Hybrid Estimation through Synergic Mode-Set Focusing

Rienmüller, Th. [*] **Bayoudh, M.** [**] **Hofbaur, M.W.** [*]
Travé-Massuyès, L. [**]

[*] *Institute of Automation and Control, Graz University of Technology, 8010 Graz, Austria, (e-mail: {theresa.rienmueller,michael.hofbaur}@TUGraz.at).*
[**] *LAAS-CNRS, Université de Toulouse, France (e-mail: {bayoudh,louise}@laas.fr)*

14 April 2009

Abstract: Tracking the evolution of hybrid systems from partial observations means tracking the continuously-valued state evolution and the interleaved discrete mode changes. Existing estimation schemes suffer from the exponential blow up of the number of hypotheses to be tracked and fall into suboptimal methods. On the other hand, hybrid parity-based mode estimation ignores the continuous state. This paper proposes a novel scheme that uses this latter as a mode focusing procedure and then applies hybrid estimation on the resulting reduced number of hypotheses. The advantages of the mixed method are on both sides: it boosts the mode identification time and the convergence of continuous state estimation.

Keywords: Hybrid Systems, Fault Detection, Diagnosis, Monitoring, Discrete-Event Dynamic Systems

1. INTRODUCTION

Many modern artifacts, mobile robotic devices, space probes, or production plants exhibit complex patterns of behavior in order to satisfy the high demand on performance and durability. Key for the physical system's operation is a sophisticated control and automation scheme that actuates the evolution of the system through its many modes of operation and reacts appropriately whenever faults occur. A detailed knowledge of the current mode of operation/failure and the current state of the physical entities that capture the continuous evolution of the physical system are important prerequisites for this automation/control task. However, it is almost always the case that the current mode and the full continuous state is not directly observable/measurable so that the missing information has to be inferred from the known actuation, available measurements and a mathematical model of the physical system.

Hybrid Systems theory provides a modeling paradigm that integrates both, the continuously-valued state evolution and the interleaved discrete mode changes in a comprehensive manner. Using such a model to track the complex evolution of a physical system requires, in theory, to consider every possible mode sequence with its associated continuous evolution. This requires one to perform both, the mode estimation and the continuous state estimation (filtering) in an interwoven form. It is easy to see that the demand to consider all estimation hypotheses is computationally infeasible due to the resulting (exponential) complexity. As a consequence, many sub-optimal estimation schemes were proposed in the literature, for example, the wide field of multi-model filtering (Ackerson and Fu [1970], Blom and Bar-Shalom [1988], Li and Bar-Shalom [1996]), particle filtering methods (de Freitas [2002], Verma et al. [2004], Narasimhan et al. [2004]) or recently developed hybrid estimation methods (Hofbaur and Williams [2002], Benazera et al. [2002], Narasimhan and Biswas [2002]) that can deal with complex systems that evolve according to a large number of modes ($l > 1{,}000$).

As mentioned above, one has to consider together the discrete estimation task that operates on the mode evolution structure of the hybrid model and the continuous estimation task that utilizes the mode-dependent continuous model. Discrete-continuous coupling is the major source of computational complexity of the hybrid estimation task.

In order to un-couple the two estimation tasks we propose to *perform continuous estimation in two (redundant) ways.* Firstly, we apply a parity-space based diagnosis technique (Bayoudh et al. [2008b]) that operates on the continuously valued input/output data. It provides a (mode) consistency information whilst ignoring the continuous state. This supplies a fine-grain abstraction of the continuous evolution that can be used for mode estimation through a discrete-event diagnoser. This diagnoser deduces a mode estimate in the form of a focused set of possible modes. In a second stage we perform an additional continuous estimation scheme that provides the neglected continuous state estimate through a traditional filtering based hybrid estimation technique (Hofbaur and Williams [2002]) that

[⋆] This work was supported in part by the Austrian Science Fund (FWF) under contract P20041-N15 and through Alenia Space, France.

can now operate on a significantly smaller set (ideally a singleton set) of possible modes.

