Hybrid Estimation through Synergetic Mode-Set Filtering

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)*

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 does not estimate the continuous state. This paper proposes a novel scheme that uses the latter method as a mode focusing procedure for hybrid estimation so that we can significantly reduce the number of hypotheses that a hybrid estimator has to consider. The advantages of this synergetic method are on both sides: it boosts the mode identification time and the convergence of continuous state estimation.

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 models. 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 hybrid 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 can now operate on a significantly smaller (ideally 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

$$S = (\zeta, Q, \Sigma, T, C, (Q_0, \zeta_0)), \tag{1}$$

where:

- ζ is the set of continuous [1] variables, which includes n_u (exogenous) *input variables* $\{u_1, \ldots, u_{n_u}\} =: \mathbf{u}$, n_x

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

[1] continuously valued to be precise

state variables $\{x_1, \ldots, x_{n_x}\} =: \mathbf{x}$ that capture the dynamic evolution of the automaton (for example in discrete time), 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 or failure 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

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

and an algebraic equation that defines the measurements through

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

where $\mathbf{x}_k, \mathbf{u}_k$ and $\mathbf{y}_k$ denote the valuation of the state, input and output variables 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 $\mathbf{N}_i = [\text{diag}(\mathbf{n}_i), \mathbf{0}]$, $\mathbf{M}_i = [\mathbf{0}, \text{diag}(\mathbf{m}_i)]$ 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*.

2.1 ILLUSTRATIVE EXAMPLE

Our framework is illustrated on the basis of a hybrid system with six modes of operation $Q = \{q_1, \ldots, q_6\}$. We define mode transitions through three observable events $\Sigma_O = \{o_1, o_2, o_3\}$ and seven unobservable events $uo_1, \ldots, uo_7$ as depicted in Fig. 1.

The underlying continuous dynamics are of form (2 - 3) and defined through

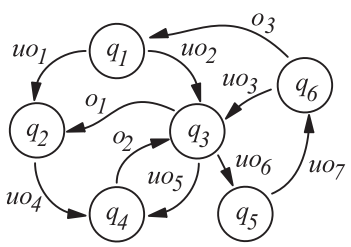

Fig. 1. Discrete automaton with six operational modes.

$$\mathbf{A}_1 = \begin{bmatrix} 0.7 & 0 & 0 \\ 0 & 0.7 & -0.1 \\ 0 & -0.1 & 0.1 \end{bmatrix} \quad \mathbf{A}_2 = \begin{bmatrix} -0.5 & 4 & 0 \\ 0 & 0.6 & 0 \\ 6 & 0 & 0.8 \end{bmatrix}$$

$$\mathbf{A}_3 = \begin{bmatrix} 0.3 & -0.3 & 0 \\ 0 & 0.6 & 0 \\ -0.3 & 0 & 0.9 \end{bmatrix} \quad \mathbf{A}_4 = \begin{bmatrix} 0.6 & -0.3 & 0 \\ 0.3 & 0.6 & 0 \\ -0.6 & 0 & 0.9 \end{bmatrix}$$

$$\mathbf{A}_5 = \begin{bmatrix} 0.3 & -0.3 & 0.3 \\ -0.3 & 0.6 & 0.3 \\ -0.3 & 0 & 0.9 \end{bmatrix} \quad \mathbf{A}_6 = \begin{bmatrix} 0.5 & 2 & 0 \\ 0 & 0.3 & 0 \\ 2 & 0 & 0.4 \end{bmatrix}$$

$$\mathbf{B}_1 = \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix} \quad \mathbf{B}_2 = \mathbf{B}_3 = \mathbf{B}_6 = \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} \quad \mathbf{B}_4 = \begin{bmatrix} 2 \\ 2 \\ 0 \end{bmatrix}$$

$$\mathbf{B}_5 = \begin{bmatrix} 1 \\ 2 \\ 1 \end{bmatrix} \quad \mathbf{C}_1 = \begin{bmatrix} 1 & 1 & 0 \\ 1 & 0 & 0 \end{bmatrix}$$

