# Monitoring and Diagnosis of Hybrid Systems Using Particle Filtering Methods

Xenofon Koutsoukos, James Kurien, and Feng Zhao
Palo Alto Research Center
3333 Coyote Hill Road
Palo Alto, CA 94304, USA
koutsouk,jkurien,zhao@parc.com

**Abstract**

Embedded systems are composed of a large number of components that interact with the physical world via a set of sensors and actuators, have their own computational capabilities, and communicate with each other via a wired or wireless network. Diagnostic systems for such applications must address new challenges caused by the distribution of resources, the networking environment, and the tight coupling between the computational and the physical worlds. Our approach is to move from centralized, discrete or continuous techniques toward a distributed, hybrid diagnosis architecture. Monitoring and diagnosis of any dynamical system depend crucially on the ability to estimate the system state given the observations. Estimation for hybrid systems is particularly challenging because it requires keeping track of multiple models and the transitions between them. This paper presents a particle filtering based estimation algorithm that addresses the challenge of the interaction between continuous and discrete dynamics in hybrid systems. The hybrid estimation methodology has been demonstrated on a rocket propulsion system.

## 1 Introduction

Our diagnostic research is motivated by existing and emerging applications of embedded systems. In such systems the physical plant is composed of a large number of distributed nodes, each of which performs a moderate amount of computation, collaborates with other nodes via a wired or wireless network, and is embedded in the physical world via a set of sensors and actuators. Examples include complex electromechanical systems with embedded controllers [18] and smart matter systems [11]. Such systems can be best represented by hybrid models and present a number of interesting new challenges for diagnostic systems.

Model-based diagnostic techniques are usually based upon a logical framework for diagnosis [3] and are thus discrete. As such, they cannot resolve between and often cannot even detect failures that manifest as small continuous variations in the plant's behavior, nor can they provide sufficient resolution to enable compensatory control of continuous degradations in the plant. These limitations render such discrete techniques ill-suited for diagnosis and control of many embedded systems, as demonstrated in practical applications [6]. Current FDI techniques [5] model continuous behavior, but cannot address the hybrid behavior exhibited by many physical systems, for example continuous processes coupled with digital controllers.

They are also typically computationally expensive in that they rely on computing statistics over raw sensor signals in order to form a diagnosis. They are therefore practical for a relatively small number of fault hypotheses.

Monitoring and diagnosis of a dynamical system depend crucially on the ability to estimate the system state given the observations. Estimation for hybrid system is particularly challenging because keeping track of multiple models and the autonomous transitions between them is computationally very expensive. Simple extension of conventional estimation techniques, like the Kalman filter, leads to algorithms that require tracking of all possible trajectories and therefore, are exponential in the number of time steps. Approximation by Gaussians is often used to collapse the distributions for each trajectory resulting in poor performance. A related approach to our work based on banks of extended Kalman filters has been presented in [9] where only trajectories with high confidence probability are traced. A related methodology that uses both discrete and continuous observers based on finite state machines and linear systems has been proposed in [1]. Sequential Monte Carlo (or particle filtering) methods can support process densities that contain both continuous and discrete dynamics and have been explored for hybrid diagnosis in [16]. However, autonomous transitions between modes triggered by the continuous dynamics have not been considered. Particle filtering has been applied also for a class of hybrid systems modeled by dynamic Bayesian networks in [12] where the autonomous transitions between discrete states are only defined using the so-called softmax conditional probability distributions. Hybrid diagnosis based on timed discrete-event representations has been studied also in [15]. In these methodologies, the continuous state is quantized and discrete methods are used. A fault modeling and diagnosis approach for hybrid systems based on qualitative representation of the fault hypotheses has been presented in [13]. A Bayesian approach for mode estimation of hybrid systems has been presented in [18] and has been demonstrated for monitoring and diagnosis of electromechanical systems. This approach uses continuous measurements to compute appropriate likehihood functions, but it is based on a temporal discrete event model of the system dynamics.

Our approach is to move from centralized, discrete or continuous techniques toward a distributed, hybrid diagnosis architecture; see [14] for details. In this paper, we focus on the problem of hybrid estimation and we present a particle filtering algorithm that address the challenge of the interaction between continuous and discrete dynamics. We show how we can estimate autonomous transitions based on complex guard conditions and we describe how we can improve the performance and robustness of the algorithm by using guard conditions that cover the state space of the system. We illustrate the algorithm for the state estimation of a two-tank system. We also demonstrate the application of the approach to the cryogenic propulsion system of a NASA experimental vehicle (X34).

The paper is organized as follows. In the remainder of this section, we briefly present our diagnostic architecture to explain the significance of hybrid estimation techniques in monitoring and diagnosis of embedded systems. In section 2, we describe our model and the hybrid estimation problem and in section 3, we present our particle filtering algorithm.

Section 4 demonstrates the application of the algorithm to a rocket propulsion system. The final section briefly discusses our current implementation of the estimation algorithm and provides directions for future work.

## The Diagnostic Architecture

The challenge of diagnosing hybrid systems is that they have both complex, hybrid dynamics and a relatively large number of components that can interact in a system-wide manner. Qualitative techniques perform diagnostic inference involving multiple components in a computationally efficient manner, but they are limited by the low-resolution introduced by discretization of the continuous variables. Hybrid estimation techniques produce high-resolution state estimates that can distinguish between these failures, but they are computationally expensive and can be only used to detect faults that can be described by detailed analytical models. Thus diagnosis of hybrid systems suggests collaboration between qualitative diagnosis and hybrid estimation. Figure 1 illustrates our conceptual hybrid diagnosis architecture for integrating these two techniques. Given a physical plant such as the propulsion system of a spacecraft described in Section 4, we make use of models at two different levels of abstraction, a qualitative model and a hybrid model of the plant. By using both qualitative diagnosis and hybrid estimation and fault detection, we leverage the speed of the qualitative diagnoser and the resolution of the hybrid model.

