Moving Horizon Estimation with a Huber Penalty Function for Robust Pose Estimation of Tethered Airplanes

K. Geebelen, A. Wagner, S. Gros, J. Swevers and M. Diehl

***Abstract*— This paper presents a Moving Horizon Estimator (MHE) for estimating the position and orientation (pose) of moving objects, in particular for tracking tethered airplanes for Airborne Wind Energy systems. In this application, absolute pose measurements are captured by a marker based stereo vision system. These measurements are fused with measurements of the angular velocity and linear acceleration from an Inertial Measurement Unit (IMU). In our MHE, the IMU measurements in the intervals between camera frames are modeled as samples of a superposition of orthonormal polynomial basis functions. This results in a MHE formulation that requires fewer optimization variables, enabling faster solution. In order to achieve robustness to marker detection errors, a formulation based on the Huber penalty is also presented. We show that our robust MHE formulation outperforms a MHE formulation using the ℓ_2-norm and a traditional extended Kalman filter.**

***Keywords* : sensor fusion, moving horizon estimation, pose estimation, IMU measurements**

I. INTRODUCTION

In the search for sustainable alternatives to today's nuclear and fossil fuel based energy production, it has become clear that wind power will play a major role in meeting future energy demands. Since theoretically the amount of power that can be extracted is proportional to the wind velocity cubed [15], and wind velocity increases with altitude, there is a strong incentive to extract power from wind at higher altitudes than can be reached by traditional horizontal axis wind turbines.

In order to capture this high energy wind, a new class of energy systems collectively know as Airborne Wind Energy systems are under development at several start-up companies and universities. The technique we consider, first proposed in [15], uses tethered airplanes that fly under automatic control at high velocities in a crosswind fashion. The resulting flight control problem is gaining increasing interest from the control community, due to its difficulty and novelty [5], [6], [10], [21]. While most of these papers designed controllers for power production, the airplanes must also be able to launch and land automatically. One of the proposed launch techniques involves accelerating the plane to flight speeds while it is fixed at the end of a rotating arm, or *carousel*, and then reeling out the tether so that the plane can gain altitude. Early in this reel-out phase, the plane is continuously flying aerobatic maneuvers (i.e. a horizontal loop) close to the ground, making precise and robust control an absolute necessity.

The sensors that are available during the early reel-out phase include an IMU and a stereo vision system. The IMU provides measurements of linear acceleration and angular velocity. Since these measurements are measured in the coordinate frame of the plane (to be estimated), we describe these measurements as *relative measurements*. While the cameras provide *absolute measurements* that are easily related to the world coordinate frame, they are captured at a much lower frequency than the IMU and, alone, would not be fast enough to directly capture the dynamics of the system[1]. Furthermore, the vision system has the potential to be influenced by illumination and objects in the scene and may occasionally produce faulty measurements (or *outliers*).

Algorithms based on extended and unscented Kalman filters, particle filters and MHEs [14], [7], [17], [16], [19], [20], [8], [18] have been developed to perform sensor fusion to estimate position and orientation. In particular, [9] applied the multiplicative Extended Kalman Filter (EKF) to our system and will be used as the baseline for our simulations.

Contributions of the paper: A specific MHE is used to robustly fuse the IMU measurements with the camera measurements, providing a drift-free estimate of the position and orientation of the airplane at a high sampling frequency. In our MHE, the IMU measurements in the intervals between camera frames are modeled as samples of a superposition of orthonormal polynomial basis functions. This results in a MHE formulation that requires far fewer optimization variables, enabling faster solution. In order to achieve robustness to marker detection errors, a formulation incorporating the Huber penalty function is also presented. We show in simulations that our robust MHE formulation outperform an EKF in case of outliers in the camera measurements.

The paper is organized as follows: Section II describes the estimation problem, Section III shows how to discretize the problem, Section IV introduces least-squares and robust MHE formulations, Sections V and VI show simulation and experimental results respectively, and Section VII concludes the paper.

K. Geebelen and Prof. J. Swevers are with the Faculty of Mechanical Engineering, division PMA KU Leuven, Celestijnenlaan 300b - box 2420, B-3001 Leuven-Heverlee, Belgium. kurt.geebelen@mech.kuleuven.be

K. Geebelen, A. Wagner, S. Gros and Prof. M. Diehl are with the Faculty of Electrical Engineering, division SCD-Sista, KU Leuven, Kasteelpark Arenberg 10 - box 2446 , B-3001 Leuven-Heverlee, Belgium.

[1]At higher altitudes it is expected that an on-board Global Positioning System will be relied on for absolute position measurements, but it's accuracy and speed are insufficient during launch.

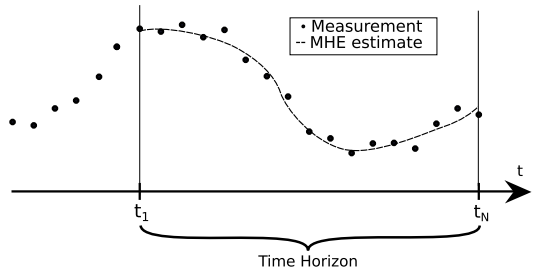


Fig. 1: The Moving Horizon Estimation approach minimizes the measurement and model noise over a time horizon that contains the most recent measurements.

