}$ and:

$$R'(p,r) = \begin{cases} R & p = r \\ 0 & p \ne r. \end{cases} \tag{19}$$

We use the following notation regarding matrix products. Repeated right matrix products are denoted:

$$\Big(\prod_{i=1}^{N} A(m_i)\Big) = A(m_1)A(m_2)\dots A(m_{N-1})A(m_N), \tag{20}$$

while repeated left matrix products are denoted:

$$\Big(A(m_i)\prod_{i=1}^{N}\Big) = A(m_N)A(m_{N-1})\dots A(m_2)A(m_1). \tag{21}$$

In principle, therefore, we can use the bound in (8) as an optimization criterion for active hybrid estimation. This is in contrast to the exact value (5), which cannot be evaluated in closed form.

However the new criterion requires evaluating $\mathcal{O}(N_{seqs}^2)$ terms, where N_{seqs} is the number of mode sequences being considered. There are D^h possible mode sequences over a horizon of h time steps, which means that even for a relatively short time horizon and a modest number of possible modes, evaluating (8) is intractable. In Section IV we overcome this problem using a principled pruning approach.

IV. Considering a Subset of Possible Mode Sequences

Since it is intractable to evaluate every term in (8), we cannot use this bound in an optimization technique such as Sequential Quadratic Programming. Instead, we find a looser bound that is tractable. In this section, we describe an approach for finding the tightest possible such bound while evaluating a fixed number of the terms in (8). The key idea is to replace the terms that come from mode sequences that are *a priori* unlikely to contribute to the probability of pruning with a looser upper bound that does not depend on the control inputs. Since these terms do not depend on the control inputs, they do not need to be evaluated in the optimization.

A. A Looser Bound on the Probability of Pruning

In deriving the looser bound, we make use of Theorem 1, which we prove in the Appendix:

$$\frac{|\frac{\Sigma_i + \Sigma_j}{2}|}{\sqrt{|\Sigma_i||\Sigma_j|}} \ge 1, \tag{22}$$

where Σ_i and Σ_j are symmetric and positive definite. We now define:

$$B_{ij}(\mathbf{u}) = P(H_i)^{\frac{1}{2}}P(H_j)^{\frac{1}{2}}e^{-k(i,j)}, \tag{23}$$

and write the bound on the probability of pruning (5) using this definition:

$$P(pruning) \le \sum_i \sum_{j>i} B_{ij}(\mathbf{u}) \tag{24}$$

Following on from the result in (22), and the definition in (9), it is clear that $k(i,j) \ge 0$, and hence $e^{-k(i,j)} \le 1$. This yields the following bound on B_{ij}:

$$C_{ij} = P(H_i)^{\frac{1}{2}}P(H_j)^{\frac{1}{2}} \ge B_{ij}(\mathbf{u}) \tag{25}$$

Note that the bound C_{ij} does not depend on the control inputs $\mathbf{u}$. By replacing $B_{ij}(\mathbf{u})$ terms in (24) with C_{ij} terms, we obtain a looser upper bound on the probability of pruning the true mode sequence.

$$p(prune) \le \sum_{i\in S}\sum_{j>i, j\in S} B_{ij}(\mathbf{u}) + \sum_{i\notin S}\sum_{j>i, j\notin S} C_{ij} \tag{26}$$

Here S is the set of s mode sequences for which the full bound is calculated as a function of $\mathbf{u}$. We would like the tightest such bound for a given size of S. We achieve this as follows.