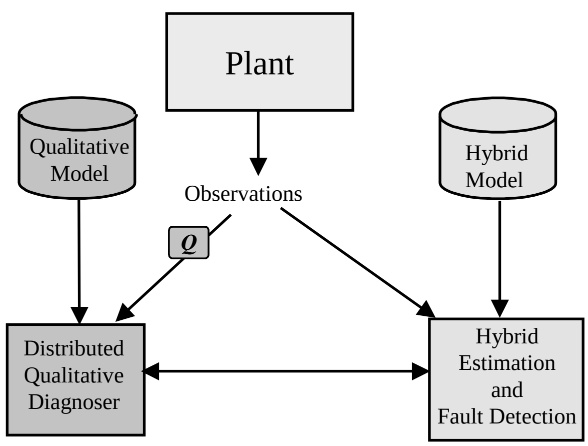


Figure 1: The PARC diagnosis architecture

The qualitative model provides a discrete abstraction of the plant model by first discretizing the range of each variable representing the system into a discrete, finite domain. For example, temperatures in a cryogenic propulsion system might be discretized into the range { *low*, *high*} where *low* describes temperatures in the range of liquid oxygen (-184 degrees) and *high* describes temperatures around the ambient air temperature. The qualitative model specifies the states of the system and an abstraction of the dynamics. Given a set of observations in the discrete space, these qualitative relationships are sufficient to very quickly rule out many possible states of the plant as inconsistent with observations and yield the most likely, consistent states. Additional observations might further reduce the set of consistent diagnoses. However, due to the loss of resolution introduced when the plant model is discretized, the qualitative diagnoser will not be able to distinguish failures that can be detected with a hybrid model.

The hybrid model represents analytically the physical phenomena that govern the dynamic evolution of the plant. For example, the continuous states of the propulsion system include the temperature and mass of the gas and liquid oxygen inside the tank and their evolution is governed by analytical equations derived by mass and energy conservation laws. It also relates the sensor measurements with the state of the system and models process and measurement noise. Controlled and autonomous events that affect the evolution of the system are also modeled. Hybrid estimation is the task of computing the most likely trajectory of the state given the observations. Estimation of the unobservable states is necessary to detect failures that are caused by subtle component degradation, for example, a leakage in a pneumatic valve.

# 2 Estimation of Hybrid Systems

Hybrid systems contain interacting discrete and continuous dynamics. The discrete dynamics are usually described by discrete event models with a finite state space. Every discrete state (or mode) corresponds to a unique differential/difference equation that governs the continuous dynamics. Mode transitions may occur either upon receiving an external control command or when the continuous state satisfies certain guard conditions. Mode transitions that depend on the continuous behavior of the system are called autonomous. The main idea in our algorithm is to focus on the mode transitions that cover most of the probability space. Of course, the probability of each mode transition changes dynamically based on the continuous behavior of the system and has to be recomputed at every time step.

**Definition 2.1** *A hybrid system is described by* $H = (Q, X, \Sigma, I, Inv, E, f)$ *where* $Q$ *is a finite set of discrete states or* modes *of the system,* $X \subseteq \Re^n$ *is the continuous state space,* $\Sigma$ *is a finite set of transition labels or* events*,* $I \subseteq Q \times X$ *is the set of initial conditions,* $Inv : Q \to 2^X$ *is the invariant associated with each mode* $q$*,* $E \subset Q \times X \times \Sigma \times Q \times X$ *is the set of discrete transitions, and* $f : Q \times X \to X$ *is the flow condition for every mode.*

The state of the hybrid system is described by $s = (q, x)$. The state can change either by a discrete transition or by a time delay. A discrete (or mode) transition may change both the mode and the continuous state, while a time delay changes only the continuous state according to the flow condition. Each transition consists of a source mode $q_i$, a target mode $q_j$, a labeling event $\sigma$ (denoted as $q_i \xrightarrow{\sigma} q_j$), a guard set $G_{ij} \subset X$, and a reset map $R_{ij}(x) = x'$. If the condition described by the guard is satisfied, then the transition can fire. Upon firing of the transition, the continuous state may be reset according to the reset map.

**Example** We consider a simple tank system to illustrate the approach. The system consists of two identical cylindrical tanks that are connected by a pipe at level $h$, as shown in figure 2. We denote by $h_1$ and $h_2$ the water levels in tanks 1 and 2 respectively. The input flow $Q_{in}$ is provided by a pump and it is described by

$$Q_{in} = V_{in} k_{in} u(t),$$

where $V_{in} \in \{0, 1\}$ represents a valve that can be used to turn on or off the pump, $k_{in}$ is a linear gain, and $u(t)$ is the input signal representing the flow at the pump. The flow $Q_a$ between the two tanks is controlled by a valve $V_a$. An outlet valve $V_{out}$ located at the bottom of tank 2 is used to empty the tank. Tank 2 is equipped with a sensor that measures the output flow which is described by

$$Q_{out} = V_{out} k_{out} \sqrt{\rho g h_2} + \xi(t) \tag{2.1}$$

where $V_{out} \in \{0, 1\}$ represents the outlet valve, $k_{out}$ is a linear gain, $\rho$ is the density of the water, $g$ is the gravity constant, and $\xi(t)$ is measurement noise.

