

A Computationally Efficient Moving Horizon Estimator for Ultra-Wideband Localization on Small Quadrotors

Pfeiffer, Sven; Wagter, Christophe De; Croon, Guido C.H.E.De

DOI

[10.1109/LRA.2021.3095519](https://doi.org/10.1109/LRA.2021.3095519)

Publication date

2021

Document Version

Final published version

Published in

IEEE Robotics and Automation Letters

Citation (APA)

Pfeiffer, S., Wagter, C. D., & Croon, G. C. H. E. D. (2021). A Computationally Efficient Moving Horizon Estimator for Ultra-Wideband Localization on Small Quadrotors. *IEEE Robotics and Automation Letters*, 6(4), 6725-6732. Article 9478211. <https://doi.org/10.1109/LRA.2021.3095519>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Green Open Access added to TU Delft Institutional Repository

'You share, we take care!' - Taverne project

<https://www.openaccess.nl/en/you-share-we-take-care>

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.

A Computationally Efficient Moving Horizon Estimator for Ultra-Wideband Localization on Small Quadrotors

Sven Pfeiffer[✉], Christophe de Wagter[✉], and Guido C.H.E. de Croon[✉], *Member, IEEE*

Abstract—We present a computationally efficient moving horizon estimator that allows for real-time localization using Ultra-Wideband measurements on small quadrotors. The estimator uses only a single iteration of a simple gradient descent method to optimize the state estimate based on past measurements, while using random sample consensus to reject outliers. We compare our algorithm to a state-of-the-art Extended Kalman Filter and show its advantages when dealing with heavy-tailed noise, which is frequently encountered in Ultra-Wideband ranging. Furthermore, we analyze the algorithm's performance when reducing the number of beacons for measurements and we implement the code on a 30 g Crazyflie drone, to show its ability to run on computationally limited devices.

Index Terms—Sensor fusion, localization, aerial systems: perception and autonomy, optimization and optimal control.

I. INTRODUCTION

OVER the past few years, the continuing miniaturization of computational hardware has advanced the development of small, agile quadrotors. This new generation of flying robots performs tasks in environments, that were previously inaccessible to them due to size and safety constraints. Most notably, this concerns indoor environments, such as warehouses, greenhouses or industrial facilities, which present numerous new potential use-cases.

A key component required for autonomous operation in these environments is a robust and accurate localization system. Ultra-Wideband (UWB) technology has been gaining a lot of attention for indoor positioning, as it offers high accuracy at an affordable cost. UWB systems allow for the localization of a mobile tag, by exchanging messages with a set of static beacons at known locations in the environment. Although a position can be calculated from a sufficient number of independent UWB measurements through multilateration, better results can be achieved when using additional filtering. Commonly, a state estimator that can fuse the UWB measurements with data collected by the on-board inertial measurement unit (IMU) is employed.

Manuscript received February 23, 2021; accepted June 17, 2021. Date of publication July 8, 2021; date of current version July 22, 2021. This letter was recommended for publication by Associate Editor J. Guivant and Editor S. Behnke upon evaluation of the reviewers' comments. This work was supported by Royal Brinkman. (*Corresponding author: Sven Pfeiffer.*)

The authors are with the Faculty of Aerospace Engineering, Delft University of Technology, 2628CD Delft, The Netherlands (e-mail: s.u.pfeiffer@tudelft.nl; c.dewagter@tudelft.nl; guido.de.croon@gmail.com).

Digital Object Identifier 10.1109/LRA.2021.3095519




Fig. 1. Moving horizon estimators can optimize the evolution of a drone's trajectory based on a limited number of previous UWB measurements. We propose a novel formulation for UWB localization that is efficient enough to run on a 30 g Crazyflie drone and performs robustly in scenarios with very noisy or limited measurements. In the plot, ξ stands for the state and y for a measurement of the state.

For linear systems with Gaussian noise, the Kalman Filter yields an optimal estimate in the sense that it minimizes a weighted two-norm of the expected value of the estimation error [1]. Different formulations have been developed to approximate this result for more general systems, while retaining the computationally efficient, iterative structure of the Kalman Filter. The Extended Kalman Filter (EKF) [1] or Unscented Kalman Filter (UKF) [2] are often used to address non-linearities, but fail to address the issues that can arise from uncertain system models and unknown or non-Gaussian noise. Since UWB measurements often exhibit heavy-tailed noise due to multi-pathing and non-line-of-sight (NLOS) conditions [3], [4], this can be a significant shortcoming.

For highly non-linear problems and systems with non-Gaussian noise, Moving Horizon Estimation (MHE) yields better results at the cost of requiring more computational power.

MHE aims at finding a sequence of states, that minimizes the noise and disturbances required to explain the observed measurements over a moving time horizon [5]. A multitude of options exist for the formulation of the cost-function to this minimization problem, but most commonly, a batch least-squares formulation is used [6]. In addition, past data can be included in the cost-function through an “arrival cost”. Unfortunately, the requirement of solving an optimization problem at every time step currently make MHE impractical for real-time applications on computationally limited devices or for systems with fast dynamics.

