2015 IEEE 54th Annual Conference on Decision and Control (CDC)
December 15-18, 2015. Osaka, Japan

Switching data-processing methods for feedback control: Breaking the speed versus accuracy trade-off

E.P. van Horssen, D. Antunes, W.P.M.H. Heemels

***Abstract*— In many digitally controlled applications, such as vision-based and data-intensive control, the choice of the data-processing method for distilling (control-relevant) state information is non-trivial and important. While accurate processing methods require significant processing time, limiting the closed-loop control rate, faster methods introduce larger inaccuracies in the distilled state information. This leads to a trade-off between speed and accuracy, which can be considered off-line or on-line. In this paper, we propose an on-line strategy to switch between different data-processing methods in real-time in order to improve closed-loop performance for linear systems.**

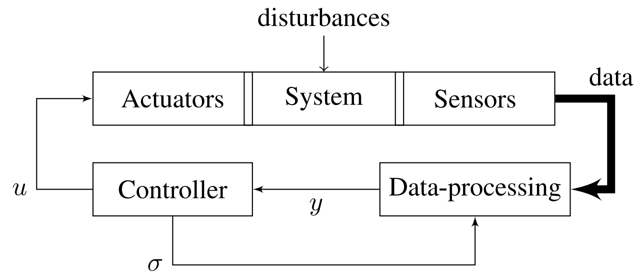

Fig. 1. Control loop with data-processing

I. Introduction

Several control applications require pre-processing of non-trivial amounts of measurement data to distill the information relevant for feedback. A prime example is vision-based control, where one or several images with millions of pixels must be translated into relevant information such as a robot position or attitude. The choice of processing method that is best for control purposes is then an important design feature.

Such data-intensive control applications require the data-processing method not only to provide accurate information, but also to be fast, in order to mitigate delays in the loop and enable a hight closed-loop control rate. These two desired features are typically conflicting. For example, in the context of vision-based control, processing a large number of features leads in general to more accurate information, but requires more processing time.

A standard solution in data-intensive applications is to select a *single* data-processing method with a reasonable trade-off between the accuracy and speed, which maximizes the control performance. This is motivated by the fact that most current control design methods assume *fixed* sensor characteristics, modeled as sensor noise (accuracy) and a measurement delay which limits the fastest sampling period (speed). However, this design choice procedure is always limited by the mentioned speed-accuracy trade-off.

In this paper, we propose to tackle this trade-off by deciding on-line which data-processing method to use in the context of sampled-data control for linear continuous-time plants with stochastic state disturbances [1]. Control-relevant information is obtained through data-processing of measurement data by either a slow method with negligible measurement noise or a fast method with non-negligible measurement noise. The challenge is to decide on-line which data-processing method to use, and which control input to apply to the plant in order to optimize closed-loop performance (see Figure 1). Performance will be measured through quadratic discounted and average cost criteria (compare [2]–[4]). We take a suboptimal approach, since the *co-design* [2], [5], [6] problem is typically combinatorial and computationally intractable in its full generality. The control input is computed based on a combination of a standard Kalman estimator-predictor and an LQR controller for state estimation and actuation (e.g. [7]), respectively. The switching (i.e. the selection of the processing method) is obtained using approximate dynamic programming algorithms, and, in particular, exploits rollout ideas [2], [8]. The works of [9]–[14] consider related scheduling and control problems.

The main result shows that the proposed switching and control policy yields a better performance than using the optimal control policy corresponding to a single data-processing method. The switching condition manifests itself as an ellipsoidal separation of the state space and is therefore easy to implement on-line.

In Section II, the problem formulation is given. Section III explains the proposed methodology, which is formalized in Section IV and the main result is stated. A numerical example in Section V illustrates the benefits of the proposed method for a second-order system.

II. Problem formulation

This section describes the plant, cost criterion, and the measurement and actuation methods used, leading to the problem formulation.

The authors are with the Control System Technology Group, Department of Mechanical Engineering, Eindhoven University of Technology, P.O. Box 513, 5600 MB Eindhoven, The Netherlands, e-mail: {e.p.v.horssen, d.antunes, w.p.m.h.heemels}@tue.nl

This work is supported by the Dutch Technology Foundation (STW) under grant number 12697 "Control based on Data-Intensive Sensing," and by the Innovational Research Incentives Scheme under the VICI grant "Wireless control systems: A new frontier in automation" (No. 11382) awarded by NWO (The Netherlands Organisation for Scientific Research) and STW (Dutch Technology Foundation).

A. Plant and performance criterion

Consider a continuous-time plant modeled by the following stochastic differential equation

$$dx_C = (A_C x_C + B_C u_C)dt + B_\omega d\omega, \;\; x_C(0) = x_0, \;\; t \in \mathbb{R}_{\geq 0}, \tag{1}$$

where $x_C(t) \in \mathbb{R}^{n_x}$ is the state and $u_C(t) \in \mathbb{R}^{n_u}$ is the control input at time $t \in \mathbb{R}_{\geq 0}$, and ω is an n_w-dimensional Wiener process with incremental covariance $I_{n_w} dt$ (cf. [7]). We assume that (A_C, B_C) is controllable and B_C has full rank. The initial condition x_0 is assumed to be known.