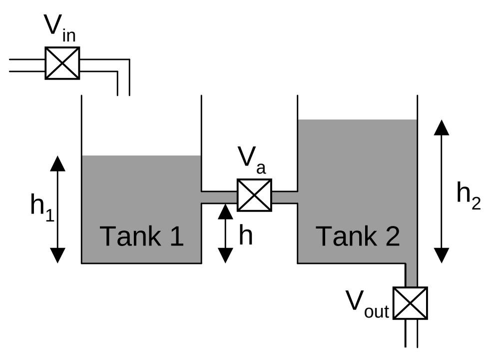


Figure 2: Two-tank system

The dynamic evolution of the system is described by

$$\begin{aligned} \dot{h}_1 &= \frac{1}{A}(Q_{in} - Q_a) \\ \dot{h}_2 &= \frac{1}{A}(Q_a - Q_{out}) \end{aligned}$$

where $A$ is the section of each cylindrical tank. There are four modes of behavior for the flow $Q_a$ which depend on the water levels $h_1$ and $h_2$ as follows:

$$Q_a = \begin{cases} 0, & \text{if } h_1 < h \text{ and } h_2 < h \\ V_a k_a \sqrt{\rho g (h_1 - h)}, & \text{if } h_1 > h \text{ and } h_2 < h \\ V_a k_a \sqrt{\rho g (h_2 - h)}, & \text{if } h_1 < h \text{ and } h_2 > h \\ \text{sign}(h_1 - h_2) V_a k_a \sqrt{\rho g |h_1 - h_2|}, & \text{if } h_1 > h \text{ and } h_2 > h \end{cases} \tag{2.2}$$

where $V_a \in \{0, 1\}$ and $k_a$ is a linear gain. The evolution of the continuous state $x = [h_1, h_2]^T$ can be described by

$$\dot{x} = f_q(x(t), u(t)) + \nu(t) \tag{2.3}$$

where $q \in \{1, 2, 3, 4\}$ is the discrete mode as described in (2.2) and $\nu(t)$ is assumed to be process noise. For every mode, we have a set of ordinary differential equations and the system transitions between modes based on $x$ as described by (2.2). Clearly, these transitions are autonomous since they depend on the continuous behavior of the system as described by the guard conditions. Figure 3 shows a typical simulation of the system for $x_0 = [.2, .75]^T$ for a pulse input. The estimation objective is to predict the hybrid state from the input and output flows.

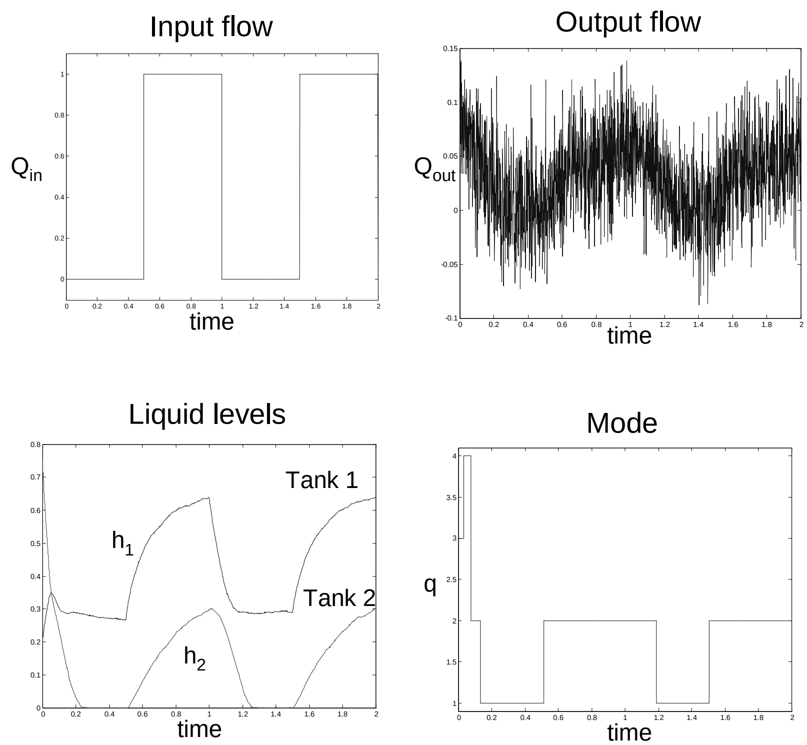


Figure 3: Simulation for the two-tank system ($x_0 = [.2, .75]^T$).

In the hybrid system literature, it is often assumed that the state is directly observable. However, in real-world applications, the state has to be reconstructed from the observations. In this paper, we follow a Bayesian state estimation approach using a discrete-time representation of the system dynamics. The continuous dynamics of the system can be described, using zero-order hold sampling for example, by the discrete-time model

$$\begin{aligned} x_{t+1} &= f_q(x_t, u_t) + \nu_t \\ y_t &= g_q(x_t) + \xi_t \end{aligned}$$

where $\nu_t$ and $\xi_t$ denote process and measurement noise respectively. The evolution of the discrete state can be described by the transition function

$$q_{t+1} = \delta(q_t, \sigma_t, x_t).$$

A discrete transition occurs when either the controller issues an appropriate command or when the continuous state satisfies the guard of the transition. In the tank system, for example, a command that opens a valve causes a controlled transition. An autonomous transition occurs when the water level exceeds $h$. We assume that control commands are issued asynchronously and therefore, $\sigma_t$ can be a null event.