The difference between the bound C_{ij} and the bound $B_{ij}(\mathbf{u})$ is largest when the control inputs $\mathbf{u}$ drive the value of $B_{ij}(\mathbf{u})$ to zero, at which point the difference is $P(H_i)^{\frac{1}{2}}P(H_j)^{\frac{1}{2}}$. Since we do not have knowledge of $\mathbf{u}$ when choosing which mode sequences to include in S, we assume this *worst case* difference between C_{ij} and B_{ij}. In order to find the tightest bound of the form (26), we therefore include in S the mode sequences that give

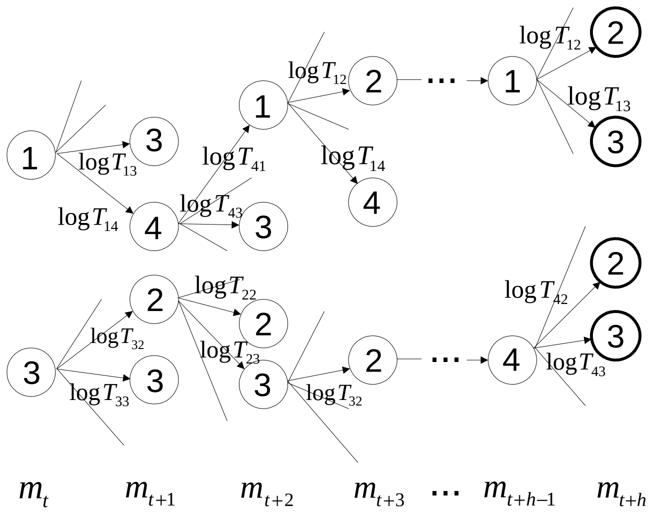

Fig. 3. Search tree for enumeration of *a priori* most likely mode sequences. At the far left there are k nodes corresponding to the mode sequences tracked by the hybrid estimator at time t. Nodes at the far right correspond to a full assignment to the modes $m_{1:t+h}$. Each arc value is the logarithm of the transition probability, and the value of a node is found by summing the value of the arcs leading to the node. Nodes are expanded in best-first order until s goal nodes are expanded. In this manner, the most likely mode sequences can be enumerated without evaluating the prior likelihood of each mode sequence, which is intractable. In this diagram $k = 2$ and $s = 4$.

the largest $P(H_i)^{\frac{1}{2}}P(H_j)^{\frac{1}{2}}$. We would like to find S to maximize:

$$L = \sum_{i \in S} \sum_{j>i, j \in S} p(H_i)^{\frac{1}{2}} p(H_j)^{\frac{1}{2}} \qquad (27)$$

Intuitively, this means that optimization will concentrate on reducing terms where the control inputs can have the greatest effect on the probability of pruning the true mode sequence, and will ignore terms that do not contribute significantly to the probability of pruning the true mode sequence. It can be seen that in order to maximize L, the set S must be chosen to contain the s hypotheses with the greatest prior probability $p(H_i)$; replacing any $p(H_i)$ value with a lower one can only reduce, or have no effect on, terms in the summation (27).

We have therefore shown how to derive a tractable upper bound on the probability of pruning the true mode sequence that involves a fixed number s of mode sequences for which the observation statistics (17) and (18) need to be calculated. By choosing the s most likely mode sequences, we achieve the tightest such bound.

B. Mode Sequence Enumeration as Best-First Search

Choosing the s most likely mode sequences is a challenging problem in itself given an exponential number of possible sequences. Prior work has, however, shown that this problem can be posed as a tree search problem[7]. This enables the best s mode sequences to be found efficiently using a best-first informed search approach[18]. We now describe how this can be applied to our problem. Fig. 3 shows part of the search tree for the most likely mode sequence enumeration problem. The tree has k initial nodes that correspond to each of the mode sequences stored by the hybrid estimator at time t. In the diagram, $k = 2$. The graph search approach aims to find the s mode sequences $m_{1:t+h}$ with the highest prior probability $p(m_{1:t+h}|\mathbf{y}_{1:t})$. Making use of the Markov assumption in Section II, we can write this as:

$$\begin{aligned} p(m_{1:t+h}|\mathbf{y}_{1:t}) &= p(m_{1:t}) \prod_{i=t+1}^{t+h} p(m_i|m_{i-1}) \\ &= p(m_{1:t}) \prod_{i=t+1}^{t+h} T_{m_i m_{i-1}}, \end{aligned} \qquad (28)$$

where T_{xy} is defined in (2). We can equivalently maximize the logarithm of this probability:

$$\log p(m_{1:t+h}|\mathbf{y}_{1:t}) = \log p(m_{1:t}|\mathbf{y}_{1:t}) + \sum_{i=t+1}^{t+h} \log T_{m_i m_{i-1}} \qquad (29)$$

