

A New Filter for Hybrid Systems and its Applications to Robust Attitude Estimation

Pedro Santana*, Renato Lopes†, Bruno Amui†, Geovany Borges†, João Ishihara†, and Brian Williams*

Abstract—Fault diagnosis and recovery are essential tools for the development of autonomous agents that can operate in hazardous environments. This can be effectively approached from a model-based perspective, where sensor faults are explicitly taken into account in a hybrid model with switching dynamics. However, practical hybrid filters are required to manage an exponential growth in the number of discrete mode sequences, also known as hypotheses. Inspired by an attitude estimation application for a quadrotor UAV with faulty sensors, this paper introduces the IP-MHMF, a novel filter for hybrid systems that generalizes the well-known IMM and introduces a more informed hypothesis-pruning step than previous algorithms. By performing hypothesis pruning on corrected rather than predicted hypothesis probabilities, the IP-MHMF is capable of much more aggressive pruning strategies that significantly reduce its computational load, while improving its estimation performance. Our numerical results on data from a real robotic platform show that the IP-MHMF outperforms state-of-the-art hybrid filters and the traditional EKF on an attitude estimation application with faulty magnetometer measurements.

I. INTRODUCTION

Navigation and 3D localization are essential tasks for robotic systems [1], particularly for operation in outdoor and uncontrolled environments. Reliable pose estimates are usually obtained by combining data from a range of different sensors by means of filtering algorithms. In the case where sensors operate normally, it is reasonable to assume that “two sensors are better than one” [2], i.e., one should always expect a more refined belief about the hidden state of the system the more sensor information is provided. However, the reliability of this approach is much more susceptible to sensor failure and might be too strong of an assumption for robots built with cheap, commercially available components. This might also be true even for highly reliable robot hardware operating under severe disturbances, such as in the vicinity of electric power transmission lines. Given that faulty measurements from a single sensor may be sufficient to degrade the performance of the entire localization system, being able to quickly detect sensor faults and recover from them becomes a problem of paramount importance [3].

Detecting failures and anomalous behavior for dynamical systems has been a topic of great interest for many years, as

demonstrated by the survey in [4]. In its simplest instantiation, anomaly detection can be achieved by means of a gating strategy, where a fault flag is raised whenever a sensor output falls outside its “allowed” range. However, not only it is difficult to determine what these “allowed” ranges should be, but it might also be the case that a sensor is failing while still operating within its normality bounds. In this case, a fault can only be detected by combining the information from the current belief state with the data coming from all the other sensors. Therefore, in this work we adopt the approach of [3], [5] and model the data fusion problem with faulty sensors as an instance of hybrid filtering. In our formulation, the discrete portion of the hybrid state (a.k.a. *mode*) corresponds to the mode of operation of each sensor, while the continuous portion of the state represents the pose quantities (position and orientation) that we are trying to estimate.

Robust attitude estimation is one important instance of the general problem of hybrid estimation. When estimating the state of a hybrid system, one must deal with the combinatorial explosion in the number of possible sequences of discrete hidden states, also known as *hypotheses* [3]. Given that the number of hypotheses tends to grow exponentially with time, a complete enumeration is often impractical. Therefore, it is necessary to settle for suboptimal approaches that manage the growth in the number of hypotheses considered, while not significantly degrading the filter’s performance. The hypothesis-mixing step introduced by the Interacting Multiple Model (IMM) filter [6], [7] represents a milestone in hypothesis-management methods for hybrid filters. By merging (“mixing”) the estimates coming from a bank of Kalman filters, the IMM exhibits linear computational complexity (in terms of the number of modes), while attaining performance levels of filters with quadratic complexity, rendering it still one of the best choices in terms of cost and efficiency [8]. However, the IMM’s fixed-depth merging approach can be very restrictive, hence recent work has extended this algorithm in order to improve its performance. For instance, the M³H [9] and the MHMF [10] are shown to outperform IMM by leveraging the idea of variable hypothesis-merging depth, originally proposed for the Generalized Pseudo Bayes (GPB) filter [11]. While the M³H adopts a maximum likelihood approach for the selection of hypotheses, the MHMF extends the mixing step of the IMM in order to allow merging depths greater than one. Moreover, both algorithms incorporate a pruning step that discards hypotheses with probabilities less than some user-defined threshold. While this helps reduce the computational cost of the filter, by eliminating highly unlikely hypotheses, it can lead the hybrid filter to instability,

* Pedro Santana and Brian Williams are with the Model-based and Embedded Robotics Systems (MERS) group, Massachusetts Institute of Technology, Cambridge, MA, USA {psantana, williams}@mit.edu

† Renato Lopes, Bruno Amui, Geovany Borges, and João Ishihara are with Laboratório de Automação e Robótica (LARA), Departamento de Engenharia Elétrica, Universidade de Brasília, Brasília, DF, Brazil rvlopes@unb.br, brunoamui@gmail.com, gaborges@unb.br, ishihara@ene.unb.br

as discussed in Section III.