II. ESTIMATION PROBLEM

The Kalman filter is the most commonly used tool for state estimation and sensor fusion, and it is optimal in an ℓ_2 sense for linear systems with Gaussian noise. In case of non-linear dynamics, an extended or unscented Kalman filter can be used, but there are no optimality results for most nonlinear systems. Another approach comes in the form of a MHE. A MHE estimates the system state by solving an optimal control problem online. In a typical MHE, the following optimization problem is solved at each sampling time:

$$\begin{aligned} \underset{x(\cdot),w(\cdot),v(\cdot)}{\text{minimize}} \quad & \int_{t=0}^{T} \| w(t) \|^2_{Q^{-1}} + \| y(t) - h(x,t) \|^2_{R^{-1}} \, dt \\ \text{subject to} \quad & \dot{x} = f(x,w,t), \qquad t \in [t_1, t_N], \end{aligned} \tag{1}$$

where w is a noise term introduced to model the model uncertainty, x is the system state, y represents the measurements, Q and R are the covariance matrices of the model and measurement noise respectively, T is the considered time horizon, f is a model of the considered system, and h is the measurement function. The problem thus minimizes a weighted sum of the noise over a moving time horizon, which is illustrated in Fig. 1. When a new measurement arrives, the horizon is shifted one measurement and the new problem is formulated and solved again.

The nature of our sensors introduces a bit more complexity into the formulation: since the IMU measurements are gathered at a much higher sampling rate than the camera measurements, we use a discrete formulation for the camera measurements and a continuous time formulation for the IMU measurements. This is illustrated in Fig. 2. Also, the statistics of the absolute measurements are far from Gaussian as there might be outliers due to false marker detections, so we might use non-quadratic penalties. With these changes, the optimization problem to be solved becomes

$$\begin{aligned} \underset{x(\cdot),w(\cdot),b}{\text{minimize}} \quad & \sum_{k=1}^{N_{\text{abs}}} \Psi(y_{\text{abs},k} - h_{\text{abs}}(x(t_k)) \\ & + \sum_{k=1}^{N_{\text{abs}}-1} \int_{t_k}^{t_{k+1}} \| w(t) - y_{\text{rel}}(t) \|^2_{Q^{-1}} \, dt \\ & + \| b - b_{\text{prev}} \|^2_B \\ \text{subject to} \quad & \dot{x}(t) = f(x(t), w(t), b), \qquad t \in \{t_1, \cdots, t_{N_{\text{abs}}}\} \\ & g(x(t_1)) = 0, \end{aligned} \tag{2}$$

where, in the objective function, Ψ is a penalty function, $x = \begin{bmatrix} p, & v, & q \end{bmatrix}^T$ is a state vector consisting of the position, the velocity and the quaternion representing the orientation of the airplane, $w = \begin{bmatrix} a, & \omega \end{bmatrix}^T$ are the estimates for acceleration and angular velocity, $y_{\text{rel}} = \begin{bmatrix} a_{\text{meas}}, & \omega_{\text{meas}} \end{bmatrix}^T$ are the relative measurements (i.e. coming from the IMU), y_{abs} are the absolute measurements (i.e. the locations of the markers in the images from the stereo vision system), N_{abs} is the number of absolute measurements in the horizon, Q is the covariance matrix of the noise on the relative measurements, $b = \begin{bmatrix} b_a & b_\omega \end{bmatrix}^T$ is the estimate for the bias on the acceleration and angular velocity measurements, b_{prev} is a previous estimate of the bias, and B is a weighting matrix. The bias b is assumed constant over the horizon, since the time scale at which it varies is much bigger than the length of the horizon.

The first term in Problem (2) is thus the mismatch between the predicted and actual absolute measurements, the second term is the mismatch between the predicted and actual relative measurements, and the last term slows the evolution of the bias estimation. Strategies for updating b_{prev} and B are discussed in e.g. [12] but are not the subject of this paper. The absolute measurement function $h_{\text{abs}}(x)$ is a pinhole camera model that maps the state x to the position of the markers in the images. Note that the image position of each marker depends on both the position and orientation of the airplane. In the constraints, $f(x(t), w(t), b)$ is a kinematic model and $g(x) = \| q \|_2^2 - 1$ enforces the quaternion norm constraint in the beginning of the horizon. The kinematic model ensures that q stays of unity norm throughout the horizon. The kinematic model $\dot{x} = f(x,w,b)$ is given by

$$f(x,w,b) = \begin{bmatrix} \dot{p} \\ \dot{v} \\ \dot{q} \end{bmatrix} = \begin{bmatrix} v \\ \text{Rot}(q)(a - b_a) - \begin{bmatrix} 0 & 0 & 1 \end{bmatrix}^T G \\ \frac{1}{2}\Omega(\omega - b_\omega) q \end{bmatrix}, \tag{3}$$

where $\text{Rot}(q)$ is the rotation matrix corresponding to the quaternion q that describes the transformation from body frame coordinates to reference frame coordinates, G is the gravitational constant, and Ω is the quaternion matrix corresponding to the quaternion vector $\begin{bmatrix} \omega - b_\omega & 0 \end{bmatrix}^{\text{T}}$ [4].