We measure the performance using the *discounted* cost given by

$$J_C^d(x_0) := \int_0^\infty \mathbb{E}[e^{-\alpha_C t} g_C(x_C(t), u_C(t))]dt, \tag{2}$$

or the *average* cost

$$J_C^a := \lim_{T\to\infty} \frac{1}{T} \int_0^T \mathbb{E}[g_C(x_C(t), u_C(t))]dt, \tag{3}$$

where $g_C(x,u) := x^\intercal Q_C x + u^\intercal R_C u$, with positive definite matrices Q_C and R_C, and $\alpha_C \in \mathbb{R}_{\geq 0}$.

B. Model of the data processing methods

A digital control platform, which incorporates a data-processing unit and a control unit, is used to measure and control the plant. At *sampling times* $t_\ell, \ell \in \mathbb{N}$, with $t_0 = 0$, a new sample of raw data pertaining to the plant is taken. At this time, a data-processing method must be selected to distill the information relevant for feedback. For simplicity, we consider that only two data-processing methods are available. Furthermore, we assume that the data-processing methods provide an estimate, in general noisy, of the full state $x(t_\ell)$.

The two data-processing methods are modeled in terms of the processing time (speed) and the noise characteristics, i.e. the covariance of the state estimate (accuracy). Method 1 is called the 'slow' method, and requires a time τ_1 to process data. When the slow method is used at time t_ℓ it will make the data $x(t_\ell) + \nu_\ell$ available for feedback at time $t_\ell + \tau_1$, where ν_ℓ is a zero-mean random variable with covariance Φ_1^ν. Method 2 is called the 'fast' method and requires a time τ_2 to process data. It makes the data $x(t_\ell) + \nu_\ell$ available for feedback at time $t_\ell + \tau_2$, where ν_ℓ is a zero-mean random variable with covariance Φ_2^ν. The fact that Method 1 requires more time to process the raw data but is more accurate than Method 2 is captured by assuming

$$\tau_1 > \tau_2, \quad \Phi_1^\nu \prec \Phi_2^\nu.$$

For simplicity, we assume that $\Phi_1^\nu = 0$.

Immediately after completion of the data-processing, new raw data can be obtained, and processed by selecting one of the two available methods. Let $\sigma_\ell \in \{1,2\}, \ell \in \mathbb{N}$, indicate the *processing method* selected at time t_ℓ. Assuming that other delays, e.g. due to the selection process or raw data acquisition, are negligible, we have

$$t_{\ell+1} = t_\ell + \tau_{\sigma_\ell}. \tag{4}$$

We denote the processed measurement by

$$y_{\ell+1} = x_C(t_\ell) + \nu_\ell, \quad \ell \in \mathbb{N}. \tag{5}$$

The noise variables ν_ℓ, $\ell \in \mathbb{N}$, are assumed to be zero-mean random with covariance $\Phi_{\sigma_\ell}^\nu$, and independent and identically distributed for all $\ell \in \mathbb{N}$.

We will assume that Method 1 used alone (i.e. $\sigma_\ell = 1$, for all $\ell \in \mathbb{N}$) would lead to a better performance in terms of criteria (2) or (3) than applying Method 2 all the time (i.e. $\sigma_\ell = 2$, for all $\ell \in \mathbb{N}$). The case when using Method 2 all the time would lead to a better performance is briefly discussed at the end of Section IV in Remark 2.

C. Problem statement

Since measurements become available for feedback at times t_ℓ (pertaining to raw data acquired at time $t_{\ell-1}$), it is reasonable to assume that control updates based on this new information also occur only at times t_ℓ. A zero-order hold (ZOH) actuation strategy is implemented such that the actuation signal is held constant between sampling times. Hence, u_C is a staircase signal. We have

$$u_C(t) = u_C(t_\ell), \quad \text{for all } t \in [t_\ell, t_{\ell+1}). \tag{6}$$

Moreover, also at times t_ℓ, new raw data is acquired and a switching decision must be taken pertaining to which data processing method to use in the interval $[t_\ell, t_{\ell+1})$.

As a result, we can state the problem we are interested in as follows: Find a *switching and control policy* (co-design problem), i.e., a sequence $\pi = (\mu_0, \mu_1, \ldots)$ of multivariate functions $\mu_\ell := (\mu_\ell^u, \mu_\ell^\sigma)$, for $\mu_\ell^\sigma : (\{1,2\} \times \mathbb{R}^{n_u} \times \mathbb{R}^{n_x})^\ell \times \mathbb{R}^{n_x} \to \{1,2\}$ and $\mu_\ell^u : (\{1,2\} \times \mathbb{R}^{n_u} \times \mathbb{R}^{n_x})^\ell \times \mathbb{R}^{n_x} \to \mathbb{R}^{n_u}$, that determine the switching and actuation inputs at sampling times t_ℓ

$$(\sigma_\ell, u_C(t_\ell)) = \mu_\ell(\mathsf{I}_\ell), \quad \ell \in \mathbb{N}, \tag{7}$$

based on the information vector

$$\mathsf{I}_\ell := (\mathsf{I}_{\ell-1}, \sigma_{\ell-1}, u_C(t_{\ell-1}), y_\ell), \quad \ell \in \mathbb{N}, \quad \mathsf{I}_0 := (x_0),$$

available to the control platform at decision times t_ℓ.

