

Switching data-processing methods in a control loop: Trade-off between delay and probability of data acquisition

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

* *Control System Technology Group, 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).*

Abstract: Many control applications, such as vision-based control, require data-processing methods to distill sensor information. This data-processing introduces several undesired effects in the control loop, such as delays, the probability of not acquiring information, and measurement inaccuracies. Often, these effects obey a trade-off. For example, the probability of acquiring control-relevant information, related to the probability of data-loss, is typically higher if a larger delay is allowed. While a single processing method with a reasonable trade-off is typically selected, we propose instead a solution to switch between data-processing methods with different delays and corresponding data-loss probabilities. We prove that the proposed method achieves a better LQG-type performance when compared to the individual methods. A simulation considering a second-order system illustrates the advantages of the proposed method.

Keywords: Switched systems, Self-triggered control, Optimal control, Data-processing, Data acquisition, Time-delay, Probability of information loss

1. INTRODUCTION

There is nowadays a growing industrial interest in modeling data-processing units in the control loop, rather than assuming that these are ideal or account for the worst case. This is especially relevant in high-end applications and in data-intensive applications, such as big data and image-/vision-based control. In such applications, the data-processing element is a non-trivial component that converts large quantities of measurement data to control-relevant information. Typically, data-processing units are not limited to only one processing method, but many are available. Several characteristics can be included in the models of those processing methods, such as delay, accuracy of information, and the probability of acquiring information, often obeying trade-offs between the characteristics. Typically, once the characteristics of the processing methods have been identified, a single method with reasonable trade-offs is selected for implementation.

Recently, in van Horssen et al. (2015), we proposed to switch between data-processing algorithms on-line to improve closed-loop performance. In van Horssen et al. (2015), we have considered the trade-off between speed, modeled by the processing rate at which a given method can run, and accuracy, modeled by the noise characteristics (covariance matrix) of the processed data. In this paper, we tackle another important trade-off present in selecting data-processing algorithms, namely the trade-off between processing delay and probability of data-loss, i.e. the probability of acquiring control-relevant information from the data. In fact, if more processing time is allowed, the probability that useful information is obtained typically becomes higher, at the cost of a larger delay. Compared

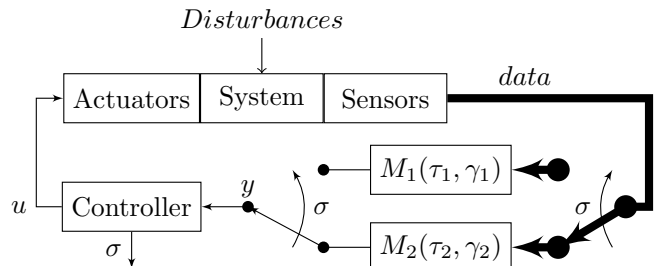


Fig. 1. Control loop with two data-processing algorithms

to the previous work, several new challenges are addressed in the design of the switching policy for this new trade-off and are described shortly.

We consider the interconnection of a physical system with sensors and actuators, a data-processing unit consisting of several data-processing methods, and a digital controller in a feedback structure, as depicted in Fig. 1. The data-processing algorithms, or methods, acquire the data from the same sensors, but the way they produce control-relevant output varies. Each data-processing method is characterized by the incurred processing delay and the probability of having a correctly processed measurement. Only one processing method is allowed to be active at any given time, which is consistent with typical limitations on processing power. The goal is to design a switching and control policy to achieve a better closed-loop performance than the typical approach of selecting and implementing only the fixed method with best performance. Although only two methods are depicted, the results presented in this paper are valid for a higher number of processing methods.

Building upon available results in optimal control theory (Åström (1970); Schenato et al. (2007)) and ideas from (approximate) dynamic programming (Bertsekas (2005); van Horssen et al. (2015); Antunes et al. (2012)), we establish a switching and actuation policy that achieves improved closed-loop performance compared to the typical non-switching approach. To guarantee this improvement, a condition is used to establish the best non-switching policy, which is different from the assumption made in van Horssen et al. (2015) (see Remark 2). Since high-tech systems are typically resource-constrained, several suggestions are made towards efficient implementation of the method. Monte Carlo simulations for motion control of a mass-spring-damper system illustrate the effectiveness of the method and the achievable performance improvement.

