Preprints of the 3rd IFAC Workshop on
Distributed Estimation and Control in Networked Systems
September 14-15, 2012. Santa Barbara, CA, USA

Event-Based State Estimation with Switching Static-Gain Observers ⋆

Sebastian Trimpe *

* *Institute for Dynamic Systems and Control (IDSC), ETH Zurich, Sonneggstr. 3, 8092 Zurich, Switzerland (e-mail: strimpe@ethz.ch).*

Abstract: An event-based state estimation problem is considered where the state of a dynamic system is observed from multiple distributed sensors that sporadically transmit their measurements to a remote estimator over a common bus. The common bus allows each sensor to run a copy of the remote estimator and to make the triggering decision based on this estimate: a measurement is transmitted only if its prediction by the estimator deviates by more than a tunable threshold. The event-based estimator is a switching observer that mimics a Luenberger observer with full communication of all measurements. It is proven that the difference between the event-based estimator and its full communication counterpart is bounded. The reduction of average sensor communication rates achieved by using the event-based state estimator for feedback control is demonstrated in experiments on a balancing cube.

Keywords: Event-based estimation, distributed estimation, state observers, networked system.

1. INTRODUCTION

Figure 1 illustrates the state estimation problem that is considered in this paper. The state of a dynamic system is estimated from measurements of multiple sensors distributed along the system. The sensors have computing capability, and decide when to broadcast their measurement over a common bus network that all sensors and the estimator are connected to. Since the system state is not necessarily observable from a single sensor alone, communication from different sensors is required. The objective is to maintain an estimate of the full system state while, at the same time, reducing the load on the communication network.

An event-based strategy to addressing this problem is explained in Fig. 2, which depicts a single sensor node of Fig. 1. The key idea is to implement, on each sensor, a copy of the remote estimator and an appropriate transmit decision rule (event generator): a measurement is transmitted *only when it is required to meet a certain estimation performance*. The event-based state estimator consists of the state estimator itself and the event generator.

The common bus (or broadcast network) is a key enabling feature of the architecture in Fig. 1. It allows all sensors to also listen to the other sensors' data and, hence, to generate the estimate $\check{x}(k)$ locally at no additional communication cost. This architecture allows for an effective implementation of a *distributed event triggering* scheme (i.e. one where the triggering depends only on local data), since the quantity of interest (the estimate $\check{x}(k)$) is actually available locally.

The input signal $u(k)$ is assumed to be accessible at the sensor nodes for making the estimator updates. For example, $u(k)$ may be an a-priori known reference input,

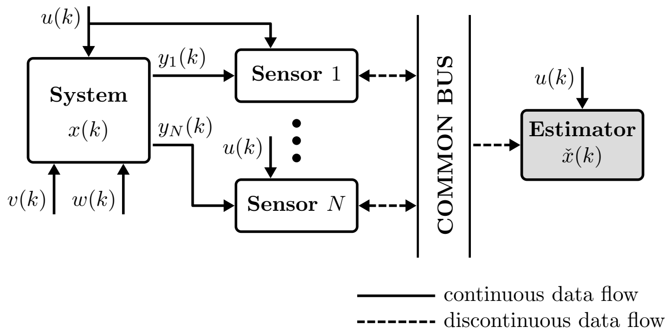

Fig. 1. Distributed state estimation problem. The state $x(k)$ of a dynamical system is observed by measurements $y_i(k)$ of N sensors. The system is driven by a known input $u(k)$ and unknown process and measurement disturbances $v(k)$ and $w(k)$. The sensor nodes are connected to each other and to a remote estimator via a common bus. Each sensor decides when to broadcast its local measurement over the bus. The remote estimator generates an estimate $\check{x}(k)$ of the state $x(k)$ based on the received data.

or $u(k) = 0$ for an autonomous system. When $u(k)$ is computed from a feedback control law, it needs to be communicated to the sensor nodes. In that case, continuously communicating $u(k)$, but only sporadically transmitting the measurements $y_i(k)$ may be beneficial if there are fewer inputs than measurements. Furthermore, with knowledge of the inputs, the state estimation problem (which is the focus of this paper) can be treated without consideration of the feedback control law. This reduces the complexity of the design problem.

Different measures are conceivable for the event generation (i.e. for the decision to transmit a measurement). Trimpe

⋆ This work was funded by the Swiss National Science Foundation.