III. Proposed Switching and Control Policy

In this section, we start by providing an overview of the main idea behind the proposed switching and control policy. This will entail (i) designing a state estimator; (ii) specifying the control input; and (iii) specifying the switching policy. Each of these topics is addressed subsequently.

A. Overview of the method

From the problem formulation, we see that a given policy may result in decision times unknown at the initial time. Our policy is to fix a priori the switching *decision times* to be

$$s_k = k\tau_1, \quad k \in \mathbb{N}$$

and assume that the processing times are related by an integer number, i.e.

$$\frac{\tau_1}{\tau_2} = n_\tau, \quad n_\tau \in \mathbb{N}_{>1}. \tag{8}$$

This can always be achieved by tuning the processing methods or adding small waiting times.

At each time $s_k, k \in \mathbb{N}$, our policy decides either to use the slow method or use n_τ times the fast method. Selecting the slow method means that the control input at time s_k is held constant in the interval $[s_k, s_{k+1})$, but an accurate measurement will become available at time s_{k+1}. Selecting to use n_τ times the fast method means that the control input is updated more frequently, at times $s_k + h\tau_2$, $h \in \{0, 1, \dots, n_\tau - 1\}$, based on the n_τ corresponding measurements. Note that the measurement available at time s_{k+1} is then inaccurate (noisy).

The switching decision at time s_k is *conceptually* defined as the one that would optimize the expected quadratic performance cost (2) or (3) while supposing that at times $s_{k+1}, s_{k+2}, \dots$ the slow method would be picked and the standard optimal LQG control input policy would be used. This can be seen as a rollout method, in the context of dynamic programming algorithms [8] (for which the base policy is to select the slow method after the optimization horizon), and is repeated in a receding horizon fashion, i.e., the same procedure is repeated at time s_{k+1}, s_{k+2}, $\dots$.

B. State-estimation

The state and (delayed) measurements at discrete times $t_\ell, \ell \in \mathbb{N}$, are described by

$$\xi_{\ell+1} = A_{\sigma_\ell}\xi_\ell + B_{\sigma_\ell}u_\ell + \omega_\ell, \quad \ell \in \mathbb{N}, \tag{9}$$

$$y_{\ell+1} = \xi_\ell + \nu_\ell, \tag{10}$$

where $\xi_\ell := x_C(t_\ell) \in \mathbb{R}^{n_x}$ and $u_\ell := u_C(t_\ell) \in \mathbb{R}^{n_u}$ are the state and control input, respectively, at discrete time $\ell \in \mathbb{N}$. Moreover,

$$A_m := e^{A_C\tau_m}, \quad B_m := \int_0^{\tau_m} e^{A_C s}B_C ds, \quad m \in \{1, 2\}.$$

We assume that the sampling periods (processing times) τ_1 and τ_2 are non-pathological, which implies that (A_1, B_1) and (A_2, B_2) are controllable[1]. The disturbances ω and ν are sequences of zero-mean independent random vectors, $\omega_\ell \in \mathbb{R}^{n_\omega}$ and $\nu_\ell \in \mathbb{R}^{n_\nu}$, respectively, with covariances $\mathbb{E}[\omega_\ell(\omega_\ell)^\intercal] = \Phi^\omega_{\sigma_\ell}$ and $\mathbb{E}[\nu_\ell(\nu_\ell)^\intercal] = \Phi^\nu_{\sigma_\ell}$, for all $\ell \in \mathbb{N}$, with $\Phi^\omega_m := \int_0^{\tau_m} e^{A_C s}B_\omega B_\omega^\intercal e^{A_C^\intercal s} ds$, $m \in \{1, 2\}$.

As a result, a predicted estimate $\bar{\xi}_\ell$ of the state ξ_ℓ from the information available until and including time t_ℓ can be given by

$$\bar{\xi}_\ell = A_{\sigma(\ell-1)}\bar{\xi}_{\ell-1} + B_{\sigma(\ell-1)}u_{\ell-1} + L_{\ell-1}(y_\ell - \bar{\xi}_{\ell-1}), \tag{11}$$

using the time-varying Kalman filter which initial condition chosen as $\bar{\xi}_0 = \xi_0 = x_0$. The gains L_ℓ, for $\ell \in \mathbb{N}$, of the estimator-predictor can be computed [7] by

$$L_\ell = A_{\sigma_\ell}\Theta_\ell(\Theta_\ell + \Phi^\nu_{\sigma_\ell})^{-1} \tag{12}$$

$$\Theta_{\ell+1} = (A_{\sigma_\ell} - L_\ell)\Theta_\ell(A_{\sigma_\ell} - L_\ell)^\intercal + L_\ell\Phi^\nu_{\sigma_\ell}L_\ell^\intercal + \Phi^\omega_{\sigma_\ell}, \tag{13}$$

[1]The sampling period τ_m is non-pathological if A_C does not have two eigenvalues with equal real parts and imaginary parts that differ by an integer multiple of $\frac{2\pi}{\tau_m}$ (cf. [1, p. 45]).

where Θ_ℓ is the covariance of the state estimate with initial condition $\Theta_0 = 0$. For a fixed switching policy, the infinite-horizon solution $\overline{\Theta}_m$, $m \in \{1, 2\}$, is the stationary solution to the Ricatti equation (13) when $\sigma_\ell = m$, for all $\ell \in \mathbb{N}$, i.e. it is the solution to corresponding the discrete-time algebraic Riccati equation (DARE). The stationary gain $\overline{L}_m, m \in \{1, 2\}$ is then given by (12).