2. HYBRID MODEL

We define the model of a hybrid system through a hybrid automaton that combines a discrete event system (automaton) with continuous dynamics in spirit of the definitions in Henzinger [1996], Hofbaur and Williams [2004], and Bayoudh et al. [2008b] through the tuple

$$S = (\zeta, Q, \Sigma, T, C, (Q_0, \zeta_0)), \quad (1)$$

where:

- ζ is the set of continuous variables, which includes n_u (exogenous) *input variables* $\{u_1, \ldots, u_{n_u}\} =: \mathbf{u}$, n_x *state variables* $\{x_1, \ldots, x_{n_x}\} =: \mathbf{x}$ that capture the dynamic evolution of the automaton, and n_y *output variables* $\{y_1, \ldots, y_{n_y}\} =: \mathbf{y}$ that represent the continuous measurements.
- $Q = \{q_1, \ldots, q_l\}$ is the set of discrete states. Each state $q_i \in Q$ represents a mode of operation, possibly a failure mode, of the system.
- Σ is the set of events. Events correspond to command value switches, spontaneous mode changes and fault events. The subset $\Sigma_O \in \Sigma$ denotes the observable events. Without loss of generality, we assume that fault events are unobservable.
- T is the transition function $T : Q \times \Sigma \to Q$ that captures the discrete evolution of the model.
- C represents the set of system constraints that link the continuous variables. It represents a set of (ordinary) differential/difference equations along with algebraic equations for each mode $q_i \in Q$ and thus defines the continuously-valued evolution of the automaton.
- $(Q_0, \zeta_0) \subset Q \times \zeta$ specifies the initial state information

When dealing with the discretely-valued part of the hybrid automaton we denote the associated discrete event system (DES) through $M := (Q, \Sigma, T, Q_0)$. Analogously, we denote the underlying continuously-valued part through the (multi-mode) system $\Xi := (\zeta, Q, C, \zeta_0)$. For the scope of this paper, we use a discrete-time linear model with sampling period T_s that associates each mode $q_i \in Q$ with a difference equation

$$\mathbf{x}_{k+1} = \mathbf{A}_i \mathbf{x}_k + \mathbf{B}_i \mathbf{u}_k + \mathbf{N}_i \mathbf{v}_k \quad (2)$$

and an algebraic equation that defines the measurements through

$$\mathbf{y}_k = \mathbf{C}_i \mathbf{x}_k + \mathbf{D}_i \mathbf{u}_k + \mathbf{M}_i \mathbf{v}_k, \quad (3)$$

where $\mathbf{x}_k, \mathbf{u}_k$ and $\mathbf{y}_k$ denote the valuation of the state, input and output variable at time $t = kT_s$, respectively. The variable $\mathbf{v} := [v_1, \ldots, v_{n_x+n_y}]^T$ defines state noise $(v_1, \ldots, v_{n_x})$ and measurement noise $(v_{n_x+1}, \ldots, v_{n_x+n_y})$ through *bounded, zero mean* noise with $|v_{h,k|} \leq 1$. The possibly mode-specific magnitude of the disturbances is specified through the scaling vectors $\mathbf{n}_i$ and $\mathbf{m}_i$ that define the noise matrices $\mathbf{N}_i = [\mathrm{diag}(\mathbf{n}_i), \mathbf{0}]$, $\mathbf{M}_i = [\mathbf{0}, \mathrm{diag}(\mathbf{m}_i)]$ which select and scale the appropriate fraction of $\mathbf{v}$.

We use the hybrid system's assumption that mode changes take place infrequently and instantly, i.e. the mode evolves, compared to the continuously-valued evolution at a slower rate. As a consequence it is legitimate to assume that only one mode change takes place within one sampling period. In terms of estimation we restrict this assumption even further and assume that an event takes place at a particular sampling time-point. As a consequence, we will use the discrete variable e to track the *observable events.* Its valuation at a time-step k will be denoted through $e_k \in \{\Sigma_O, \epsilon\}$, where ϵ stands for *no observable event.*

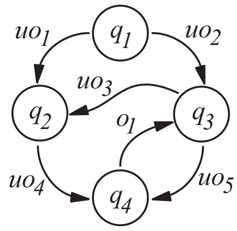

Fig. 1. Discrete automaton with four operational modes and six possible mode transitions.

2.1 Illustrative Example

Our framework is illustrated on the basis of a hybrid system with four modes of operation $Q = \{q_1, q_2, q_3, q_4\}$. We define mode transitions through one observable event $\Sigma_O = \{o_1\}$ and five unobservable events uo_1, uo_2, uo_3 uo_4 and uo_5 as depicted in Figure 1. The underlying continuous dynamics in equations (2 - 3) are defined through

$$A_1 = \begin{bmatrix} 0.7 & 0 & 0 \\ 0 & 0.7 & -0.1 \\ 0 & -0.1 & 0.1 \end{bmatrix} \; A_2 = \begin{bmatrix} -0.5 & 4 & 0 \\ 0 & 0.6 & 0 \\ 6 & 0 & 0.8 \end{bmatrix}$$