$$\mathbf{C}_2 = \mathbf{C}_3 = \mathbf{C}_4 = \mathbf{C}_6 = \begin{bmatrix} 1 & 0 & 1 \\ 0 & 1 & 1 \end{bmatrix} \quad \mathbf{C}_5 = \begin{bmatrix} 1 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix}$$

$$\mathbf{D}_1 = \mathbf{D}_2 = \mathbf{D}_3 = \mathbf{D}_4 = \mathbf{D}_5 = \mathbf{D}_6 = \begin{bmatrix} 1 \\ 0 \end{bmatrix}$$

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 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 $\{\mathbf{u}_0, \ldots, \mathbf{u}_k\}$, compute an estimate of the hybrid state $\hat{\mathbf{x}}_{h,k}$ that combines the discrete mode estimate $\hat{\mathbf{x}}_{d,k} = q_i \in Q$ and the continuous state estimate $\hat{\mathbf{x}}_{c,k}$ as tuple $\langle \hat{\mathbf{x}}_{d,k}, \hat{\mathbf{x}}_{c,k} \rangle$ 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 Hybrid estimation suboptimal Methods

Early solutions to the hybrid estimation problem such as the multi-model IMM algorithm (Blom and Bar-Shalom [1988]) track hypotheses, i.e. mode sequences and their associated continuous evolution, over a limited number of time-steps only

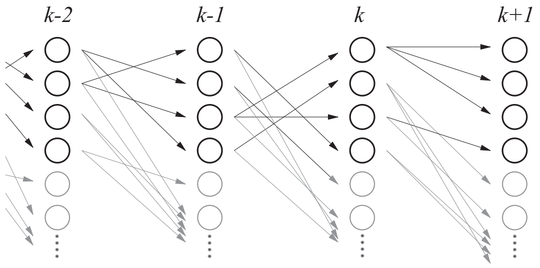

Fig. 2. Beam-search procedure to focus on the $\lambda = 4$ leading hypotheses

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 hypotheses' estimate and the observations but might also include prior transition probability information, if available.

Our hybrid estimation algorithm (hME, see Hofbaur and Williams [2002, 2004]) uses the likelihood measure b to guide a best-first search process onto the set of most likely hypotheses. Starting from a set of λ hybrid estimates $\hat{\mathbf{x}}^{(1)}_{h,k-1}, \ldots, \hat{\mathbf{x}}^{(\lambda)}_{h,k-1}$ that are ordered according to their likelihood $b^{(i)}_{k-1}, i = 1 \ldots, \lambda$ we use the discrete event system M of our model to compute and rank the valid mode transitions $\hat{\mathbf{x}}^{(i)}_{d,k-1} = q_\nu \rightarrow \hat{\mathbf{x}}^{(j)}_{d,k} = q_\iota$.

A consecutive continuous filtering operation that uses the continuous system Ξ of our model provides an associated continuous state estimate $\hat{\mathbf{x}}^{(i)}_{c,k}$ and the associated likelihood $b^{(i)}_{k}$ for each hypothesis i. As a result, we obtain an any-time any-space algorithm that uses a focused search strategy to efficiently the leading hypotheses compute at each time-step and avoids to compute the majority of hypotheses with low likelihood. The resulting overall behavior of our hME algorithm can be seen as a beam-search procedure as shown in Fig. 2.

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

$$U_k := [\mathbf{u}^T_{k-p}, \ldots, \mathbf{u}^T_k]^T, \;\; Y_k := [\mathbf{y}^T_{k-p}, \ldots, \mathbf{y}^T_k]^T,$$
$$V_k := [\mathbf{v}^T_{k-p}, \ldots, \mathbf{v}^T_k]^T,$$

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

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

We filter this vector Bayoudh et al. [2008b] to eliminate transients and obtain a *mode signature* with dedicated zero elements for each mode of operation. Discrete events Σ^{Sig} are then generated upon mode 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).