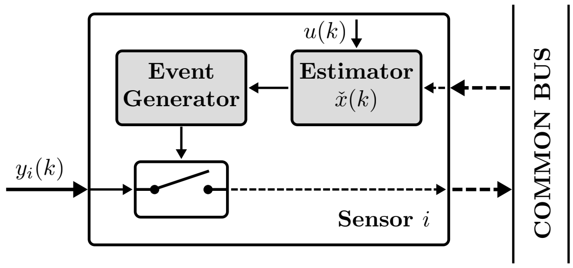

Fig. 2. A single sensor node. Each sensor node listens to all data communicated from other sensors on the common bus. Therefore, it has access to the same data set as the remote estimator of Fig. 1 and can implement a copy of the same. Based on this estimate, the event generator decides when to broadcast the local measurement.

and D'Andrea [2011a] use a constant threshold on the difference of an actual measurement and its prediction by the estimator as condition to trigger an event. We refer to such a triggering rule, which depends on (realtime) measurement data, as *measurement-based triggering*. This is also the approach taken herein. In contrast, a triggering condition that is based on the estimator variance is referred to as *variance-based triggering* (considered in Trimpe and D'Andrea [2011b]). Measurement-based triggering allows for the estimator to react to events in real time, whereas transmit schedules can typically be computed off-line with variance-based triggering.

In contrast to the previously mentioned work, where the state estimator is a time-varying Kalman filter, we use a switching Luenberger-type observer herein. The design is based on a Luenberger observer for the case of synchronous communication of all measurements (referred to as the *full communication state estimator* (FCSE)). The *event-based state estimator* (EBSE) updates its estimate with the available measurements using the corresponding sub-blocks of the static FCSE gain matrix. Hence, the estimator switches between different pre-computed static gains. Compared to a time-varying Kalman filter, where the filter gain is recomputed at every time step from a few matrix multiplications, the approach taken herein is less computationally demanding.

The main contributions of this paper are: the proposal of the EBSE as a switching static-gain observer combined with an appropriate event trigger; proof of an upper bound on the difference between the EBSE and the FCSE; and demonstration of the EBSE performance in experiments on an unstable networked control system (the Balancing Cube; see www.cube.ethz.ch for a video).

This paper is organized as follows: after a brief review of related work and an introduction of notation below, the estimation problem is stated in Sec. 2. The event-based state estimator is derived in Sec. 3, and its stability is analyzed in Sec. 4. Experimental results on the Balancing Cube testbed are given in Sec. 5, and the paper concludes with remarks in Sec. 6.

Related Work. The approach to event-based *state estimation* taken herein is conceptually related to the approach for event-based *control* by Lunze and Lehmann [2010] for a centralized design, and by Stöcker et al. [2012] for a decentralized problem. Therein, the authors design an event triggering rule such that the difference between the state of a reference system with continuous feedback and the state of the event-based control system is bounded. Here, the FCSE is the reference estimator, and event triggers are designed such that the difference of the EBSE to the FCSE is bounded.

Event-based strategies are a popular means of ensuring an efficient use of the communication resource in control, estimation, and optimization problems in networked control systems (see Lemmon [2011] for an overview). For a single sensor and a single estimator node, event-based state estimation problems have been studied by several researchers (see Lemmon [2011] and references therein). A distributed estimation problem related to the one herein is addressed by Weimer et al. [2012]. The authors design communication policies for wireless sensor nodes that may either transmit information to a central estimator, listen to information from the central estimator, or be turned off. When either sensor or estimator transmit data, all data since the last update is sent; hence, the load per packet is variable whereas it is fixed for the approach herein.

Notation. For a vector $v \in \mathbb{R}^n$ and $q \in [1, \infty]$, $\|v\|_q$ (or simply $\|v\|$) denotes the vector Hölder norm of v (see Bernstein [2005])

$$\|v\| = \|v\|_q = \begin{cases} \left(\sum_{j=1}^{n} |v_j|^q \right)^{1/q} & \text{for } 1 \leq q < \infty \\ \max\limits_{j \in \{1,\dots,n\}} |v_j| & \text{for } q = \infty. \end{cases} \tag{1}$$

For a matrix A, $\|A\|_q$ (or simply $\|A\|$) denotes the matrix norm of A induced by the chosen vector norm. For a vector-valued sequence $v = \{v(0), v(1), v(2), \dots\}$, $\|v\|_\infty$ denotes the ℓ^∞ norm of v (see Callier and Desoer [1991])

$$\|v\|_\infty := \sup_{k \geq 0} \|v(k)\|,$$

where $\|v(k)\|$ is the chosen vector norm.