**Problem 2.1** *The hybrid estimation problem is to compute the most likely hybrid state $s_t = (q_t, x_t)$ given the observation sequence $Y_t = (y_0, y_1, \ldots, y_t)$ the sequence of continuous control inputs $U_t = (u_0, u_1, \ldots, u_t)$, and the history of control events $(\sigma_1, \sigma_2, \ldots)$ up to time $t$.*

The most challenging aspect of every hybrid estimation algorithm is how to monitor the autonomous mode transitions and use the appropriate mode $q$ for updating the estimate of the continuous state $x$. The probability of mode transitions triggered by control commands can be usually computed by discrete estimation techniques based, for example, on hidden Markov models. Let's focus on autonomous transitions and define the mode transition probability matrix with elements

$$T_{ij}(x_{t-1}) = p(q_t = j | x_{t-1}, q_{t-1} = i), \; i, j = 1, \ldots, |Q|.$$

Let $G_{ij}$ be the guard corresponding to the transition from mode $i$ to mode $j$. Assuming that the system is at mode $q_i$ and that the probability of the transition $q_i \to q_j$ is equal to the probability the guard $G_{ij}$ is satisfied, we have

$$T_{ij}(x_{t-1}) = \int_{G_{ij}} p(x_{t-1} | Y_{t-1}, U_{t-1}, q_{t-1} = i) dx_{t-1} \tag{2.4}$$

where $p(x_{t-1} | Y_{t-1}, U_{t-1}, q_{t-1} = i)$ is the conditional density of the continuous state at time $t-1$. The above integral represents the probability of switching from mode $q_i$ to mode $q_j$. The general idea of our estimation algorithm is that at every time step we evaluate the transition probability matrix based on the estimate of the continuous state. Then, we focus on the most likely modes and we update the continuous estimate by conditioning our belief on the new measurements using the corresponding flow conditions. Our current implementation is based on a particle filtering approach described in Section 3. This approach allows the efficient computation of the transition probabilities using Monte Carlo methods. The transition probabilities are then used to dynamically assign particles to the discrete modes, thus focusing on the most likely transitions. Before presenting the hybrid estimation algorithm, we discuss how we can improve the performance of the algorithm by transforming the guard conditions so that they form a cover of the state space.

## Robustness of Hybrid Estimation

The probability of occurrence of the autonomous transitions is represented by the transition probability matrix that can be computed at every time step as a function of the continuous state. The estimation algorithm will be robust if small changes in the continuous state do not result in large changes in the probabilities $T_{ij}$. Practically, it is desirable to (1) avoid chattering phenomena, where the probability mass oscillates between modes at every time step, and (2) allow enough time after a mode change for the transient to converge to the steady state behavior for that particular mode. These aspects of the algorithm can be considerably improved by transforming the guard conditions so that they form a cover of the continuous state space as explained in the following.

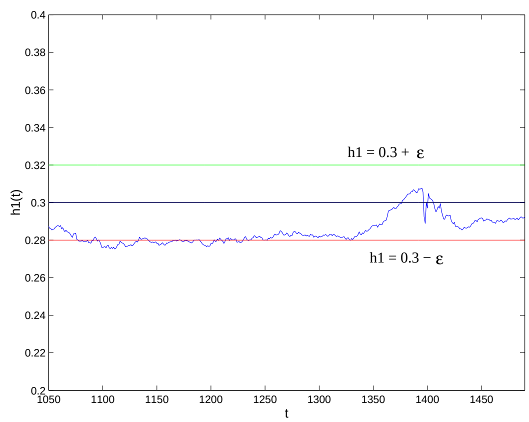


Figure 4: Guard conditions that cover the state space

Figure 4 shows the estimated water level $h_1(t)$ for the two-tank example. Let's assume that $h = .3$. The system switches from $q_1$ to $q_2$ if $h_1(t) > .3$ and from $q_2$ to $q_1$ if $h_1(t) < .3$. The estimation algorithm returns a probability distribution over possible continuous states that approximates the actual state $x(t) = [h_1, h_2]^T$ at every time step. If the transition probability matrix $T$ is computed using the original guard conditions, the performance of the algorithm is degraded by the fast switching around $t = 1380$ (ms) and leads to chattering between modes $q_1$ and $q_2$. While the most likely discrete state oscillates between $q_1$ and $q_2$, the estimation of the continuous state is unreliable.

Hybrid estimation can be considerably improved by transforming the guard conditions to form a cover of the state space as illustrated in Figure 4. The transition $q_1 \to q_2$ occurs if $h_1(t) > h + \epsilon$. Similarly, the transition $q_2 \to q_1$ occurs if $h_1(t) < h - \epsilon$. The small variations of the state around $h_1(t) = h - \epsilon$, for example, will not trigger any transitions since the system is not in mode $q_2$. The design parameter $\epsilon$ depends on the process and measurement noise. The transition probability matrix can be represented by the transformed guard conditions by equation (2.4). It should be noted that the continuity of analog-to-digital maps based on covers of the state space has been studied using small topologies in [17].

# 3 Particle Filtering for Hybrid State Estimation

In the following, we briefly describe an algorithm for hybrid estimation based on particle filtering. In particle filters, complex integrals as that of equation (2.4) are computed efficiently by approximating the belief state by finitely many samples. General process densities that can represent the interaction between discrete and continuous dynamics in hybrid systems can be used in an efficient manner. Detailed descriptions of particle filtering methods for estimation of dynamical systems can be found in [4]. Our approach is similar to algorithms with mixed-state and automatic model switching that have been successfully applied for tracking of motion boundaries in video images [10, 2].