$$A_3 = \begin{bmatrix} 0.3 & -0.3 & 0 \\ 0 & 0.6 & 0 \\ -0.3 & 0 & 0.9 \end{bmatrix} \; A_4 = \begin{bmatrix} 0.6 & -0.3 & 0 \\ 0.3 & 0.6 & 0 \\ -0.6 & 0 & 0.9 \end{bmatrix}$$

$$B_1 = \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix} \; B_2 = B_3 = \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} \; B_4 = \begin{bmatrix} 2 \\ 2 \\ 0 \end{bmatrix}$$

$$C_1 = \begin{bmatrix} 1 & 1 & 0 \\ 1 & 0 & 0 \end{bmatrix} \; C_2 = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix}$$

$$C_3 = \begin{bmatrix} 1 & 0 & 1 \\ 0 & 1 & 1 \end{bmatrix} \; C_4 = \begin{bmatrix} 1 & 0 & 1 \\ 0 & 1 & 1 \end{bmatrix}$$

$$D_1 = D_2 = D_3 = D_4 = \begin{bmatrix} 1 \\ 0 \end{bmatrix}$$

For illustrative purposes of the individual algorithmic aspects we first use the model without noise ($\mathbf{N}_i = \mathbf{0}, \mathbf{M}_i = \mathbf{0}$). However, the final analysis is conducted with non-zero bounded noise.

3. HYBRID ESTIMATION

Hybrid estimation reconstructs the mode of operation and its associated continuous state at each time-step k :

Hybrid Estimation Problem: Given a hybrid model S, the discrete-time sequence of noisy (continuous) observations $\{\mathbf{y}_1, \ldots, \mathbf{y}_k\}$, the sequence of observable events $\{e_0, \ldots, e_k\}$ and the actuated control inputs $\{\mathbf{u}_0, \ldots, \mathbf{u}_k\}$, estimate the hybrid state that is comprised of the mode of operation $q_i \in Q$ and the continuous state $\mathbf{x}_k$ for time-step k.

We cannot fully observe the mode evolution of the automaton, nor do we usually know the initial mode exactly ($Q_0 \subset Q$ is not necessarily a singleton). As a consequence, a hybrid estimator has to consider all possible evolutions that are conform with the actuation and observations.

3.1 Suboptimal Hybrid Estimation

Early solutions to the hybrid estimation problem such as the multi-model IMM algorithm (Blom and Bar-Shalom [1988]) track hypotheses over a limited number of time-steps only and merge the continuous estimates according to a measure of likelihood. This likelihood is mostly drawn from the continuous filters and expresses the level of agreement between the estimate and the observations but might also include prior transition probability information, if available.

Our hybrid estimation algorithm (hME, see Hofbaur and Williams [2004]) uses the likelihood measure to *focus* on the set of most likely hypotheses. As a result, we obtain an any-time any-space algorithm that uses a focused search strategy to efficiently compute the leading hypotheses filtering out the majority of hypotheses with low likelihood.

Even though we were able to show that our approach can successfully handle systems with a large (l > 100,000) number of modes, we can say that every additional method to focus on possible hypotheses improves the estimation result and, of course, reduces the associated computational effort.

3.2 Mode Estimation through Parity-Space Methods

Our recent work on hybrid systems diagnosis (Bayoudh et al. [2008b]) provides an alternative approach that, because we were mostly concerned about diagnosis, concentrates on the mode estimate only. It uses a parity-space method that we extended to hybrid systems.

In a first step, we derive for every mode $q_i \in Q$ of the hybrid system S a set of Analytical Redundancy Relations (ARRs) that relate the continuous inputs $\mathbf{u}_{k-p}, \ldots, \mathbf{u}_k$ with the observable continuous outputs $\mathbf{y}_{k-p}, \ldots, \mathbf{y}_k$ over a time-window of length $p+1$. Selecting p appropriately (typically $p \leq n_x$) allows us to eliminate any dependencies upon the system's continuous state $\mathbf{x}$. This standard procedure from FDI Gertler [1991] can be summarized for a particular mode q_i of our hybrid system (2-3) as follows:

If we stack the input, output and noise according to

$$U_k := [\mathbf{u}_{k-p}^T, \ldots, \mathbf{u}_k^T]^T, \quad Y_k := [\mathbf{y}_{k-p}^T, \ldots, \mathbf{y}_k^T]^T,$$
$$V_k := [\mathbf{v}_{k-p}^T, \ldots, \mathbf{v}_k^T]^T,$$

we can re-write (2-3) and obtain for the continuous evolution in mode q_i

$$Y_k = \mathbf{O}_i \mathbf{x}_{k-p} + \mathbf{L}(\mathbf{A}_i, \mathbf{B}_i, \mathbf{C}_i, \mathbf{D}_i) U_k + \mathbf{L}(\mathbf{A}_i, \mathbf{N}_i, \mathbf{C}_i, \mathbf{M}_i) V_k, \tag{4}$$

with the matrices

$$\mathbf{O}_i := \begin{bmatrix} \mathbf{C}_i \\ \mathbf{C}_i \mathbf{A}_i \\ \vdots \\ \mathbf{C}_i \mathbf{A}_i^p \end{bmatrix} \tag{5}$$

and

$$\mathbf{L}(\mathbf{A}, \mathbf{B}, \mathbf{C}, \mathbf{D}) := \begin{bmatrix} \mathbf{D} & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{CB} & \mathbf{D} & \ddots & \vdots \\ \vdots & \ddots & \ddots & \mathbf{0} \\ \mathbf{CA}^{p-1}\mathbf{B} & \cdots & \mathbf{CB} & \mathbf{D} \end{bmatrix}. \tag{6}$$

For a sufficiently large p, there always exists a matrix $\mathbf{\Omega}_i$ that is orthogonal to the matrix $\mathbf{O}_i$, i.e. $\mathbf{\Omega}_i \mathbf{O}_i = \mathbf{0}$, so that we can eliminate the state $\mathbf{x}_{k-p}$ in (4) through left-hand multiplication with $\mathbf{\Omega}_i$. Hence, we can define the residual vector

$$\mathbf{r}_{i,k} := \mathbf{\Omega}_i Y_k - \mathbf{\Omega}_i \mathbf{L}(\mathbf{A}_i, \mathbf{B}_i, \mathbf{C}_i, \mathbf{D}_i) U_k. \tag{7}$$

In a noise-free environment we have to check the ARR consistency simply through $\mathbf{r}_{i,k} = [r_{i1,k}, \ldots, r_{im_i,k}]^T = \mathbf{0}$. However, if we include bounded noise as in our model (2-3) we can compute bounds ε_{ij} on the individual residuals $r_{ij,k}$ through the matrix

$$\mathbf{W} := \mathbf{\Omega}_i \mathbf{L}(\mathbf{A}_i, \mathbf{N}_i, \mathbf{C}_i, \mathbf{M}_i) \tag{8}$$

that captures the influence of the disturbances within the observation window of length $p+1$.

With this information, we can write the consistency check

$$\tilde{r}_{ij,k} := \begin{cases} 0 & \text{if } |r_{ij,k}| \leq \varepsilon_{ij} \\ 1 & \text{otherwise} \end{cases}, \quad j = 1, \ldots, m_i \tag{9}$$

and obtain a boolean residual vector for mode q_i at time-step k as

$$\tilde{\mathbf{r}}_{i,k} := [\tilde{r}_{i1,k}, \ldots, \tilde{r}_{im_i,k}]^T. \tag{10}$$

To extend this rather standard ARR approach to multi-mode systems, we proposed in (Bayoudh et al. [2008b]) to use the residuals for all $l = |Q|$ modes of the automaton concurrently, i.e. we *combine* all l residual vectors to form

$$\tilde{\mathbf{r}}_k := [\tilde{\mathbf{r}}_{1,k}^T, \ldots, \tilde{\mathbf{r}}_{l,k}^T]^T \tag{11}$$

We filter this vector to eliminate transients and obtain a *mode signature* with dedicated zero elements for each mode of operation. Additional discrete events Σ^{Sig} are then generated upon signature changes that provide additional evidence about mode transitions on the basis of the continuous evolution. We now extend the discrete-event system part of our hybrid automaton with the additional events Σ^{Sig} and use this extended discrete-event system

$$\tilde{M} := (Q, \{\Sigma, \Sigma^{Sig}\}, T, Q_0) \tag{12}$$

to derive a discrete event diagnoser (Sampath et al. [1995]) for our hybrid model. This diagnoser provides a mode estimate that takes both, the discrete and continuous evolution of the hybrid automaton into account. In Bayoudh et al. [2008a], we provided a detailed discussion on hybrid systems diagnosability analysis on the basis of (12). Figure 2 shows the associated architecture with the multi-mode residual/signature generator, consecutive filtering to deal with transients of the computed residuals $r_{ij,k}$ and the DES-diagnoser that provides the mode estimate.