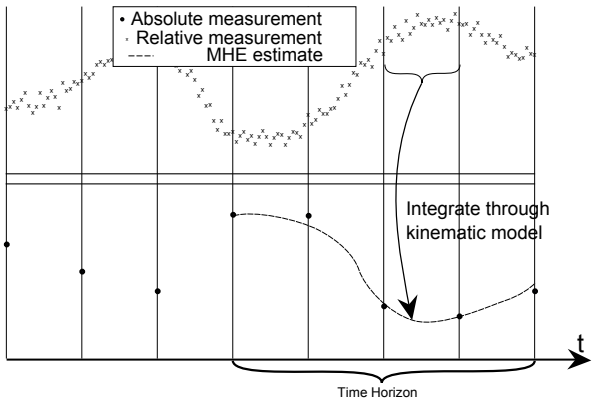


Fig. 2: Illustration of the measurements.

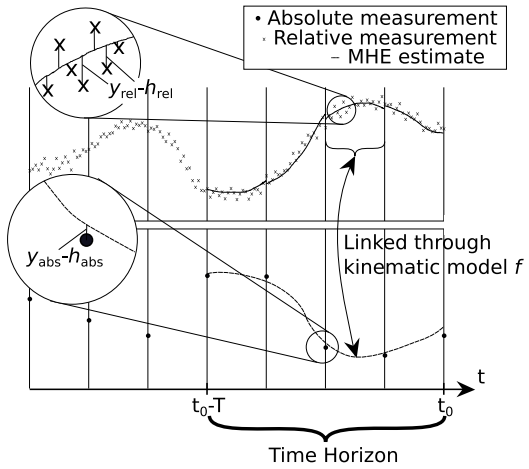


Fig. 3: The relative measurements are projected on a polynomial basis and then used in MHE. These polynomials are also affected by the absolute measurements through the kinematic model.

III. DISCRETIZATION OF CONTINUOUS PROBLEM

Problem (2) is still an infinite optimal control problem. One way to deal with this would be to discretize the unknown input w with the same grid as the actual relative measurements y_{rel}. In our context this would result in a very large-scale optimzation problem, since the relative measurements are densely sampled.

Motivated by this, we have developed an approach in which the relative measurements are approximated by a polynomial basis of degree m. Fig. 3 illustrates this. The polynomial coefficients, denoted by u, are then the decision variables for the optimization problem. Each set of relative measurements between two absolute pose measurements is represented by a polynomial. In our case, we thus have a polynomial for a_x, a_y, a_z, ω_x, ω_y and ω_z for each frame in the horizon. A frame in the context of our MHE formulation is defined to be the window between two absolute measurements. The parametrization for the estimates for acceleration and angular velocity in time will thus become

$$\tilde{w}(t;u_k) = \begin{bmatrix} \tilde{a}_x(t) \\ \tilde{a}_y(t) \\ \tilde{a}_z(t) \\ \tilde{\omega}_x(t) \\ \tilde{\omega}_y(t) \\ \tilde{\omega}_z(t) \end{bmatrix} = \sum_{p=0}^{m} L_p(t-t_k) u_k^{(p)}, \tag{4}$$

where L_p is the polynomial of order p in the polynomial basis, and $u_k^{(p)} = \begin{bmatrix} u_1^{(p)} & u_2^{(p)} & u_3^{(p)} & u_4^{(p)} & u_5^{(p)} & u_6^{(p)} \end{bmatrix}^{\mathrm{T}}$ are the coefficients for the polynomial of order p for frame k. The indices $\{1,\ldots,6\}$ indicate the type of relative measurement (i.e. acceleration or angular velocity in x,y or z direction).

To discretize $f(x(t),w(t),b)$ we use a multiple shooting parametrization with multiple shooting nodes on each time where there is an absolute measurement. To go from $x(t_k)$ to $x(t_{k+1})$, the kinematic model (3) is numerically integrated using the polynomials for frame k as input. Note that the integrator can take much larger steps than the sampling rate of the IMU. The state, integrated between two multiple shooting nodes, is denoted by $\tilde{x}(t;x_k,u_k,b)$. The MHE problem thus becomes:

$$\begin{aligned} \underset{\substack{x_1,\ldots,x_{N_{\text{abs}}} \\ u_1,\ldots,u_{N_{\text{abs}}-1} \\ b}}{\text{minimize}} \quad & \textstyle\sum_{k=1}^{N_{\text{abs}}} \Psi(y_{\text{abs},k} - h_{\text{abs}}(x_k)) \\ & + \textstyle\sum_{k=1}^{N_{\text{abs}}-1} \int_{t_k}^{t_{k+1}} \| \tilde{w}(t-t_k,u_k) - y_{\text{rel}}(t) \|_{Q^{-1}}^2 \, dt \\ & + \| b - b_{\text{prev}} \|_B^2 \\ \text{subject to} \quad & x_{k+1} = \tilde{x}(t_{k+1};x_k,u_k,b),\ k \in \{1,\ldots,N_{\text{abs}}-1\} \\ & g(x_1) = 0, \end{aligned} \tag{5}$$

where $u_k = \begin{bmatrix} u_k^{(0)^{\mathrm{T}}} & u_k^{(1)^{\mathrm{T}}} & \ldots & u_k^{(m)^{\mathrm{T}}} \end{bmatrix}^{\mathrm{T}}$. In practice, due to the high but finite sampling rate of the IMU, we replace the integrals in the objective function by sums. The corresponding term in Problem (5) becomes

$$\begin{aligned} & \int_{t_k}^{t_{k+1}} \| \tilde{w}(t-t_k,u_k) - y_{\text{rel}}(t) \|_{Q^{-1}}^2 \, dt \\ \approx & \sum_{j=0}^{N_{\text{rel}}-1} \| \tilde{w}(t_k + j\triangle t, u_k) - y_{\text{rel}}(t_k + j\triangle t) \|_{\bar{Q}^{-1}}^2, \end{aligned} \tag{6}$$

where $\triangle t$ is the sampling period of the relative measurements and where N_{rel} is the number of relative measurements between two absolute measurements and $\bar{Q}$ the covariance of each single relative measurement as instructed by the IMU producer. $\bar{Q}$ is assumed diagonal and constant in time, meaning the noise of the relative measurements is uncorrelated in time with a constant variance, as is usually the case. It can thus be written as

$$\bar{Q} = \begin{bmatrix} \sigma_1^2 & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & \sigma_6^2 \end{bmatrix}. \tag{7}$$