The cost of each initial node in the search tree is the probability of the corresponding mode sequence $m_{1:t}^{(i)}$ given the observations up to time t. This probability is provided by the hybrid estimator running up to time t. Each intermediate node in the graph corresponds to a partial assignment to the modes $m_{1:t+h}$, and a goal node corresponds to a full assignment. The cost of a node is the sum of the arcs from the initial node to that node. We can therefore explore the graph using best-first search[18] to find the set of s goal nodes with the largest log-probability given in (29). In this manner, the most likely mode sequences can be enumerated without evaluating the prior likelihood of each mode sequence, which is intractable.

V. SUMMARY

The new active hybrid estimation approach can be summarized as follows:

- Hybrid estimation calculates the probabilities of the k tracked mode sequences, as well as the distribution over the continuous state conditioned on each mode sequence.
- Starting from the k tracked mode sequences, active estimation enumerates the s most likely future mode sequences over the horizon $t+1, \dots, t+h$ using best-first search.
- Active estimation forms a cost function involving only terms corresponding to the s most likely future mode sequences:

$$J = \sum_{i \in S} \sum_{j>i, j \in S} P(H_i)^{\frac{1}{2}} P(H_j)^{\frac{1}{2}} e^{-k(i,j)} \qquad (30)$$

- Active estimation optimizes the finite sequence of control inputs $\mathbf{u}$ to minimize the cost function J, subject to constraints, using Sequential Quadratic Programming.
- Optimized control inputs are executed, while hybrid estimation continues to estimate the hybrid state.

Since control inputs are designed subject to hard constraints, we can ensure that:

1) A control task, defined in terms of the expected state, is achieved.

2) The expected system state remains within a linearization region.
3) Actuator saturation limits are not violated.

In this sense, the active hybrid estimation approach uses constraints to perform control, while optimizing with regard to estimation.

VI. Simulation Results

In this section we present simulation results that demonstrate the new active hybrid estimation approach with a satellite fault detection scenario.

Satellite dynamics, relative to a nominal circular orbit, can be expressed in linear form using Hill's equations[19], assuming that disturbances do not act on the spacecraft and that deviations from the nominal orbit remain small (hundreds of meters). The system state is defined in a rotating reference frame, where the axes correspong to *radial*, *in-track* and *out-of-plane* motion. For the simulations given here we consider the motion in the radial and in-track directions only; this motion is decoupled from the out-of-plane motion.

We assume a 90 minute low-Earth orbit, time steps of 60 seconds, and a planning horizon of 10 time steps. Observations of radial and in-track velocity are made, with additive zero-mean Gaussian white noise of standard deviation $0.1mm/s$. There are thrusters that act to change the radial and in-track velocities. The input $\mathbf{u}_t$ is defined such that $\left[\mathbf{u}_t\right]_1$ is an impulsive velocity change in the radial direction while $\left[\mathbf{u}_t\right]_2$ is an impulsive velocity change in the in-track direction. In order to model imprecision in thrusting, we assume zero-mean Gaussian white process noise, with standard deviation 0.5% of the maximum thruster effort.

The satellite's mean initial state is at rest, at zero displacement from the nominal orbit. The standard deviation of the initial state is $0.1m$ in displacement and $0.1mm/s$ in velocity in both in-track and radial directions. In all of the Gaussian distributions we assume diagonal covariance matrices, i.e. zero correlation between different uncertain variables.

