

Reduced Communication State Estimation for Control of an Unstable Networked Control System

Sebastian Trimpe and Raffaello D'Andrea

Abstract—A state estimation method is presented that allows the designer to trade off estimator performance for communication bandwidth in a networked control system. The method is based on a time-varying Kalman filter and a communication decision rule for each sensor: a sensor measurement is transmitted and used to update the Kalman filter if its associated prediction variance exceeds a certain tolerable bound. The resulting equation for the estimation error variance is deterministic, which enables its off-line analysis. If a periodic solution to the variance equation is found, it facilitates a straight-forward implementation of the communication decision: each sensor transmits its measurements with a fixed periodic sequence. This state estimation method is applied in the feedback control system of a cube balancing on one of its edges. Six rotating bodies on the cube stabilize the system and constitute the agents in the networked control system: each one is equipped with local actuation, sensing, and computation, and the agents share their sensor data over a broadcast network. Experimental results compare the performance of the reduced communication state estimation algorithm to a Kalman filter with full measurements.

I. INTRODUCTION

This paper considers the problem of estimating the state of a dynamic system from multiple distributed sensors, while at the same time seeking to reduce the number of sensor measurements that serve as input to the estimator algorithm. While the number and the arrangement of sensors is considered as given, the sensors' transmit rates are variable. Since reducing the set of sensor data generally decreases the estimator performance (provided the reduced sensor data is not defective), a designer would thus be able to trade off between estimator performance and communication bandwidth.

Networked control systems (NCSs) are an example of the described scenario, where transmitting a sensor measurement is associated with a certain cost. In NCSs, a multi-purpose communication network is shared by multiple control, sensor, and actuator units, [1]. Accordingly, a sensor node transmitting its measurement means that the other units cannot use the network without increasing load-induced delays. In wireless sensor networks, reducing the amount of transmitted data often reduces the energy consumption on the sensor nodes, [2].

We approach the problem of reduced communication state estimation by using a (standard) time-varying Kalman filter combined with a constraint on the usage of sensor measurements: a particular sensor measurement is used to

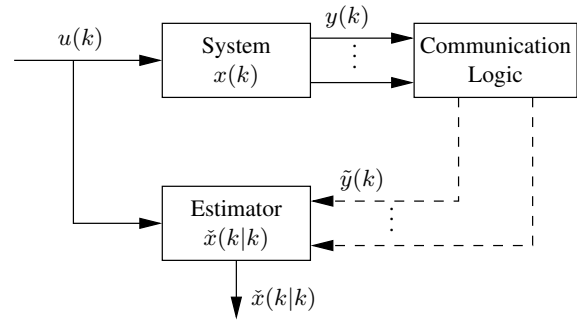


Fig. 1. The design problem: state estimator and communication logic are designed in order to estimate the process state $x(k)$ from a reduced number of measurements. Solid lines denote continuous data flow (i.e. transmission at every discrete time step k), dashed lines denote discontinuous flow of data.

update the estimator if its associated prediction error variance exceeds a certain tolerable bound. Hence, a measurement is only used when it is required to meet a certain estimation performance. The bounds represent tuning parameters that allow one to trade off communication bandwidth for estimation performance. The constraints define transmission rules at each sensor.

The estimator design problem is depicted in Fig. 1. The system state $x(k)$ is estimated at discrete time instants k from a subset of the measurements $y(k)$. A communication logic block selects the subset $\tilde{y}(k)$ from the full measurement vector $y(k)$ and sends the data over a network to the state estimator. We assume an ideal communication network, where the transmission of measurements is instantaneous and no data is lost. The physical representation of the communication logic block may be a sensor with computation capabilities, a network agent (possibly itself running a state estimator), or it may simply represent sensors with different transmit rates. It is assumed that the system inputs $u(k)$ are known to the state estimator.

Actively reducing the transmission of data to a remote estimator is known as *controlled communication*, [1]. Such algorithms have been proposed in [3]–[6], for example. In previous work [6], the sending decision is based on real time measurement data. The method presented here differs in that the sending decision is based on the estimation error variance. Since the variance can be computed off-line, this approach offers a tractable solution. Specifically, if a periodic solution of the estimation variance evolution is found, it corresponds to a periodic sending sequence for each sensor, and hence allows for a straight-forward implementation of

S. Trimpe and R. D'Andrea are with the Institute for Dynamic Systems and Control (IDSC), ETH Zurich, Sonneggstr. 3, 8092 Zurich, Switzerland {strimpe, rdandrea}@ethz.ch.




Fig. 2. The Balancing Cube (edge length 1.2m) is an example of a networked control system: six rotating bodies, each having sensors, actuation, and computational unit, share information over a network to balance the cube on its edge.

the resulting communication logic. The periodic sending sequence can be computed in advance and fixed for each sensor. For two heterogeneous sensors, variance-based and periodic scheduling of sensor transmissions have also been studied in [7].

The state estimator is a time-varying Kalman filter that handles the arrival of measurements at different rates. The proposed technique can therefore also be regarded as a method for designing a multi-rate Kalman filter, where the update rates are determined based on upper bounds on the tolerable error variance. The resulting state estimator switches periodically between different modes defined by the set of measurements available at an update step. Switching or periodic state estimators have been studied, for example, in [8]–[10], and references therein.