We can define a function

$$h_{\text{rel},r}(u_{k,r},t_k)_r = \begin{bmatrix} \tilde{w}(t_k;u_{k,r})_r \\ \tilde{w}(t_k+\triangle t;u_{k,r})_r \\ \tilde{w}(t_k+2\triangle t;u_{k,r})_r \\ \vdots \\ \tilde{w}(t_k+(N_{\text{rel}}-1)\triangle t;u_{k,r})_r \end{bmatrix}. \tag{8}$$

The index $r \in \{1,\ldots,6\}$ refers to the type of the relative measurement that is represented (i.e. acceleration or angular velocity in one of the directions). The function $h_{\text{rel},r}$ thus evaluates the polynomial for a specific type of relative measurement at all time points there is a relative measurement. It is linear in the polynomial coefficients, and can be written as

$$h_{\text{rel},r}(u_{k,r},t_k)_r = \Gamma u_{k,r}, \tag{9}$$

where Γ is a matrix that corresponds to the polynomial basis, and $u_{k,r}$ refers to the polynomial coefficients for frame k

of the relative measurements denoted by r. Eq. (6) hence becomes

$$\sum_{r=1}^{6} \sigma_r^{-2} \parallel \Gamma u_{k,r} - y_{\text{rel},k,r}^{\text{vec}} \parallel_2^2, \tag{10}$$

where $y_{\text{rel},k,r}^{\text{vec}}$ is a vector containing all relative measurements in one frame for the specific type of relative measurements denoted by r. It can be shown by expansion of terms that Eq. (10) is equivalent to

$$\sum_{r=1}^{6} \sigma_r^{-2} \parallel \Gamma u_{k,r} - y_{\text{rel},k,r}^{\text{vec}} \parallel_2^2 = \sum_{r=1}^{6} \sigma_r^{-2} \parallel u_{k,r} - \tilde{u}_{k,r} \parallel_{\Gamma^{\text{T}}\Gamma}^2 + C, \tag{11}$$

where C is a constant term that has no influence on the optimization, and where

$$\tilde{u}_{k,r} = (\Gamma^{\text{T}}\Gamma)^{-1}\Gamma^{\text{T}} y_{\text{rel},k,r}^{\text{vec}}, \tag{12}$$

which is the solution of the unconstrained linear least squares problem. By using this formulation, the relative measurements can be discarded after having solved the unconstrained linear least squares problem by the matrix-vector multiplication (12), which saves memory and computation time of the objective function. This solution will also be used to initialize the MHE.

The polynomials that we choose for our MHE are the '*Gram polynomials*' [2]. They are orthonormal on an equidistant grid on the interval $[-1,1]$, i.e. they satisfy the property

$$\sum_{k=0}^{N_{\text{rel}}} p_i(x_k) p_j(x_k) = \begin{cases} 1, & i = j \\ 0, & i \neq j \end{cases}, \quad x_k = -1 + \frac{2k}{N_{\text{rel}}}, \tag{13}$$

and thus form a matrix Γ that is orthonormal, meaning that $\Gamma^{\text{T}}\Gamma = I$, the identity matrix. Fig. 4 shows the Gram polynomials up to order 4 for a grid containing 50 nodes. Eq. (12) now simplifies to

$$\tilde{u}_{k,r} = \Gamma^{\text{T}} y_{\text{rel},k,r}^{\text{vec}}, \tag{14}$$

and Eq. (11) becomes

$$\sum_{r=1}^{6} \sigma_r^{-2} \parallel u_{k,r} - \tilde{u}_{k,r} \parallel_2^2 . \tag{15}$$

By choosing an orthonormal polynomial basis, Problem (5) can now be formulated as

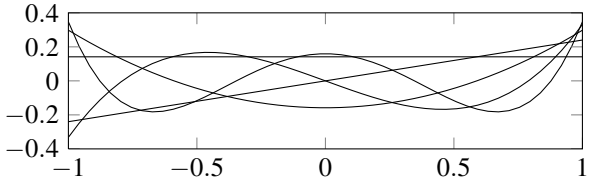


Fig. 4: The Gram polynomials up to order 4 for a grid of 50 nodes.

$$\begin{aligned} \underset{\substack{x_1,\dots,x_{N_{\text{abs}}} \\ u_1,\dots,u_{N_{\text{abs}}-1} \\ b}}{\text{minimize}} \quad & \textstyle\sum_{k=1}^{N_{\text{abs}}} \Psi(y_{\text{abs},k} - h_{\text{abs}}(x_k)) \\ & + \textstyle\sum_{k=1}^{N_{\text{abs}}-1} \sum_{r=1}^{6} \sigma_r^{-2} \parallel u_{k,r} - \tilde{u}_{k,r} \parallel_2^2 \\ & + \parallel b - b_{\text{prev}} \parallel_B^2 \\ \text{subject to} \quad & x_{k+1} = \tilde{x}(t_{k+1}; x_k, u_k, b),\ k \in \{1,\dots,N_{\text{abs}}\} \\ & g(x_1) = 0. \end{aligned} \tag{16}$$

Problem (16) should be solved every time a new absolute measurement is available. After the problem is solved, the state is propagated by integrating the kinematic model (Eq. (3)) using the IMU measurements, corrected with the bias estimate b, as input to this model. This is done until a new absolute measurement is available. In navigation, this is known as '*dead reckoning*'.