Our work is motivated by the accuracy and performance required by attitude estimation for a quadrotor unmanned aerial vehicle (UAV) designed to inspect electrical power transmission lines. We propose the Informed Pruning, Multiple-Hypotheses Mixing Filter (IP-MHMF), a new filtering algorithm for hybrid systems. It improves upon the prior art by introducing a more informed hypothesis-pruning step, which allows it to retain the computational speed gains of [9], [10] while preventing filter divergence. Its performance was evaluated on a nonlinear attitude estimation application featuring real sensor data with intermittent faults. It was compared against the MHMF, the M³H, and a traditional strategy employing an Extended Kalman filter (EKF) in terms of computational cost and robustness to sensor faults.

This paper is organized as follows. Section II describes hybrid filtering and multiple-hypotheses tracking, followed by the formulation of the IP-MHMF in Section III. Section IV presents the model of the attitude estimation system used to generate the numerical results in Section V. Finally, concluding remarks are presented in Section VI.

II. PROBLEM FORMULATION

This work was motivated by the problem of stabilization of a quadrotor UAV designed to operate in environments with strong magnetic disturbances. Equipped with a sonar, a three-axis magnetometer, and a 6 DOF IMU (three-axial accelerometer and gyro), the quadrotor's localization system relies heavily on its magnetometer to correct its orientation estimates. However, local magnetic field distortions tend to cause the magnetometer to occasionally indicate erroneous orientation changes even if the quadrotor remains static. Moreover, its low-cost magnetic sensor occasionally yields spurious readings for unknown reasons.

In this context, a hybrid approach is appropriate. Hybrid systems denote a class of dynamical systems whose behavior combines continuous and discrete state variables [12]. The discrete variables usually denote the system's operating mode and define how the continuous state evolves. In this paper, we focus on an important subclass of hybrid systems [13], [14] known as multiple model (MM) systems, where the discrete mode is used to index a family of different dynamical models. Hence, the state evolution and output models of an MM system are allowed to change over time. An MM system can be described as

$$x_k = f_{m_k}(x_{k-1}, u_{k-1}, w_{k-1}), \quad (1)$$

$$y_k = h_{m_k}(x_k, v_k), k \in \mathbb{N}, \quad (2)$$

where $x_k \in \mathbb{R}^{n_x}$ is the continuous state vector; $m_k \in \mathbb{M} \triangleq \{1, 2, \dots, M\}$ is the discrete mode of the system; $f_{m_k} : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \times \mathbb{R}^{n_w} \rightarrow \mathbb{R}^{n_x}$ is a possibly nonlinear mode-dependent process evolution function; $h_{m_k} : \mathbb{R}^{n_x} \times \mathbb{R}^{n_v} \rightarrow \mathbb{R}^{n_y}$ and $y_k \in \mathbb{R}^{n_y}$ are the mode-dependent measurement function and measurement vector, respectively; $u_{k-1} \in \mathbb{R}^{n_u}$ is the input vector; and $v_{k-1} \in \mathbb{R}^{n_v}$ and $w_{k-1} \in \mathbb{R}^{n_w}$ are independent white noise processes. As observed in (1)-(2), the discrete mode m_k defines a set of M different state-space evolution and

measurement functions f_{m_k} and h_{m_k} describing the system's dynamics. The parameter m_k is assumed to follow a Markov chain with an unknown initial probability vector $p(m_0)$ and an unknown transition probability matrix (TPM)

$$\Pi_k = \{\pi_{i,j}\}, \pi_{i,j} = \Pr\{m_k = j | m_{k-1} = i\}, i, j \in \mathbb{M}, \forall k \in \mathbb{N}. \quad (3)$$

Assuming that neither x_k nor m_k are directly measurable, the hybrid stochastic filtering problem aims to estimate the joint *a posteriori* probability density function

$$p(x_k, m_k | y_{1:k}) = p(x_k | m_k, y_{1:k}) \Pr(m_k | y_{1:k}) \quad (4)$$

of x_k and m_k , based on a sequence $y_{1:k} = \{y_1, y_2, \dots, y_k\}$ of noise-corrupted measurements generated according to (2). Considering the MM system described in (1)-(3) and given a sequence of output measurements $y_{1:k}$, plus initial conditions \hat{x}_0 , \hat{P}_0 , $\hat{p}(m_0)$, and $\hat{\Pi}_0$, one wishes to obtain:

- 1) \hat{x}_k , the estimated minimum variance state vector;
- 2) \hat{P}_k , the estimation error covariance matrix;
- 3) $\hat{p}(m_k)$, the estimated mode probability vector;
- 4) $\hat{\Pi}_k$, the estimated TPM.