4. SYNERGETIC HYBRID ESTIMATION

Mode estimation through parity space methods, as described in section 3.2, provides good mode estimates. If one is also interested in the continuous state estimate, one has to perform an additional state filtering operation. It is tempting to directly use the mode estimate from the parity space estimator to specify the mode of operation for the filtering process. However, this leads to an overall hybrid estimation scheme with unsatisfactory behavior. The parity space method assumes that the system operates at a constant mode over the last p time-steps. As a consequence, we obtain a valid mode estimate after at least p time-steps following the actual transition. In fact, the mode estimation requires even additional time since we pre-filter the residuals in order to avoid wrong mode predictions due to noise in the system. A continuous filter on the basis of this mode estimate would therefore continue for at least p time-steps with the wrong mode estimate causing a potential divergence of the continuous estimate with a consecutive convergence phase once the continuous estimator switched to the right mode again. As a consequence, we can obtain problematic continuous estimation performance in the vicinity of mode transitions. Furthermore, the DES diagnoser-based mode estimator would fail to track fast mode changes with several mode changes within the observation window of p time-steps.

Our first idea was therefore to combine the parity-space mode estimator with our hME algorithm to improve the hybrid estimation quality whenever mode changes occur in the system Rienmüller et al. [2009]. We can use the parity-space method to detect mode changes whenever some of its residuals go astray. Whenever this happens, we give control to hME to track all possible hypotheses from that time-point onwards. Once the mode estimator settles upon a new mode estimate, we can pick the associated hybrid estimates from hME and continue onwards with this, highly focused, set (ideally singleton) of hypotheses.

Our mode estimator builds upon the careful interplay between a set of Analytic Redundancy Relations and a DES-diagnoser Bayoudh et al. [2008b]. Combining this approach with a hybrid estimator for estimating both the mode of operation/failure and the continuous state has some drawbacks. First, the DES-diagnoser that is very useful for diagnosis analysis purposes is of high complexity Sampath et al. [1995], which imposes computational-time and space limitations. Furthermore, it requires the evaluation of the ARRs of all modes at every sampling time, which is problematic in systems with many modes. Another drawback of this approach arises whenever the system has mutually un-diagnosable modes, i.e. modes that share the same input/output behavior. The modes q_3 and q_5 of our example can illustrate this fact. Their models represent two different state-space realizations of a common dynamic system. As a consequence, the residuals for both modes share the same matrices $\mathbf{\Omega}_3 = \mathbf{\Omega}_5$ and $\mathbf{L}(\mathbf{A}_3, \mathbf{B}_3, \mathbf{C}_3, \mathbf{D}_3) = \mathbf{L}(\mathbf{A}_5, \mathbf{B}_5, \mathbf{C}_5, \mathbf{D}_5)$ so that we cannot use the residuals to generate an additional event for the unobservable mode transition $uo_6 : q_3 \rightarrow q_5$. As a consequence, we obtain a DES diagnoser that does not consider the mode q_5 at all but traverses from mode q_3 directly to mode q_6. A continuous state estimation would therefore again require a significant adaption phase since it misses the q_5 state evolution.

We therefore propose to use a simpler *residual-based mode estimator* (rME) instead. This mode estimator is used to *focus* the hybrid estimation result of hME upon those hypotheses that operate at modes that are consistent with the mode estimators' ARR constraints. On the other hand hME reduces the set of modes so that rME operates on a subset of the system's modes. As a consequence, we can apply this combined method to systems with a significantly larger number of modes. However, the correct operation relies on a careful interplay between hME and rME. With respect to hME, rME runs according to two modes of operation :

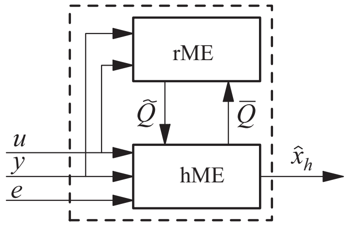

Fig. 3. Hybrid estimation architecture.