The satellite can switch between four different modes, which are used to model three single-point failures:

1) Mode 1: Nominal operation
2) Mode 2: Radial velocity sensor failure
3) Mode 3: In-track velocity sensor failure
4) Mode 4: Radial thruster failure

In the case of sensor failure, we assume that the observation is zero-mean Gaussian noise, while in the case of a thruster failure, we assume that the thruster exerts no control effort. The transition matrix as defined in (2) is:

$$T = \begin{bmatrix} 0.989 & 0.01 & 0.001 & 0.01 \\ 0.0 & 1.0 & 0.0 & 0.0 \\ 0.0 & 0.0 & 1.0 & 0.0 \\ 0.0 & 0.0 & 0.0 & 1.0 \end{bmatrix} \qquad (31)$$

Note that we assume that once a fault occurs it persists indefinitely.

For the examples given, active estimation optimizes with regard to the 20 most likely mode sequences, i.e. $s = 20$. However when evaluating the performance of the method, we report the upper bound on the probability of pruning the true mode sequence calculated by taking into account all mode sequences.

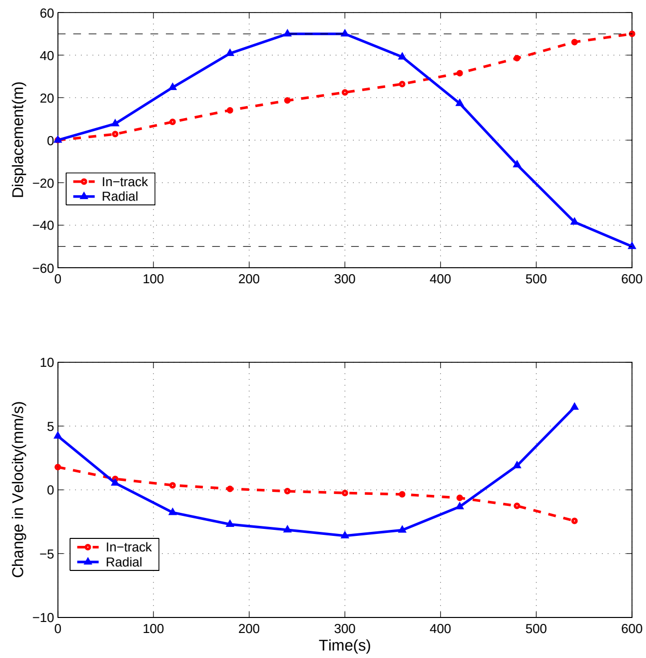

Fig. 4. Typical input sequence designed by active hybrid estimation. The expected in-track and radial displacements can be at most $\pm 50m$. The optimized sequence makes full use of the available displacement and yields an upper bound on the probability of pruning the true mode sequence of 0.12.

A. Box-Constrained Displacement

In this section we use active hybrid estimation to design a control sequence to optimally detect failures while ensuring that the expected spacecraft position remains within a box region. The maximum actuator effort was constrained to be $50mm/s$. Fig. 4 shows a typical input sequence designed by the active hybrid estimation approach for a maximum expected displacement of $50m$. This sequence gives an upper bound on the probability of pruning of 0.12. Fig. 5 shows how the upper bound on the probability of error varies with the size of the box allowed for the maneuver. Increasing the space for the maneuver decreases the probability of pruning the true mode sequence.

B. Displacement Maneuver

In this section we use active hybrid estimation to detect failures optimally while carring out an in-track displacement maneuver. In this maneuver, the expected system state must go from an expected in-track displacement of zero to an in-track displacement of $20m$ (conditioned on there being no actuator failure). By using constrained optimization, the active hybrid estimation approach is able to ensure that this goal is achieved, while optimizing to detect failures.

Fig. 6 shows a typical sequence designed by the method. This yields an upper bound on the probability of pruning of 0.10. For comparison we also show a fuel-optimal sequence.