Because mode transitions are potentially unobservable, there will be M^k possible mode sequences (hypotheses) at the k -th time instant if all elements in the TPM are nonzero, even if m_0 is known. Let \mathcal{I}_k be the set of all possible mode sequences for m_k . Each particular sequence $I_k^{(i)} \in \mathcal{I}_k$ is the i -th hypothesis at the k -th time instant. Because of the exponential growth in the number of hypotheses, it is impossible to implement an optimal estimator for (1)-(3) in practice because the limited memory and computation [11]. Therefore, suboptimal approaches, which merge similar hypotheses or exclude unlikely ones, must be adopted in practical MM estimators. The IP-MHMF algorithm presented in the next section leverages additional information about state predictions in order to perform a more informed hypotheses pruning step, therefore retaining the low computational cost of previous approaches while making the filter empirically more robust to divergence.

III. INFORMED PRUNING MHMF

This section contains the main contribution of this work, namely the IP-MHMF. This hybrid stochastic filter generalizes the well-known IMM algorithm and has improved stability when compared to the M³H and the MHMF while exhibiting relatively low computational cost.

Similar to the MHMF, the IP-MHMF improves its performance with respect to the IMM by allowing hypothesis-merging depths $d \geq 1$. Choosing $d > 1$ tends to improve the quality of the estimates, but has the undesirable effect of increasing the number of considered hypotheses at any given point in time. Therefore, as in the M³H and the MHMF, the IP-MHMF includes a hypothesis-pruning step that eliminates hypotheses with probabilities below a user-defined threshold α . Pruning hypotheses is beneficial because it reduces the number of necessary calculations and prevents the filter estimates from being degraded during the hypothesis-merging step [15]. However, the elimination of hypotheses may cause

stability and estimate degradation issues, as was empirically observed during filtering tests using the MHMF and the M³H with faulty sensor data.

In both the M³H and the MHMF, the pruning step is executed immediately before the filter-dependent prediction step, in order to minimize the number of EKFs that are used to track different hypotheses. However, problems may occur when the system changes to a mode with a low transition probability, such as the fault mode for sensors that work properly during most of the time. In this case, all hypotheses that consider sensor faults might be pruned, because the probability prediction step that uses the TPM may reduce their probability below the pruning threshold. Ignoring the possibility of sensor faults leads to the incorporation of strongly corrupted sensor data, causing degradation of filter estimates and potentially leading to filter instability. One simple method to circumvent the erroneous hypothesis-pruning problem is to increase the pruning threshold. However, this method eventually leads to the consideration of many hypotheses, which increases the computational load. Hence, the solution proposed in the IP-MHMF is to not rely on a finely-tuned pruning threshold α to ensure stability, while keeping the computational load to a minimum. Instead, the IP-MHMF performs an informed pruning step that leverages additional information from state predictions, in order to correct the probability estimates for each one of its tracked hypotheses. Unlike the M³H and the MHMF, which prune hypotheses based on their predicted probabilities, the IP-MHMF uses α as a pruning threshold on corrected hypothesis probability estimates, therefore greatly reducing the possibility of the true system hypothesis being eliminated by the pruning step, and hence preventing divergence. For Kalman filters and its variants, changes in the number of hypotheses during the prediction step do not significantly affect the overall computational load, given that the most expensive operation is the correction step involving matrix inversions [16].

Furthermore, estimators for Markovian switching systems typically assume previous knowledge of the TPM Π , which is rarely the case [17]. Hence, the IP-MHMF explicitly includes an on-line TPM estimation step based on system outputs.

A. Algorithm

The IP-MHMF estimates the hybrid state of (1)-(2) by tracking multiple hypotheses between consecutive merging steps. Based on its current estimate of the TPM computed by the *Quasi-Bayesian* TPM estimation algorithm in [17], it is able to predict the probability of each hypothesis (Step *i*). These are then used to merge the estimates in an IMM-like fashion whenever the merging depth is attained (Step *ii*). The state estimates for each hypothesis are propagated forward in time by single-model filters in Step *iii*, which are subsequently used in Step *iv* to correct the hypotheses' probabilities. Step *v* then prunes unlikely hypotheses, i.e., the ones whose corrected probabilities fall below a user-defined threshold. Since the IP-MHMF uses corrected probabilities in order to perform its pruning in Step *v*, it is much more effective than the MHMF and the M³H at determining

the subset of its tracked hypotheses that are likely given the sensor output data, therefore avoiding divergence. The estimates of the remaining hypotheses are corrected with sensor data in Step *vi*, followed by the generation of the IP-MHMF outputs in Step *vii*. Finally, TPM estimate is updated using the *Quasi-Bayesian* algorithm in Step *viii*.