4. HYBRID ESTIMATION THROUGH MODE-SET FOCUSING

The mode set estimate from Section 3.2 is now used to focus hybrid estimation from Section 3.1 which will then provide the full hybrid estimate, i.e. the mode of operation together with the continuous state estimate. Ideally, our

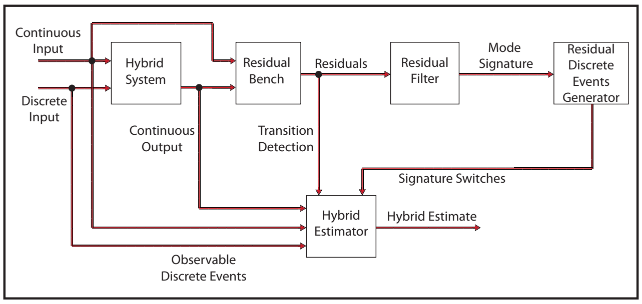

Fig. 2. Hybrid Diagnosis Scheme

ARR-based mode estimator deduces a single mode estimate for the current mode of operation/failure. In such case, this mode can be simply used and the associated continuous filter applied for state estimation. However, uncertain initial mode information Q_0 might require some time until the mode estimator obtains enough evidence through several signature changes to arrive at a unique mode estimate. A limited amount of sensors may also lead to reduced diagnosability so that the mode estimator won't be able to discriminate among several modes.This discrimination problem applies to successive multi-model filters or traditional hybrid estimators as well. This is simply due to the fact that a well designed parity-space diagnoser can always be replaced by a corresponding filter-based diagnoser (Gertler [1991]). However, it is shown further on that it is still worthwhile to assemble both in a coherent architecture. Such a rather un-intuitive approach will highly contribute to the estimation quality, both in terms of prompt mode change identification and also in terms of the estimation quality for the continuous state estimate.

4.1 Mode-Change Detection and Identification

The residuals and the discrete events Σ^{sig} generated from them can detect and identify previously unobservable mode changes. Mode-change detection can be drawn from abrupt residual changes almost instantly. Mode identification, however, requires additional time. The residuals need at least p time-steps to settle after a mode change since the ARR constraint (7) assumes single-mode dynamics within the observation window. Figure 3 shows the adaption process for a mode transition $q_2 \rightarrow q_4$ at time-step $k = 180$ with window-length $p = 2$ and filter length $\tau_{filter} = 3$.

The consecutive filter postpones the mode identification additionally for another τ_f time-steps (we used $\tau_f = n_x$ in our example). A single continuous filter would thus

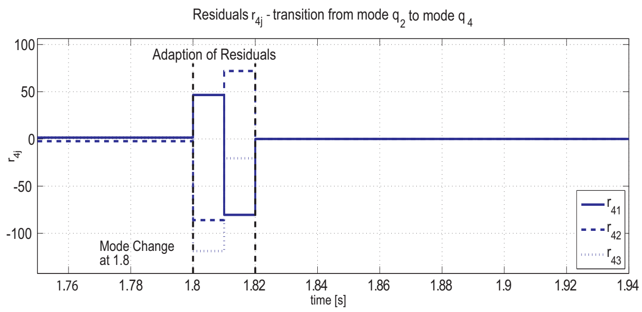

Fig. 3. Residual of mode q_4: $\mathbf{r}_4 = [r_{41}, r_{42}, r_{43}]^T$

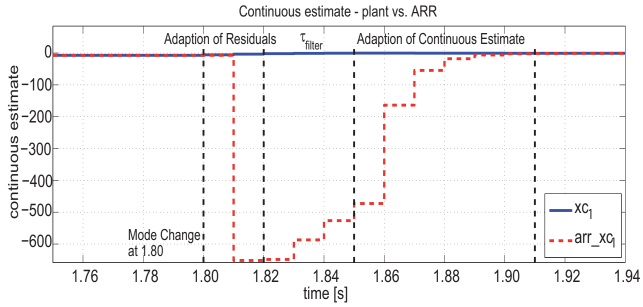

Fig. 4. Continuous state deviation due to the mode identification delay for a mode change $q_2 \rightarrow q_4$

estimate at the old mode of operation during this period of length $(p + \tau_f)T_s$. This can cause a dramatic increase of the estimation error for the continuous state. As a consequence, the filter needs additional time to adapt the continuous state once the mode estimator identifies the correct mode. In our example (see Fig. 4), the error is dramatically big around $t = 1.82s$ as the estimation of xc_1 drops down below -600, whereas the actual state is around $xc_1 \approx -4$ (see also Figure 5, solid line). At time-step $t = 1.85s$ the residuals are settled and the continuous state is estimated using the filter for mode q_4. Still, the filter needs another 6 time-steps until the continuous estimate complies with the actual state of the system.