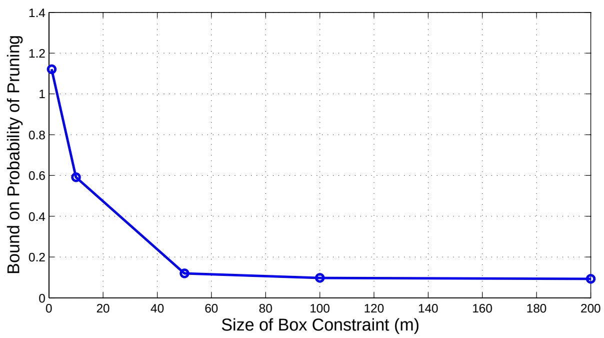

Fig. 5. Variation of probability of pruning with size of box allowed for active estimation maneuver. As more room is allowed for the maneuver, the upper bound on the probability of pruning the true mode sequence decreases. Note that for very small boxes, the value of the bound is greater than 1, which means that it loses meaning as a bound on the true probability of pruning.

This gives an upper bound on the probability of pruning of 0.87. Hence the active estimation approach significantly reduces the probability of the hybrid estimator losing the true mode sequence while achieving the same objective as the fuel-optimal sequence.

VII. Conclusion

This paper introduced a novel method for active state estimation in Switching Linear Dynamic Systems. The method designs finite sequences of control inputs that reduce the probability of pruning the true mode sequence while ensuring that a given control task is achieved. The key innovation was to derive a tractable upper bound on the probability of pruning the true mode sequence. Simulation results demonstrated the approach using a satellite failure scenario.

VIII. Discussion

One of the key assumptions used in creating the new approach is that pruning is carried out at the end of the finite horizon control sequence, rather than at every time step. In most hybrid estimation schemes, pruning occurs at every time step, however. While it would be possible, and in fact simpler, to formulate the control design problem for one time step consistent with such a scheme, we did not two so for two reasons. First, in the discrete-time formulation, in many systems the effect of control inputs at time t is not manifested in the observations until some time after $t+1$. Hence a design approach that only takes into account how the inputs at time t will affect the observations at $t+1$ will be severely limited. Second, by constraining the system state over a long horizon, the optimization has much greater latitude in designing a powerful control sequence for the purposes of estimation; the system state can be driven far from its initial value, while being brought back to its goal value by the end of the time horizon. Hence by considering the probability of pruning over a horizon, rather than over a single time step, the approach yields sequences that are more powerful with regard to estimation.

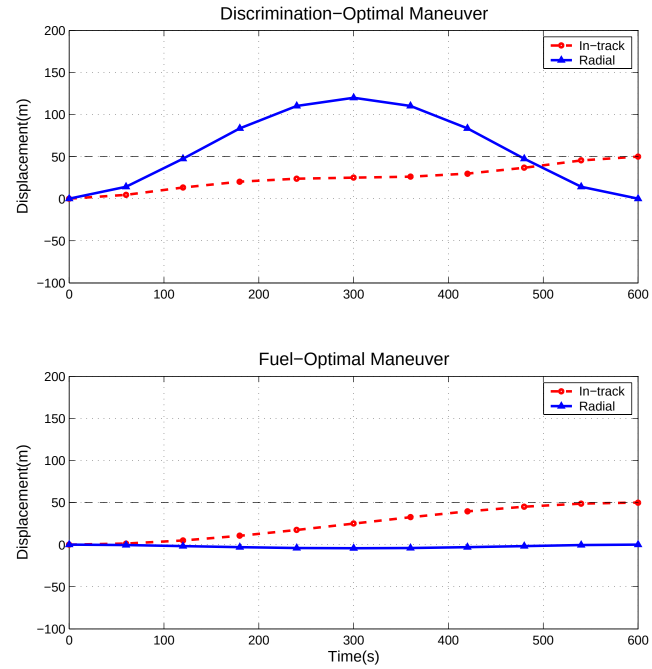

Fig. 6. Typical displacement maneuver. The expected in-track displacement is constrained to change from zero to $50m$, while the expected radial displacement must be zero at the end of the maneuver. Top: the spacecraft trajectory designed by the active hybrid estimation apprach to be optimal with regard to estimation. Bottom: the fuel-optimal spacecraft trajectory. The estimation-optimal trajectory gives an upper bound on the probability of pruning of 0.10, while the fuel optimal trajectory has an upper bound of 0.87. Note that the fuel-optimal trajectory has very little radial motion, while the discrimination-optimal trajectory uses a large radial motion to detect failures in the radial thruster and velocity sensor.