IP-MHMF Let $\hat{x}_k^{(i)}$ and $\hat{P}_k^{(i)}$, $i \in \{1, 2, \dots, M^{q_k}\}$, $q_k \in \{1, 2, \dots, d\}$, be the state vector and associated covariance matrix, respectively, that correspond to the filter that tracks the system hypothesis $I_k^{(i)} \in \mathcal{I}_k$, where \mathcal{I}_k is the set of all possible hypotheses at the k -th sample instant and d is the hypothesis-merging depth. Let $0 \leq \alpha < 1$ be the hypothesis-pruning threshold. Furthermore, denote $n(\mathcal{I}_k)$ as the total number of elements in \mathcal{I}_k and $m_k^{(i)} \in \mathbb{M}$ as the current mode for hypothesis $I_k^{(i)}$. Define $\hat{p}^{(i)}(\mathcal{I}_k) = \Pr(I_k^{(i)} | y_{1:k})$, $\hat{p}^{(i)}(m_k) = \Pr(m_k = i | y_{1:k})$, and assume the following initial conditions:

$$\hat{p}^{(i)}(\mathcal{I}_0), \hat{x}_0^{(i)}, \hat{P}_0^{(i)}, \hat{\Pi}_0, q_0=1, i \in \{1, 2, \dots, M\},$$

the hybrid data fusion algorithm is expressed as follows:

i Hypotheses probability prediction

$$\bar{p}^{(i)}(\mathcal{I}_k) = \hat{\pi}_{a,b}(k-1) \hat{p}^{(j)}(\mathcal{I}_{k-1}), \quad a = m_{k-1}^{(j)}, b = m_k^{(i)},$$

where $i \in \{1, \dots, n(\mathcal{I}_k)\}$, $j \in \{1, \dots, n(\mathcal{I}_{k-1})\}$.

ii Initial conditions for the current step

- If $n(I_k) \leq M$ or $q_k < d$

$$\hat{x}_k^{(i)} = \hat{x}_{k-1}^{(j)}, \quad \hat{P}_k^{(i)} = \hat{P}_{k-1}^{(j)}, \quad q_k = q_{k-1} + 1,$$

which implies that $I_k^{(i)}$ was obtained from $I_{k-1}^{(j)}$.

- Else

$$\bar{p}^{(i)}(m_k) = \sum_{j=1}^{n(\mathcal{I}_k)} \Pr(m_k = i | I_k^{(j)}, y_{1:k-1}) \bar{p}^{(j)}(\mathcal{I}_k),$$

$$\hat{x}_{k-1}^{(i)} = \sum_{j=1}^M \frac{\hat{\pi}_{j,i}(k-1) \hat{p}^{(j)}(m_{k-1}) r_{k-1}^{(j)}}{\bar{p}^{(i)}(m_k)},$$

$$r_{k-1}^{(j)} = \sum_{l=1}^{n(\mathcal{I}_{k-1})} \hat{x}_{k-1}^{(l)} \hat{\Pr}(I_{k-1}^{(l)} | m_{k-1} = j, y_{1:k-1}),$$

$$\hat{P}_{k-1}^{(i)} = \sum_{j=1}^M \frac{\hat{\pi}_{j,i}(k-1) \hat{p}^{(j)}(m_{k-1}) [\Delta_{k-1}^{(j)} + \delta(i, j)]}{\bar{p}^{(i)}(m_k)},$$

$$\Delta_{k-1}^{(j)} = \sum_{l=1}^{n(\mathcal{I}_{k-1})} \hat{P}_{k-1}^{(l)} \hat{\Pr}(I_{k-1}^{(l)} | m_{k-1} = j, y_{1:k-1}),$$

$$\delta(i, j) = \left(r_{k-1}^{(j)} - \hat{x}_{k-1}^{(i)} \right) (\cdot)^T, \quad q_k = 1,$$

where

$$\hat{\Pr}(I_{k-1}^{(l)} | m_{k-1} = j, y_{1:k-1}) = \frac{\Pr(m_{k-1} = j | I_{k-1}^{(l)}, y_{1:k-1}) \Pr(I_{k-1}^{(l)} | y_{1:k-1})}{\Pr(m_{k-1} = j | y_{1:k-1})}. \quad (5)$$

iii (Filter-dependent) prediction step

$$\left(\hat{x}_{k-1}^{(i)}, \hat{P}_{k-1}^{(i)} \right) \xrightarrow{\text{Prediction}} \left(\bar{x}_k^{(i)}, \bar{P}_k^{(i)} \right). \quad (6)$$

iv Hypotheses probability correction

$$\hat{p}^{(i)}(\mathcal{I}_k) = \frac{p(y_k | I_k^{(i)}, \hat{\Pi}_{k-1}, y_{1:k-1}) \bar{p}^{(i)}(\mathcal{I}_k)}{c_i},$$

$$\gamma_p = \sum_{j=1}^{n(\mathcal{I}_k)} \hat{p}^{(j)}(\mathcal{I}_k),$$

$$\hat{p}(\mathcal{J}_k) = [\hat{p}^{(1)}(\mathcal{J}_k) \dots \hat{p}^{(n(\mathcal{J}_k))}(\mathcal{J}_k)]^T (1/\gamma_p),$$

where c_i is a normalizing constant.

Hypotheses pruning

Eliminate hypotheses $I_k^{(i)}$ with $\frac{\hat{p}^{(i)}(\mathcal{J}_k)}{\sum_{j=1}^{n(\mathcal{J}_k)} \hat{p}^{(j)}(\mathcal{J}_k)} \leq \alpha$,

normalize all $\hat{p}^{(i)}(\mathcal{J}_k)$, and update $n(\mathcal{J}_k)$ accordingly.