Let $\{s_{t-1}^{(k)}, w_{t-1}^{(k)}, k = 1, \dots, N\}$ denote the sample set at time $t-1$ where $s_{t-1}^{(k)} = (q_{t-1}^{(k)}, x_{t-1}^{(k)})$ is the $k^{th}$ sample of the hybrid state and $w_{t-1}^{(k)}$ its probability weight. The estimation algorithm consists of the following steps:

**1. Initialization** $t = 0$.

i. sample $s_0^{(k)}, k = 1, 2, \dots, N$ from $p(q_0), p(x_0)$ and set $t = 1$.

**2. Prediction**

i. apply $p(s_t|s_{t-1}^{(k)})$ to compute each $\tilde{s}_t^{(k)}$.

ii. evaluate the importance weights $w_t^{(k)} = p(y_t|\tilde{s}_t^{(k)})$.

iii. normalize the weights.

**3. Resampling**

i. resample $N$ particles $s_t^{(k)}$ from $\tilde{s}_t^{(k)}$.

ii. set $t \leftarrow t + 1$ and go to step 2.

Consider that at time $t$ the prediction $p(q_{t-1}, x_{t-1}|Y_{t-1}, U_{t-1})$ is represented by the sample set $\{q_{t-1}^{(k)}, x_{t-1}^{(k)}, w_{t-1}^{(k)}, k = 1, \dots, N\}$. The mode transition probabilities can be computed by

$$T_{ij}(x_{t-1}) = \begin{cases} \dfrac{\sum_{k \in \hat{G}_{ij}} w_{t-1}^{(k)}}{\sum_{k \in \hat{I}} w_{t-1}^{(k)}} & i \neq j \\ 1 - \sum_{\ell \neq i} T_{i\ell}(q_{t-1}, x_{t-1}) & i = j \end{cases} \tag{3.5}$$

where $k \in \hat{G}_{ij} \Leftrightarrow q_{t-1}^{(k)} = i \ \wedge \ x_{t-1}^{(k)} \in G_{ij}$ and $k \in \hat{I} \Leftrightarrow q_{t-1}^{(k)} = i$. Let $(q_{t-1}^{(k)}, x_{t-1}^{(k)}, w_{t-1}^{(k)})$ be the $k^{th}$ particle and assume $q_{t-1}^{(k)} = i$, then we sample from the $i^{th}$ row of the mode transition probability matrix $[T_{i1}, T_{i2}, \dots, T_{i|Q|}]$ to select the $k^{th}$ sample $q_t^{(k)}$ for the discrete mode. Suppose that $q_t^{(k)} = j$, then we sample from the density $p_{ij}(x_t|x_{t-1}^{(k)}) = p(x_t|x_{t-1}^{(k)}, q_{t-1} = i, q_t = j)$ to compute the $k^{th}$ sample $x_t^{(k)}$ for the continuous state. Next, we compute that importance weights, normalize, reinforce the predicted state using the observations, and resample the particles as described in the above algorithm.

The mode of the system is computed using the particles as the most likely mode at every time step and the continuous state is computed using only particles from the most likely mode, that is

$$\hat{q}_t = \arg\max_i \sum_{k \in \hat{Q}_i} w_t^{(k)} \tag{3.6}$$

and

$$\hat{x}_t = \frac{\sum_{k\in\hat{Q}} w_t^{(k)} x_t^{(k)}}{\sum_{k\in\hat{Q}} w_t^{(k)}} \tag{3.7}$$

where $\hat{Q}_i = \{k | q_t^{(k)} = i\}$ and $\hat{Q} = \{k | q_t^{(k)} = \hat{q}_t\}$.

**Example** In the following, we describe the application of the estimation algorithm to the two-tank system shown in figure 2. The observation history is generated using a Matlab/Simulink model of the system. The hybrid estimation algorithm estimates the the mode of the system and the water levels given the input and output flow (shown in figure 3). We assume that the initial state is described by $q_0 = 3$ and $x_0 \sim \mathcal{N}(\mu, P)$, with $\mu = [.2, .75]^T$ and $P = \text{diag}(.1, .1)$. We sample $N = 100$ particles from the distribution of the initial conditions to get

$$\{s_0^{(k)} = (q_0^{(k)}, x_0^{(k)}), w_0^{(k)} = 1/N, k = 1, \ldots, N\}.$$

Given the sample set $\{q_{t-1}^{(k)}, x_{t-1}^{(k)}, w_{t-1}^{(k)}, k = 1, \ldots, N\}$, we compute the transition probability matrix using equation (3.5). Let the $k^{th}$ particle be $(q_{t-1}^{(k)}, x_{t-1}^{(k)}, w_{t-1}^{(k)})$, and assume $q_{t-1}^{(k)} = i$, then we sample from the $i^{th}$ row of the mode transition probability matrix to select the $k^{th}$ sample $q_t^{(k)}$ for the discrete mode. Suppose that $q_t^{(k)} = j$, then we sample from $p_{ij}(x_t | x_{t-1}^{(k)}, u_{t-1}) = p(x_t | x_{t-1}^{(k)}, u_{t-1}, q_{t-1} = i, q_t = j)$ to compute the $k^{th}$ sample $x_t^{(k)}$ for the continuous state. The density $p_{ij}(x_t | x_{t-1}^{(k)}, u_{t-1})$ is computed using the flow condition of the mode $j$ from equation (2.3) by assuming zero-mean Gaussian process noise. Next, we set $\tilde{s}_t^{(k)} = (q_t^{(k)}, x_t^{(k)})$ and we evaluate the importance weights using the likelihood function $w_t^{(k)} = p(y_t | \tilde{s}_t^{(k)})$ derived from (2.1) by assuming zero-mean Gaussian measurement noise. Finally, we resample the $N$ particles to multiply particles with high importance weights and eliminate particles with low importance weights [7]. The estimated hybrid state of the system, computed using (3.6) and (3.7), is shown in figure 5. As it can be seen by comparing the simulation in figure 3 and the estimation results, although there is initially some estimation error due to the uncertainty of the initially condition, we are able to track the hybrid state of the system.