IV. CHOICES FOR PENALTY FUNCTION

The penalty function Ψ for the absolute measurements is not yet specified. In this paper, we consider two choices for Ψ: a least-squares penalty function using the ℓ_2-norm and a robust penalty function using the Huber penalty function.

A. Least Squares Estimation

A first choice for the penalty function Ψ for the residues of the absolute measurements is the ℓ_2-norm:

$$\Psi_{\text{LS}}(y) = \parallel y \parallel_{R^{-1}}^2, \tag{17}$$

where R is the covariance matrix for the absolute measurements. Problem (16) hence becomes

$$\begin{aligned} \underset{\substack{x_1,\dots,x_{N_{\text{abs}}} \\ u_1,\dots,u_{N_{\text{abs}-1}} \\ b}}{\text{minimize}} \quad & \textstyle\sum_{k=1}^{N_{\text{abs}}} \parallel y_{\text{abs},k} - h_{\text{abs}}(x_k) \parallel_{R^{-1}}^2 \\ & + \textstyle\sum_{k=1}^{N_{\text{abs}}-1} \sum_{r=1}^{6} \sigma^{-2} \parallel u_{k,r} - \tilde{u}_{k,r} \parallel_2^2 \\ & + \parallel b - b_{\text{prev}} \parallel_B^2 \\ \text{subject to} \quad & x_{k+1} = \tilde{x}(t_{k+1}; x_k, u_k, b),\ k \in \{1,,\cdots,N_{\text{abs}}\} \\ & g(x_1) = 0. \end{aligned} \tag{18}$$

B. Robust Estimation

The algorithm that detects the marker positions in the image does not always converge to the correct solution due to disturbances present in the image. This causes outliers present in the absolute measurements. To increase the robustness of the MHE to such outliers, the Huber penalty function can be used as penalty function for the absolute measurements. The Huber penalty function is given by

$$\Psi_{\text{Huber}}(x) = \begin{cases} x^2 & \mid x \mid \leq \rho \\ \rho(2 \mid x \mid -\rho) & \mid x \mid > \rho. \end{cases} \tag{19}$$

It is thus equal to an ℓ_2-norm for small errors, but acts like the ℓ_1-norm as soon as the errors are bigger than a certain threshold, and thus is much less sensitive to outliers than the ℓ_2-norm. In an optimization framework, the Huber penalty function can be formulated as (see [3])

$$\begin{aligned} \Psi_{\text{Huber}}\left(\tfrac{x}{\sigma}\right) = \quad \underset{\mu,\nu}{\min} \quad & \mu^2 + 2\rho\nu \\ \text{s.t.} \quad & -\mu - \nu \leq \tfrac{x}{\sigma} \leq \mu + \nu \\ & 0 \leq \mu \leq \rho \\ & 0 \leq \nu, \end{aligned} \tag{20}$$

where v and w are slack variables. Thus, the overall optimization problem can be formulated as an NLP with three continuously differentiable problem functions. For each residual two slack variables are needed. The threshold ρ should be large enough that normal noise levels lie within the ℓ_2-region, but not so large that outliers lie within the ℓ_1-region. A good choice for ρ is e.g. 3, which means that residuals smaller than 3 times the standard deviation of the normal noise lie within the ℓ_2-region. For a Gaussian distribution of the noise, the ℓ_1-region of the Huber penalty function will thus be active only about 0.3% of the time when there is a normal measurement.

Applying the Huber penalty function to the term for the absolute measurements, and using the formulation with slack variables, results in the following optimization problem:

$$
\begin{aligned}
\underset{\substack{x_1,\ldots,x_{N_{\text{abs}}}\\ u_1,\ldots,u_{N_{\text{abs}}-1}\\ \mu_1,\ldots,\mu_{N_{\text{abs}}-1}\\ \nu_1,\ldots,\nu_{N_{\text{abs}}-1}\\ b}}{\text{minimize}} \quad & \textstyle\sum_{k=1}^{N_{\text{abs}}}(\mu_k^2+2\rho\nu_k) \\
& \textstyle+\sum_{k=1}^{N_{\text{abs}}-1}\sum_{r=1}^{6}\sigma^{-2}\parallel u_{k,r}-\tilde{u}_{k,r}\parallel_2^2 \\
& +\parallel b-b_{\text{prev}}\parallel_B^2 \\
\text{subject to} \quad & x_{k+1}=\tilde{x}(t_{k+1};x_k,u_k,b),\ k\in\{1,\ldots,N_{\text{abs}}\} \\
& g(x_1)=0 \\
& -\mu_k-\nu_k\le R^{-1/2}(y_{\text{abs},k}-h_{\text{abs}}(x_k))\le\mu_k+\nu_k \\
& 0\le\mu_k\le\rho \\
& 0\le\nu_k, \qquad\qquad k\in\{1,\ldots,N_{\text{abs}}\},
\end{aligned}
\tag{21}
$$

where μ_k and ν_k are vectors of slack variables.