The Balancing Cube¹ shown in Fig. 2 and 3 serves as the testbed to demonstrate the presented method. The cube balances on one of its edges through the action of six rotating bodies on its inner faces. The rotating bodies carry a motor, sensors, a computer, and a battery. Their computers share the local sensor data over a Controller Area Network (CAN). The cube therefore represents an example of a networked control system with the rotating bodies being its agents. The reduced communication state estimation scheme of Fig. 1 is applied by implementing a copy of the same state estimator on each agent. Being a broadcast network, the CAN ensures the consistency of the state estimates in the network.

This paper is organized as follows: The equations for the time-varying Kalman filter and associated constraints on the usage of sensor measurements are presented in Sec. II. Section III treats the special case of a periodic solution to the estimation error variance iteration and the obtained simplification of the communication logic. Experimental

¹For a video of the Balancing Cube, please refer to the project website <http://www.cube.ethz.ch>.




Fig. 3. Rendering of the Balancing Cube, shown in the same orientation as in Fig. 2. The cube has six rotating arms, one on each face.

results from the Balancing Cube testbed are given in Sec. IV. The paper concludes with remarks in Sec. V.

II. REDUCED COMMUNICATION STATE ESTIMATOR

We consider the stochastic linear time-invariant system

$$x(k) = Ax(k-1) + Bu(k-1) + v(k-1) \quad (1)$$

$$y(k) = Cx(k) + w(k), \quad (2)$$

where k is the discrete time index; $x(k), v(k) \in \mathbb{R}^n$; $u(k) \in \mathbb{R}^m$; $y(k), w(k) \in \mathbb{R}^p$; and all matrices are of corresponding dimensions. The process noise, the measurement noise, and the initial state $x(0)$ are assumed mutually independent, Gaussian distributed with $v(k) \sim \mathcal{N}(0, Q)$, $w(k) \sim \mathcal{N}(0, R)$, and $x(0) \sim \mathcal{N}(x_0, P_0)$, where $\mathcal{N}(m, V)$ denotes a normally distributed random variable with mean m and covariance matrix V . Furthermore, the pair (A, C) is assumed detectable, (A, Q) stabilizable, and R diagonal. The latter assumption means that the measurement noise is mutually independent for any two sensors considered, which is often the case in practice. The presented state estimation method can, however, be readily extended to the case of block diagonal R by sending blocks of correlated measurements at once.

Throughout this paper j is used to index a single measurement, i.e. an element of the vector y . Accordingly, C_j denotes the j th row of C and R_{jj} the j th diagonal element of R . We use the index set $J(k)$, a subset of $\{1, \dots, p\}$, to denote a selection of measurements at time k . The notation $[C_j]_{j \in J(k)}$ is used to denote the matrix constructed from stacking the rows C_j for all $j \in J(k)$; and $\text{diag}[R_{jj}]_{j \in J(k)}$ denotes the diagonal matrix with entries R_{jj} , for $j \in J(k)$, on its diagonal.

It is well-known that the optimal state estimator for the system (1), (2) with full measurements ($J(k) = \{1, \dots, p\}$) is the Kalman filter, which is restated in Sec. II-A. The constraints on the usage of measurements are set up in Sec. II-B and the corresponding Kalman filter equations for the reduced set of measurements ($J(k) \subseteq \{1, \dots, p\}$) are derived.

A. Kalman Filter with Full Set of Measurements

The Kalman filter recursions for the system (1), (2) can be written as

$$\hat{x}(k|k-1) = A\hat{x}(k-1|k-1) + Bu(k-1) \quad (3)$$

$$P(k|k-1) = AP(k-1|k-1)A^T + Q \quad (4)$$

$$K(k) = P(k|k-1)C^T(CP(k|k-1)C^T + R)^{-1} \quad (5)$$

$$\hat{x}(k|k) = \hat{x}(k|k-1) + K(k)(y(k) - C\hat{x}(k|k-1)) \quad (6)$$

$$P(k|k) = (I - K(k)C)P(k|k-1) \quad (7)$$

with the following meaning of the Kalman filter variables

$$\hat{x}(k|k-1) = \mathbb{E}[x(k)|\mathcal{Y}(k-1), \mathcal{U}(k-1)]$$

$$\hat{x}(k|k) = \mathbb{E}[x(k)|\mathcal{Y}(k), \mathcal{U}(k)]$$

$$P(k|k-1) = \text{Var}[x(k)|\mathcal{Y}(k-1), \mathcal{U}(k-1)]$$

$$P(k|k) = \text{Var}[x(k)|\mathcal{Y}(k), \mathcal{U}(k)],$$

where $\mathbb{E}[\cdot]$ denotes the conditional expected value, $\text{Var}[\cdot]$ the conditional variance, and \mathcal{Y} and \mathcal{U} denote the sets of measurements and inputs up to time k , i.e.

$$\mathcal{Y}(k) = \{y(l) \mid 0 \leq l \leq k\}$$

$$\mathcal{U}(k) = \{u(l) \mid 0 \leq l \leq k\}.$$

The filter is initialized by $\hat{x}(0|0) = x_0$ and $P(0|0) = P_0$. To reflect the fact that for the filter (3)–(7) all sensor measurements $y(k)$ are communicated, it is referred to below as *full communication Kalman filter*.