Appendix

In this section we prove Theorem 1, which was used in deriving a looser upper bound on the probability of pruning in Section IV. We make use of the following inequality:

Lemma 1: If A, B are positive-definite symmetric $n \times n$ matrices then,

$$|\mathrm{A}+\mathrm{B}| \geq |\mathrm{A}| + |\mathrm{B}|\,, \tag{32}$$

holds for all $n > 0$.

Proof: It is sufficient to prove this statement when A is diagonal (with positive components), as we can always find a basis where this is true.

First note that when $n = 1$, A and B are just positive real numbers, and (32) is clearly true. Now let us suppose that (32) is true when $n = k$, for some $k > 0$.

Let $\mathrm{A}^{(k+1)}$ be a diagonal $(k+1) \times (k+1)$ matrix with positive components λ_i, and $\mathrm{B}^{(k+1)}$ be a symmetric matrix of the same dimensions, with components $b_{ij}^{(k+1)}$. Then

$$\begin{aligned}\left|\mathrm{A}^{(k+1)}+\mathrm{B}^{(k+1)}\right| = &\left(\lambda_1 + b_{11}^{(k+1)}\right)\left|\mathrm{M}'^{(k)}_{11}\right| \\ &+ \left(b_{12}^{(k+1)}\right)\left|\mathrm{M}'^{(k)}_{12}\right| + \cdots \\ &+ \left(b_{1,k+1}^{(k+1)}\right)\left|\mathrm{M}'^{(k)}_{1,k+1}\right|\end{aligned} \tag{33}$$

where $\mathrm{M'}^{(k)}_{ij}$ is the $k \times k$ matrix formed by removing the i^{th} row and j^{th} column from the matrix $\left(\mathrm{A}^{(k+1)} + \mathrm{B}^{(k+1)}\right)$.

We can now apply (32) to $\left|\mathrm{M'}^{(k)}_{ij}\right|$, as we have assumed (32) to be true in k dimensions. After some thought about the form of the determinant of A with the i^{th} row and j^{th} column removed, and defining $\mathrm{M}^{(k)}_{ij}$ to be B with the i^{th} row and j^{th} column removed, we find

$$\mathrm{M'}^{(k)}_{11} \geq \lambda_2\lambda_3 \cdots \lambda_{k+1} + \left|\mathrm{M}^{(k)}_{11}\right| \tag{34}$$

$$\mathrm{M'}^{(k)}_{1m} \geq \left|\mathrm{M}^{(k)}_{1m}\right| \tag{35}$$

where $2 \leq m \leq k+1$. Applying these inequalities to (33), and using the expansion of $\left|\mathrm{B}^{(k+1)}\right|$ in terms of $\mathrm{M}^{(k)}_{ij}$, we find that

$$\left|\mathrm{A}^{(k+1)} + \mathrm{B}^{(k+1)}\right| \geq \left|\mathrm{A}^{(k+1)}\right| + \left|\mathrm{B}^{(k+1)}\right|. \tag{36}$$

Hence we have shown that if (32) is true for $n = k$, it is also true for $n = k+1$. Therefore, by induction, it is true for all $n > 0$. ■

We now use this lemma to prove the following theorem:

Theorem 1: Any two symmetric positive-definite $n \times n$ real matrices Σ_0, Σ_1 satisfy the following inequality:

$$\frac{\left|\frac{\Sigma_0+\Sigma_1}{2}\right|}{\sqrt{|\Sigma_0|\,|\Sigma_1|}} \geq 1. \tag{37}$$

Proof: Using Lemma 1 in the last line,

$$|\Sigma_0|\,|\Sigma_1| = \left(\frac{|\Sigma_0| + |\Sigma_1|}{2}\right)^2 - \left(\frac{|\Sigma_0| - |\Sigma_1|}{2}\right)^2 \tag{38}$$

$$\leq \left(\frac{|\Sigma_0| + |\Sigma_1|}{2}\right)^2 \tag{39}$$

$$\leq \left|\frac{\Sigma_0 + \Sigma_1}{2}\right|^2. \tag{40}$$

■

References

[1] R. Dearden and D. Clancy, "Particle filters for real-time fault detection in planetary rovers," in *Proceedings of the 13th International Workshop on Principles of Diagnosis (DX02)*, May 2002, pp. 1–6.