C. Policy for the control input

We show in this subsection that the feedback policy for the proposed rollout method is given by

$$u_\ell = -K_\ell\bar{\xi}_\ell, \tag{14}$$

where the control gains K_ℓ for the period k (i.e. for $t_\ell \in [s_k, s_{k+1})$) are determined at the decision time s_k, independent of the measurements. For the slow method this corresponds to one gain, whereas for the fast method this corresponds to multiple gains.

To provide the expressions for the gains K_ℓ, we start by discretizing the cost function (2) with a *fixed sampling period* τ_m, $m \in \{1, 2\}$, corresponding to either the slow or the fast method. We focus for now on the discounted cost (2) and discuss the average cost in Remark 1. It can be shown (using for example the arguments in [7]) that, for $\xi_0 = x_0$, this discounted cost (2), apart from an additive constant factor, is given by

$$J^d_m(\xi_0) := \sum_{l=0}^{\infty} \mathbb{E}[\alpha^l_m g(\xi_l, u_l, m)], \quad m \in \{1, 2\}, \tag{15}$$

where $\alpha_m := e^{-\alpha_C\tau_m}$, $g(\xi, u, m) := \xi^\intercal Q_m\xi + 2\xi^\intercal S_m u + u^\intercal R_m u$, and, where

$$\begin{bmatrix} Q_m & S_m \\ (\star)^\intercal & R_m \end{bmatrix} := \int_0^{\tau_m} e^{-\alpha_C s} e^{\begin{bmatrix} A_C & B_C \\ 0 & 0 \end{bmatrix}^\intercal s} \begin{bmatrix} Q_C & 0 \\ 0 & R_C \end{bmatrix} e^{\begin{bmatrix} A_C & B_C \\ 0 & 0 \end{bmatrix} s} ds.$$

In accordance with the proposed rollout algorithm (see Section III-A), the control input obtained in the interval $[s_k, s_{k+1})$ results from the choice O of one of two options $O \in \{opt1, opt2\}$ at time s_k.

In the first option, $O = opt1$, the slow method is selected at time s_k supposing that it will also be selected at times $s_{k+1}, s_{k+2} \dots$. The discounted cost in this case is simply given by (15) when $m = 1$, which is then to be optimized by the control inputs. Thus, the control inputs follow from the standard optimal control policy for the discrete-time LQG problem resulting from minimizing (15) for (9) when $m = 1$ and $\sigma_\ell = 1$ for every $\ell \in \mathbb{N}$ (taken into account our assumption on controllability and positive definiteness for the cost). The control input of this policy is given by (see, e.g., [7], [8])

$$u_\ell = -\bar{K}\bar{\xi}_\ell,$$

where

$$\bar{K} = \bar{G}^{-1}(B_1^\intercal\bar{P}A_1 + (S_1)^\intercal)$$

and $\bar{P}$ and $\bar{G}$ result from the solution to the algebraic Riccati equation

$$\begin{aligned} \bar{P} &= \alpha_1 A_1^\intercal\bar{P}A_1 + Q_1 - \bar{K}^\intercal\bar{G}\bar{K} \\ \bar{G} &= (R_1 + \alpha_1 B_1^\intercal\bar{P}B_1). \end{aligned} \tag{16}$$

In the second option, $O = opt2$, the fast method is selected at the n_τ times $s_k, s_k+\tau_2, \dots, s_k+(n_\tau-1)\tau_2$ assuming that the slow method will be selected at times $s_{k+1}, s_{k+2} \dots$. The optimal control policy can still be derived from standard LQG arguments for time-varying systems, since the scheduling decisions are fixed. At times $s_k, s_k+\tau_2, \dots, s_k+(n_\tau-1)\tau_2$ this policy is given by

$$u_{\ell+h} = -\underline{K}_h \bar{\xi}_{\ell+h}, \quad h \in \{0, \dots, n_\tau - 1\},$$

where the control gains $\underline{K}_h$ are given by the solution to

$$\underline{G}_h = R_2 + \alpha_2 B_2^\intercal \underline{P}_{h+1} B_2 \tag{17}$$
$$\underline{K}_h = (\underline{G}_h)^{-1}(\alpha_2 B_2^\intercal \underline{P}_{h+1} A_2 + (S_2)^\intercal) \tag{18}$$
$$\underline{P}_h = (Q_2 + \alpha_2 A_2^\intercal \underline{P}_{h+1} A_2) - \underline{K}_h^\intercal \underline{G}_h \underline{K}_h \tag{19}$$

solved backward, as in dynamic programming, for $h \in \{0, \dots, n_\tau - 1\}$ starting from $\underline{P}_{n_\tau} = \bar{P}$, where $\bar{P}$ is the solution to (16).