In this work, we present a computationally efficient Moving Horizon Estimator with RANSAC [7] outlier rejection that can perform real-time localization with UWB on a small quadrotor with limited computational power. We compare it with a state-of-the-art EKF in simulated noise environments and show its advantages when dealing with heavy-tailed noise, which is observed on UWB measurements when multi-pathing and NLOS effects occur. We also show that the MHE is still able to perform state estimation when the number of UWB beacons in the environment is drastically reduced.

In Section II we will present an overview of the related work. Section III outlines the design of our computationally efficient moving horizon estimator, while Section IV will cover its implementation for position estimation with UWB. We will then compare our MHE with a state-of-the-art EKF in Section V. Finally we will conclude on our findings in Section VI.

II. RELATED WORK

Shortly after the original paper on the Kalman Filter appeared in 1960 [8], the concept of limited memory filters was introduced by Andrew Jazwinski to address cases in which the Kalman Filter diverges due to modelling errors [9]. Several strategies have since been proposed to reduce the high computational cost of such limited memory filters.

The Unbiased Finite Impulse Response (UFIR) Filter was developed with a computationally efficient iterative structure like the Kalman Filter. It offers robustness with respect to unknown noise and initial position, but still relies on linearization when dealing with non-linear system models. Furthermore, the UFIR filter is only unbiased and not optimal [10].

Moving Horizon Estimation (MHE) on the other hand solves an online optimization problem and can therefore accomodate non-linear system models. However, the solution of these optimization problems can quickly become expensive, which has mostly limited the use of MHE to slow processes, e.g., in the chemical industry. There are however approaches to improve the computational complexity of moving horizon estimation. Zavala, Laird and Biegler [11] have used sensitivity analysis to formulate and solve a reference optimization problem based on previous measurements, while waiting for a new measurement to arrive. Once the new measurement is known, the solution to the real problem can be calculated more quickly. Kühl *et al.* [12] have demonstrated a scheme that only relies on a single iteration of a sequential quadratic programming method per time step. Both papers demonstrate performance

on examples from the chemical industry, which deal with more states but involve significantly slower dynamics than quadrotors.

For this work, the use of computationally more efficient gradient descent methods was investigated, which were proven to be stable by Alessandri and Gaggero [13]. Similar to Kühl *et al.* they were able to show, that also with simple gradient descent minimization, a single iteration at every time step is sufficient to get a stable state estimate.

On autonomous flying robots, MHE is rarely used due to the before mentioned computational complexity. Girrbach *et al.* have investigated the influence of horizon length when fusing GNSS and IMU with MHE, but only perform comparisons offline and without addressing the challenges of computational cost [14]. Shuo *et al.* formulate the MHE problem with position measurements and a simplified model in such a way, that there is an analytical solution to the optimization problem. This allows them to use their MHE on a 72 g quadrotor in a drone racing context [15].

Instead of using MHE, state estimation with UWB measurements on quadrotors is usually based on variations of the Kalman Filter. The Crazyflie drone which we use as our testing platform in this work for example uses an EKF for fusing data from a wide variety of on-board sensors [16]. As an alternative to Kalman Filters, Xu *et al.* [17], [18] recently started to investigate different forms of UFIR filters for the purpose of fusing UWB and IMU on drones.

In our work, we are approaching the problem of sensor fusion on quadrotors by using a moving horizon estimator, which uses a single iteration gradient descent algorithm to optimize its state-estimate. We will show that this reduces the computational load of MHE sufficiently, while providing accurate position estimates on computationally limited drones.

III. MOVING HORIZON ESTIMATOR

A. Problem Formulation

Let us consider a non-linear, discrete-time dynamic system as described by (1):

$$\begin{cases} \xi_{t+1} = f(\xi_t, u_t) + w_t \\ y_t = h(\xi_t, u_t) + v_t \end{cases} \quad (1)$$

where $t = 0, 1, \dots$ is the time step, $\xi_t \in \mathbb{R}^n$ is the state vector, $u_t \in \mathbb{R}^p$ is the input vector, and $y_t \in \mathbb{R}^m$ is the measurement vector. Finally $w_t \in \mathbb{R}^n$ and $v_t \in \mathbb{R}^m$ are the process and measurement noise vectors, all at time step t respectively.

The classical MHE problem can be formulated, as finding the state $\hat{\xi}_{t-N|t}$, that minimizes the cost function $J(\hat{\xi}_{t-N|t})$, where $\hat{\xi}_{t-N|t}$ is the estimated state at time step $t - N$, based on information available at time step t , N being the length of the moving horizon. The cost function typically includes an arrival cost, which takes into account a prior estimate of the state, $\bar{\xi}_{t-N|t}$, the sum of squares of the process noise and the sum of squares of the measurement noise. In this work, we simplify the problem to an output noise MHE, which neglects the process noise terms. This simplification is necessary to reduce the size of the optimization problem sufficiently for it to be solved on

an embedded system. It removes the ability of the estimator to adjust inputs to better fit the measurements and is therefore likely to reduce the accuracy of the estimate. Since it does however not impact the convergence of the estimate, this can be an acceptable trade-off for many systems.