4.2 Mode-Set Estimation

One way to reduce the mode identification delay would be to use shorter observation windows $p < n_x$ whenever one has more than one measurement ($n_y > 1$). However, this might lead to diagnosability problems since mode signatures are not necessarily distinct anymore. Our example provides equivalent signatures for mode q_3 and q_4 for $p = 1$ (Bayoudh et al. [2008b]). As a consequence, we obtain a mode set $\bar{Q}_k \subset Q$ as mode estimate at time-step k. The idea is to leave the final decision about mode consistency to the consecutive multi-mode/hybrid estimator, for example our hME algorithm, that also provides the associated continuous state estimate. The focused mode set highly contributes to the computational effort of the consecutive estimator.

4.3 Mode-set focused Hybrid Estimation

A careful interaction between the ARR-based mode estimator and a successive hybrid estimator can significantly contribute to the overall hybrid estimation quality. In particular at time-steps in the vicinity of mode changes. As mentioned before, we can detect mode changes almost immediately, whereas the mode estimator settles upon the mode of operation after several time steps only. Thus, the mode estimator can withdraw its predicted mode or mode-set upon mode-change detection. This allows the hybrid estimator to consider every possible hypothesis starting from the previous, possibly unique, estimate. The consequence is twofold:

- continuous state estimates for all hypotheses under consideration with their associated likelihood values are immediately obtained. Likelihoods allow us to indicate the most likely estimation hypothesis and

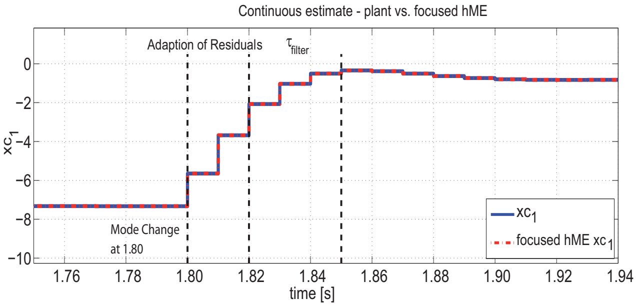

Fig. 5. Mode-set focused hybrid estimation for the mode change $q_2 \rightarrow q_4$

thus possibly identify the mode of operation even before the mode estimator provides new evidence as well.

- the adaptation delay of the continuous estimate that we saw earlier is avoided by tracking multiple hypotheses. The mode estimator simply provides additional evidence that allows the hybrid estimator to focus upon the correct estimation hypothesis, thus it immediately provides the correct continuous estimate!

Figure 5 shows this improved estimation result for the mode transition that was previously analysed in Figure 4. The continuous estimate $focused\ hME\ xc_1$ (dotted line) follows the actual behavior of the system xc_1(continuous line) immediately (graphs are superposed).

Of course, an analogous interaction can be used whenever the mode estimator settles upon a focused set of modes. This additional evidence can easily be included in the multi-mode estimation scheme as additional focusing method that contributes to the estimation quality but also reduces the computational effort that is necessary for hybrid estimation.

5. EXAMPLE CONT.

For our final experiments we considered the example of Section 2.1 with four modes of operation and a sampling period of $T_s = 0.01s$. To fully demonstrate our improved algorithm our example is now supplemented with additional noise that acts through

$$N_i = \begin{bmatrix} 0.025 & 0 & 0 & 0 & 0 \\ 0 & 0.025 & 0 & 0 & 0 \\ 0 & 0 & 0.025 & 0 & 0 \end{bmatrix}, \quad i = 1, \ldots, 4$$

$$M_i = \begin{bmatrix} 0 & 0 & 0 & 0.025 & 0 \\ 0 & 0 & 0 & 0 & 0.025 \end{bmatrix}, \qquad i = 1, \ldots, 4\ .$$

To point out the improvements of our new hybrid estimation scheme in correctly estimating both the continuous and the discrete behavior of the system, we performed several test-runs defining mode sequences for the system under consideration and compared the estimation results to the output of an ARR based approach with a window-length of $p = 2$ and $\tau_f = 3$ and a Rao-Backwellised Particle Filter with $n = 1000$ particles.