The control input policy, in continuous-time, can be written in the interval $t \in [s_k, s_{k+1})$ as

$$u_C(t) = \begin{cases} -\bar{K}\bar{\xi}_\ell, & \text{if } \sigma_\ell = 1 \text{ for } t_\ell = s_k, \\ -\underline{K}_h \bar{\xi}_{\ell+h}, & \text{if } \sigma_\ell = 2 \text{ for } t_\ell = s_k \text{ and} \\ & \quad t \in [s_k + h\tau_2, s_k + (h+1)\tau_2), \end{cases} \tag{20}$$

where $h \in \{0, \dots, n_\tau - 1\}$. The choice of the switching signal σ_ℓ is formalized in the next section.

D. Switching policy

According to the proposed rollout method, the data processing method to be selected at time s_k is the one which would minimize the quadratic discounted cost if that method were to be applied in the interval $[s_k, s_{k+1})$, and that the slow method would be selected afterwards as a rollout 'base' policy.

Let $J_O^d(\xi_\ell)$ denote the cost-to-go according to (2) from state ξ_ℓ at time t_ℓ for switching choice $O \in \{opt1, opt2\}$. Since the state ξ_ℓ is not available at time t_ℓ, our switching decision is based on the expectation of the cost-to-go[2] with respect to the available information, i.e. expressed in terms of state estimate $\bar{\xi}_\ell$ and its variance Θ_ℓ (which are sufficient statistics [8, Ch.5]), for each option $O \in \{opt1, opt2\}$, denoted as

$$\bar{J}_O^d(\bar{\xi}_\ell, \Theta_\ell) := \mathbb{E}\left[\, J_O^d(\xi_\ell) \mid \mathsf{I}_\ell \,\right], \tag{21}$$

at each decision time, i.e. when $t_\ell = s_k$ for any $k \in \mathbb{N}$. This is then repeated for every time $s_k, k \in \mathbb{N}$ in a receding horizon fashion.

For the first option, the cost-to-go (21) is now given by

$$\bar{J}_{opt1}^d(\bar{\xi}_\ell, \Theta_\ell) = \bar{\xi}_\ell^\intercal \bar{P} \bar{\xi}_\ell + \chi_{opt1}(\Theta_\ell), \tag{22}$$

where

$$\begin{aligned}\chi_{opt1}(\Theta_\ell) &= tr(\Theta_0^k(\bar{P} + \bar{K}^\intercal \bar{G} \bar{K})) \\ &\quad + \tfrac{\alpha_1}{1-\alpha_1} tr(\Phi_1^\omega \bar{P} + \bar{\Theta}_1 \bar{K}^\intercal \bar{G} \bar{K})\end{aligned} \tag{23}$$

[2]The discount up to time t_ℓ does not affect the decision and is omitted.

with $\Theta_0^k = \Theta_\ell$.

For the second option, the cost-to-go (21) at decision time $t_\ell = s_k$ is given by

$$\bar{J}_{opt2}^d(\bar{\xi}_\ell, \Theta_\ell) = \bar{\xi}_\ell^\intercal \underline{P}_0 \bar{\xi}_\ell + \chi_{opt2}(\Theta_\ell), \tag{24}$$

where

$$\begin{aligned}\chi_{opt2}(\Theta_\ell) &= tr(\Theta_0^k(\underline{P}_0 + \underline{K}_0^\intercal \underline{G}_0 \underline{K}_0)) \\ &\quad + \sum_{h=1}^{n_\tau - 1} [\alpha_2^h tr(\Phi_2^\omega \underline{P}_h + \Theta_h^k \underline{K}_h^\intercal \underline{G}_h \underline{K}_h)] \\ &\quad + \alpha_1 tr(\Phi_1^\omega \bar{P} + \Theta_{n_\tau}^k \bar{K}^\intercal \bar{G} \bar{K}) \\ &\quad + \tfrac{\alpha_1^2}{1-\alpha_1} tr(\Phi_1^\omega \bar{P} + \bar{\Theta}_1 \bar{K}^\intercal \bar{G} \bar{K}),\end{aligned} \tag{25}$$

and where Θ_h^k for $h \in \{0, \dots, n_\tau\}$ follows from (13) with initial condition $\Theta_0^k = \Theta_\ell$.

For $opt1$ and $opt2$ the costs are given by (22) and (24), respectively. The switching signal, in accordance with Section III-A, is then given by

$$\sigma_\ell = \overline{\sigma}_k \quad \text{for all } t_\ell \in [s_k, s_{k+1}), \tag{26}$$

where the proposed switching decisions $\overline{\sigma}_k \in \{1, 2\}$, for times t_ℓ which correspond to decision times s_k, i.e. at $t_\ell = s_k$, are given by

$$\begin{aligned}\overline{\sigma}_k &:= \arg \min_{O \in \{opt1, opt2\}} \bar{J}_O^d(\bar{\xi}_\ell, \Theta_\ell) \\ &= \arg \min_{O \in \{opt1, opt2\}} \bar{\xi}_\ell^\intercal \Pi_O \bar{\xi}_\ell + \chi_O(\Theta_\ell) \\ &= \arg \min_{O \in \{opt1, opt2\}} \bar{\xi}_\ell^\intercal \Pi_O \bar{\xi}_\ell + \eta_O(\Theta_\ell),\end{aligned} \tag{27}$$