- a *focusing-mode* when it has settled upon a set of mode estimates that can be used to focus the set of hypotheses of hME,
- a *bypass-mode* during the transitory phase following mode transition in which no focus can be provided. The overall estimator will thus behave like standard hME that tracks even fast mode evolutions with several mode changes within p time steps.

Fig. 3 illustrates this basic interaction between the hybrid estimator hME and the mode estimator rME.

4.1 Residual-based Mode Estimator (rME)

The mode estimator rME maintains the ARR for the modes Q of the hybrid model, but evaluates ARRs to compute the associated residuals only for a reduced mode set $\bar{Q} \subseteq Q$ that it receives as an input (from hME). It utilizes the computed residuals according to its two modes of operation.

In the main *focusing-mode* it uses the computed and filtered residuals to select the reduced set of consistent modes $\tilde{Q}$. In addition it uses the residuals for the modes $\bar{Q}$ as indicators for a mode-transition event.

rME changes to the *bypass-mode* upon a transition detection where it defines the mode-set output simply through $\tilde{Q} = \bar{Q}$. It continues to track the filtered residuals for the modes $\bar{Q}$ and switches back to the *focusing-mode* as soon as the residuals provide evidence that the system settled at a new mode of operation again.

Analytic redundancy relations of each mode may be precomputed or generated on the fly from the corresponding dynamic model whenever $\bar{Q}$ contains new modes. If a subset of modes arises from the piecewise linearization of a global nonlinear model, a trade-off strategy may be to pre-compute the non linear analytic redundancy relations and instantiate them with their corresponding linear counterpart whenever the corresponding mode enters into play Bayoudh et al. [2009].

4.2 Focused hME

hME receives as input $\tilde{Q}$, which is generated by rME as the current set of consistent modes. hME hence uses $\tilde{Q}$ to discard hypotheses with modes that are not in this set as indicated in Fig 4. hME then provides the continuous state estimate for each hypothesis corresponding to the focused mode set $\tilde{Q}$.

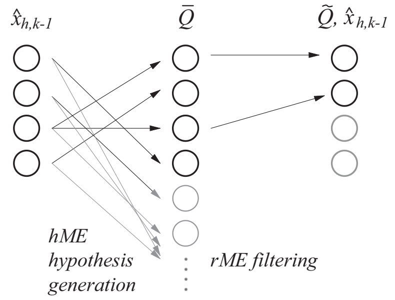

Fig. 4. Hybrid estimation computational steps.

This mechanism may significantly reduce the number of hypotheses followed by the hybrid estimator hME, avoiding the otherwise inevitable exponential blow-up.

4.3 Merging Hypotheses

Whenever rME reduces the mode set $\bar{Q}$, obtained by the hME, to a single mode $\tilde{Q} = q_i$, we can focus onto hypotheses at this particular mode, i. e. all hypotheses assuming a different mode can be discarded. Moreover, a settled mode estimate of rME indicates that the hybrid system evolved at this particular mode for at least p time-steps. So we have to focus on hypotheses that not just share the mode q_i in their estimate at time-step k but also in their p predecessor estimates. Typically, only a small set of hypotheses will satisfy this condition. Nevertheless, we can still end up with more than one hybrid estimate for time-step k. The following operation of hME can be twofold. (i) we can decide to keep all hypotheses independently and continue the estimation process with all of them, or (ii) we can merge the continuous estimate into a single inclusive continuous estimate and thus, continue hybrid estimation with a perfectly focused single estimate for time-step k. The example section 5 will illustrate both modes of operation.