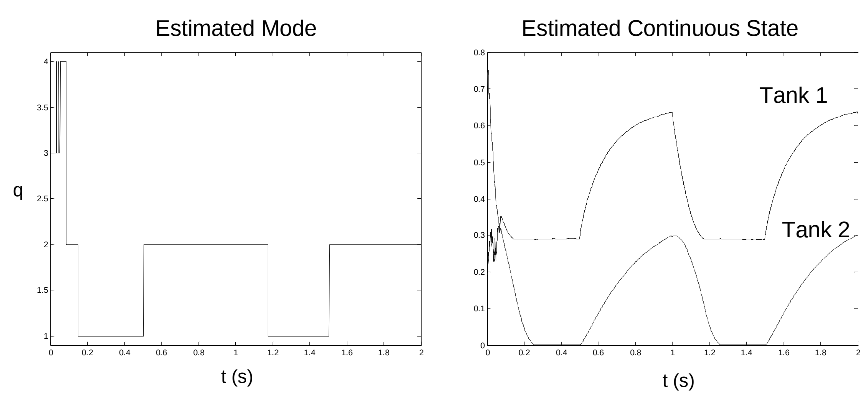


Figure 5: Estimated mode and continuous state for the tank system

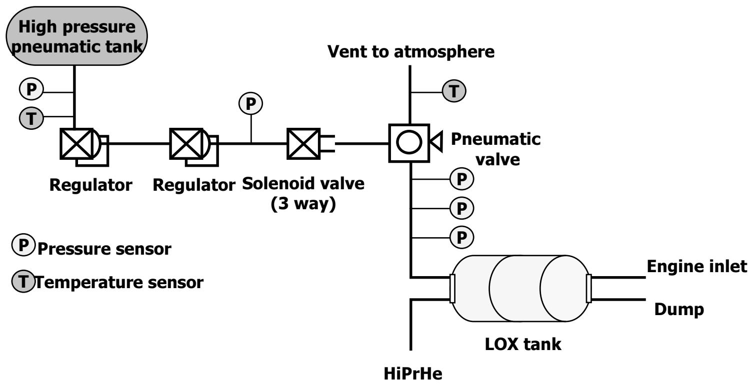


Figure 6: Liquid oxygen (LOX) tank system

# 4 The Propulsion System Domain

In order to illustrate the use of the hybrid estimation approach in our diagnostic architecture, we present an application related to rocket propulsion. All space launch vehicles that reach Earth orbit do so by carrying large quantities of oxygen which is combined with a fuel and burned to produce thrust. The oxygen is stored in the form of liquid oxygen (LOX) at a temperature several hundred degrees below that of the launch environment.

Figure 6 illustrates the LOX venting system for the X-34, an experimental, rocket-powered vehicle designed for NASA. When the pneumatic valve is open, the LOX tank can vent to the atmosphere. The vehicle's control system does not directly actuate the pneumatic valve. Instead, the pneumatic valve opens when it is pressurized by the pneumatic system to its left. The pneumatic tank and regulators provide high pressure gas to the solenoid valve. When the control system opens the solenoid valve, the pneumatic valve is pressurized and opens. There are a wide variety of failures possible within this system. The pneumatic valve might fail to open because either of the valves is stuck closed, either of the regulators are too low, or the pneumatic tank is leaking. The pneumatic valve might stay open because it is stuck or the solenoid valve is stuck open, and similarly might open or close more slowly than originally anticipated. The LOX tank may also lose mass because the pneumatic valve is leaking, the LOX tank is leaking, or components downstream of the Engine Inlet, Dump or HiPrHe lines (not shown) are leaking. A slowly actuating pneumatic valve might be compensated for by the control system, whereas a leaking LOX tank is not recoverable and a potential safety hazard. Detection and diagnosis rely heavily upon estimating the mass of LOX and gaseous oxygen (GOX) in the tank, a task complicated by the fact that the propulsion system is described by a $10^{th}$ order hybrid system with nonlinear dynamics and both commanded (venting, not-venting) and autonomous (boiling, not boiling) transitions.

The particle filtering algorithm presented in section 3 is used for fault detection using an observer-like scheme as shown in figure 7. The particle filter algorithm plays the role of a hybrid observer which is computing the most likely discrete mode $\hat{q}$ and continuous state $\hat{x}$ and is generating the expected output $\hat{y}$ based on the plant model. The residual signal

$r_t = y_t - \hat{y}_t$ is then thresholded, after low-pass filtering, to detect possible failures. Fault detection and isolation is performed by considering both the residual $r_t$ and the mode $\hat{q}_t$. For example, the observer may not be able to perfectly track fast transients after each mode transition and therefore, the residual exceeding the threshold immediately after a mode transition does not necessarily correspond to a fault. Also information about the modes for which the discrepancy is present can be used for fault isolation. A leakage in a pneumatic valve, for example, will cause a discrepancy only if the valve is closed. Our diagnostic system is using a feature extraction algorithm and a neurofuzzy classifier to compute the probability of the fault hypotheses based on the residual signals generated by the hybrid estimation algorithm; details are out of the scope of this paper. In the following, we present simulation results for the propulsion system for two scenarios (1) normal behavior, and (2) leakage in the pneumatic valve.