While the ideas of the present paper and of van Horssen et al. (2015) are, to our best knowledge, novel, there is some related work available in literature. Switching approaches to schedule measurements from different sensors have been considered previously, e.g. for sensor data scheduling (Wu et al. (2013); Kouchiyama and Ohmori (2010); Molin and Hirche (2009); Leong et al. (2015)). The recent self-triggered (Araújo et al. (2011); Antunes et al. (2012); Gommans et al. (2014)) and event-triggered (Wu et al. (2013); Rabi et al. (2008); Molin and Hirche (2010)) approaches to schedule transmissions in a control loop also exploit switching to improve the control performance. The relation between delay and information loss was addressed using a different approach in Demirel et al. (2015). Dropouts in the optimal control context are addressed in Schenato et al. (2007). The problem of selecting which part of the data is most relevant is considered in sensor management Hero and Cochran (2011), sensor fusion, and sensor selection literature. An alternative tool to construct switching and control policies is the embedding method (Benigea and DeCarlo (2005); Vasudevan et al. (2013)).

With respect to van Horssen et al. (2015), the present paper addresses in particular the following issues. No restricting relation is assumed between the delays of the different processing methods, leading to aperiodic sampling (see Remark 1). The new result allows asynchronous decision intervals, i.e. future decisions instances are not fixed in time (see Remark 3). Uncertainty of acquiring useful processing results inhibits regular innovation of the state information.

The problem formulation is given in Section 2. Section 3 explains the proposed methodology, provides the main result, and gives details on the implementation. A numerical example in Section 4 illustrates the benefits of the proposed method for a second-order system.

2. PROBLEM FORMULATION

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

2.1 Plant and performance criterion

Let a linear stochastic system be described by the differential equation

$$\frac{d}{dt}x_C(t) = A_C x_C(t) + B_C u_C(t) + B_\omega \frac{d\omega}{dt}, \quad (1)$$

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_ω -dimensional Wiener process with incremental covariance $I_{n_\omega} dt$ (cf. Åström (1970)). We assume that (A_C, B_C) is controllable and B_C has full rank. The initial condition is a Gaussian random vector $x_C(0) \sim \mathcal{N}(\bar{x}_0, \Phi^{x_0})$.

Performance of the system is measured by the average cost function, as in the linear quadratic Gaussian (LQG) framework, and is described by

$$J_C^a := \limsup_{T \rightarrow \infty} \mathbb{E} \left[\frac{1}{T} \int_0^T g_C(x_C(t), u_C(t)) dt \right], \quad (2)$$

where $g_C(x, u) := x^\top Q_C x + u^\top R_C u$, with positive semi-definite and positive definite matrices Q_C and R_C , respectively. Additionally, we assume that the pair $(A_C, Q_C^{\frac{1}{2}})$ is observable.

2.2 Measurements from data-processing

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 $\sigma_\ell \in \mathcal{M}$ is activated to distill information that is relevant for feedback control. In this section, we assume, for simplicity, that $\mathcal{M} = \{M_1, M_2\}$ and we will refer to the processing methods by their indices $\{1, 2\}$. Furthermore, only one method may be active at a given time.

After a certain method-dependent delay incurred by the choice of processing method

$$\tau_\ell := \bar{\tau}_{\sigma_\ell} = \begin{cases} \bar{\tau}_1, & \text{if } \sigma_\ell = 1, \\ \bar{\tau}_2, & \text{if } \sigma_\ell = 2, \end{cases} \quad \bar{\tau}_{\sigma_\ell} \in \mathbb{R}_{\geq 0}, \quad (3)$$

the system provides new information y_ℓ to the controller.