V. SIMULATION RESULTS

The sensors that are used for the experiments, and that are also used for the simulations, consists of a stereo vision system and an IMU. The stereo vision system consists of two cameras that are mounted on the rotating carousel. These cameras observe three markers mounted on the airplane. This provides absolute measurements of the position and orientation of the airplane if the position and orientation of the cameras is known w.r.t. a certain reference frame. The cameras normally operate at a sampling frequency of 15 Hz. The IMU is mounted inside the airplane, and measures the linear acceleration and angular velocity of the airplane in all 3 directions at a sampling frequency of 800 Hz.

The MHE is implemented using CasADi [1]. It is solved by the SQP method from CasADi, allowing full convergence of each problem. The horizon includes 7 absolute measurements. In the simulation, there are 50 relative measurements in between 2 absolute measurements. The horizon hence spans 0.375 s. Gaussian noise with a standard deviation of 5 pixels was added to the absolute measurements, as well as occasional outliers. For the relative measurements, Gaussian noise with a standard deviation of 0.1 m/s^2 and 0.57 deg/s was added for the acceleration and angular velocity measurements respectively. These values come from the specifications of the IMU, and were also checked by taking the standard deviation of a time series of static measurements, and were doubled for the simulations. Outliers are added to the absolute measurements with a 5 % chance that there is an outlier on one of the marker positions in one camera image. An outlier means that the position of the marker is random within the image. Polynomials of degree 2 are used to approximate the relative measurements. Both the MHEs using the ℓ_2-norm and the Huber penalty function for the absolute measurements are implemented and tested for data with and without outliers. A comparison with a multiplicative Extended Kalman Filter (EKF) is also made. For details of this filter, refer to [9].

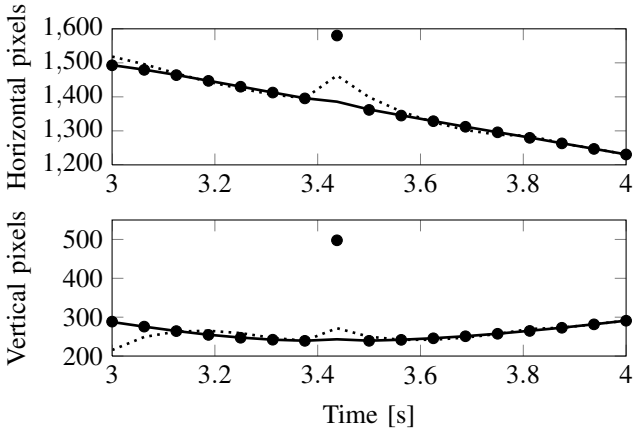


Fig. 5: Camera image reprojection of MHE state estimate for one marker on one image. The dots denote the actual measurements. The dotted line is the reprojected ℓ_2-estimate, the full line is the reprojected estimate using the Huber penalty function.

In case there are no outliers present in the measurements, all estimators have the same estimation performance and converge to the solution. Fig. 5 shows the reprojection for one marker in time in case there are outliers present in the measurements for both the ℓ_2-norm and the Huber penalty function. There is an outlier at a time of about 3.45 s. It is clear from the figure that the Huber penalty function (full line) gives a better estimate than the ℓ_2-norm (dotted line) in case of an outlier. Once the outlier is out of the horizon, the performance of both MHEs is the same. Each time the horizon hits an outlier, the state estimate is perturbed, but as the horizon window passes over the outlier this perturbation fades away.

Fig. 6 shows the corresponding error in the estimation of position and orientation for both moving horizon estimators and for the EKF. For the error in orientation, the equivalent angle of rotation of the relative orientation between the estimate and the solution is taken. This simulation experiment was repeated 100 times, using noise with the same characteristics. Table I shows the average and maximum estimation error over all runs. The maximum position and orientation error for the EKF is 0.69 m and 0.75 rad, for the ℓ_2-norm approach 0.54 m and 0.18 rad, and for the approach with the Huber penalty function, is

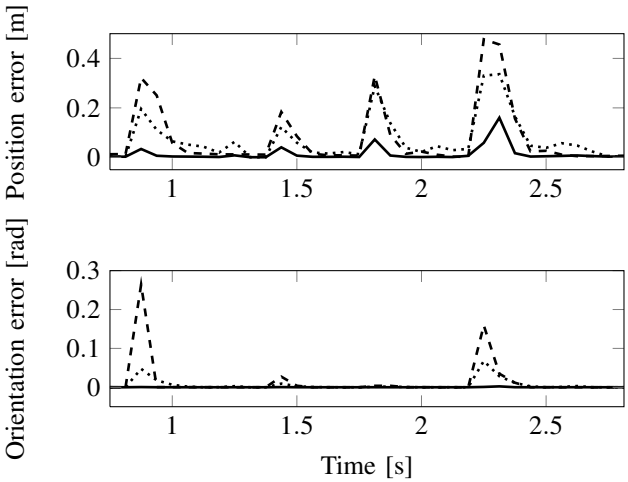


Fig. 6: Error in position and orientation estimates for one run. The dashed line is the EKF estimate, the dotted line is the MHE estimate using the ℓ_2-norm, the full line is the MHE estimate using the Huber penalty function.