It is well-known (cf. e.g. [11]) that the time-varying Kalman filter (3)–(7) is the *optimal* state estimator for the considered problem class. Optimality, in this case, means that the Kalman filter keeps track of the entire distribution of the state $x(k)$ conditioned on all measurements and inputs up to time k .

Under the condition that the pair (A, C) is detectable and (A, Q) is stabilizable, the Kalman filter recursion for the prediction variance $P(k|k-1)$ converges to a positive semi-definite matrix \bar{P} , [11], i.e.

$$\lim_{k \rightarrow \infty} P(k|k-1) = \bar{P} \geq 0,$$

which satisfies the discrete algebraic Riccati equation (DARE)

$$\bar{P} = A\bar{P}A^T + Q - A\bar{P}C^T(C\bar{P}C^T + R)^{-1}C\bar{P}A^T. \quad (8)$$

B. Kalman Filter with Reduced Set of Measurements

In order to reduce communication requirements, we now seek a state estimator for the system (1) that receives a time-varying number $\tilde{p}(k) \leq p$ of measurements

$$\tilde{y}(k) = \tilde{C}(k)x(k) + \tilde{w}(k), \quad (9)$$

where $\tilde{y}(k), \tilde{w}(k) \in \mathbb{R}^{\tilde{p}(k)}$, $\tilde{w}(k) \sim \mathcal{N}(0, \tilde{R}(k))$. Notice that $\tilde{y}(k), \tilde{w}(k), \tilde{C}(k) \in \mathbb{R}^{\tilde{p}(k) \times n}$, and $\tilde{R}(k) \in \mathbb{R}^{\tilde{p}(k) \times \tilde{p}(k)}$ have time varying dimensions, which includes the case $\tilde{p}(k) = 0$; that is, at time k there is no measurement available at the estimator. In order to avoid special treatment of this case, we use the convention that $y(k) = \emptyset$, and the measurement

update step in the Kalman filter below is omitted in case no measurement is available at time k .

We next state the Kalman filter equations for the system (1), (9) and then make precise how the measurements (and hence the matrices $\tilde{C}(k)$ and $\tilde{R}(k)$) are selected at each time step.

For any given sequence of $\{\tilde{C}(k)\}_k$ and $\{\tilde{R}(k)\}_k$, the time-varying Kalman filter

$$\tilde{x}(k|k-1) = A\tilde{x}(k-1|k-1) + Bu(k-1) \quad (10)$$

$$\tilde{P}(k|k-1) = A\tilde{P}(k-1|k-1)A^T + Q \quad (11)$$

$$\begin{aligned} \tilde{K}(k) = & \tilde{P}(k|k-1)\tilde{C}^T(k) \\ & \cdot (\tilde{C}(k)\tilde{P}(k|k-1)\tilde{C}^T(k) + \tilde{R}(k))^{-1} \end{aligned} \quad (12)$$

$$\tilde{x}(k|k) = \tilde{x}(k|k-1) + \tilde{K}(k)(y(k) - \tilde{C}(k)\tilde{x}(k|k-1)) \quad (13)$$

$$\tilde{P}(k|k) = (I - \tilde{K}(k)\tilde{C}(k))\tilde{P}(k|k-1) \quad (14)$$

is the optimal estimator for the system (1), (9) (cf. [11]). The estimator keeps track of the state distribution conditioned on all measurements $\tilde{y}(k)$ and inputs up to time k . Hence,

$$\tilde{x}(k|k-1) = \mathbb{E}[x(k)|\tilde{\mathcal{Y}}(k-1), \mathcal{U}(k-1)]$$

$$\tilde{x}(k|k) = \mathbb{E}[x(k)|\tilde{\mathcal{Y}}(k), \mathcal{U}(k)]$$

$$\tilde{P}(k|k-1) = \text{Var}[x(k)|\tilde{\mathcal{Y}}(k-1), \mathcal{U}(k-1)]$$

$$\tilde{P}(k|k) = \text{Var}[x(k)|\tilde{\mathcal{Y}}(k), \mathcal{U}(k)],$$

where $\tilde{\mathcal{Y}}(k)$ denotes the collection of measurements $\tilde{y}(k)$,

$$\tilde{\mathcal{Y}}(k) = \{\tilde{y}(l) \mid 0 \leq l \leq k\}.$$

Among all possible sequences $\{\tilde{C}(k)\}_k$ and $\{\tilde{R}(k)\}_k$, we now wish to choose those that correspond to a reduced set of measurements $J(k) \subseteq \{1, \dots, p\}$. Following the idea outlined in the introduction, the estimator update uses only those measurements whose prediction variance exceeds a certain bound. Since the resulting estimator hence relies on a subset of all measurements $\mathcal{Y}(k)$, its estimation variance is greater than the variance of the full measurement Kalman filter (3)–(7). Therefore, the prediction variance of measurement $y_j(k)$,

$$\text{Var}[y_j(k) | \tilde{\mathcal{Y}}(k-1), \mathcal{U}(k-1)] = C_j\tilde{P}(k|k-1)C_j^T + R_{jj} \quad (15)$$

is compared to its steady-state counterpart of the full measurements filter,

$$\bar{P}_{y_j} := \lim_{k \rightarrow \infty} \text{Var}[y_j(k) | \mathcal{Y}(k-1), \mathcal{U}(k-1)] = C_j\bar{P}C_j^T + R_{jj}. \quad (16)$$