Figure 6 shows an exemplary mode trace with the associated mode estimates of our focused hybrid estimation (fhe) in comparison to the pure ARR-based mode estimator (a) and the Rao-Backwellised Particle Filter (b). In this experiment, we start with the initial state $Q_0 = q_1$ and command the system to follow the discrete trajectory (time in s):

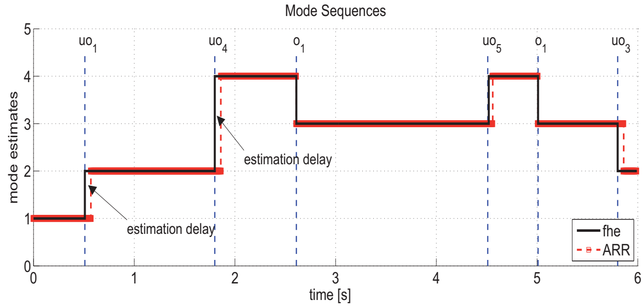

(a) Comparison of mode identification delays - mode set focused hybrid estimation vs ARR

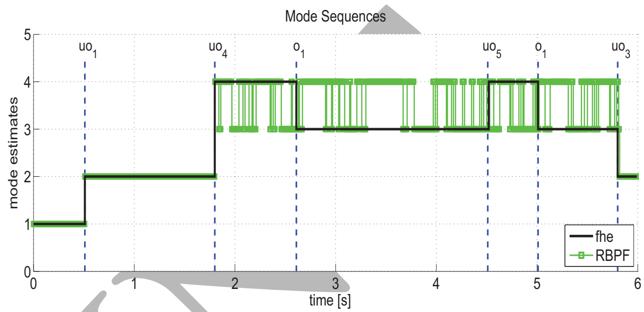

(b) Comparison of mode estimates - mode set focused hybrid estimation vs RBPF

Fig. 6. Mode sequence and mode estimate

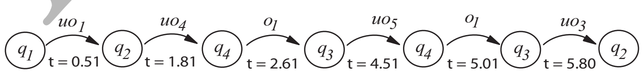

Our focused hybrid estimation approach follows the mode changes of the plant within a delay of at most one time-step, whereas the ARR approach needs $p + \tau_f = 5$ time-steps to identify the correct mode. The RBPF reacts to a mode change immediately, however it is not always able to discriminate between mode 3 and mode 4. The result of this behavior is again a considerable increase of the estimation error while the estimator is presuming the wrong mode. In Figure 7 a detail of the appropriate state estimates of the continuous state x_3 before and after a transition from mode q_3 to q_4 is depicted. Our approach identifies the correct mode one step after the transition takes place at $t = 4.52$ and the continuous estimate follows the plant immediately.

6. CONCLUSION

A novel scheme for hybrid estimation that de-couples mode and continuous state estimation through performing two redundant schemes is demonstrated. It is based on a parity-space based estimation for mode estimation and a consecutive multi-model filter based continuous estimation. This approach enables us to detect and identify mode changes / faults in the system with small delays and to focus the computationally complex hybrid filtering task onto few estimation hypotheses. The assessment of how well this algorithm performs compared to existing methods

(Rao-Backwellised KF (RBPF), standard IMM and hME without ARR based focusing) remains to be done through empirical testing on a set of case studies.

Hybrid systems with many modes require a slightly modified approach since it is difficult to compute the full signature that contains all modes at all time-steps. We are currently investigating approaches where we compute filters and ARRs on demand, similar to our on-line filter decomposition and deduction scheme (Hofbaur and Williams [2004]), and use partial signatures to guide the mode estimation. Another line of research investigates a distributed version of our hybrid estimator where several estimators work concurrently and cooperatively on the hybrid estimation task for complex multi-component systems.

Our ultimate goal is to provide a good hybrid estimate for the system. Sometimes, however, it can be possible that the current operation of an artifact does not reveal enough information so that an estimator can discriminate among several estimation candidates (e.g. due to insufficient excitation whenever it operates at a specific operational mode/point). As a consequence, it is important for us to develop a technique that can maintain several estimation hypotheses, both in terms of the discrete (mode) estimate and the continuous state estimate. In order to refine the estimate, one has to *actively* excite the system in a manner that satisfies the operational goal but also that reveals enough information to guide the mode discrimination. In Bayoudh et al. [2008c] it was shown that the discrete event diagnoser of our proposed parity-space diagnosis technique provides the essential basis to guide active hybrid diagnosis among the modes of operation. The supplemented continuous estimation scheme and its interaction with the parity-space diagnoser that we propose in this paper provides the other ingredient, the associated continuous state estimate, so that an *active hybrid diagnoser* can compute a suitable combined continuous/discrete actuation that leads to an un-ambiguous diagnosis/estimation for the system whilst as well maintaining at best its operational goal.