Estimator	EKF	MHE(ℓ_2)	MHE(Huber)
Mean position error [mm]	93.2	72.6	6.6
Max position error [mm]	687.2	540.8	65.8
Mean orientation error [mrad]	39.9	13.2	0.1
Max orientation error [mrad]	745.9	182.9	1.3

TABLE I: Average estimation error over 100 runs for EKF and MHE using Huber penalty function and ℓ_2-norm.

only about 0.066 m and 0.0013 rad. The maximum position error for the robust MHE is thus 10 times lower than the maximum position error of the EKF, and about 8 times lower that the maximum position error of the MHE using the ℓ_2-norm. The robust MHE thus performs much better than the EKF and the MHE using the ℓ_2-norm.

VI. RESULTS WITH EXPERIMENTAL DATA

The MHE was run off line on several batches of experimental data. In the experiments, the IMU was running at 800 Hz, and the cameras were running at 7.5 Hz. Fig. 7 shows a time series of absolute measurements for the position of one marker in one camera image (dots) and the corresponding reprojected state estimates (full line) from the robust MHE. Because the current experiments are done in a darkened room, the marker detection algorithm is always successful and therefore there are no outliers in the data. The variation of the measurements is due to actual periodic movement of the airplane. The figure shows the reprojection errors are reasonably small. Because there are no outliers present in the data, the performance of the ℓ_2-norm and the Huber norm is the same. It is expected that there will be more outliers in the data once the experiments are done outdoors, where there is disturbance from ambient light. In this case, the Huber norm will be better suited than the ℓ_2-norm, as was shown in the simulation results.

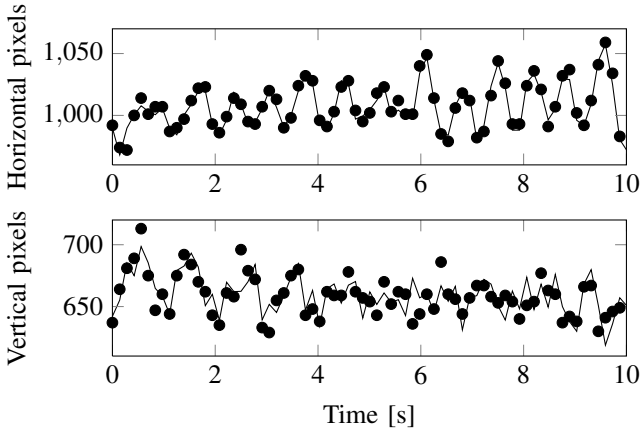


Fig. 7: Camera image reprojection of state estimate for one marker on one image. The dots denote the measurements. The line is the estimate of the robust MHE.

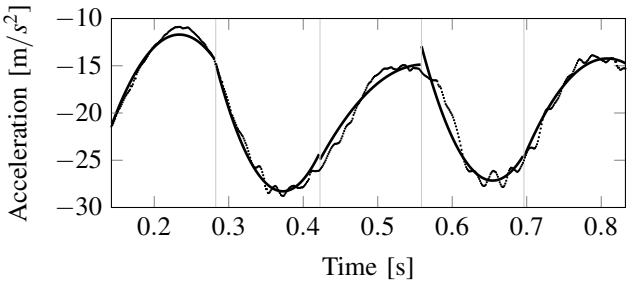


Fig. 8: Acceleration measurements (dotted line) and polynomials that fit this acceleration measurement (full line) for four frames of one MHE solution.

Fig. 8 shows one component of the acceleration measurements (dotted line) in combination with the polynomials that try to fit this acceleration component (full line). Five frames of one MHE solution are shown. Note that we are not aiming for the optimal fit of the polynomial on these measurements, but for an optimal fit in combination with the camera measurements. Polynomials of degree 2 are used to approximate the IMU measurements. There are jumps in the polynomials (e.g. at $t = 0.53$s) because there is no continuity constraint on the polynomials in between frames. Note that since the polynomials are modelling derivatives of the state, the state will not be discontinuous. A motivation for these discontinuities is that in reality, accelerations can be discontinuous. For the polynomials representing the angular velocities, you could consider including continuity constraints. From the IMU measurements, we can see that the airplane is oscillating at a frequency of about 4 Hz. The polynomials seem to capture this oscillation quite well. On top of this frequency of 4 Hz, the IMU also measures an oscillation at 20 Hz. This oscillation is not due to a rigid body movement of the airplane, but due to structural modes. The polynomials filter out these high frequency components, and thus act as a low pass filter.

VII. CONCLUSION

We have demonstrated a moving horizon estimator that is able to fuse IMU measurements and measurements from a stereo vision system, and can reliably estimate of the pose of a tethered airplane while it performs high acceleration (up to 10g) aerobatic maneuvers . The estimator approximates the IMU measurements with polynomials to reduce the number of decision variables in the optimization problem. By limiting the degree of these polynomials, our formulation is also able to simultaneously filter out high frequency vibrations present in the IMU measurements. The estimator is made robust against outliers present in the camera measurements by using the Huber penalty function instead of the ℓ_2-norm. A comparison of both MHE formulations is performed in simulation, showing that the Huber penalty function performs much better than the ℓ_2-norm when outliers are present in the measurements. In experiments on data recorded on our test apparatus, there were only slight differences between the filters, but it is expected that once experiments are done in a less controlled environment there will be outliers present in the data, and the Huber penalty function will be needed. Future work will focus on development of a real-time implementation of the MHE, involving techniques such as automatic code generation similar to that which has previously been done for nonlinear model predictive control [11] using *real-time algorithms* for MHE [13].