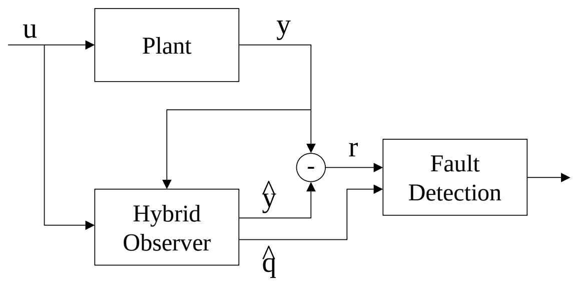


Figure 7: Fault detection using hybrid observer

## Simulation Results

We have tested extensively the particle filtering algorithm for state estimation of the propulsion system domain. Here, we present simulation results for a subsystem of the propulsion domain consisting of the LOX tank and the pneumatic valve. This subsystem interacts with the high pressure pneumatic tank only via the command that controls the solenoid valve that in turn, drives the pneumatic valve. This subsystem is best described as a hybrid system with autonomous transitions between the discrete modes corresponding to the oxygen boiling or not boiling in the tank. The continuous dynamics of the subsystem are described by a set of $4^{th}$ order nonlinear differential equations that are discretized using a sampling period $T = 100ms$. The discrete modes correspond to the oxygen boiling or not, which is determined by a nonlinear guard of the form $P_{sat} \geq P_{GOX}$ where the saturation pressure $P_{sat}$ is approximated using a $5^{th}$ order polynomial of the LOX temperature, while the GOX pressure $P_{GOX}$ dependes on the GOX mass and temperature using the ideal gas law. The outputs are the GOX pressure and temperature and are contaminated with Gaussian noise.
*Normal behavior.* We have demonstrated that the algorithm can track the state in the case when there are no faults in the system. The continuous states corresponding to the LOX and GOX masses are shown in figure 8. The expected venting pressure, as computed using

the estimated state, is plotted versus the actual venting pressure and the discrete mode are shown in figure 9. The simulation was performed using $N = 100$ particles, requiring approximately $2500s$ for a time horizon of $9000s$ in a PC workstation using MATLAB.

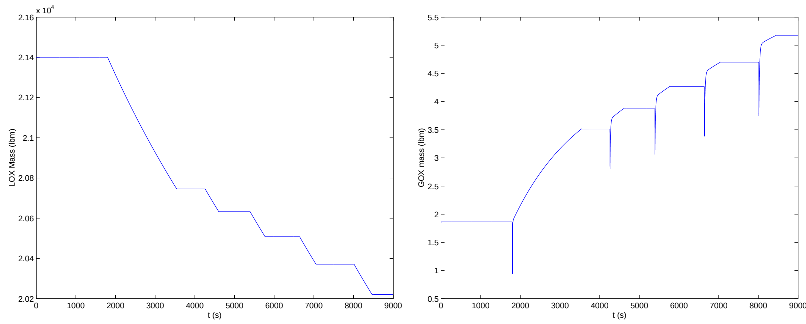


Figure 8: LOX and GOX mass

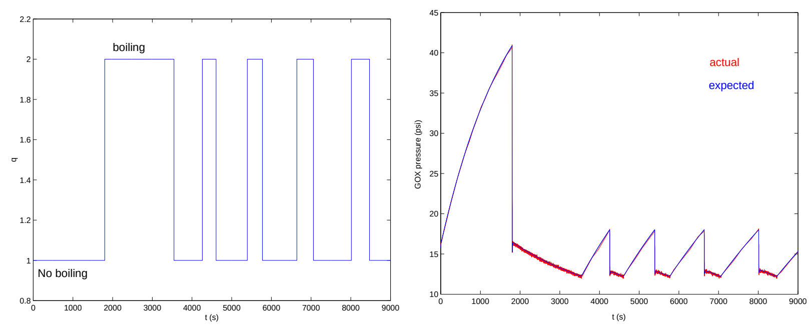


Figure 9: Discrete mode and venting pressure

*Pneumatic valve leakage.* The estimation algorithm can be used also to detect continuous faults such as leakage in the pneumatic valve. The valve leakage was simulated by including an additive term in the equation that represents the flow balance when the pneumatic valve is closed. Figure 10 shows the expected and the actual venting pressure. The estimated discrete mode and the residual signal computed as the difference between the actual GOX pressure and the expected are also shown. Whenever there is no boiling then the actual pressure is less than the expected one and a fault is detected.