2. ESTIMATION PROBLEM FORMULATION

Consider the linear time-invariant (LTI) system

$$x(k) = A\,x(k-1) + B\,u(k-1) + v(k-1) \tag{2}$$

$$y_1(k) = C_1\,x(k) + w_1(k) \tag{3}$$

$$\vdots$$

$$y_N(k) = C_N\,x(k) + w_N(k), \tag{4}$$

where $k \geq 1$ is the discrete-time index, $x(k) \in \mathbb{R}^n$ is the system state, $u(k) \in \mathbb{R}^{n_u}$ the control input, $y_i(k) \in \mathbb{R}^{p_i}$, $i \in \{1, \dots, N\}$, are measurements by N sensors, $v(k) \in \mathbb{R}^n$, $w_i(k) \in \mathbb{R}^{p_i}$, $i \in \{1, \dots, N\}$, are disturbances, and all matrices are of corresponding dimensions. We use $y(k)$ to denote the vector that combines all measurements; that is,

$$y(k) := \begin{bmatrix} y_1(k) \\ \vdots \\ y_N(k) \end{bmatrix} = \underbrace{\begin{bmatrix} C_1 \\ \vdots \\ C_N \end{bmatrix}}_{=:C} x(k) + \underbrace{\begin{bmatrix} w_1(k) \\ \vdots \\ w_N(k) \end{bmatrix}}_{=:w(k)}$$

$$= Cx(k) + w(k).$$

Hence, $y(k), w(k) \in \mathbb{R}^p$ with $p := \sum_{i=1}^{N} p_i$. We assume that (A, C) is detectable. Notice that (A, C_i) is not assumed to be detectable; that is, the system state is not necessarily detectable from any individual sensor alone.

Notice that no assumption on the characteristics of the disturbances $v(k)$ and $w(k)$ is made. For example, $v(k)$ and $w(k)$ may be random variables with known statistics in a stochastic setting; or they may be bounded disturbances in a deterministic setting.

2.1 Full Communication State Estimator

Next, we introduce the FCSE, which uses the full measurement vector $y(k)$ at every time step and serves as a reference to the EBSE design later:

$$\hat{x}(k|k-1) = A\,\hat{x}(k-1|k-1) + B\,u(k-1) \tag{5}$$

$$\hat{x}(k|k) = \hat{x}(k|k-1) + L\,\big(y(k) - C\,\hat{x}(k|k-1)\big), \tag{6}$$

with the static estimator gain L. The estimator is initialized with some $\hat{x}(0) \in \mathbb{R}^n$. It generates an estimate $\hat{x}(k|k)$ of the state $x(k)$ based on all past measurements up to, and including, $y(k)$. For ease of notation, we write $\hat{x}(k) := \hat{x}(k|k)$.

The estimator gain L is designed such that $(I - LC)A$ is stable (i.e. all eigenvalues have magnitude less than one). If (A, C) is detectable, such an L is guaranteed to exist (see Åström and Wittenmark [1997]). It can be designed, for instance, via pole placement (see Åström and Wittenmark [1997]) or as the steady-state solution of the Kalman filter (see Anderson and Moore [2005]).

Let $\hat{e}(k) := x(k) - \hat{x}(k)$ denote the estimation error of the FCSE. The error evolves according to

$$\hat{e}(k) = (I-LC)A\hat{e}(k-1) + (I-LC)v(k-1) - Lw(k). \tag{7}$$

In a state estimation scenario where $v(k)$ and $w(k)$ are bounded disturbances, the stability of $(I - LC)A$ implies that the estimation error is also bounded. If $v(k)$ and $w(k)$ are independent random variables with finite variance, (7) ensures that the estimation error variance is bounded.

2.2 Problem Statement

An EBSE is sought that approximates the estimate $\hat{x}(k)$ of the FCSE (5)–(6) up to a guaranteed bound, but uses fewer measurements. The EBSE consists of an event generator and a state estimator (with state $\check{x}(k)$) as symbolized by the blocks in Fig. 2.

The sensor nodes and the remote estimator are assumed to be synchronized in time, and transmission via the communication network is assumed to be instantaneous and without data loss. Hence, the state estimates $\check{x}(k)$ on all nodes are assumed identical.

3. EVENT-BASED STATE ESTIMATOR

In this section we present the EBSE, which addresses the problem stated above.

Event Generator. The event generator on sensor i decides at every time step k whether or not to transmit the local measurement $y_i(k)$. The measurement $y_i(k)$ is transmitted whenever a prediction of that measurement based on the previous estimate $\check{x}(k-1)$ is off by more than a tolerable threshold.

Without any information on $v(k)$ and $w_i(k)$, a prediction $\check{y}_i(k)$ of the measurement $y_i(k)$ may be obtained from (2)–(4) by setting $v(k) = 0$ and $w_i(k) = 0$; that is,

$$\check{y}_i(k) := C_i\big(A\check{x}(k-1) + Bu(k-1)\big).$$

Using $\check{x}(k|k-1) := A\check{x}(k-1) + Bu(k-1)$, the employed event triggering rule is

$$\text{transmit } y_i(k) \quad \Leftrightarrow \quad \|y_i(k) - C_i\check{x}(k|k-1)\| \geq \delta_i, \tag{8}$$

with threshold parameters $\delta_i \geq 0$, $i \in \{1, \ldots, N\}$. Tuning δ_i allows the designer to trade off each sensor's frequency of events (and, hence, the communication rate) for estimator performance. For notational convenience, we denote $\delta = (\delta_1, \ldots, \delta_N) \in \mathbb{R}^N$ the vector of threshold parameters δ_i.