(Filter-dependent) correction step

$$(\bar{x}_k^{(i)}, \bar{P}_k^{(i)}) \xrightarrow{\text{Correction}} (\hat{x}_k^{(i)}, \hat{P}_k^{(i)}). \quad (7)$$

Estimate generation

$$\begin{aligned} \hat{x}_k &= \sum_{i=1}^{n(\mathcal{J}_k)} \hat{p}^{(i)}(\mathcal{J}_k) \hat{x}_k^{(i)}, \\ \hat{P}_k &= \sum_{i=1}^{n(\mathcal{J}_k)} \hat{p}^{(i)}(\mathcal{J}_k) \left[\hat{P}_k^{(i)} + (\hat{x}_k^{(i)} - \hat{x}_k)(\cdot)^T \right], \\ \hat{p}^{(i)}(m_k) &= \sum_{j=1}^{n(\mathcal{J}_k)} P(m_k=i|I_k^{(j)}, y_{1:k}) \hat{p}^{(j)}(\mathcal{J}_k), \\ \hat{p}(m_k) &= \left[\hat{p}^{(1)}(m_k) \dots \hat{p}^{(M)}(m_k) \right]^T. \end{aligned}$$

TPM update: $\hat{\Pi}_{k-1} \xrightarrow{\text{Algorithm [17]}} \hat{\Pi}_k$.

The term $I_{k-1}^{(l)}$ is one particular hypothesis that is tracked between two merging steps. In (5), $\hat{\Pr}(m_{k-1}=j|I_{k-1}^{(l)}, y_{1:k-1})$ is either 0 or 1 depending on which mode corresponds to $I_{k-1}^{(l)}$. The IMM's mixing step is a particular case of Step *ii* with $d=1$. No details are provided in (6) and (7) because these steps vary depending on the filter chosen to track each hypothesis. For example, if (1)-(2) are linear, the KF is a reasonable choice. Since attitude estimation is nonlinear, the results in Section V were obtained using the EKF.

IV. ATTITUDE ESTIMATION

Similar to [5], this section leverages hybrid modeling to address the problem of embedded sensor faults. This modeling enables the IP-MHMF to track the UAV's attitude in conditions where the EKF fails. Because there are no sensor measurements of the aircraft movement in the XY plane, translation is not discussed here. For the following equations, consider the coordinate frames in Fig. 1.

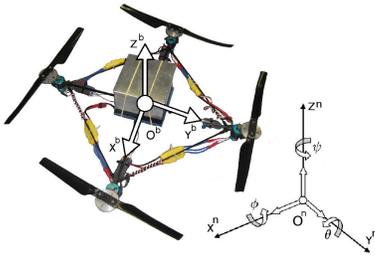


Fig. 1. Body (b) and reference (n) coordinate frames. The rotation angles between the two frames are called roll (ϕ), pitch (θ), and yaw (ψ).

A. Prediction

Let $q_n^b = [q_0 \ q_1 \ q_2 \ q_3]^T$, $\|q_n^b\|=1$, be the quaternion that represents the orientation of b with respect to n . The equation

that relates q_n^b to its corresponding rotation matrix is

$$C_n^b = \begin{bmatrix} q_0^2 + q_1^2 - q_2^2 - q_3^2 & 2(q_1 q_2 - q_0 q_3) & 2(q_1 q_3 + q_0 q_2) \\ 2(q_1 q_2 + q_0 q_3) & q_0^2 - q_1^2 + q_2^2 - q_3^2 & 2(q_2 q_3 - q_0 q_1) \\ 2(q_1 q_3 - q_0 q_2) & 2(q_2 q_3 + q_0 q_1) & q_0^2 - q_1^2 - q_2^2 + q_3^2 \end{bmatrix}. \quad (8)$$

During rotation, gyros measure angular rates ω_x , ω_y , and ω_z around the axes X^b , Y^b , and Z^b , respectively (Fig. 1). The discrete-time quaternion evolution model is given by [18]

$$\begin{aligned} q_{n,k}^b &= \left[\mathbb{I}_{4 \times 4} \cos\left(\frac{\delta}{2}\right) - W \tau \frac{\sin\left(\frac{\delta}{2}\right)}{\delta} \right] q_{n,k-1}^b, \\ W &= \begin{bmatrix} 0 & \omega_x & \omega_y & \omega_z \\ -\omega_x & 0 & -\omega_z & \omega_y \\ -\omega_y & \omega_z & 0 & -\omega_x \\ -\omega_z & -\omega_y & \omega_x & 0 \end{bmatrix}, \delta = \left(\sqrt{\omega_x^2 + \omega_y^2 + \omega_z^2} \right) \tau, \end{aligned} \quad (9)$$

where τ denotes the sampling period and the subscript $k \in \mathbb{N}$ denotes the sample taken at instant $k\tau$.