4.4 Resetting the algorithm

Because of the truncation operation of hME - the algorithm considers only the *leading* set of hypotheses and discards all other hypotheses - it is possible that the true mode is not contained in $\bar{Q}$ and thus no mode hypothesis survives the rME test, i.e. $\tilde{Q} = \emptyset$. In this case it is necessary to reset the algorithm. Whenever the total number of modes is moderately large, we can start this resetting process with performing rME on $\bar{Q} = Q$ and using the thus obtained set $\tilde{Q}$ to initiate new hybrid estimates for hME to continue estimation onwards. Whenever we deal with multi-component systems it is possible to focus on the subpart of the system onto which hME lost its track and thus re-initiate only onto a fraction of the mode set $\bar{Q} \subset Q$ Hofbaur and Williams [2004].

5. ILLUSTRATIVE EXAMPLE CONT.

For our final experiments we considered the example of Section 2.1 with six modes of operation and a sampling period of $T_s = 0.01s$. We first examined the results of continuous estimation for a pure ARR-based mode estimator with a single continuous filter for each mode q_i. As mentioned before in section 4 a parity based method provides a valid mode estimate after at least p time-steps only. The consecutive filter postpones the mode identification additionally for another τ_{filter} time-steps.

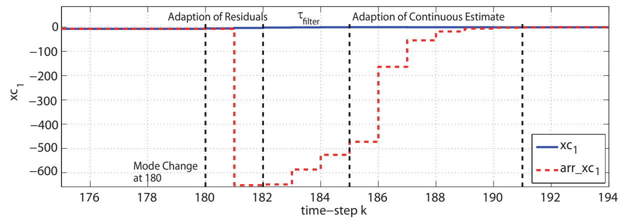

(a) Continuous state deviation due to the mode identification delay for a mode change $q_2 \rightarrow q_4$

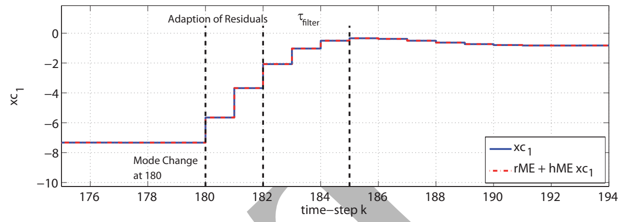

(b) Improved estimation result of the ARR focused hME

Fig. 5. Comparison of continuous estimation for a purely parity based method and the ARR focused hME.

A single continuous filter would thus estimate at the old mode of operation during this period. 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. Fig. 5a shows the adaption process on the estimation of the first state component xc_1 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 error is dramatically big around $k = 182$ as the estimation of xc_1 drops down below -600.

Using the new approach the rME switches to the *bypass mode* upon a transition detection, defining the new mode set simply through $\tilde{Q} = \bar{Q}$. This allows the hME to consider every possible hypothesis starting from the previous, possibly unique, estimate. As a consequence continuous state estimates for all hypotheses under consideration with their associated likelihood values are immediately obtained. The adaptation delay of the continuous estimate can thus be reduced. Fig. 5b depicts this improved estimation result. The continuous estimate arr_x_1 (dotted line) follows the actual behavior of the system (continuous line) immediately (graphs are superposed).

5.1 Merging Hypotheses

We then compared the different techniques from Section 4.3 - (i) continuing with all valid estimates or (ii) merging of identical mode hypotheses. In many cases the additional hypotheses can be discarded because their preceding hypotheses assumed a mode $q_i \notin \tilde{Q}$ within p time steps before the ARRs are settled. However, when mode changes happen faster than the rME can follow them or, as in our example mode q_3 and q_5, several modes cannot be distinguished by the residual equations, the hME is forced to consider every possible hypothesis in the usual way. This is important as we would otherwise loose track of the continuous estimate. As a result, several hypotheses may develop that end up with the same mode estimate, but with a different mode history and thus a slightly different continuous estimate. Whenever the rME gains certainty about the current mode again, all wrong hypotheses can be discarded and the hME would continue the estimation process starting with all hypotheses assuming the correct mode. This process leads to

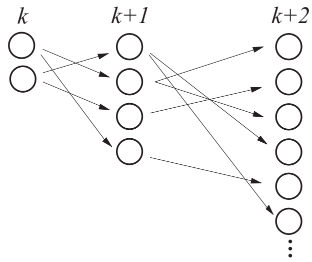

(a) Hybrid Mode Estimation starting with two hypotheses at time step k

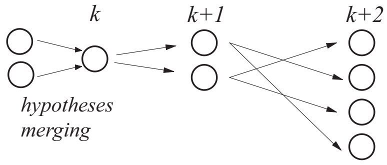

(b) Hybrid Mode Estimation after merging of two hypotheses.