Accordingly, we use the following rule to decide if a single measurement $y_j(k)$ is transmitted for use in the Kalman filter update:

transmit $y_j(k) \Leftrightarrow$

$$\frac{\text{Var}[y_j(k) | \tilde{\mathcal{Y}}(k-1), \mathcal{U}(k-1)] - \bar{P}_{y_j}}{\bar{P}_{y_j}} \geq \delta_j,$$

which, with (15) and (16), simplifies to

$$\text{transmit } y_j(k) \Leftrightarrow C_j(\tilde{P}(k|k-1) - \bar{P})C_j^T \geq \delta_j\bar{P}_{y_j}. \quad (17)$$

The tuning parameters δ_j capture the tolerable normalized deviation of each sensor's measurement prediction variance from the full communication, steady-state variance. For example, $\delta_j = 0$ means no deviation, $\delta_j = 1$ means deviation of \bar{P}_{y_j} , etc. Clearly, if $\delta_j = 0$ for all sensors, the reduced communication Kalman filter (10)–(14) is equivalent to the full communication filter (3)–(7).

Using the transmit rule (17), the index set $J(k)$ of all measurements used in the estimator at time k is

$$J(k) = \{j \mid 0 \leq j \leq p, C_j(\check{P}(k|k-1) - \bar{P})C_j^T \geq \delta_j \bar{P}_{y_j}\}, \quad (18)$$

and the corresponding time-dependent output and measurement noise variance matrices are

$$\check{C}(k) = [C_j]_{j \in J(k)} \quad (19)$$

$$\check{R}(k) = \text{diag}[R_{j_j}]_{j \in J(k)}. \quad (20)$$

The sequences $\{\check{C}(k)\}_k$ and $\{\check{R}(k)\}_k$ are well defined by (18), (19), (20) and knowledge of $\check{P}(k|k-1)$. Together, the equations (11), (12), (14), (18), (19), and (20) provide recursive update equations for obtaining the sequences $\check{P}(k|k-1)$, $\check{P}(k|k)$, $\check{C}(k)$, $\check{R}(k)$ from the problem data (A, C, R, Q, P_0) , and the tuning parameters δ_j . Note that this is fundamentally different from approaches such as [6], where the decision whether to use a measurement in the estimator update depends on the actual measurement data $y(k)$. If the decision depends on real time data, the Kalman filter variables $\check{P}(k|k-1)$ and $\check{P}(k|k)$ become random variables themselves, whereas with the presented method, they can be computed off-line from the problem data.

Since the Kalman filter (10)–(14) is the optimal state estimator for any sequences $\{\check{C}(k)\}_k$ and $\{\check{R}(k)\}_k$, it is also optimal for those sequences given by (18), (19), (20). In other words, given the constraints (17) expressing the objective to use “valuable” measurements only, the Kalman filter (10)–(14) is the optimal state estimator. It is referred to below as *reduced communication Kalman filter*.

To later study the evolution of the estimation error variance, the equations (11), (12), and (14) are combined to

$$\begin{aligned} \check{P}(k+1) &= A\check{P}(k)A^T + Q - A\check{P}(k)\check{C}^T(\check{P}(k)) \\ &\quad \cdot \left(\check{C}(\check{P}(k))\check{P}(k)\check{C}^T(\check{P}(k)) + \check{R}(\check{P}(k)) \right)^{-1} \\ &\quad \cdot \check{C}(\check{P}(k))\check{P}(k)A^T \\ &=: \mathcal{G}(\check{P}(k)), \end{aligned} \quad (21)$$

where the short-hand $\check{P}(k) := \check{P}(k|k-1)$ is used; $\check{C}(\check{P}(k)) := \check{C}(k)$ and $\check{R}(\check{P}(k)) := \check{R}(k)$ have been introduced to emphasize their dependence on $\check{P}(k)$ according to (18), (19), (20); and $\mathcal{G}(\cdot)$ denotes the map of $\check{P}(k)$ to $\check{P}(k+1)$.

III. PERIODIC SOLUTIONS

The Kalman filter derived in Sec. II-B can readily be implemented as a means to manage the communication rate (measured as the number of measurements per time unit) for the problem of Fig. 1. The state estimator block is given by

(10)–(14) and the communication logic by (17). However, in view of the fact that the Kalman filter iteration (21) can be computed off-line, it may be beneficial to analyze the iteration for a given choice of threshold parameters δ_j . If a periodic solution is found, it gives rise to a simplified implementation of the measurement sending decisions.

For (A, C) detectable and (A, Q) stabilizable, the variance of the full communication Kalman filter (3)–(7) converges to the unique solution $\bar{P} \geq 0$ of the the DARE (8). Clearly, it cannot be generally expected that the reduced communication Kalman filter (10)–(14) converges to the same steady-state solution. We illustrate this point with the following example: if A is unstable and δ_j is chosen large, then if $\check{P}(k)$ starts close to \bar{P} , $J(k) = \emptyset$, and $\check{P}(k)$ will grow according to $\check{P}(k+1) = A\check{P}(k)A^T + Q$; hence, \bar{P} is not a solution.