# 5 Conclusions

Monitoring and diagnosis of embedded systems depends crucially on the ability to estimate the hidden hybrid state from the available measurements. In this paper, we have presented a particle filtering based method and demonstrate the algorithm using a two-tank example and a rocket propulsion system. The algorithm can be applied in the case of autonomous transitions, nonlinear system dynamics, and non-Gaussian noise. Performance characterization of the algorithm is an important and open problem. Convergence of the algorithm

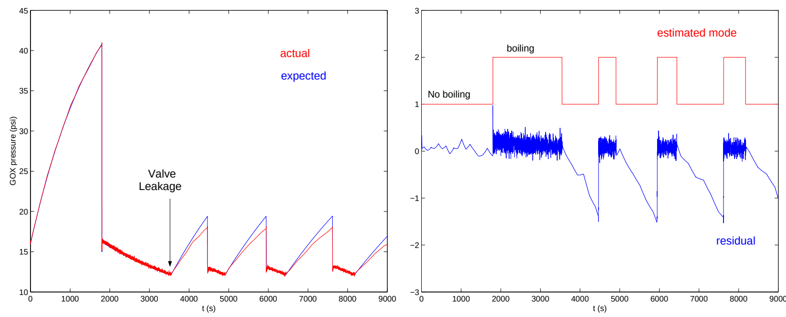


Figure 10: Expected vs. actual venting pressure, residual, and estimated discrete mode in the case of leakage

depends crucially on the number of particles that, in turn, depends on the dimension of the continuous state space and number of discrete modes. We have observed that the time interval between discrete transitions also affects the performance of the algorithm. Currently, we address some of these problems by increasing the number of particles and/or assigning a small number of particles at every mode even if the measurements indicate that some of the modes are not probable. Theoretical aspects regarding the performance characterization of the approach are subjects of current and future research.

**Acknowledgement**: This work is supported in part by the Defense Advanced Research Projects Agency (DARPA) under contract F33615-99-C3611. We thank NASA Ames Research for providing the simulation software for the cryogenic propulsion system.

# References

[1] A. Balluchi, L. Benvenuti, M. D. Benedetto, and A. Sangiovanni-Vincentelli. Design of observers for hybrid systems. In C. Tomlin and M. Greenstreet, editors, *Hybrid Systems: Computation and Control (HSCC'02)*, Vol. 2289, *LNCS*, 76–89. Springer-Verlag, 2002.

[2] M. Black and D. Fleet. Probabilistic detection and tracking of motion boundaries. *International Journal of Computer Vision*, 38(3):231–245, 2000.

[3] J. de Kleer and B. C. Williams. Diagnosing multiple faults. *Artificial Intelligence*, 32(1):97–130, 1987. Reprinted in [8].

[4] A. Doucet, N. D. Freitas, and N. Gordon, editors. *Sequential Monte Carlo Methods in Practice*. Statistics for Engineering and Information Science. Springer, 2001.

[5] P. Frank, S. Ding, and B. Köppen-Seliger. Current developments in the theory of FDI. In *4th IFAC Symp. SAFEPROCESS*, pages 16–27, 2000.

[6] C. Goodrich and J. Kurien. Continous measurements and quantitative constraints - challenge problems for discrete modeling techniques. In *Proc. of iSAIRAS-2001*, 2001.

[7] N. Gordon, D. Salmond, and A. Smith. Novel approach to nonlinear/non-gausian bayesian state estimation. *IEE-Proceedings-F*, 140(2):107–113, 1993.

[8] W. Hamscher, L. Console, and J. de Kleer. *Readings in Model-Based Diagnosis*. Morgan Kaufmann, San Mateo, CA, 1992.

[9] M. Hofbaur and B. Williams. Mode estimation of probabilistic hybrid systems. In C. Tomlin and M. Greenstreet, editors, *Hybrid Systems: Computation and Control (HSCC'02)*, Vol. 2289, *LNCS*, 253–266. Springer-Verlag, 2002.

[10] M. Isard and A. Blake. A mixed-state condensation tracker with automatic model switching. In *Proc. of the 6th International Conference on Computer Vision*, pages 107–112, 1998.

[11] W. Jackson, M. Fromherz, D. Biegelsen, J. Reich, and D. Goldberg. Constrained optimization based control of real time large scale systems: Airject movement object system. In *Proc. of the 40th IEEE Conference on Decision and Control*, pages 4717–4720, Orlando, FL, December 2001.

[12] D. Koller and U. Lerner. Sampling in factored dynamic systems. In Doucet et al. [4], pages 445–464.

[13] X. Koutsoukos, F. Zhao, H. Haussecker, J. Reich, and P. Cheung. Fault modeling for monitoring and diagnosis of sensor-rich hybrid systems. In *Proc. of the 40th IEEE Conference on Decision and Control*, pages 793–801, Orlando, FL, December 2001.

[14] J. Kurien, X. Koutsoukos, and F. Zhao. Distributed diagnosis of networked hybrid systems. In *AAAI Spring Symposium on Information Refinement and Revision for Decision Making: Modeling for Diagnostics, Prognostics, and Prediction*, pages 37–44, Stanford, CA, March 2002.

[15] J. Lunze. Diagnosis of quantised systems by means of timed discrete-event representations. In N. Lynch and B. Krogh, editors, *Hybrid Systems: Computation and Control*, Vol. 1790, *LNCS*, 258–271. Springer, 2000.

[16] S. McIlraith, G. Biswas, D. Clancy, and V. Gupta. Hybrid systems diagnosis. In N. Lynch and B. Krogh, editors, *Hybrid Systems: Computation and Control*, Vol. 1790, *LNCS*, 282–295. Springer, 2000.

[17] A. Nerode and W. Kohn. Models for hybrid systems: Automata, topologies, controllability, observability. In R. L. Grossman, A. Nerode, A. P. Ravn, and H. Rischel, editors, *Hybrid Systems*, Vol. 736, *LNCS*, 317–356. Springer-Verlag, 1993.

[18] F. Zhao, X. Koutsoukos, H. Haussecker, J. Reich, P. Cheung, and C. Picardi. Distributed monitoring of hybrid systems: A model-directed approach. In *Proc. IJCAI'2001*, pages 557–564, Seattle, WA, 2001.