Fig. 6. Comparison of the number of mode hypotheses generated by the hME (a) with and (b) without hypotheses merging.

an unnecessary blow up of the hypotheses tree. We thus had the idea of merging those hypotheses so that the estimation process then continues from one single hypothesis onwards, which reduces the number of hypotheses considerably without showing any drawbacks in terms of continuous state estimation. For this purpose we merge the intervals of the different hypotheses and sum up their probabilities. Hybrid estimation now continues with a perfectly focused single estimate $\hat{\mathbf{x}}_{d,k}^{(1)} = q_i$. Of course, the hME estimation can spread over several hypotheses for consecutive time-steps again, whenever rME/$\bar{Q}$ does specify more than one possible mode (e.g. q_3 and q_5 in our example) or a mode transition takes place(see Fig. 6).

5.2 Comparison of the new ARR focused hME with the standard hME

For our final experiments we supplement our example with additional noise that is specified through

$$N_i = \begin{bmatrix} 0.1 & 0 & 0 & 0 & 0 \\ 0 & 0.1 & 0 & 0 & 0 \\ 0 & 0 & 0.1 & 0 & 0 \end{bmatrix} \qquad (13)$$
$$M_i = \begin{bmatrix} 0 & 0 & 0 & 0.125 & 0 \\ 0 & 0 & 0 & 0 & 0.125 \end{bmatrix}.$$

Furthermore, the continuous input is specified by a sequence of random numbers with $|\mathbf{u}_k| \leq 2$ and the hME operates with $\lambda = 24$ leading hypotheses. Of course, hME considers more hypotheses until it finds the leading set. To point out the improvements of our new hybrid estimation scheme in correctly estimating both the continuous and the discrete behavior of the system while reducing the number of considered hypotheses significantly, we performed several test-runs with different mode sequences and compared the estimation results to the output of a standard hME.

Fig. 7 shows an exemplary mode trace with the associated mode estimates of our ARR focused hybrid estimation in comparison to the standard hME. In this case we start with the initial state $Q_0 = q_1$ and command the system to follow the discrete trajectory:

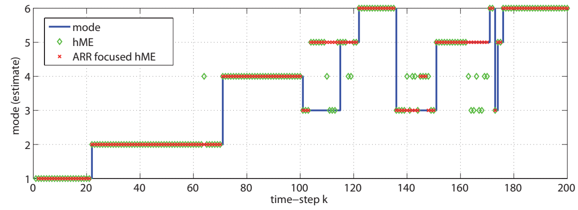

(a) Mode Estimate of the leading hypothesis: $\hat{\mathbf{x}}_{d,k}^{(1)}$

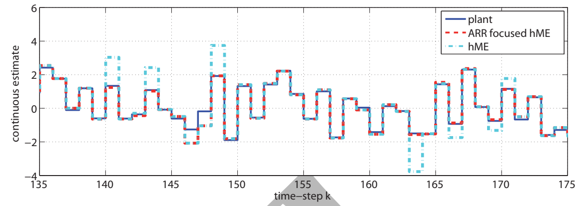

(b) Appropriate continuous estimate of the leading hypothesis: $\hat{\mathbf{x}}_{c,k}^{(1)}$

Fig. 7. Comparison of the hybrid estimates of standard hME and ARR focused hME.