Since the measurement frequency of UWB varies unpredictably, some measurements in $\mathbf{y}_k$ might be outdated. We can accommodate for this fact, by writing the cost function in terms of individual measurements $y_i, i = 0, \dots, m$, and including the binary switching sequences $\theta_{i,k}$, which are equal to 1 if a new measurement arrives and 0 otherwise [19].

$$\theta_{i,k} = \begin{cases} 1 & \text{if } y_{i,k} \neq y_{i,k-1} \\ 0 & \text{otherwise} \end{cases} \quad (2)$$

To express the cost function in term of the state at the beginning of the horizon, $\bar{\xi}_{t-N|t}$, it is helpful to define the composite prediction and measurement equations at time step k :

$$\mathcal{F}_k(\bar{\xi}_{t-N|t}) = \mathbf{f}^{u_k} \circ \dots \circ \mathbf{f}^{u_{t-N}}(\bar{\xi}_{t-N|t}) \quad (3)$$

$$\mathcal{H}_{i,k}(\bar{\xi}_{t-N|t}) = h_i^{u_k} \circ \mathcal{F}_{k-1}(\bar{\xi}_{t-N|t}) \quad (4)$$

For compactness, we omit the input vectors $\mathbf{u}_{t-N}$ to $\mathbf{u}_k$ in the composite functions, as the output error structure of the MHE does not allow them to change anyway. For the composing functions we note the corresponding input vector in the superscript. This results in the final cost function shown in (5).

$$J_t(\hat{\xi}_{t-N|t}) = \mu \|\hat{\xi}_{t-N|t} - \bar{\xi}_{t-N|t}\|^2 + \sum_{k=t-N}^t \sum_{i=1}^m \left(\theta_{i,k} \cdot \|y_{i,k} - \mathcal{H}_{i,k}(\hat{\xi}_{t-N|t})\|^2 \right) \quad (5)$$

B. Computationally Efficient MHE

To solve the MHE problem efficiently, we use a single iteration of the gradient method at every time step [13]. The prior estimate $\bar{\xi}_{t-N|t}$ is obtained by applying the prediction equations on the last available state estimate.

$$\hat{\xi}_{t-N|t} = \bar{\xi}_{t-N|t} - \alpha \nabla J_t(\bar{\xi}_{t-N|t}) \quad (6)$$

$$\bar{\xi}_{t-N|t} = \mathbf{f}(\hat{\xi}_{t-N-1|t-1}, \mathbf{u}_{t-N-1}) \quad (7)$$

If the number of states n becomes large, the numerical calculation of $\nabla J_t(\bar{\xi}_{t-N|t})$ becomes very expensive. This is due to the fact, that every evaluation of $J_t(\bar{\xi}_{t-N|t})$ requires the prediction of the state vector's evolution over the complete horizon. We are therefore using an iterative, analytical solution based on the chain rule, which only requires a single prediction cycle over the full horizon.

In the analytical expression of the gradient of the cost function, the first term (arrival cost) disappears since we evaluate the gradient at $\bar{\xi}_{t-N|t}$. What remains is a sum of terms that stem from the individual measurements:

$$\nabla J_t(\bar{\xi}_{t-N|t}) = -2 \sum_{k=t-N}^t \sum_{i=1}^m$$

$$[\theta_{i,k} (y_{i,k} - \mathcal{H}_{i,k}(\bar{\xi}_{t-N|t})) \nabla \mathcal{H}_{i,k}(\bar{\xi}_{t-N|t})] \quad (8)$$

By stepping through the full horizon, we can calculate the sequence of prior estimates $\bar{\xi}_{k|t}, k = t-N, \dots, t$ in an iterative manner, which makes calculating the predicted measurement $\mathcal{H}_{i,k}(\bar{\xi}_{t-N|t}) = h_i(\bar{\xi}_{k|t}, \mathbf{u}_k)$ trivial.

To calculate the gradient of $\mathcal{H}_{i,k}$, we make use of the chain rule to split the equation into two factors: The gradient based on the current predicted state and input, $\nabla h_i(\bar{\xi}_{k|t}, \mathbf{u}_k)$, and the Jacobian of the composite prediction equation $D\mathcal{F}_{k-1}(\bar{\xi}_{t-N|t})$:

$$\nabla \mathcal{H}_{i,k}(\bar{\xi}_{t-N|t}) = (D\mathcal{F}_{k-1}(\bar{\xi}_{t-N|t}))^T \cdot \nabla h_i(\bar{\xi}_{k|t}, \mathbf{u}_k) \quad (9)$$

$$D\mathcal{F}_{k-1}(\bar{\xi}_{t-N|t}) = D\mathbf{f}(\bar{\xi}_{k-1|t}, \mathbf{u}_{k-1}) \cdot D\mathbf{f}(\bar{\xi}_{k-2|t}, \mathbf{u}_{k-2}) \cdot \dots \cdot D\mathbf{f}(\bar{\xi}_{t-N|t}, \mathbf{u}_{t-N}) \quad (10)$$

Similar to the predicted state, the Jacobians of the prediction equation can be calculated in an iterative manner by premultiplying the previous value with the Jacobian evaluated at the current prediction. In the end, calculating the Jacobian of the cost function therefore requires the following computations for each time step k in the horizon (i.e. N times):