The Kalman filter iteration (21) may, however, have periodic solutions. An example of this is shown in the next section. A periodic solution of the prediction variance corresponds to fixed, periodic sending sequences for the sensors given by (17), and hence provides a straight-forward way for implementing the reduced communication state estimator. The definition of a periodic solution, and an immediate property that is useful for practical implementation, are given in the following:

Definition 1: A symmetric positive definite matrix \check{P} is called a κ -periodic solution to (21) if $\check{P} = \mathcal{G}^\kappa(\check{P})$, where \mathcal{G}^κ denotes the κ times application of \mathcal{G} .

Proposition 1: Let \check{P} be a κ -periodic solution. If $\check{P}(1) = \check{P}$, then $\forall k \check{P}(k) = \mathcal{G}^{\text{mod}(k-1, \kappa)}(\check{P})$.

Proof: From the definition of a κ -periodic solution it follows that $\forall m \in \mathbb{N}, \check{P}(m\kappa + 1) = \check{P}$. Furthermore, $\mathcal{G}^i(\check{P})$ for $i = 0, 1, \dots, \kappa - 1$ are also κ -periodic solutions. Hence, $\forall m \in \mathbb{N}$ and $\forall i \in \{0, 1, \dots, \kappa - 1\}, \check{P}(m\kappa + i + 1) = \mathcal{G}^i(\check{P})$, from which the claim follows. ■

A. Seeking a Periodic Solution

One practical way to find a periodic solution is to simply simulate the Kalman filter iteration (21) initialized with $\check{P}(0) = \bar{P}$, and observe if a periodic solution exists for some κ . This practical approach is pursued in Sec. IV.

The question of existence of a periodic solution is an interesting theoretical question, but is, however, beyond the scope of this paper. Even in the absence of a known periodic solution, one may be able to approximate a switching sequence with a periodic solution and analyze in advance if it performs satisfactorily.

B. State Estimator with Periodic Sending

This section addresses how the communication logic and estimator blocks of Fig. 1 can be implemented when a periodic solution for the estimator variance is known. This implementation is used in the experimental demonstration in Sec. IV.

Communication logic. With a known κ -periodic solution \check{P} , the implementation of the sending decision becomes particularly straight-forward. One simply has to store the

sending sequence $\{\gamma_j(k)\}_k$ for each sensor j over κ steps; that is, defining for $k = 1, \dots, \kappa$,

$$\gamma_j(k) := \begin{cases} 1 & \text{if } C_j(\mathcal{G}^{k-1}(\tilde{P}) - \bar{P})C_j^T \geq \delta_j \bar{P}_{y_j} \\ 0 & \text{otherwise,} \end{cases}$$

and using Proposition 1, the transmit decision (17) becomes

$$\text{transmit } y_j(k) \Leftrightarrow \gamma_j(\text{mod}(k-1, \kappa) + 1) = 1, \quad (22)$$

which is simply a check of a binary condition.

State estimator. The state estimator is given by the Kalman filter (10)–(14), which handles the arrival of varying numbers of measurements. In an ideal network, the variance and gain matrices can be computed off-line. However, in order to cope with non-idealities in physical networks such as imperfect synchronization and communication delays, the estimator node checks which measurements have arrived at every step k , builds the output and measurement noise matrices of output equation (9), and performs the Kalman filter update steps (10)–(14).

IV. EXPERIMENTAL DEMONSTRATION ON THE BALANCING CUBE

We applied the reduced communication state estimator to a networked control system with unstable dynamics. In this section we present the experimental results, and compare the closed-loop performance of the reduced communication estimator to that of the full communication Kalman filter.

The testbed for the estimation algorithm is the Balancing Cube – a dynamic sculpture that can balance on any of its edges or corners through the action of six rotating bodies located on its inner faces, see Fig. 2 and 3. Each rotating body is rigidly mounted to the cube structure and we refer to the body together with its housing as a *module*. Each module is identically equipped with local actuation, sensors, power, and a computation unit. Sensor data can be shared between the modules over a broadcast network. Even though the cube can balance on its corners (as has been shown in [12]), for the purpose of this work, it balances on one of its edges.

The experimental setup is the same as the one presented in [6], and we therefore keep the description of the system in Sec. IV-A to the essentials. Further details and, in particular, an explanation of the feedback controller design (a static gain LQR controller) may be found in the mentioned reference. The design of the reduced communication state estimator is addressed in Sec. IV-B and experimental results of its application on the Balancing Cube are shown in Sec. IV-C.

A. System Description and Linear Model

The active building blocks of the Balancing Cube are the six rotating modules on its faces. Each one is actuated by a DC motor, which tracks velocity commands by a local high gain feedback controller. The angular position of a module relative to the cube body is measured by an absolute encoder. An inertial measurement unit (IMU) with tri-axis accelerometer and tri-axis rate gyroscope (gyro) is mounted on each face of the cube and associated with a module. A single-board computer (SBC) on each module reads data

from the local encoder and IMU and issues commands to the motor. The SBCs exchange data with each other over a Controller Area Network (CAN), whose wires run through slip rings and along the cube structure. The low-level CAN protocols allow each module to broadcast its local measurements to all other modules on the network.