with

$$\Pi_{opt1} := \bar{P}, \quad \Pi_{opt2} := \underline{P}_0,$$

and

$$\begin{aligned}\eta_{opt1}(\Theta_\ell) &= tr(\Theta_0^k(\bar{P} + \bar{K}^\intercal \bar{G} \bar{K})) \\ &\quad + \alpha_1 tr(\Phi_1^\omega \bar{P} + \bar{\Theta}_1 \bar{K}^\intercal \bar{G} \bar{K}), \\ \eta_{opt2}(\Theta_\ell) &= tr(\Theta_0^k(\underline{P}_0 + \underline{K}_0^\intercal \underline{G}_0 \underline{K}_0)) \\ &\quad + \sum_{h=1}^{n_\tau - 1} [\alpha_2^h tr(\Phi_2^\omega \underline{P}_h + \Theta_h^k \underline{K}_h^\intercal \underline{G}_h \underline{K}_h)] \\ &\quad + \alpha_1 tr(\Phi_1^\omega \bar{P} + \Theta_{n_\tau}^k \bar{K}^\intercal \bar{G} \bar{K}),\end{aligned}$$

the unequal parts of $\chi_{opt1}(\Theta_\ell)$ and $\chi_{opt1}(\Theta_\ell)$, respectively.

Remark 1: (Average cost) The average cost (3) can be described in terms of the discretized system for a fixed sampling period τ_m, $m \in \{1, 2\}$, as follows

$$J_m^a := \lim_{\kappa \to \infty} \frac{1}{\tau_m \kappa} \sum_{k=0}^{\kappa - 1} \mathbb{E}[g(\xi_k, u_k, m)]. \tag{28}$$

and the optimal cost-to-go for fixed data processing methods $m \in \{1, 2\}$ is given by

$$J_m^a = \tfrac{1}{\tau_m} tr(\Phi_m^\omega \bar{P}_m + \bar{\Theta}_m \bar{K}_m^\intercal \bar{G}_m \bar{K}_m), \tag{29}$$

using the stationary solutions to the Ricatti equations (13) and (19) for the system discretized with period τ_m. Note that our proposed policy (11), (20), (27) is also valid for the case where $\alpha_C = 0$, and consequently $\alpha_m = 1$, $m \in \{1, 2\}$. We shall consider this policy in the average cost case.

IV. Performance analysis

We are now in the position to present our main result. Let $J_\pi^d(\xi_0)$, J_π^a denote the discounted cost (2) and average cost (3), respectively, when the proposed policy, characterized by (11), (20), (27) is applied to the plant (1).

Theorem 1: Consider the plant (1) with two possible sensor data-processing methods providing measurements (5) at times (4), and a digital controller, which estimates the state at sampling times $t_\ell, \ell \in \mathbb{N}$, according to (11), provides the control input according to (20), and selects on-line the data processing method at times $s_k = k\tau_1$, $k \in \mathbb{N}$ according to (27). Then

$$J_\pi^d(\xi_0) \leq J_1^d(\xi_0), \quad \text{for all } \xi_0 \in \mathbb{R}^n,$$

and

$$J_\pi^a \leq J_1^a.$$

Moreover,

$$\Pi_2 \preceq \Pi_1,$$

and therefore the switching policy (27), at switching decision time $t_\ell = s_k$, can be written as

$$\overline{\sigma}_k = \begin{cases} 1, & \text{if } \overline{\xi}_\ell^\mathsf{T} \Omega \overline{\xi}_\ell \leq \eta_2(\Theta_\ell) - \eta_1(\Theta_\ell), \\ 2, & \text{otherwise,} \end{cases}$$

with constant positive semidefinite matrix $\Omega := \Pi_1 - \Pi_2$. □

The theorem states that both the discounted and average costs of the proposed policy are not larger than the corresponding costs of applying the slow method at every time step. The proof is omitted for brevity.

The fact that $\Pi_2 \preceq \Pi_1$ implies that the fast data processing method is selected if the state estimate lies outside an ellipsoid in the state space, and the slow method is selected otherwise. Note that if $\overline{\sigma}_k = 1$, then $\Theta_\ell = \Phi_1^\omega$ at $t_\ell = s_{k+1}$, i.e. the variance (and therefore the ellipsoid) 'resets'.

Remark 2: In this paper, we have assumed that the fixed policy for Method 1 is inherently better than the fixed policy for Method 2. Therefore, we selected Method 1 as the base policy. If Method 2 would be better, we would choose Method 2 as the base policy. The proposed methodology can be adapted to account for this.

Remark 3: Note that our policy can also be used if x_0 is a random variable by taking $\overline{\xi}_0 = \mathbb{E}[x_0]$ and $\Theta_0 = \mathbb{E}[x_0 x_0^\mathsf{T}]$.