- Forward prediction of the prior estimate

$$\bar{\xi}_{k|t} = \mathbf{f}(\bar{\xi}_{k-1|t}, \mathbf{u}_{k-1}) \quad (11)$$

- Calculation of the $n \times n$ Jacobian of the prediction function $\mathbf{f}$ at $t = k-1$

$$D\mathbf{f}(\bar{\xi}_{k-1|t}, \mathbf{u}_{k-1}) = \frac{\partial \mathbf{f}}{\partial \xi}(\bar{\xi}_{k-1|t}, \mathbf{u}_{k-1}) \quad (12)$$

- Calculation of the composite Jacobian of $\mathcal{F}_{k-1}$ by multiplication of two $n \times n$ matrices (chain rule)

$$D\mathcal{F}_{k-1}(\bar{\xi}_{t-N|t}) = D\mathbf{f}(\bar{\xi}_{k-1|t}, \mathbf{u}_{k-1}) \cdot D\mathcal{F}_{k-2}(\bar{\xi}_{t-N|t}) \quad (13)$$

- Up to m measurement predictions for measurements that arrived at $t = k$,

$$h_i(\bar{\xi}_{k|t}, \mathbf{u}_k) \quad (14)$$

- Calculation of up to m measurement gradients for measurements that arrived at $t = k$,

$$\nabla h_i(\bar{\xi}_{k|t}, \mathbf{u}_k) = \frac{\partial h_i}{\partial \xi}(\bar{\xi}_{k|t}, \mathbf{u}_k) \quad (15)$$

- Calculation of up to m composite measurement gradients by multiplication of an $n \times n$ matrix with a size n vector, (9)
- Up to m subtractions, $m \cdot n$ multiplications and $m \cdot n$ additions to form the full measurement sum at $t = k$, (8)

To complete the estimation process, the gradient descent step needs to be performed (6), which requires n multiplications and n subtractions. Finally, a forward prediction of $\bar{\xi}_{t-N|t}$ must be performed to actually know the current state $\hat{\xi}_{t|t}$, which requires k evaluations of the prediction equation $\mathbf{f}$.

From these steps, the most expensive calculation in terms of computational complexity is the multiplication of two $n \times n$ matrices, which behaves as $\mathcal{O}(n^3)$ without using specialized algorithms. Neglecting lower order terms, our MHE therefore ends up with a computational complexity of $\mathcal{O}(N \cdot n^3)$. With respect to the number of states, our MHE therefore scales just as well as the EKF and the increase in complexity only comes from the number of time steps in the moving horizon.

Algorithm 1: Computationally efficient MHE.

```

input :  $u_t, y_t$ 
output:  $\hat{\xi}_t$ 
1  $N_t \leftarrow 0$ 
2  $\bar{u} \leftarrow [], \bar{y} \leftarrow []$ 
3  $\bar{\xi}_{\text{prior}} \leftarrow \xi_0$ 
4 while true do
5   Update the moving window
6   if  $N_t < N_{\text{max}}$  then
7      $\bar{u}[N_t] \leftarrow u_t$ 
8      $\bar{y}[N_t] \leftarrow y_t$ 
9      $N_t \leftarrow N_t + 1$ 
10  else
11     $\bar{\xi}_{\text{prior}} \leftarrow f(\bar{\xi}_{t-N}, \bar{u}[0])$ 
12    shift_left ( $\bar{u}$ )
13    shift_left ( $\bar{y}$ )
14     $\bar{u}[N_t] \leftarrow u_t$ 
15     $\bar{y}[N_t] \leftarrow y_t$ 
16  Calculate the gradient of the cost function
17   $\xi \leftarrow \bar{\xi}_{\text{prior}}$ 
18   $DF \leftarrow \mathcal{I}_{n \times n}$ 
19   $M_\Sigma \leftarrow 0$ 
20  for  $k \in [0, \dots, N_t]$  do
21    if  $k \neq 0$  then
22       $DF \leftarrow Df(\xi, \bar{u}[k-1]) \cdot DF$ 
23       $\xi \leftarrow f(\xi, \bar{u}[k-1])$ 
24    for  $i \in [0, \dots, m-1]$  do
25       $\mathcal{H}_{i,k} = h_i(\xi, \bar{u}[k])$ 
26       $\nabla \mathcal{H}_{i,k} = DF^T \cdot \nabla h_i(\xi, \bar{u}[k])$ 
27       $M_\Sigma \leftarrow M_\Sigma - \theta_{i,k} (\bar{y}[k][i] - \mathcal{H}_{i,k}) \cdot \nabla \mathcal{H}_{i,k}$ 
28  Calculate the new estimate
29   $\nabla J \leftarrow -2 \cdot M_\Sigma$ 
30   $\hat{\xi}_{t-N} \leftarrow \bar{\xi}_{\text{prior}} - \alpha \nabla J$ 
31   $\hat{\xi}_t \leftarrow \hat{\xi}_{t-N}$ 
32  for  $k \in [0, \dots, N_t]$  do
33     $\hat{\xi}_t \leftarrow f(\hat{\xi}_t, \bar{u}[k])$ 
34  yield ( $\hat{\xi}_t$ )

```

C. RANSAC Outlier Rejection

The occurrence of outliers in UWB measurements is a problem for any state estimator, as they can have a big impact on the estimated state. To deal with outliers, other authors have used the Mahalanobis distance [16], [17], which weighs the measurement error with its variance. To eliminate any

reliance on the knowledge of error statistics, we have decided to use a random sample consensus (RANSAC) scheme instead.