The networked control architecture of the Balancing Cube is shown in Fig. 4. The modules with local actuation, sensing, and computation constitute the agents of the NCS (the terms *module* and *agent* are therefore used synonymously below). The broadcast protocol ensures that all agents receive the same data from the network (if one agent sends data, it is received by *all* other agents).

For the purpose of demonstrating the reduced communication state estimation technique, only two sensors are used per module: the absolute encoder and the rate gyro measurement that is parallel to the axis of rotation of the cube.

A linear discrete-time model of the Balancing Cube with sampling time $T_s = 1/60$ s is given by

$$x(k) = Ax(k-1) + B_1 u(k-1) + B_2 u(k-2) + v(k-1), \quad (23)$$

$$y(k) = Cx(k) + w(k), \quad (24)$$

where $u(k)$ are the velocity commands issued to the motor and originating from a stabilizing feedback controller. The model captures the dynamics of the cube about the equilibrium configuration shown in Fig. 2 and 3. The matrices of the state space model may be found in [6]; the states and outputs of the system are summarized in Table I.

The special structure of the state update equation (23) with the additional delayed input $u(k-2)$ is due to an approximation of the module velocity states as the previously issued velocity commands. This approximation is legitimate due to the high gain inner velocity feedback on the motors, which ensures fast command tracking. It reduces the state dimension, and hence, the complexity of the state estimation problem. The Kalman filter equations (3) and (10) are adapted accordingly by adding the extra input.

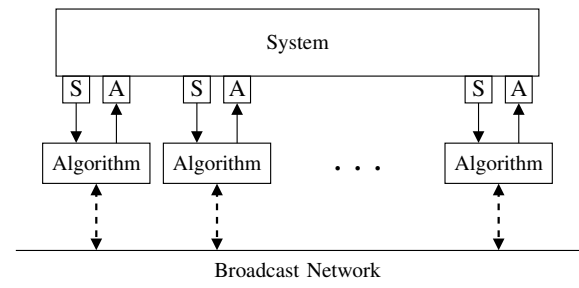


Fig. 4. The networked control architecture testbed: the blocks A and S denote actuator and sensor units; the *Algorithm* block runs estimation and control algorithms as well as the communication logic. An actuator and a sensor unit together with the associated *Algorithm* block are considered as an agent of the NCS. Solid lines denote continuous and dashed lines discontinuous data flow.

TABLE I
STATES AND MEASUREMENTS OF THE BALANCING CUBE MODEL.

state	physical meaning	meas.	sensor
x_1	angle module 1	y_1	encoder module 1
x_2	angle module 2	y_2	rate gyro module 1
x_3	angle module 3	y_3	encoder module 2
x_4	angle module 4	y_4	rate gyro module 2
x_5	angle module 5	y_5	encoder module 3
x_6	angle module 6
x_7	cube angle	y_{11}	encoder module 6
x_8	cube ang. vel.	y_{12}	rate gyro module 6

B. Design and Implementation of the Reduced Communication State Estimator

We first design the full communication Kalman filter (3)–(7) for the system (23), (24). We treat the noise variance matrices as tuning parameters of the estimator. The following values provide acceptable performance in experiments:

$$Q = \text{diag}([1 \ 1 \ 1 \ 1 \ 1 \ 1 \ 0.01 \ 1]) \quad (25)$$

$$R = \text{diag}([0.1 \ 1 \ 0.1 \ 1 \ 0.1 \ 1 \ 0.1 \ 1 \ 0.1 \ 1 \ 0.1 \ 1]). \quad (26)$$

The full communication Kalman filter provides an upper bound on the achievable performance for the reduced communication counterpart; the performance of the two is compared in Sec. IV-C. The solution of the DARE (8) for full communication and the parameters (25) and (26) is²

$$\bar{P} = \begin{bmatrix} 1.09 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1.09 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1.09 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1.09 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1.09 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1.09 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0.72 & 0.13 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0.13 & 1.17 \end{bmatrix}.$$

The only additional tuning parameters required for the reduced communication Kalman filter are the threshold parameters δ_j in (17), which were chosen as $\delta_j = 40$ for the absolute encoder measurements ($j = 1, 3, \dots, 11$), and $\delta_j = 2$ for the rate gyro measurements ($j = 2, 4, \dots, 12$).

In order to extract a periodic solution for the Kalman filter variance iteration, (21) is simulated³ with initial value $\tilde{P}(0) = \bar{P}$. From the simulation data (shown for some diagonal elements of $\tilde{P}(k)$ in Fig. 5), a periodic solution with $\kappa = 50$ can be identified. The fixed point iteration $\tilde{P}(k+1) = \mathcal{G}^{50}(\tilde{P}(k))$ converges to the κ -periodic solution

$$\tilde{P} = \begin{bmatrix} 1.10 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1.10 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1.10 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 4.10 & 0 & 0 & 0 & -0.01 \\ 0 & 0 & 0 & 0 & 1.10 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 4.10 & 0 & 0.01 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0.75 & 0.14 \\ 0 & 0 & 0 & -0.01 & 0 & 0.01 & 0.14 & 1.19 \end{bmatrix}.$$

²For easier reading, the elements of \bar{P} and \tilde{P} (below) are rounded to two decimal places.