B. Correction

The accelerometer and the magnetometer are used to correct attitude estimates. Their readings relate to the true attitude of the system according to the model

$$f_k^b = \left(C_{n,k}^b \right)^T (a_k^n - g_E^n) + \varepsilon_{f,k}, \varepsilon_{f,k} \sim N(0, R_{\varepsilon_{f,k}}), \quad (10)$$

$$m_{mag,k}^b = \left(C_{n,k}^b \right)^T m_E^n + \varepsilon_{m,k}, \varepsilon_{m,k} \sim N(0, R_{\varepsilon_{m,k}}), \quad (11)$$

where f_k^b is the specific force measurement of the accelerometer in b ; $m_{mag,k}^b$ is the magnetometer reading; $C_{n,k}^b$ is as in (8); a_k^n is the body acceleration in n ; g_E^n and m_E^n are the local gravitational and magnetic fields, respectively; and $\varepsilon_{i,k}$ and $i \in \{f, m\}$, model sensor noise. For attitude correction, it is important that a_k^n be minimized in (10) so that it reduces to

$$f_k^b \approx \left(C_{n,k}^b \right)^T (-g_E^n) + \varepsilon_{f,k}, \quad (12)$$

creating a correspondence between attitude and the gravity vector g_E^n in the reference frame b . The approximation (12) is reasonable for non-acrobatic flights.

C. Modeling magnetometer faults

The magnetometer is important in the quadrotor attitude estimation, because it is the only sensor responsible for correcting the yaw angle estimates of the aircraft. The accelerometer measurements in (12) remain unchanged when the UAV rotates along its vertical axis because the local gravitational field g_E^n and Z^n are collinear. Therefore, it is important to mitigate all sources of disturbances that act over this sensor, because reliable yaw estimates are necessary to control the UAV's heading during flight. Following the hybrid-system approach in [5], the attitude estimation instruments of the aircraft are modeled as an MM system with two distinct modes. The first mode corresponds to the equation

$$\frac{m_{mag,k}^b}{\|m_{mag,k}^b\|} = \left(C_{n,k}^b \right)^T \frac{m_E^n}{\|m_E^n\|} + \varepsilon_{m,k}^{norm}, \varepsilon_{m,k}^{norm} \sim N(0, R_{\varepsilon_{m,k}^{norm}}), \quad (13)$$

and models normalized magnetometer measurements during nominal operation, that is, when the sensor readings are analogous to the body's orientation in space. The second mode seeks to mathematically represent the spurious readings and local magnetic field distortions observed during the

calibration tests. Because these disturbances do not seem to have a clear correlation with the vehicle’s state, the magnetometer faults are modeled as

$$\frac{m_{mag,k}^b}{\|m_{mag,k}^b\|} = 0.5 + \varepsilon_{m,k}^{fault}, \varepsilon_{m,k} \sim N\left(0, R_{\varepsilon_{m,k}}^{(2)}\right), \quad (14)$$

where the 0.5 value represents half of the range of the normalized readings. The noise covariance matrices for different magnetometer operating modes are $R_{\varepsilon_{m,k}}^{(1)} = \mathbb{I} \times 4.4e-3$ and $R_{\varepsilon_{m,k}}^{(2)} = \mathbb{I} \times 0.1$.

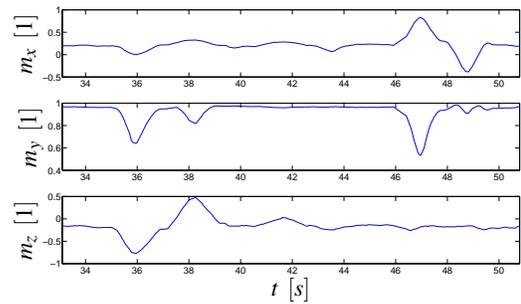
V. EXPERIMENTAL RESULTS

This section compares the performance of four nonlinear filters that were used to estimate the attitude of a quadrotor based on real inertial and magnetic sensor measurements. The classical solution using a single EKF was compared against the MHMF, the M³H filter, and the IP-MHMF, all benefiting from the hybrid modeling presented in Section IV-C to handle magnetometer faults. All MM filters used a bank of EKFs to track each hypothesis of the system.

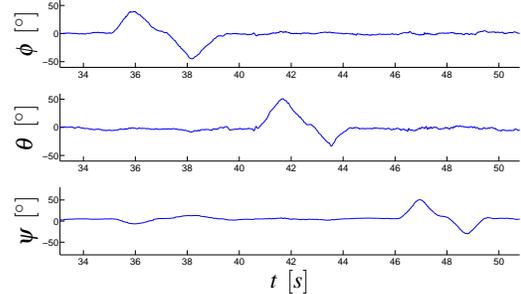
Figure 2 shows the results of the first experiment performed to verify whether the EKF and the MM filters were capable of performing accurate attitude estimation according to the model in Section IV in the absence of sensor faults. The experiment consisted of three manual rotations of the quadrotor around each of its coordinated axes by approximately 45° in both directions. Figure 2(a) shows the actual magnetometer measurements collected during the experiment, which were used to generate the attitude estimates in Fig. 2(b). The latter only shows the results for the EKF, but the curves were virtually identical for all other filters. This experiment served to demonstrate that the attitude estimation algorithms were correctly implemented and served as a basis of comparison for the next experiment.