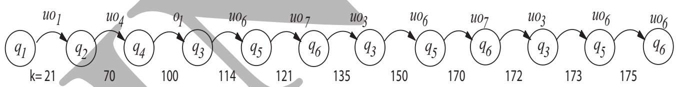

In terms of estimation quality both algorithms show a similar behavior. Modes q_1, q_2 and q_6 can be distinguished by both the hME and the rME. However, as mentioned before, modes q_3 and q_5 share the same matrices $\mathbf{\Omega}_3 = \mathbf{\Omega}_5$ so that the rME only can tell us that the system is in either of the two modes. Additionally, as the continuous behavior of mode q_4 is very similar to that of mode q_3, the residuals of q_4 sometimes lie within the bounds ϵ_{4j} although the system is in mode q_3 or q_5. As a consequence hypotheses that estimate q_4 survive the rME test in some cases (see for example $k = 145$ up to $k = 147$). Most of the incorrect mode classifications are due to the mutual undiagnosability of q_3 and q_5, in particular as hME provides the mode estimate through the most likely hypothesis. Still, the total number of incorrect mode estimates of the new approach is below the number of wrong estimates of the standard hME. However, the biggest advantage of the ARR focused hME in comparison to the hME is a *significant reduction of the number of hypotheses that the algorithm considers* for continuous filtering during its operation. Fig. 8 shows the number of hypotheses at every time step k for the ARR focused hME with and without the hypotheses merging operation and the standard hME. Especially when the system is in one of the critical modes (q_3, q_4 and q_5), the number of nodes in the search tree explodes for the hME, while our approach can keep the number of hypotheses at a moderate level. Furthermore, using the merging operation, the number of hypotheses is reduced perfectly to a single estimate and increases only in the vicinity of mode transitions or in modes with ambiguous residuals. This, of course, has dramatic implications on the overall runtime behavior of our algorithm.

6. CONCLUSIONS AND FUTURE WORKS

A novel scheme for hybrid estimation that uses a mode focusing procedure to reduce the number of hypotheses that have to be considered is proposed. It is based on the interaction of a hybrid estimator (hME) and a parity-space based mode estimator (rME). Among a set of modes corresponding to current hME

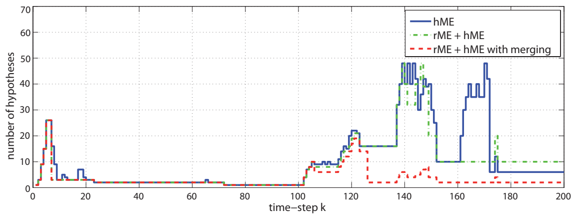

Fig. 8. Comparison of the number of considered hypotheses at each time-step k for the standard hME algorithm and the new ARR focused approach (rME und hME).

hypotheses, rME is able to identify the subset of modes that are consistent with observations, hence providing this focused set of modes as input to hME. This mechanism may significantly reduce the number of hypotheses followed by the hybrid estimator, avoiding the otherwise inevitable exponential blow-up.

Compared to our previous work Rienmüller et al. [2009] that used a mode estimator built upon an interplay between a set of ARRs and a DES diagnoser, the presently proposed mode estimator is computationally simpler as it only uses the ARRs. The proposed rME and hME interconnection avoids the problem of mutually un-diagnosable modes that are just not seen by the DES diagnoser. In addition, it allows us to compute filters and ARRs on demand, similar to our on-line filter decomposition and deduction scheme (Hofbaur and Williams [2004]).

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 - mode and continuous state - 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). As a consequence, it is important for us to develop a technique that can maintain several hybrid estimation hypotheses. 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 DES 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 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.

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.

M. Bayoudh, L. Trav-Massuys, and X. OLIVE. Diagnosis of a class of non linear hybrid systems by on-line instantiation of parameterized analytic redundancy relations. In *20th International Workshop on Principles of Diagnosis DX'09*, 2009.

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 *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 *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.

Th. Rienmüller, M. Mayoudh, M. Hofbaur, and L. Travé-Massuyès. Hybrid estimation through synergic mode-set focusing. In *Accepted for publication at the IFAC SAFEPROCESS Symposium*, 2009.

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.