³The files to run the simulation may be requested from the first author or downloaded at <http://www.idsc.ethz.ch/Research/DAndrea/Cube/downloads>.




Fig. 5. Simulation result of the reduced communication Kalman filter iteration (21) after 1000 steps. Shown are some diagonal elements of $\tilde{P}(k)$. The solution is periodic with $\kappa = 50$.

The corresponding fixed transmit sequences in (22) are

$$\begin{aligned} \gamma_i(k) &= \mathbf{1}_{50}(k), & \text{for } i = 1, 3, 5, 9 \text{ (encoder)} \\ \gamma_i(k) &= \mathbf{1}_{47}(k), & \text{for } i = 7, 11 \text{ (encoder)} \\ \gamma_i(k) &= \mathbf{1}_5(k) + \mathbf{1}_{10}(k) + \dots + \mathbf{1}_{50}(k), & \text{for } i = 2, 4, \dots, 12 \text{ (gyro)}, \end{aligned} \quad (27)$$

for $k = 1, \dots, 50$ and with

$$\mathbf{1}_{\bar{k}}(k) := \begin{cases} 1 & \text{if } \bar{k} = k \\ 0 & \text{otherwise} \end{cases}.$$

Hence, the absolute encoders transmit their measurements once every 50 steps and the gyros once every five steps.

The reduced communication state estimator is implemented on each module of the Balancing Cube as described in Sec. III-B. Using the decision rule (22), each module decides at every time step whether or not to transmit a measurement for its associated sensors. Furthermore, each module gathers all measurements $\tilde{y}(k)$ that have arrived over the network, constructs the matrices of the corresponding output equation (9), and updates its state estimate according to (10)–(14).

To keep the exposition of the state estimation method simple, local sensor measurements are used to update the estimator subject to the same constraint (22) as is used for

the transmit decision (that is, a local sensor measurement is used to update the estimate if and only if it is transmitted), even though there is no communication cost involved in using the local measurements at every step (this approach is pursued in [6]). Since, furthermore, the network is a broadcast network, all agents have access to the same sensor data $\tilde{y}(k)$ and therefore run a copy of the same state estimator (10)–(14). This also implies that each agent can compute all agents' control inputs $u(k)$, which makes their exchange unnecessary.

C. Experimental Results

The presented state estimator is used in the feedback control system of the Balancing Cube. The control performance and communication rates of the reduced communication estimator are compared in experiments to those of the full communication estimator: their state estimates are used as input to the same state feedback controller. The control objective is the stabilization of the system about the equilibrium $x(k) = 0$.

In order to evaluate the experimental control performance, we use a truth model that is based on the nonlinear state estimation method for the Balancing Cube presented in [12], and augmented with further non-causal post-processing. For this purpose, all sensor data (including, in particular, the accelerometer data) is recorded and the truth model state $x^{\text{truth}}(k)$ is obtained in post-processing. The estimate of the cube tilt obtained from this method has been verified with a camera-based motion capture system (cf. results in [12]) and has proven to work well on the cube.

The same measures for control performance and communication rate are used as in [6]. The performance \mathcal{P} of the control system is measured as the root mean square (RMS) value of the system state,

$$\mathcal{P} := \sqrt{\frac{1}{K} \sum_{k=1}^K (\bar{x}^{\text{truth}}(k))^T \bar{x}^{\text{truth}}(k)},$$

for data of length K and where $\bar{x}^{\text{truth}}(k)$ is the full state vector that also includes the angular velocities of the modules. The *communication rate* $\mathcal{R}_j(k)$ of sensor j is defined as the moving average of transmissions over the last M steps; that is,

$$\mathcal{R}_j(k) := \frac{\text{number of } y_j(k) \text{ transmits in } [k-M+1, k]}{M}.$$

Furthermore, the *total communication rate* \mathcal{R} is defined as

$$\mathcal{R} := \frac{1}{p} \sum_{i=1}^p \left(\frac{1}{K} \sum_{k=1}^K \mathcal{R}_i(k) \right).$$

The communication rates $\mathcal{R}_j(k)$ and \mathcal{R} are in the interval $[0, 1]$ by definition. In particular, $\mathcal{R} = 1$ corresponds to the case where at each time step all data is exchanged between the agents, while $\mathcal{R} = 0$ means no data is exchanged. For $M > \kappa$ the rates of the reduced communication Kalman filter are constant: $\mathcal{R}_j(k) = 0.02$ for the encoder measurements

TABLE II
EXPERIMENTAL COMMUNICATION AND PERFORMANCE MEASURES.

	\mathcal{R}	\mathcal{P}
full communication Kalman filter	1.000	0.2095
reduced communication Kalman filter	0.110	0.3855

($j = 1, 3, \dots, 11$) and $\mathcal{R}_j(k) = 0.2$ for the rate gyro measurements ($j = 2, 4, \dots, 12$).

Experiment: balancing about an equilibrium. The cube was balanced in two separate experiments: one using the full communication Kalman filter, the other using the reduced communication Kalman filter for control. The experimental data presented below originates from module 1 (the other modules' estimates are essentially the same except for small deviations caused by imperfections of the physical communication network such as delays).

The obtained performance \mathcal{P} and total communication rate \mathcal{R} for experimental runs of 2 minutes are shown in Table II. As expected, reducing the number of measurements negatively affects the control performance. However, the performance decrease is less than a factor of 2, while only 11% of the total measurement data was used.