V. Simulation

Here, we show simulation results for a second-order system. In particular, we consider the system (1) given by

$$A_C = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}, \; B_C = \begin{bmatrix} 0 \\ 1 \end{bmatrix}, \; B_\omega = \begin{bmatrix} 0.81 \\ 0.30 \end{bmatrix}, \tag{30}$$

and a cost given by (3) and

$$Q_C = \begin{bmatrix} 0.1 & 0 \\ 0 & 10 \end{bmatrix}, \; R_C = 1. \tag{31}$$

We define the base sampling time to be $\tau_s = 1$, and $n_\tau = 3$ for the sampling time ratio in (8). Hence, the system is discretized with two sampling periods $\tau_1 = \tau_s$ and $\tau_2 = \tau_s/3$ to get (9)-(10), which correspond to two data-processing algorithms {Alg1, Alg2}. The measurement noise on Alg2 is given by

$$\Phi_2^\nu = \begin{bmatrix} 9 & 0 \\ 0 & 9 \end{bmatrix}. \tag{32}$$

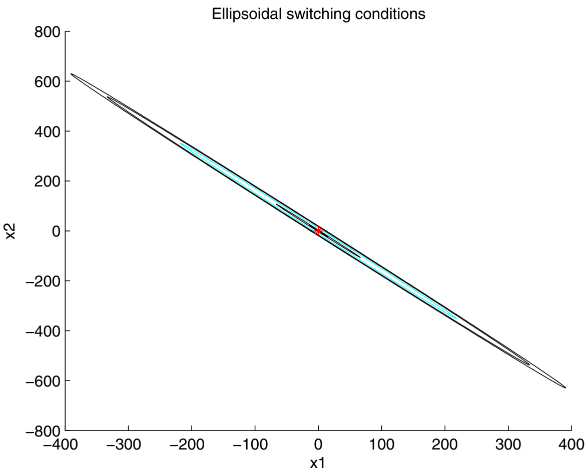

Fig. 2. Ellipsoid conditions

The expected average costs of the fixed policies can be computed (from (29)) to be $J_1^a = 3.1319 \cdot 10^3$ and $J_2^a = 3.1774 \cdot 10^3$, i.e. Alg2 performs 1.45% worse. This shows that persistent use of Alg2 would perform worse than using only Alg1, which is consistent with the assumptions made in our problem set-up.

Figure 2 shows ellipsoids representing the switching condition, based on the average cost, for various values of the covariance Θ_l of the state estimate. The outer black ellipsoid is for the case $\Theta_\ell = \Phi_1^\omega$ at $s_k = t_\ell$, i.e., the previous measurement was taken using Alg1. The inner black ellipsoids are further iterations if Alg2 is used multiple times consecutively. When Alg1 is used, the ellipsoid resets to the outer black ellipsoid. The cyan-colored ellipsoid is generated from $\Theta_\ell = \bar{\Theta}_2$, i.e. the non-switching solution for Alg2. For a given system, the initial ellipsoid results from $\Theta_0 = 0$.

We run 10 Monte Carlo simulations of the system. We use a small initial condition near the origin $x_0 \approx [0, 0]^\mathsf{T}$. We take $k \in \mathbb{N}_{[0,10000]}$ decision moments, since we are interested in the average cost.

In Figure 3, the state (green), and the estimate of the state (red/blue) are shown for one simulation. The color indicates the switching decision, red is Alg2 and blue is Alg1. The black lines are part of the outer ellipsoid in Figure 2. For this case, from Figure 3, we see that outside the black ellipsoid Alg1 is never selected. Also, occasionally, it happens that Alg2 is chosen inside the ellipsoid due to the scaling of the ellipsoids, thereby explaining the occurrence of red circles inside the region with mainly blue markers. The green markers indicate the true state at the decision moments.

The disturbances occasionally push the state outside the region where using only Alg1 is optimal, i.e. outside the

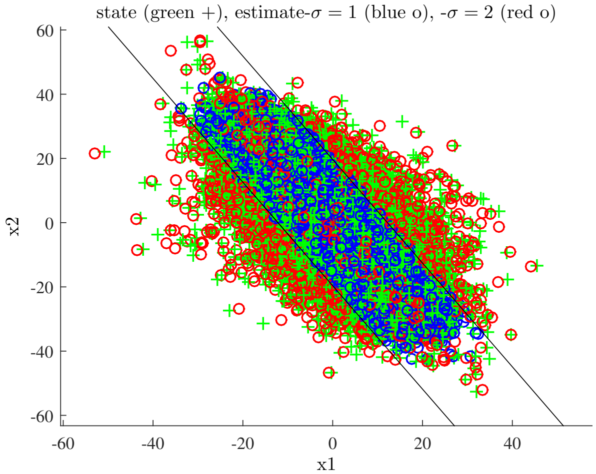

Fig. 3. State, state estimate and switching decision for one simulation

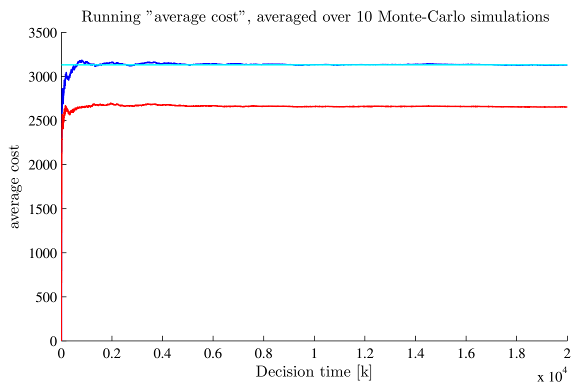

Fig. 4. Running (average) cost

black ellipsoid. This happens even though we have an initial condition near the origin and an optimal control strategy is used with Alg1.