State Estimator. Let $I(k)$ denote the tuple of indices of those sensors that transmitted their measurement at time k; that is,

$$I(k) := \big(i \mid 1 \leq i \leq N,\ \|y_i(k) - C_i\check{x}(k|k-1)\| \geq \delta_i\big). \tag{9}$$

The filter that is used to generate an estimate $\check{x}(k) := \check{x}(k|k)$ of the state $x(k)$ based on the measurements broadcast up to time k is given by:

$$\check{x}(k|k-1) = A\,\check{x}(k-1|k-1) + B\,u(k-1) \tag{10}$$

$$\check{x}(k|k) = \check{x}(k|k-1) + \sum_{i \in I(k)} L_i\big(y_i(k) - C_i\check{x}(k|k-1)\big), \tag{11}$$

where $L = [L_1, L_2, \ldots, L_N]$ with $L_i \in \mathbb{R}^{n \times p_i}$ is the decomposition of the estimator gain matrix according to the dimensions of the individual measurements. By rewriting (6) as

$$\hat{x}(k) = \hat{x}(k|k-1) + \sum_{i \in \{1,\ldots,N\}} L_i\big(y_i(k) - C_i\hat{x}(k|k-1)\big), \tag{12}$$

one can see that (11) is obtained from (6) by including only those elements in the summation where measurements are available. If, at time k, no measurement is transmitted (i.e. $I(k) = \emptyset$), (11) is to be understood such that the summation vanishes; that is, $\check{x}(k|k) = \check{x}(k|k-1)$. In order to ease the presentation, this case is not explicitly mentioned hereafter. The filter (10)–(11) and the triggering rule (8) constitute the EBSE.

We assume henceforth that the EBSE is initialized with the same value as the FCSE, i.e. $\check{x}(0) = \hat{x}(0)$. This is a reasonable assumption since we seek to mimic the FCSE with the EBSE.

Notice that the estimator gains L_i in (11) are blocks of the constant matrix L; that is, the entries can be computed off-line. This is different from the approach in Trimpe and D'Andrea [2011a], where a time-varying Kalman filter is used, and the entries of the estimator gain are recomputed at every step k. The approach herein thus has lower computational complexity.

For a fixed sequence $\{I(1), \ldots, I(k)\}$, the filter (10)–(11) is a linear filter. The index tuple $I(k)$ does, however, depend on $y(k)$ by (9); hence, (10)–(11) represent a nonlinear filter. It is a switching observer, whose switching modes correspond to the available measurements at time k.

Notice that any individual mode of the filter may be unstable ((A, C_i) is not necessarily detectable). Moreover, even if the individual modes were stable, this would not imply the stability of the switching observer, as discussed in Lunze [2000] and Böker and Lunze [2002] (Sec. 6.2).

The stability of the presented EBSE follows from the combination of (10)–(11) with the triggering condition (8). This is discussed in detail in the next section.

4. ANALYSIS

We introduce the error measures

$$e(k) := \hat{x}(k) - \check{x}(k) \quad \text{and} \tag{13}$$

$$\check{e}(k) := x(k) - \check{x}(k), \tag{14}$$

which are analyzed below. The error $e(k)$ is the difference between the state estimate of the FCSE and the EBSE. It is required to be bounded according to the problem statement in Sec. 2.2. The error $\check{e}(k)$ is the estimation error of the EBSE (defined analogously to $\hat{e}(k)$ for the FCSE).

Difference of the EBSE to the FCSE. From (5), (10), (11), (12), and (13), we get

$$\begin{aligned} e(k) &= \hat{x}(k) - \check{x}(k) \\ &= A\hat{x}(k-1) + Bu(k-1) - A\check{x}(k-1) - Bu(k-1) \\ &\quad + \sum_{i\in\{1,\dots,N\}} L_i\big(y_i(k) - C_iA\hat{x}(k-1) - C_iBu(k-1)\big) \\ &\quad - \sum_{i\in\{1,\dots,N\}} L_i\big(y_i(k) - C_iA\check{x}(k-1) - C_iBu(k-1)\big) \\ &\quad + \sum_{i\in\bar{I}(k)} L_i\big(y_i(k) - C_iA\check{x}(k-1) - C_iBu(k-1)\big), \end{aligned}$$

where

$$\begin{aligned} \bar{I}(k) &:= (1,\dots,N) \setminus I(k) \\ &= \big(i \mid 1 \le i \le N, \|y_i(k) - C_i\check{x}(k|k-1)\| < \delta_i\big). \end{aligned} \tag{15}$$