To do this in an efficient manner, we split the measurement terms in (8) in two sums: A common sum and a RANSAC sum. The common sum includes measurements that are not subject to RANSAC outlier rejection (e.g., altitude measurements) while the RANSAC sum includes only a fraction of the terms from UWB measurements. By adding to multiple RANSAC sums at the same time, multiple gradients of the cost function can be computed after going through only one prediction cycle for the full horizon.

Since this removes the need to step through the full horizon several times, the only real contribution of the outlier rejection to the computational cost is from the evaluation of the different RANSAC estimates. The evaluation of every RANSAC estimate still requires an independent prediction cycle over the full horizon, to choose the best estimate.

IV. IMPLEMENTATION

In our implementation we limit ourselves to the estimation of the drone's position and velocity, while the attitude is estimated using Mahony's AHRS algorithm [20].

A. Prediction Equations

We use a slightly simplified version of the model in [16], where we represent the attitude with quaternions and neglect the dependence of the drag on the propellers' angular velocities.

$$\dot{\mathbf{x}} = \left[\mathbf{q} \otimes \begin{bmatrix} 0 \\ \boldsymbol{\rho} \end{bmatrix} \otimes \mathbf{q}^{-1} \right]_v \quad (16)$$

$$\dot{\boldsymbol{\rho}} = \frac{f}{m} \mathbf{e}_3 + (\mathbf{K}_a - [[\boldsymbol{\omega}_\times]]) \boldsymbol{\rho} - \left[\mathbf{q}^{-1} \otimes \begin{bmatrix} 0 \\ \mathbf{g} \end{bmatrix} \otimes \mathbf{q} \right]_v \quad (17)$$

In this model, our six dimensional state $\boldsymbol{\xi} = [\mathbf{x}^T, \boldsymbol{\rho}^T]^T$ encompasses the position $\mathbf{x} \in \mathbb{R}^3$ in the global frame and the local velocities $\boldsymbol{\rho} \in \mathbb{R}^3$ in the body frame. The inputs to the system are the total thrust f , the attitude quaternion $\mathbf{q} \in \mathbb{R}^4$, and the angular velocities as measured by the gyroscopes, $\boldsymbol{\omega} \in \mathbb{R}^3$. The drone's mass m and drag coefficients $\mathbf{K}_a = \text{diag}(k_\perp, k_\perp, k_\parallel)$ enter as parameters. $\otimes$ is the quaternion product and the subscript 0 and v refer to the scalar and vector portion of the quaternion respectively. Finally, $\mathbf{g}$ is the directed acceleration of gravity and $[[\boldsymbol{\omega}_\times]]$ is the skew symmetric matrix defined as

$$[[\boldsymbol{\omega}_\times]] = \begin{bmatrix} 0 & -\omega_3 & \omega_2 \\ \omega_3 & 0 & -\omega_1 \\ -\omega_2 & \omega_1 & 0 \end{bmatrix} \quad (18)$$

Since inputs are usually available at higher rates than the estimator's update rate, we accumulate these inputs between two estimation cycles and average them to reduce input noise. For the thrust and each of the gyroscope's axes, this is done by simply calculating the mean, while for the attitude quaternion, the average is found by calculating the dominant eigenvector of

the matrix Q [21]:

$$Q = \sum_i \mathbf{q}_i \mathbf{q}_i^T \quad (19)$$

B. Measurement Equations

We consider two different types of UWB measurements that can be used with the estimator. In two-way ranging (TWR), the UWB system returns the distance between the tag at its current position, and the (known) location of an anchor, $\mathbf{p}_{\text{uwb},i}$. The measurement equation, including measurement noise η_{twr} is then given by

$$y_{\text{twr}} = \|\mathbf{p}_{\text{uwb},i} - \mathbf{x}\| + \eta_{\text{twr}} \quad (20)$$

Alternatively, time-difference of arrival (TdoA) measurements can be used. In TdoA schemes, the measurement yields the difference between the distances to two different anchors:

$$y_{\text{tdoa},ij} = \|\mathbf{p}_{\text{uwb},i} - \mathbf{x}\| - \|\mathbf{p}_{\text{uwb},j} - \mathbf{x}\| + \eta_{\text{tdoa}} \quad (21)$$

Since in most indoor environments, the distribution of anchors is mostly in the xy-plane with comparatively small variations in height, altitude estimation based on UWB alone is often difficult and noisy. We therefore also include the measurement equations for dedicated altitude measurements from a barometer or from laser time-of-flight (ToF) measurements:

$$y_{\text{baro}} = x_3 + z_{\text{ref}} + \eta_{\text{baro}} \quad (22)$$

$$y_{\text{tof}} = x_3 \cdot \cos \theta \cos \phi + \eta_{\text{tof}} \quad (23)$$

Note that the ToF measurement equation must take into account the current pitch θ and current roll ϕ since the laser is not measuring the shortest distance to the ground, but the distance on the body frame's z-axis.