In Figure 4, the simulated average costs, averaged over all Monte Carlo simulations, are shown as they change with the length of the simulation. Since the average cost is defined over an infinite amount of time, we are interested in the tail of the figure where the cost stabilizes. Here, we see that our switched policy (red) performs better, i.e., it has lower average cost, than using Alg1 all the time (blue). We can see that using only Alg1 approximately converges to the theoretical value (cyan). In particular, we can determine the true improvement by comparing the stabilized values. These are given in Table I. The cost difference is ≈ 475, which amounts to a cost profit of $\approx 15\%$ compared to the theoretical cost of Alg1.

From these results, we can deduce that implementing an on-line switching strategy can be very beneficial. Note that a cost reduction of 15% is achieved, despite the fact that Alg2 performs worse (by 1.45%). Here, we have used a very simple 2D system. For more complex systems, it is expected that, by tuning parameters of the system and data-processing methods, even larger performance benefits can be realized.

	Cost	Theoretical value
Only Alg1	3128.8	3131.9
Only Alg2	3170.5	3177.4
Switched	2653.8	n.a.

TABLE I

AVERAGE COST

VI. CONCLUSION

In this paper, we proposed a switching and control policy that breaks the speed versus accuracy trade-off encountered in many data-intensive or vision-based control applications having both slow but accurate data-processing methods and fast but coarse data-processing methods to extract state information from (big) measurement data. We formally showed that a lower expected cost is achieved by our policy compared to when no switching is used. In fact, simulation results demonstrated that a profit of around 15% can already be achieved for a simple second-order system, and we expect that even larger gains in performance can be realized for more complex systems. Our results can be extended to more than two data processing methods in a straightforward manner.

REFERENCES

[1] T. Chen and B. Francis, *Optimal Sampled-Data Control Systems*. Springer-Verlag, 1995.
[2] D. Antunes and W. Heemels, "Rollout event-triggered control: Beyond periodic control performance," *Automatic Control, IEEE Transactions on*, vol. 59, no. 12, pp. 3296–3311, Dec 2014.
[3] J. Geromel, G. Deaecto, and J. Daafouz, "Suboptimal switching control consistency analysis for switched linear systems," *Automatic Control, IEEE Trans. on*, vol. 58, no. 7, pp. 1857–1861, July 2013.
[4] G. S. Deaecto, A. R. Fioravanti, and J. C. Geromel, "Suboptimal switching control consistency analysis for discrete-time switched linear systems," *European Journal of Control*, vol. 19, no. 3, pp. 214–219, 2013.
[5] S. Reimann, S. Al-Areqi, and S. Liu, "Output-based control and scheduling codesign for control systems sharing a limited resource," in *Decision and Control (CDC), 2013 IEEE 52nd Annual Conference on*, Dec 2013, pp. 4042–4047.
[6] J. Melin, M. Jungers, J. Daafouz, and C. Iung, "Performance analysis and design of dynamic output feedback control for switched systems," *International Journal of Control*, vol. 84, no. 2, pp. 253–260, 2011.
[7] K. Astrom, *Introduction to stochastic control theory*. Elsevier, 1970.
[8] D. P. Bertsekas, *Dynamic Programming and Optimal Control*. Athena Scientific, 2005, 3rd Edition. Vol. 1 and Vol. 2.
[9] A. Cervin and P. Alriksson, "Optimal on-line scheduling of multiple control tasks: A case study," in *Real-Time Systems, 2006. 18th Euromicro Conference on*, 2006, pp. 141–150.
[10] J.-W. Lee, "Infinite-horizon joint lqg synthesis of switching and feedback in discrete time," *Automatic Control, IEEE Transactions on*, vol. 54, no. 8, pp. 1945–1951, Aug 2009.
[11] J.-W. Lee and G. Dullerud, "Supervisory control and measurement scheduling for discrete-time linear systems," *Automatic Control, IEEE Transactions on*, vol. 56, no. 4, pp. 873–879, April 2011.
[12] A. Molin and S. Hirche, "On LQG joint optimal scheduling and control under communication constraints," in *Decision and Control, 2009 and 2009 28th Chinese Control Conf. CDC/CCC 2009. Proc. of the 48th IEEE Conf. on*, Dec 2009, pp. 5832–5838.
[13] B. Lincoln and B. Bernhardsson, "LQR optimization of linear system switching," *Automatic Control, IEEE Transactions on*, vol. 47, no. 10, pp. 1701–1705, Oct 2002.
[14] D. Henriksson and A. Cervin, "Optimal on-line sampling period assignment for real-time control tasks based on plant state information," in *Decision and Control, 2005 and 2005 European Control Conf. 44th IEEE Conf. on*, Dec 2005, pp. 4469–4474.