Straightforward manipulation then yields

$$\begin{aligned} e(k) &= \Big(A - \sum_{i\in\{1,\dots,N\}} L_iC_iA\Big) e(k-1) \\ &\quad + \sum_{i\in\bar{I}(k)} L_i\big(y_i(k) - C_i\check{x}(k|k-1)\big) \\ &= (I-LC)Ae(k-1) + \sum_{i\in\bar{I}(k)} L_i\big(y_i(k) - C_i\check{x}(k|k-1)\big). \end{aligned} \tag{16}$$

The dynamics of $e(k)$ are those of a stable LTI system ($(I - LC)A$ is stable) with an input which is bounded according to (15). Hence, we have the following theorem.

Theorem 1. Let all eigenvalues of $(I - LC)A$ have magnitude less than one (i.e. the error dynamics of the FCSE are stable). Then, the difference $e(k)$ between the FCSE and the EBSE is bounded for all k. In particular, there exist constants $m > 0$ and $\rho \in [0, 1)$ such that

$$\|e\|_\infty \le \frac{m}{1-\rho} \|L\| \, \|\delta\| =: e_{\max}. \tag{17}$$

Proof. Rewrite (16): for $k \ge 1$,

$$\bar{e}(k+1) = (I - LC)A\bar{e}(k) + L_{\bar{I}(k)}\Delta_{\bar{I}(k)}(k), \tag{18}$$

where $\bar{e}(k+1) := e(k)$; $\Delta_i(k) := y_i(k) - C_i\check{x}(k|k-1)$; $\Delta_{\bar{I}(k)}(k)$ denotes the matrix obtained from consecutively stacking the vectors $\Delta_i(k)$, $i \in \bar{I}(k)$, from top to bottom; and $L_{\bar{I}(k)}$ denotes the matrix from consecutively stacking L_i, $i \in \bar{I}(k)$, from left to right.

First, notice that $\bar{e}(k+1) = (I-LC)A\bar{e}(k)$ is exponentially stable by assumption. Therefore, there exist constants $m > 0$ and $\rho \in [0, 1)$ such that for all $k_0 \in \mathbb{N}$ and $k \ge k_0$,

$$\|\big((I - LC)A\big)^{k-k_0}\| \le m\,\rho^{k-k_0}, \tag{19}$$

[Callier and Desoer, 1991, p. 212–213, Def. 17, Thm. 33].

Recalling the definition of the vector norm (1), one can see that, for $1 \le q < \infty$,

$$\|\Delta_{\bar{I}(k)}(k)\|_q^q = \sum_{i\in\bar{I}(k)} \|\Delta_i(k)\|_q^q \underset{(15)}{<} \sum_{i\in\bar{I}(k)} \delta_i^q \le \sum_{i=1}^{N} \delta_i^q = \|\delta\|_q^q,$$

and, for $q = \infty$,

$$\|\Delta_{\bar{I}(k)}(k)\|_q = \max_{i\in\bar{I}(k)} \|\Delta_i(k)\|_q \underset{(15)}{<} \max_{i\in\bar{I}(k)} \delta_i \le \|\delta\|_q.$$

Hence, for $1 \le q \le \infty$, $\|\Delta_{\bar{I}(k)}(k)\| < \|\delta\|$, and

$$\sup_{k\ge1} \|\Delta_{\bar{I}(k)}(k)\| \le \sup_{k\ge1} \|\delta\| = \|\delta\|.$$

Since also $\|L_{\bar{I}(k)}\| \le \|L\|$, it follows that the input term $L_{\bar{I}(k)}\Delta_{\bar{I}(k)}(k)$ in (18) is bounded. Using these results and applying the bounded trajectories theorem [Callier and Desoer, 1991, p. 218, Thm. 75] then yields

$$\sup_{k\ge1} \|\bar{e}(k)\| \le m\|\bar{e}(1)\| + \frac{m}{1-\rho} \|L\| \, \|\delta\|.$$

Equation (17) follows by $\bar{e}(1) = e(0) = \hat{x}(0) - \check{x}(0) = 0$. □

Notice that the bound (17) holds irrespective of the representation of the disturbances $v(k)$ and $w_i(k)$ in (2)–(4). In particular, it holds for the case where the disturbances are unbounded, such as for Gaussian noise.