V. RESULTS

A. Setup

For testing and evaluation, we implemented our own algorithm and the EKF from [16] (which is used on the Crazyflie) in Python.¹ We compared their real-time position estimates at each timestep t , $\hat{\mathbf{x}}_{t|t}$, with ground-truth from an optitrack motion capture system. The evaluation was performed on real IMU and UWB data, which was collected on a Crazyflie drone at 100 Hz, and optitrack data collected in the “Cyberzoo” testing area at the Aerospace Faculty of TU Delft² (seen in Fig. 2). In total, 24 data sets were gathered on 6 different trajectories, containing either two-way ranging (TWR) or time-difference of arrival (TdoA) data from the UWB system. The different trajectories used were a square, a triangle, an octagon, an hourglass, a five-pointed star and a random sequence of points. On each trajectory, we recorded four runs, two for TWR data and two for TdoA data. The UWB setup included eight beacons, positioned roughly in the eight corners of a cube, enclosing our testing area (cf. Table I). An example for a run of the Octagon trajectory where



Fig. 2. Flight data was collected in the “Cyberzoo” at TU Delft using a Crazyflie drone and the loco positioning system by bitcraze. Groundtruth is provided by an Optitrack motion capture system.

TABLE I
UWB BEACON POSITIONS

ID	x [m]	y [m]	z [m]
0	-4.48	-4.87	1.16
1	-4.78	4.52	0.67
2	4.81	4.57	0.81
3	4.91	-4.77	0.96
4	-4.34	-5.05	2.60
5	-4.91	4.41	2.65
6	4.56	4.77	2.59
7	4.91	-4.77	2.53




Fig. 3. Example of one of the collected trajectories (Octagon) with offline performance of MHE and EKF on real TWR data from 4 beacons.

we test the estimator performance offline but on real TWR data is shown in Fig. 3.

B. Heavy-Tailed Measurement Noise

While the noise on TWR measurements can be assumed to be Gaussian in an ideal scenario, it is more realistically modelled as being heavy-tailed. This is due to multi-pathing and non-line-of-sight (NLOS) effects, which cause UWB signals to arrive with a slight delay compared to the shortest path [3], [4]. To simulate heavy-tailed noise for TWR, we adapt the noise model by Jimenez and Seco [3], which models the heavy-tailed noise as the sum of a Gaussian distribution for the LOS component, a Gamma distribution for the NLOS component, and a constant

¹<https://github.com/SUPfeiffer/uwb-simulator>

²Data available at https://doi.org/10.4121/14_827_680

for outliers. To avoid dealing with a probability density function (PDF) that does not integrate to 1, we remove the constant, and rely on getting outliers from the Gamma distribution alone. To show the effect of an increasingly non-Gaussian distribution, we multiply the Gamma distribution with a scale factor s_{ht} , which we vary between 0 and 1.25. Since a reduction in the NLOS component should also cause the mean to go towards zero, we multiply μ by the same factor, which causes the noise distribution to represent the ideal LOS condition (i.e. a pure, zero-mean Gaussian) for $s_{ht} = 0$. Finally, we divide the PDF by $(1 + s_{ht})$, which causes the integral to be equal to 1. The resulting model is given in (24).

$$f(x) = \frac{1}{1 + s_{ht}} \cdot \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{(x - s_{ht}\mu)^2}{2\sigma^2}\right) + \frac{s_{ht}}{1 + s_{ht}} \cdot \frac{\lambda^k}{(k-1)!} x^{k-1} \exp(-\lambda x) \quad (24)$$

Even though there are significant differences between individual anchors, we find that using the parameters suggested by Jimenez and Seco ($\mu = 0.1$, $\lambda = 3.5$, $k = 2$) and a standard deviation of $\sigma = 0.1$ (determined from our own data), our TWR data is best fit with a heavy-tail scale factor of $s_{ht} = 0.2$. The main differences between anchors are the need for some higher values for λ and k , but we also noted some anchors with shifted or multiple peaks in the noise distribution. These odd behaviours might originate from localized biases on different trajectories and were not modeled in this research.

Since TdoA measurements can be seen as resulting from two TWR measurements to different beacons, the heavier tail appears on both sides and is better modeled by using a Cauchy distribution. To vary the intensity of the heavy-tail, we model the noise distribution for TdoA as the weighted sum of a Cauchy distribution ($\gamma = 0.3$) and a Gaussian distribution ($\sigma = 0.3$), both centered around $\mu = 0$. The resulting noise model, shown in (25) best matches the recorded TdoA data for a relative weight of the Cauchy distribution of $r_{ht} = 0.5$.

$$f(x) = (1 - r_{ht}) \cdot \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{x^2}{2\sigma^2}\right) + r_{ht} \cdot \left[\pi\gamma \left(1 + \left(\frac{x}{\gamma}\right)^2\right) \right]^{-1} \quad (25)$$

Using these noise models, we compare the performance of MHE and EKF on a mix of real IMU data combined with simulated UWB measurements from 4 anchors (IDs 1, 3, 4, 6) that provide measurements on average every 0.05 s (TWR) or 0.1 s (TdoA) each. We run the simulation 5 times on every set of IMU data collected, for a total of 120 runs per value of s_{ht} and r_{ht} . Fig. 5 and Fig. 6 show the average RMSE of these runs for TWR and TdoA measurements respectively, with the shaded regions representing the standard error of the mean.