[2] L. Blackmore, S. Funiak, and B. C. Williams, "Combining stochastic and greedy search in hybrid estimation," in *Proc. 20th National Conference on Artificial Intelligence (AAAI05)*, Pittsburgh, PA, 2005.

[3] V. Pavlovic, J. Rehg, T.-J. Cham, and K. Murphy, "Multiple-model adaptive estimation using a residual correlation kalman filter bank," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 36, pp. 393–406, Apr. 2000.

[4] S. M. Oh, J. M. Rehg, T. Balch, and F. Dallaert, "Data-driven mcmc for learning and inference in switching linear dynamic systems," in *Proc. 20th National Conference on Artificial Intelligence (AAAI-2005)*, Pittsburgh, PA, 2005.

[5] U. Lerner and R. Parr, "Inference in hybrid networks: Theoretical limits and practical algorithms," in *Proceedings of the 17th Annual Conference on Uncertainty in Artificial Intelligence (UAI-01)*, Seattle, Washington, August 2001, pp. 310–318.

[6] C. B. Chang and M. Athans, "State estimation for discrete systems with switching parameters," *IEEE Transactions on Aerospace and Electronic Systems*, vol. AES-14, pp. 418–424, May 1978.

[7] M. Hofbaur and B. C. Williams, "Mode estimation of probabilistic hybrid systems," in *Intl. Conf. on Hybrid Systems: Computation and Control*, May 2002.

[8] A. Doucet, S. Godsill, and C. Andrieu, "On sequential Monte Carlo sampling methods for Bayesian filtering," *Statistics and Computing*, vol. 10, no. 3, pp. 197–208, July 2000.

[9] U. Lerner, R. Parr, D. Koller, and G. Biswas, "Bayesian fault detection and diagnosis in dynamic systems," in *Proc. of the 17th National Conference on A. I.*, July 2000, pp. 531–537. [Online]. Available: citeseer.nj.nec.com/lerner00bayesian.html

[10] Y. Bar-Shalom and T. E. Fortmann, *Tracking and Data Association*. Academic Press, 1988.

[11] U. Lerner, "Hybrid bayesian networks for reasoning about complex systems," Ph.D. dissertation, Stanford University, October 2002.

[12] L. Blackmore and B. C. Williams, "Finite horizon control design for optimal model discrimination," in *Proc. Control and Decision Conference 2005*, Seville, Spain, Dec. 2005.

[13] L. Blackmore and B. Williams, "Finite horizon control design for optimal discrimination between several models," Massachusetts Institute of Technology, CSAIL Technical Report CSAIL-TR-2006-013, Feb. 2006.

[14] J. Nocedal and S. Wright, *Numerical Optimization*. Springer, 1999.

[15] R. E. Kalman, "A new approach to linear filtering and prediction problems," *Transactions of the ASME–Journal of Basic Engineering*, vol. 82, pp. 35–45, Apr. 1960.

[16] J. K. Tugnait, "Detection and estimation for abruptly changing systems," *Automatica*, vol. 18, 1982.

[17] H. Blom and Y. Bar-Shalom, "The interacting multiple model algorithm for systems with Markovian switching coefficients," *IEEE Transactions on Automatic Control*, vol. 33, 1988.

[18] S. J. Russell and P. Norvig, Eds., *Artificial Intelligence: A Modern Approach*. Prentice Hall, 2003.

[19] M. H. Kaplan, *Modern Spacecraft Dynamics and Control*. New York NY: Wiley, 1976.