Estimator error. The estimation error $\check{e}(k)$ of the EBSE can be written as

$$\check{e}(k) = x(k) - \hat{x}(k) + \hat{x}(k) - \check{x}(k) = \hat{e}(k) + e(k). \tag{20}$$

Therefore, Theorem 1 can be used to deduce properties of the EBSE from properties of the FCSE. For example, if the estimation error $\hat{e}(k)$ of the FCSE is bounded, then by (20) and Theorem 1, the error $\check{e}(k)$ of the EBSE is also bounded. In general, Theorem 1 shows that for $\delta_i \to 0$, $e(k)$ becomes arbitrarily small; that is, the performance of the FCSE is recovered.

5. EXPERIMENTS

The Balancing Cube shown in Fig. 3 serves as the testbed for demonstrating the event-based state estimation method. Six rotating arms on the inner faces of the cube allow the cube to balance on any of its edges or corners. The arms (called *modules*) constitute the agents of the networked control system: each one is equipped with sensors, actuation, and a single-board computer. The computers share data over a Controller Area Network (CAN) bus. For the purpose of this paper, the cube balances on one of its edges. The experimental setup is the same as in Trimpe and D'Andrea [2011a], where more detailed descriptions of the system, the linear model, and the controller can be found.

The eight states of the system model (2) are the angles of the six modules (rotation relative to the cube structure), and the angle and angular rate of the cube about its axis of rotation. Since each module regulates its angular velocity locally with a fast feedback controller, the angular velocities of the six modules are treated as plant inputs. Two types of sensors are used on each module: an absolute encoder measuring the module angle and a rate gyroscope

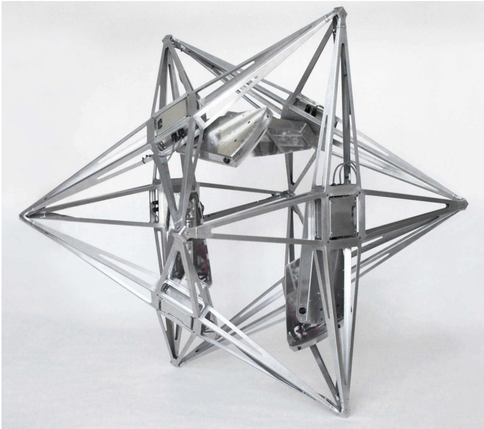

Fig. 3. The experimental testbed: six rotating modules balance a cubic structure on one of its edges. (See www.cube.ethz.ch for a video.)

measuring the angular velocity of the cube. Hence, $N = 12$ sensors are used in total. Each module is able to observe its own angle and the cube states with its local sensors, but the complete system state is not locally observable.

The gain L of the linear observer (5)–(6) is designed as the gain of a steady-state Kalman filter. The resulting eigenvalues of $(I - LC)A$ are 0.936, 0.427, and (with algebraic multiplicity 6) 0.382. For the rate gyro sensors, the transmit threshold $\delta_{\text{gyro}} = 0.004\,\text{rad/s}$ is used, which corresponds to roughly one standard deviation of the sensor noise. For the encoders, $\delta_{\text{enc}} = 0.008\,\text{rad}$ is chosen.

Every module implements a copy of the EBSE (10)–(11) as shown in Fig. 2. In addition to using the estimate $\check{x}(k)$ for the transmit decision in the event generator, the estimator feeds an LQR controller that computes the input to the local actuator (i.e. one of the elements of $u(k)$). Hence, there is feedback from the estimators on the sensor nodes to the system input $u(k)$ (not shown in Fig. 1). A detailed discussion and analysis of the distributed feedback control system is beyond the scope of this paper and a subject for future work.

The sensors are sampled, and the controller commands are updated every 16.6 ms. The control inputs are shared over the CAN bus at every time step, so that $u(k)$ is available to compute (10) on each module. The network bandwidth is sufficient to broadcast all input and measurement data within the duration of one time step. Hence, data transmission is assumed synchronous for all practical purposes.

A truth model state is computed in post-processing from all recorded sensor data, which also includes measurements from multiple accelerometers on the cube (see Trimpe and D'Andrea [2011a] for details). The performance $\mathcal{P}$ of the feedback control system is measured as the root-mean square (RMS) value of the truth model state. By $\mathcal{R}_i(k)$ we denote the average communication rate of measurement y_i. It is computed at time k as the moving average over the last 100 steps. Furthermore, $\bar{\mathcal{R}}_i$ denotes the time average of $\mathcal{R}_i(k)$ over the duration of the experiment, and $\mathcal{R}$ denotes the *average total rate* (the average of $\bar{\mathcal{R}}_i$ over all sensors i). Hence, $\mathcal{R}$ is a measure of the total communication in the network ($\mathcal{R} = 1$ means full communication, $\mathcal{R} = 0$ means no communication).

Table 1. Experimental results.