For a 30-second sequence, module 1's estimates of its own angle, module angle 4, and both cube states are shown in Fig. 6 together with the reference state $x^{\text{truth}}(k)$. The corresponding estimation error is given in Fig. 7.

V. CONCLUDING REMARKS

The approach for reducing communication requirements for state estimation in a networked control system presented herein follows the same basic idea as the approach in [6]: a sensor measurement is employed for updating a state estimate (and hence transmitted from sensor to estimator) if it is required to meet a certain estimator performance – or, loosely speaking, if the measurement cannot be predicted well enough by the state estimator. Unlike [6], however, the sensor transmission decision used in this paper is not based on the real time measurement, but on its prediction variance. This has the benefit that, just as for the standard Kalman filter, the variance evolution can be computed and analyzed off-line. The adaptation of the communication requirements to unmodeled, real time events such as external disturbances is, however, not possible. A promising approach is to combine the two approaches by augmenting fixed minimum sensor communication rates with bounds on real time prediction errors.

A periodic solution to the reduced communication Kalman filter corresponds to periodic sensor transmission rates, which gives rise to a very efficient implementation of the sensor's communication logic. The resulting state estimator is a periodic Kalman filter with possibly different sensor arrival rates. This method can therefore also be used as a design tool for a multi-rate Kalman filter, where the required sensor update rates are determined from tolerable bounds on the estimation error variance selected by the designer. The question of the existence of a periodic solution and cor-




Fig. 6. State estimates $\hat{x}(k|k)$ (blue) compared to the truth model state $x^{\text{truth}}(k)$ (red). From top to bottom: angle of module 1, angle of module 4, cube angle, and cube angular velocity.




Fig. 7. Estimation errors $e(k) = x^{\text{truth}}(k) - \hat{x}(k|k)$ for the data shown in Fig. 6. From top to bottom: angle of module 1, angle of module 4, cube angle, and cube angular velocity.

responding analysis of the reduced communication Kalman filter iteration (21) are interesting theoretical problems for future study.

In the example of Sec. IV, some of the sensors transmit at the same time step (cf. equation (27)). In order to reduce the required network capacity, it may be desirable to extend the method to avoid multiple sensors transmitting data at the same time. Regardless, the presented method reduces the *average* communication rate, which allows to use the network for data of different purpose (such as higher level adaptation) when it is not in use for state estimation.

VI. ACKNOWLEDGMENTS

The authors would like to thank the anonymous reviewers for pointing out interesting future directions.

This work was partially funded by the Swiss National Science Foundation (SNSF).

REFERENCES

- [1] J. Hespanha, P. Naghshtabrizi, and Y. Xu, "A survey of recent results in networked control systems," *Proceedings of the IEEE*, vol. 95, no. 1, pp. 138–162, Jan. 2007.
- [2] G. J. Pottie, "Wireless sensor networks," in *Information Theory Workshop*, Jun. 1998, pp. 139–140.

- [3] Y. Xu and J. Hespanha, "Optimal communication logics in networked control systems," in *Proc. of the 43rd IEEE Conference on Decision and Control*, Atlantis, Bahamas, Dec. 2004, pp. 3527–3532.
- [4] —, "Estimation under uncontrolled and controlled communications in networked control systems," in *Proc. of the 44th IEEE Conference on Decision and Control and the European Control Conference*, Seville, Spain, Dec. 2005, pp. 842–847.
- [5] J. Yook, D. Tilbury, and N. Soparkar, "Trading computation for bandwidth: reducing communication in distributed control systems using state estimators," *IEEE Transactions on Control Systems Technology*, vol. 10, no. 4, pp. 503–518, Jul. 2002.
- [6] S. Trimpe and R. D'Andrea, "An experimental demonstration of a distributed and event-based state estimation algorithm," in *Proc. of the 18th IFAC World Congress*, Milano, Italy, Aug. 2011, pp. 8811–8818.
- [7] H. Sandberg, M. Rabi, M. Skoglund, and K. H. Johansson, "Estimation over heterogeneous sensor networks," in *Proc. of the 47th IEEE Conference on Decision and Control*, Cancun, Mexico, Dec. 2008, pp. 4898–4903.
- [8] Y. Liu, "Switching observer design for uncertain nonlinear systems," *IEEE Transactions on Automatic Control*, vol. 42, no. 12, pp. 1699–1703, Dec. 1997.
- [9] G. Böker and J. Lunze, "Stability and performance of switching Kalman filters," *International Journal of Control*, vol. 75, no. 16/17, pp. 1269–1281, Nov. 2002.
- [10] D. S. Flamm, "A new shift-invariant representation for periodic linear systems," *Systems & Control Letters*, vol. 17, no. 1, pp. 9–14, 1991.
- [11] B. Anderson and J. Moore, *Optimal Filtering (Dover Books on Engineering)*. Dover Publications, Jan. 2005.
- [12] S. Trimpe and R. D'Andrea, "Accelerometer-based tilt estimation of a rigid body with only rotational degrees of freedom," in *Proc. of the IEEE International Conference on Robotics and Automation*, Anchorage, Alaska, USA, May 2010, pp. 2630–2636.