The second experiment consisted of introducing typical magnetometer faults in the magnetometer data so as to simulate a “worst-case” scenario for the attitude estimation system, as shown in Figure 3(a). There are basically two types of typical disturbances: either the magnetometer returns a short sequence of spurious readings, often with one component of its measurement vector set to 0, or it becomes biased in a fixed direction due to the presence of a strong electromagnetic field. Figure 3(b) shows the performance of the traditional attitude estimation approach using a single EKF. One can observe in the plots the presence of minor disturbances in the roll and pitch estimates, since the accelerometer helps correct them. The yaw estimates, on the other hand, became strongly degraded when compared to Fig. 2(b), which is due to the fact that they depend solely on magnetometer readings for their correction step.

This strong performance degradation is not found on the attitude estimates in Fig. 3(c), which shows the best performance for the IP-MHMF (all other MM filters yielded almost identical plots at their best performances). The roll and pitch estimates remained virtually unchanged when compared to the undisturbed case and the yaw estimates exhibited only minor perturbations compared to Fig. 2(b). Despite the



(a) Normalized magnetometer measurements.



(b) Attitude estimates for the EKF (virtually identical results for all other filters).

Fig. 2. Attitude estimates using undisturbed magnetometer measurements.

similarity in the quality of their estimates, the true distinction between the different algorithms can be seen in the results shown in Table I. It shows the computational load results for the best performance of all filters using both normal and faulty measurements, where the computational complexity for MM filters is linear in the number of tracked hypotheses. The MHMF exhibited strong performance degradation even for small values of α , due to incorrect pruning of hypotheses, and was outperformed by the IP-MHMF in all test instances. Therefore, Table I only shows the comparison between the EKF, the M³H, and the IP-MHMF.

The EKF was the lightest choice in all cases, but Fig. 3(b) clearly illustrates that its estimates are not robust to the sensor disturbances and may jeopardize the stability and safe operation of the quadrotor. Regarding the MM filters, the M³H filter performed slightly better than the IP-MHMF when operating with the measurements in Fig. 2(a). However, an important aspect should be considered when comparing these two filters: the M³H filter considers the TPM as given and requires it to be properly tuned to operate correctly. The IP-MHMF, on the other hand, assumes no initial information regarding the mode transition probabilities, and estimates $\hat{\Pi}_k$ in real time using the output sensor measurements, forcing it to conserve a larger number of hypotheses during the initial sample instants. Nevertheless, during the approximately 18 s in Fig. 3, the filters only differ during the first 60 ms. Then, both filters conserve a single hypothesis that corresponds to the correct magnetometer operation.

The second half of Table I, which corresponds to the measurements in Fig. 3(a), presents a different perspective for the comparison of the MM filters. Because the more informed hypothesis-pruning step in the IP-MHMF uses corrected

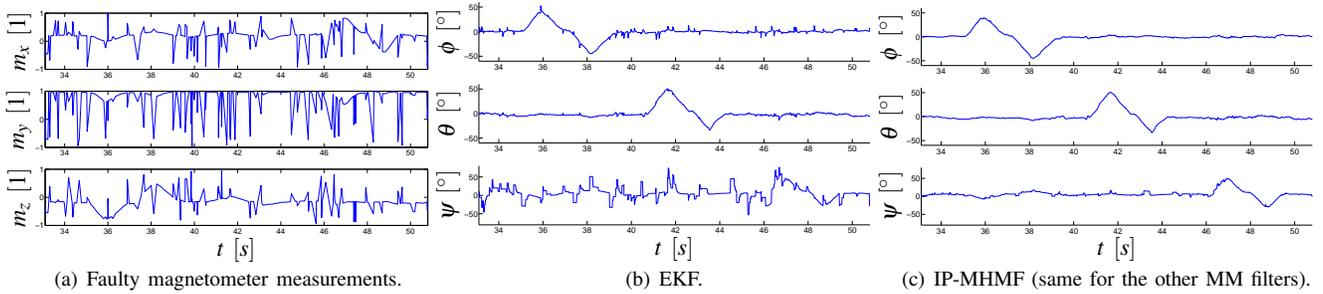


Fig. 3. Attitude estimates using faulty magnetometer measurements.

TABLE I

COMPUTATIONAL COST FOR THE BEST FILTERING PERFORMANCE.