	$\mathcal{R}$	$\mathcal{P}$
FCSE	1.000	0.183
EBSE	0.221	0.206

Table 2. Average communication rates $\bar{\mathcal{R}}_i$.

Module #	1	2	3	4	5	6
Encoder	0.0001	0.024	0	0	0.022	0
Gyro	0.429	0.406	0.384	0.472	0.453	0.459

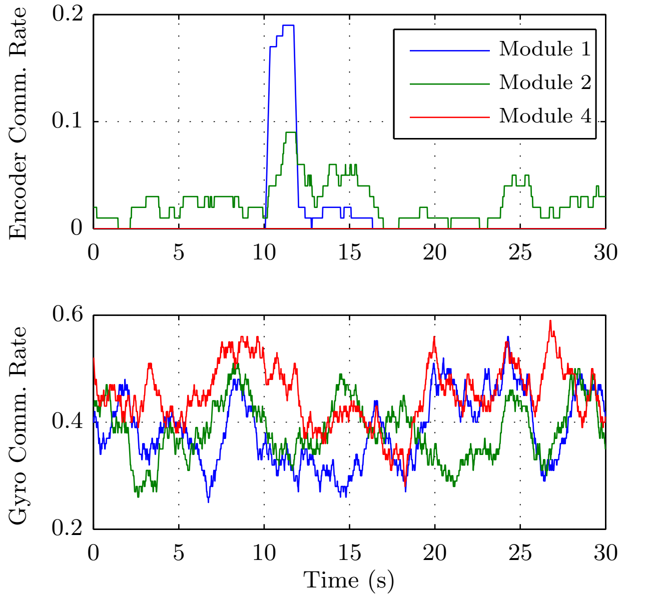

Fig. 4. Experimental communication rates. Roughly at 10 s, Module 1 was displaced by pushing it.

Table 1 shows the communication and performance measures $\mathcal{R}$ and $\mathcal{P}$ for two three-minute balancing experiments: one experiment using the EBSE and another with the FCSE. The results illustrate the expected trade-off between control performance and average communication.

In Table 2, the average communication rates $\bar{\mathcal{R}}_i$ of the individual sensors are given for the same EBSE experiment. They can be interpreted as follows: due to local regulation of the module angular velocities, the module angles can be predicted very accurately from the known velocity input; hence, little communication is required. The gyro sensors, on the other hand, observe the unstable mode of the cube; hence, sufficiently high rates are required for stabilizing the cube. The encoder communication rates of Modules 2 and 5 (on the front and back face of the cube in Fig. 3) are greater than those of the other modules, because Modules 2 and 5 move more during balancing, and their motion is affected more severely by gear backlash in the actuation. Gear backlash is not captured by the linear model (2), and its effect can therefore not be predicted by (10).

In another experiment, Module 1 was manually displaced during balancing (a clutch in the actuation mechanism allows the module to slip when pushed). The communication rates over time are shown in Fig. 4. Clearly, Module 1's encoder rate adapts to the external disturbance. While the module is being pushed, (10) cannot accurately predict the module angle, hence its communication rate goes up.

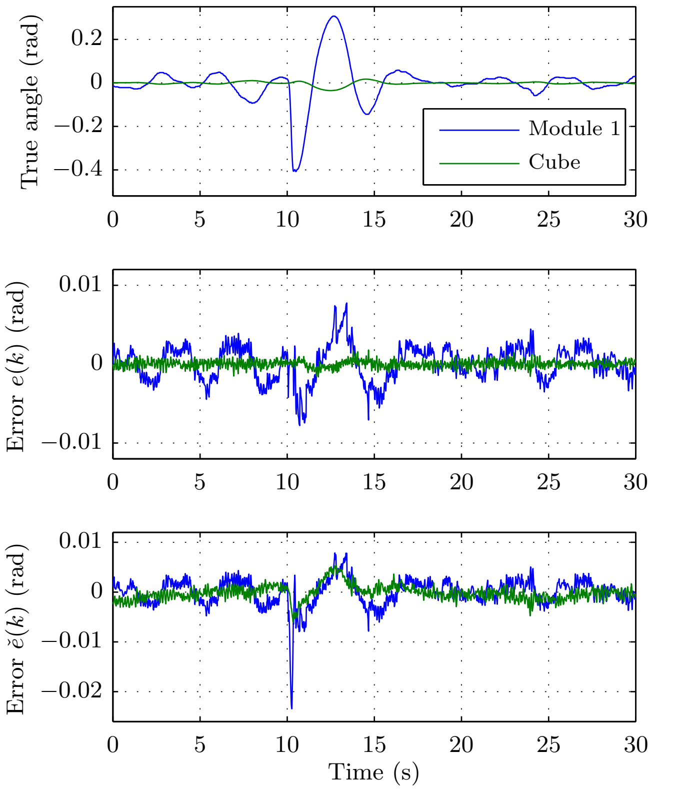

Fig. 5. Experimental estimator performance (same data sequence as in Fig. 4). The estimates of Module 1's angle and the cube angle are shown. Top graph: truth model states; middle: difference between FCSE and EBSE (the error $e(k)$); bottom: EBSE error $\breve{e}(k)$.