In the case of TWR, both estimators perform similarly well, with the EKF having a slight edge for Gaussian noise and the MHE showing an advantage for noise with a stronger heavy tail. In the case of TdoA, the MHE performs better over the full range of simulated noise, with the difference increasing at higher




Fig. 4. Probability Density Function (PDF) of the heavy-tailed noise distribution used for generating UWB measurements. The scale factor s_{ht} is used to change the influence of the heavy-tailed Gamma distribution added to the underlying Gaussian.




Fig. 5. Estimator performance on TWR measurements with increasingly heavy-tailed noise. Shaded regions represent the standard error of the mean.




Fig. 6. Estimator performance on TdoA measurements with increasingly heavy-tailed noise. Dashed lines show success rate of the estimator ($RMSE < 1$ m). Unsuccessful runs are excluded from the calculation of the mean RMSE.

values of r_{ht} . The cause for this difference between TWR and TdoA (also for the Gaussian noise) most likely stems from the larger standard deviation for TdoA, which also leads to more outliers.




Fig. 7. Estimator performance on TWR measurements with varying number of beacons. Dashed lines show the success rate (RMSE < 1 m), shaded regions represent the standard error of the mean.




Fig. 8. Estimator performance on TdoA measurements with varying number of beacons. Dashed lines show the success rate (RMSE < 1 m), shaded regions represent the standard error of the mean.

C. Removing UWB Anchors

Of particular interest in the field of UWB localization is the reduction of required fixed infrastructure, namely the number of beacons that need to be placed in the environment. To assess the performance of our estimator in applications with respect to the number of beacons, we use our full data sets (IMU and UWB) and remove the measurements from randomly selected beacons. Specifically, we run the filters with real UWB data with the number of enabled beacons ranging from 1-8 for TWR and 2-8 for TdoA (TdoA measurements require two beacons). We perform 10 runs for every number of beacons on each of the 12 corresponding (TWR or TdoA) data sets, for a total of 120 runs per number of beacons. For each individual run, we randomly select the beacons used, since the placement of beacons with respect to the trajectory was found to have a significant effect on performance for low numbers of beacons. Due to beacons losing connection on three TdoA trajectories, we only use 9 of the 12 data sets for the TdoA comparison. Results of these runs are shown in Fig. 7 and Fig. 8 respectively.

One of the downsides of the proposed gradient-descent method is that the magnitude of the cost function's gradient $\nabla J(\hat{\xi}_{t-N|t})$ varies with the number of measurements. This

means that the value of the step size α had to be slightly adjusted depending on the number of beacons.

As was to be expected, the lower number of beacons causes a significant increase in the estimation error. While the MHE still shows an advantage, the observed errors are higher than with simulated measurements. This is likely due to inaccuracies in measuring out the anchor positions, and spatially varying biases that have also been mentioned by [16]. In fact, we found that measurements from some beacons exhibit far more challenging noise than what is modeled by a heavy-tailed distribution, such as multiple peaks. For TDOA measurements, both estimators are failing quite often. Indeed, there seems to be a much larger discrepancy between simulated and real TDOA data than for TWR. This is probably due to the fact, that with a large number of beacons, the TDOA algorithm requires a lot of communication between beacons, causing the frequency of the individual measurements to go down as more beacons are present. When removing beacons after the measurements were taken, this leads to much lower measurement frequencies than otherwise expected. It is however still visible, that the reduced measurement frequency seems to have a smaller effect on the MHE's performance.

D. Computational Complexity

To compare the computational complexity, we implement our algorithm as a separate task on a Crazyflie drone which performs these computations on an STM32F405 microcontroller. This task replaces the EKF task in the original Crazyflie firmware, so that only one estimator is running at the same time.

By analyzing task dumps, we can identify how much of the available computational resources a task uses. When using the same settings as in simulation, the MHE task accounted for about 15% of the CPU load, while the EKF only accounted for about 11%. This shows the potential of the algorithm to indeed be used on computationally limited devices. With more potent microprocessors being released continuously, we believe that computationally efficient MHE algorithms can become useful for many applications, for which they were previously considered too expensive.

VI. CONCLUSION

A computationally efficient MHE algorithm for quadrotor localization with UWB was developed and tested. We showed that by simplifying the problem to an output noise MHE and using a single-iteration gradient descent method, MHE can yield good results on computationally limited devices. While these simplifications helped us in reducing the computational cost of MHE, they also pose some limitations on the applications of the algorithm.

The simplification to an output noise MHE assumes that the process model is perfectly known and does not exhibit any process noise, which is rarely the case. Even though the algorithm performed well, there is potential for improvement, especially when only few measurements are available.

Using a simple single-iteration gradient descent algorithm causes the optimization for the correct state to happen over

several time steps, and requires the step-size α to be well chosen to keep the algorithm stable. Especially when the number of measurements is varying, this can become difficult, due to the varying magnitude of the gradient of the cost function. We found that conservative values of α can yield stable results for large variations in the number of measurements, but the performance will be worse when fewer measurements are available.

We plan to address these issues in future research to further improve the accuracy of the proposed method and make it useful for a wider variety of tasks.