REFERENCES

G.A. Ackerson and K.S. Fu. On state estimation in switching environments. *IEEE Transactions on Automatic Control*, 15:10–17, 1970.

M. Bayoudh, L. Travé-Massuyès, and X. Olive. Coupling continuous and discrete event system techniques for hybrid system diagnosability analysis. In *Proceedings of the European Conference on Artificial Intelligence (ECAI 08), Patras, Greece*, pages 219–223, 2008a.

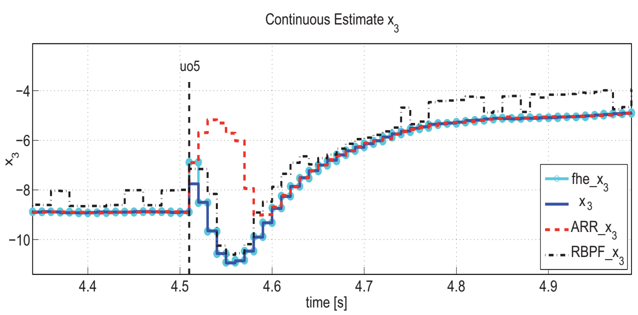

Fig. 7. Exemplary continuous state estimate (x_3)

M. Bayoudh, L. Travé-Massuyès, and X. Olive. Hybrid systems diagnosis by coupling continuous and discrete event techniques. In *Proceedings of the IFAC World Congress, Seoul, Korea*, pages 7265–7270, 2008b.

M. Bayoudh, L. Travé-Massuyès, and X. Olive. Towards active diagnosis of hybrid systems. In *Proceedings of the 19th workshop on Principles of Diagnosis (DX08), Blue Mountains, Australia*, pages 231–237, 2008c.

E. Benazera, L. Travé-Massuyès, and P. Dague. State tracking of uncertain hybrid concurrent systems. In *Proceedings of the 13th International Workshop on Principles of Diagnosis (DX02)*, pages 106–114, May 2002.

H.A.P. Blom and Y. Bar-Shalom. The interacting multiple model algorithm for systems with markovian switching coefficients. *IEEE Transactions on Automatic Control*, 33:780–783, 1988.

N. de Freitas. Rao-blackwellised particle filtering for fault diagnosis. In *Proceedings of the IEEE Aerospace Conference 2002*, volume 4, pages 1767–1772, 2002.

J. Gertler. Analytical redundancy methods in failure detection and isolation. In *Preprints of the IFAC SAFEPROCESS Symposium*, pages 9–21, 1991.

T. Henzinger. The theory of hybrid automata. In *Proceedings of the 11th Annual IEEE Symposium on Logic in Computer Science (LICS '96)*, pages 278–292, New Brunswick, New Jersey, 1996.

M. W. Hofbaur and B. C. Williams. Mode estimation of probabilistic hybrid systems. In C.J. Tomlin and M.R. Greenstreet, editors, *Hybrid Systems: Computation and Control, HSCC 2002*, volume 2289 of *Lecture Notes in Computer Science*, pages 253–266. Springer Verlag, 2002.

M. W. Hofbaur and B. C. Williams. Hybrid estimation of complex systems. *IEEE Transactions on Systems, Man, and Cybernetics - Part B: Cybernetics*, 34(5):2178–2191, October 2004.

X.R. Li and Y. Bar-Shalom. Multiple-model estimation with variable structure. *IEEE Transactions on Automatic Control*, 41:478–493, 1996.

S. Narasimhan and G. Biswas. An approach to model-based diagnosis of hybrid systems. In C.J. Tomlin and M.R. Greenstreet, editors, *Hybrid Systems: Computation and Control, HSCC 2002*, volume 2289 of *Lecture Notes in Computer Science*, pages 308–322. Springer Verlag, 2002.

S. Narasimhan, R. Dearden, and E. Benazera. Combining particle filters and consistency based approaches for monitoring and diagnosis of stochastic hybrid systems. In *Proceedings of the 15th International Workshop on Principles of Diagnosis (DX04)*, pages 123–128, 2004.

M. Sampath, R. Sengputa, S. Lafortune, K. Sinnamohideen, and D. Teneketzis. Diagnosability of discrete event systems. *IEEE Transactions on Automatic Control*, 40:1555–1575, Sept. 1995.

V. Verma, G. Gordon, R. Simmons, and S. Thrun. Real-time fault diagnosis. *IEEE Robotics and Automation Magazine*, 11(2):56–66, 2004.