For the same experiment, Fig. 5 illustrates the performance of the EBSE exemplarily for Module 1's angle and the cube angle. For $q = \infty$ in (1), the bound in (17) is $e_{\max} = 0.1162$. Clearly, the error signals $e(k)$ in Fig. 5 are well below. The conservatism of the bound (17) stems from the upper bound (19) on the state-transition matrix of the multivariate system (16).

6. CONCLUDING REMARKS

The proposed event-based state estimator for a distributed arrangement of sensors is a direct extension of well-known methods for linear state estimation with synchronous (time-sampled) measurement feedback (such as the Luenberger observer or the steady-state Kalman filter). The approach allows one to trade off estimator performance achievable with the full communication design for communication bandwidth. Experiments on the Balancing Cube illustrated the ability of the event-based control system to discriminate different sensor types and adapt the sensor communication rates to the need for feedback control.

The setup for the experiments on the Balancing Cube is essentially the same as in Trimpe and D'Andrea [2011a], but the event-based state estimator implementation herein is different: a switching Luenberger observer is used instead of a time-varying Kalman filter. Whereas a time-varying filter is computationally more expensive, it potentially has a better performance, since the filter gains adapt on-line to the set of received measurements at a time step. The experimental results (e.g. in Table 1) should, however, not directly be compared with those in Trimpe and D'Andrea [2011a], because different design parameters were chosen. An experimental comparison of the different methods is planned for future work.

Each agent on the Balancing Cube implements a copy of the event-based state estimator and uses the estimate to compute its local control input. To satisfy the assumption that the input vector $u(k)$ is available at all sensor nodes (see Fig. 1), the inputs are shared between the agents over the network. The experiments demonstrated the performance of the event-based control system under realistic conditions, where the individual estimates are not perfectly identical. The stability of the distributed feedback system with non-identical estimators is, however, not analyzed herein. Strategies for removing the requirement of the continuous exchange of the control inputs, as well as the stability analysis of the distributed feedback control system shall be addressed in future work.

ACKNOWLEDGEMENTS

The author would like to thank Raffaello D'Andrea, Jan Lunze, and Christian Stöcker for insightful discussions.

REFERENCES

B. D. O. Anderson and J. B. Moore. *Optimal Filtering*. Dover Publications, Mineola, New York, 2005.

K. J. Åström and B. Wittenmark. *Computer-controlled systems: theory and design*. Prentice Hall, 1997.

D. S. Bernstein. *Matrix mathematics: theory, facts, and formulas with applications to linear system theory*. Princeton University Press, New Jersey, 2005.

G. Böker and J. Lunze. Stability and performance of switching Kalman filters. *International Journal of Control*, 75(16/17):1269–1281, 2002.

F. M. Callier and C. A. Desoer. *Linear System Theory*. Springer-Verlag, 1991.

M. Lemmon. Event-triggered feedback in control, estimation, and optimization. In A. Bemporad, M. Heemels, and M. Johansson, editors, *Networked Control Systems*, volume 406, pages 293–358. Springer-Verlag, 2011.

J. Lunze. Ein Beispiel für den Entwurf schaltender Beobachter. *Automatisierungstechnik*, 48:556–562, 2000.

J. Lunze and D. Lehmann. A state-feedback approach to event-based control. *Automatica*, 46(1):211–215, 2010.

C. Stöcker, J. Lunze, and D. Vey. Stability analysis of interconnected event-based control loops. In *Proc. 4th IFAC Conf. on Analysis and Design of Hybrid Systems*, pages 58–63, Eindhoven, Netherlands, 2012.

S. Trimpe and R. D'Andrea. An experimental demonstration of a distributed and event-based state estimation algorithm. In *Proc. 18th IFAC World Congress*, pages 8811–8818, Milan, Italy, August 2011a.

S. Trimpe and R. D'Andrea. Reduced communication state estimation for control of an unstable networked control system. In *Proc. 50th IEEE Conf. on Decision and Control and European Control Conf.*, pages 2361–2368, Orlando, Florida, USA, 2011b.

J. Weimer, J. Araujo, and K. H. Johansson. Distributed event-triggered estimation in networked systems. In *Proc. 4th IFAC Conf. on Analysis and Design of Hybrid Systems*, pages 178–185, Eindhoven, Netherlands, 2012.