REFERENCES

- [1] D. Simon, *Optimal State Estimation*. Hoboken, NJ, USA: Wiley, 2006.
- [2] S. J. Julier and J. K. Uhlmann, "New extension of the Kalman filter to nonlinear systems," in *Proc. Signal Process., Sensor Fusion, Target Recognit. VI*, I. Kadar, Ed., Int. Soc. Opt. Photon, vol. 3068, 1997, pp. 182–193.
- [3] A. R. Jiménez and F. Seco, "Comparing ubisense, BeSpoon, and DecaWave UWB location systems: Indoor performance analysis," *IEEE Trans. Instrum. Meas.*, vol. 66, no. 8, pp. 2106–2117, Aug. 2017.
- [4] M. Kok, J. D. Hol, and T. B. Schon, "Indoor positioning using ultrawideband and inertial measurements," *IEEE Trans. Veh. Technol.*, vol. 64, no. 4, pp. 1293–1303, Apr. 2015.
- [5] D. A. Allan and J. B. Rawlings, "Moving horizon estimation," in *Proc. Handbook Model Predictive Control*, S. V. Raković and W. S. Levine, Eds. Cham, Switzerland: Springer Nature, pp. 99–144, 2019.
- [6] D. G. Robertson, J. H. Lee, and J. B. Rawlings, "A moving horizon-based approach for least-squares estimation," *AIChE J.*, vol. 42, no. 8, pp. 2209–2224, 1996.
- [7] M. A. Fischler and R. C. Bolles, "Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography," *Commun. ACM*, vol. 24, no. 6, pp. 381–395, Jun. 1981.
- [8] R. E. Kalman, "A new approach to linear filtering and prediction problems," *J. Basic Eng.*, vol. 82, no. 1, pp. 35–45, 1960.
- [9] A. H. Jazwinski, "Limited memory optimal filtering," *IEEE Trans. Automat. Contr.*, vol. AC-13, no. 5, pp. 558–563, Oct. 1968.
- [10] Y. S. Shmaliy, S. Zhao, and C. K. I. Ahn, "Unbiased finite impulse response filtering," *IEEE Control Syst. Mag.*, vol. 37, no. 5, pp. 70–89, Oct. 2017.
- [11] V. M. Zavala, C. D. Laird, and L. T. Biegler, "A fast moving horizon estimation algorithm based on nonlinear programming sensitivity," *J. Process Control*, vol. 18, no. 9, pp. 876–884, 2008.
- [12] P. Kühn, M. Diehl, T. Kraus, J. P. Schlöder, and H. G. Bock, "A real-time algorithm for moving horizon state and parameter estimation," *Comput. Chem. Eng.*, vol. 35, no. 1, pp. 71–83, 2011.
- [13] A. Alessandri and M. Gaggero, "Fast moving horizon state estimation for discrete-time systems using single and multi iteration descent methods," *IEEE Trans. Automat. Contr.*, vol. 62, no. 9, pp. 4499–4511, 2017.
- [14] F. Gierbach, J. D. Hol, G. Bellusci, and M. Diehl, "Optimization-based sensor fusion of GNSS and IMU using a moving horizon approach," *Sensors (Basel, Switzerland)*, vol. 17, no. 5, pp. 1–17, 2017.
- [15] S. Li, E. van der Horst, P. Duernay, C. De Wagter, and G. C. de Croon, "Visual model-predictive localization for computationally efficient autonomous racing of a 72-g drone," *J. Field Robot.*, vol. 37, no. 4, pp. 667–692, 2020.
- [16] M. W. Mueller, M. Hamer, and R. D'Andrea, "Fusing ultra-wideband range measurements with accelerometers and rate gyroscopes for quadcopter state estimation," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2015, pp. 1730–1736.
- [17] Y. Xu, C. K. Ahn, Y. S. Shmaliy, X. Chen, and Y. Li, "Adaptive robust INS/UWB-integrated human tracking using UFIR filter bank," *Measurement: J. Int. Meas. Confederation*, vol. 123, pp. 1–7, Jul. 2018. [Online]. Available: <https://doi.org/10.1016/j.measurement.2018.03.043>
- [18] Y. Xu, Y. S. Shmaliy, T. Shen, D. Chen, M. Sun, and Y. Zhuang, "INS / UWB-Based quadrotor localization under colored measurement noise," *IEEE Sensors J.*, vol. 21, no. 5, pp. 6384–6392, Mar. 2021.
- [19] A. Liu, W. A. Zhang, M. Z. Chen, and L. Yu, "Moving horizon estimation for mobile robots with multirate sampling," *IEEE Trans. Ind. Electron.*, vol. 64, no. 2, pp. 1457–1467, Feb. 2017.
- [20] R. Mahony, T. Hamel, and J. Pflimlin, "Nonlinear complementary filters on the special orthogonal group," *IEEE Trans. Automat. Contr.*, vol. 53, no. 5, pp. 1203–1218, Jun. 2008.
- [21] F. L. Markley, Y. Cheng, J. L. Crassidis, and Y. Oshman, "Averaging quaternions," *J. Guid., Control, Dyn.*, vol. 30, no. 4, pp. 1193–1197, 2007.