	EKF	M ³ H	IP-MHMF
Undisturbed measurements			
Mean number of hyp.	1	1.004	1.007
Max number of hyp.	1	2	4
Pruning threshold	—	30%	50%
Faulty measurements			
Mean number of hyp.	1	2.297	1.007
Max number of hyp.	1	3	4
Pruning threshold	—	3%	50%

probability estimates rather than predicted ones, one can maintain much higher pruning thresholds, which drastically reduce the computational complexity, while not affecting the filter's performance. For the M³H, however, raising the pruning threshold above 3% caused incorrect elimination of hypotheses and performance degradation. Once again, the two filters have different maximum numbers of hypotheses because the IP-MHMF conserves more hypotheses at the beginning due to the lack of information on the TPM, but the computational complexity of 4 hypotheses is no longer observed after the first 70 ms.

VI. CONCLUSIONS

This paper presented the IP-MHMF, a novel filter for hybrid multiple model systems. It generalizes the IMM's hypothesis-merging step to depths greater than one and introduces a more informed hypothesis-pruning step than previous algorithms in the literature. By leveraging state prediction information from its bank of single-model filters, the IP-MHMF performs hypotheses pruning on corrected rather than predicted probability estimates, therefore improving its performance. It also allows a much more aggressive hypothesis-pruning strategy that significantly reduces its computational load, therefore making it suitable for embedded applications with processing and energy constraints. Its usefulness was demonstrated on an attitude estimation application for a quadrotor UAV with faulty sensors.

ACKNOWLEDGMENTS

The authors would like to acknowledge the financial support of CNPq, CAPES, FAP-DF, and the Dean of Research and Graduate Studies (DPP-UnB). This research was partially funded by AFOSR grant 6926079.

REFERENCES

- [1] D. Titterton and J. Weston, *Strapdown Inertial Navigation Technology*. London: Peter Peregrinus Ltd., 1997.
- [2] P. Allen and R. Bajcsy, "Two sensors are better than one: example of vision and touch," *Proceedings of 3rd International Symposium on Robotics Research*, pp. 48–55, 1986.
- [3] M. W. Hofbauer and B. C. Williams, "Hybrid estimation of complex systems," *Systems, Man, and Cybernetics, Part B: Cybernetics, IEEE Transactions on*, vol. 34, no. 5, pp. 2178–2191, 2004.
- [4] A. S. Willsky, "A survey of design methods for failure detection in dynamic systems," *Automatica*, vol. 12, no. 6, pp. 601 – 611, 1976.
- [5] P. Santana, G. Borges, and J. Ishihara, "Hybrid data fusion for 3D localization under heavy disturbances," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, 10–15 2010, pp. 55 –60.
- [6] H. Blom, "An efficient filter for abruptly changing systems," *Proceedings of 23rd Conference on Decision and Control*, vol. 4, no. 30, pp. 656–658, December 1984.
- [7] H. Blom and Y. Bar-Shalom, "The interacting multiple model algorithm for systems with Markovian switching coefficients," *IEEE Transactions on Automatic Control*, vol. 33, no. 8, pp. 780–783, August 1988.
- [8] E. Mazar, A. Averbuch, Y. Bar-Shalom, and J. Dayan, "Interacting multiple model methods in target tracking: a survey," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 34, no. 1, pp. 103–123, January 1998.
- [9] Y. Boers and H. Driessen, "A multiple model multiple hypothesis filter for Markovian switching systems," *Automatica*, vol. 41, no. 4, pp. 709 – 716, 2005.
- [10] P. Santana, H. Menegaz, G. Borges, and J. Ishihara, "Multiple hypotheses mixing filter for hybrid Markovian switching systems," in *Proceedings of the 49th IEEE Conference on Decision and Control*, 9–11 2010, pp. 3457–3462.
- [11] G. Ackerson and K. Fu, "On state estimation in switching environments," *IEEE Transactions on Automatic Control*, vol. 15, no. 1, pp. 10–17, February 1970.
- [12] R. Goebel, R. Sanfelice, and A. Teel, "Hybrid dynamical systems," *IEEE Control Systems Magazine*, vol. 29, no. 2, pp. 28–93, April 2009.
- [13] M. J. McCourt and P. J. Antsaklis, "Stability of interconnected switched systems using qsr dissipativity with multiple supply rates," in *American Control Conference (ACC)*, Montreal, QC, June 2012, pp. 4564 – 4569.
- [14] J. Wang and T. Chen, "Online identification of switched linear output error models," in *IEEE International Symposium on Computer-Aided Control System Design (CACSD)*, Denver, CO, Sept. 2011, pp. 1379 – 1384.
- [15] X.-R. Li and Y. Bar-Shalom, "Multiple-model estimation with variable structure," *IEEE Transactions on Automatic Control*, vol. 41, no. 4, pp. 478–493, April 1996.
- [16] C. Chui and G. Chen, *Kalman Filtering with Real-Time Applications*, 3rd ed. Springer-Verlag New York, 1987.
- [17] V. Jilkov and X. Li, "Online Bayesian estimation of transition probabilities for Markovian jump systems," *IEEE Transactions on Signal Processing*, vol. 52, no. 6, pp. 1620–1630, June 2004.
- [18] R. M. Rogers, *Applied mathematics in integrated navigation systems*. AIAA Education Series, 2003.