VIII. ACKNOWLEDGMENTS

This research was supported by Research Council KUL: PFV/10/002 Optimization in Engineering Center OPTEC, GOA/10/09 MaNet and GOA/10/11 Global real- time optimal control of autonomous robots and mechatronic systems. Flemish Government: IOF/KP/SCORES4CHEM, FWO: PhD/postdoc grants and projects: G.0320.08 (convex MPC), G.0377.09 (Mechatronics MPC); IWT: PhD Grants, projects: SBO LeCoPro; Belgian Federal Science Policy Office: IUAP P7 (DYSCO, Dynamical systems, control and optimization, 2012-2017); EU: FP7-EMBOCON (ICT-248940), FP7-SADCO (MC ITN-264735), ERC ST HIGHWIND (259 166), Eurostars SMART, ACCM.

REFERENCES

[1] J. Andersson, J. Åkesson, and M. Diehl. CasADi – A symbolic package for automatic differentiation and optimal control. In S. Forth, P. Hovland, E. Phipps, J. Utke, and A. Walther, editors, *Recent Advances in Algorithmic Differentiation*, Lecture Notes in Computational Science and Engineering, Berlin, 2012. Springer.

[2] A. Björck. *Numerical Methods for Least Squares Problems*. SIAM, 1996.

[3] S. Boyd and L. Vandenberghe. *Convex Optimization*. University Press, Cambridge, 2004.

[4] H. Bruyninckx and J. De Schutter. *Introduction to Intelligent Robotics*. KU Leuven, 2001.

[5] M. Canale, L. Fagiano, and M. Milanese. Power kites for wind energy generation - fast predictive control of tethered airfoils. *IEEE Control Systems Magazine*, 27:25–38, 2007.

[6] M. Canale, L. Fagiano, and M. Milanese. High Altitude Wind Energy Generation Using Controlled Power Kites. *IEEE Transactons On Control Systems Technology*, 18:168 – 180, 2010.

[7] D. Choukroun, I.Y. Bar-Itzhack, and Y. Oshman. Optimal-request algorithm for attitude determination. *Journal of Guidance Control and Dynamics*, 27, no. 3, 2004.

[8] J.L. Crassidis, F.L. Markley, and Y. Cheng. Survey of nonlinear attitude estimation methods. *Journal of Guidance Control and Dynamics*, 30(1):12–28, 2007.

[9] K. Geebelen, H. Ahmad, M. Vukov, S. Gros, J. Swevers, and M. Diehl. An experimental test set-up for launch/recovery of an Airborne Wind Energy (AWE) system. In *Proceedings of the 2012 American Control Conference*, 2012.

[10] B. Houska and M. Diehl. Robustness and Stability Optimization of Power Generating Kite Systems in a Periodic Pumping Mode. In *Proceedings of the IEEE Multi - Conference on Systems and Control*, pages 2172–2177, Yokohama, Japan, 2010.

[11] B. Houska, H.J. Ferreau, and M. Diehl. An Auto-Generated Real-Time Iteration Algorithm for Nonlinear MPC in the Microsecond Range. *Automatica*, 47(10):2279–2285, 2011.

[12] Tom Kraus. Real-Time State And Parameter Estimation for NMPC-Based Feedback Control With Application To The Tennessee Eastman Benchmark Process. Master's thesis, University of Heidelberg, 2007.

[13] Peter Kühl, Moritz Diehl, Tom Kraus, Johannes P. Schlöder, and Hans Georg Bock. A real-time algorithm for moving horizon state and parameter estimation. *Computers & Chemical Engineering*, 35(1):71–83, January 2011.

[14] E.J. Lefferts, F.L. Markley, and M.D. Shuster. Kalman filtering for spacecraft attitude estimation. *Journal of Guidance, Control, and Dynamics*, 5(5):417–429, 1982.

[15] M.L. Loyd. Crosswind Kite Power. *Journal of Energy*, 4(3):106–111, July 1980.

[16] F.L. Markley. Attitude determination using vector observations: A fast optimal matrix algorithm. *The Journal of the Astronautical Science*, 41:261–280, 1993.

[17] F.L. Markley. Attitude Estimation or Quaternion Estimation. *Guidance, Navigation and Control Systems Engineering*, 2003.

[18] T. Ohtsuka and H.A. Fujii. Nonlinear Receding-Horizon State Estimation by Real-Time Optimization Technique. *Journal of Guidance, Control, and Dynamics*, 19(4), 1996.

[19] C. V. Rao, J. B. Rawlings, and D. Q. Mayne. Constrained state estimation for nonlinear discrete-time systems: Stability and moving horizon approximations. *IEEE Transactions on Automatic Control*, 48(2):246–258, 2003.

[20] J.B. Rawlings and B.R. Bakshi. Particle filtering and moving horizon estimation. *Computers and Chemical Engineering*, 30:1529–1541, 2006.

[21] P. Williams, B. Lansdorp, and W. Ockels. Optimal Crosswind Towing and Power Generation with Tethered Kites. *Journal of Guidance, Control, and Dynamics*, 31(1):81–92, January–February 2008.