The new information, which arrives at the controller at $t_\ell + \tau_\ell$, contains either information about the full state of the system at the sampling time, or no information at all, depending on an indicator $\gamma_\ell \in \{0, 1\}$, i.e.

$$y_\ell := \begin{cases} x_C(t_\ell), & \text{if } \gamma_\ell = 1, \\ \emptyset, & \text{if } \gamma_\ell = 0. \end{cases} \quad (4)$$

Upon information arrival, a new sample is taken, i.e. $t_{\ell+1} = t_\ell + \tau_\ell$.

Apart from the delay, the processing methods are distinguished by the probability that they will provide information. This property is modeled by the Bernoulli distribution of γ_ℓ for each method. In particular, we have that

$$\Pr(\gamma_\ell = 1 \mid \sigma_\ell) =: \bar{\gamma}_{\sigma_\ell} = \begin{cases} \bar{\gamma}_1, & \text{if } \sigma_\ell = 1, \\ \bar{\gamma}_2, & \text{if } \sigma_\ell = 2. \end{cases} \quad \bar{\gamma}_{\sigma_\ell} \in \mathbb{R}_{(0,1]}, \quad (5)$$

Typically, when the processing methods are given a shorter processing time to compute the output, they have also a higher probability of not producing an output. For two methods, this can be captured by the properties $\bar{\tau}_1 > \bar{\tau}_2$ and $\bar{\gamma}_1 > \bar{\gamma}_2$, which can be generalized to several processing methods by a proper ordering.

Remark 1. An important challenge introduced in this setting with respect to van Horssen et al. (2015) is that, by

selecting different delays immediately after the previous delay, the intervals between scheduling decisions become time-varying.

2.3 Digital control

The digital control platform applies a zero-order-hold actuation signal to the system, i.e. the signal

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

is constant between samples. Then, using Åström and Wittenmark (2013) and Åström (1970) the system (1) can be represented at the sampling times t_ℓ by

$$\begin{aligned} x_{\ell+1} &= A_{\sigma_\ell} x_\ell + B_{\sigma_\ell} u_\ell + \omega_\ell, & \omega_\ell &\sim \mathcal{N}(0, \Phi_{\sigma_\ell}^\omega), \\ y_\ell &= \begin{cases} x_\ell, & \text{if } \gamma_\ell = 1, \\ \emptyset, & \text{if } \gamma_\ell = 0. \end{cases} \end{aligned} \quad (7)$$

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

$$A_m := e^{A_C \bar{\tau}_m}, \quad B_m := \int_0^{\bar{\tau}_m} e^{A_C s} B_C ds, \quad m \in \mathcal{M}.$$

The disturbance ω is a sequence of zero-mean independent Gaussian random vectors $\omega_\ell \in \mathbb{R}^{n_\omega}$ with covariances $\mathbb{E}[\omega_\ell(\omega_\ell)^\top] = \Phi_{\sigma_\ell}^\omega$ for all $\ell \in \mathbb{N}$, with

$$\Phi_m^\omega := \int_0^{\bar{\tau}_m} e^{A_C s} B_\omega B_\omega^\top e^{A_C^\top s} ds, \quad m \in \mathcal{M}.$$

The cost function can be written in terms of the discrete-time system as

$$J_C^a = \limsup_{\ell \rightarrow \infty} \mathbb{E} \left[\frac{1}{T_\ell} \sum_{l=0}^{\ell-1} g(x_l, u_l, \tau_l) \right], \quad \text{for } T_\ell := \sum_{l=0}^{\ell-1} \tau_l, \quad (9)$$

where $g(x, u, \tau) := x^\top Q(\tau)x + 2x^\top S(\tau)u + u^\top R(\tau)u$, and, where

$$\begin{bmatrix} Q(\tau) & S(\tau) \\ (*)^\top & R(\tau) \end{bmatrix} := \int_0^\tau e^{\begin{bmatrix} A_C & B_C \\ 0 & 0 \end{bmatrix} s} \begin{bmatrix} Q_C & 0 \\ 0 & R_C \end{bmatrix} e^{\begin{bmatrix} A_C & B_C \\ 0 & 0 \end{bmatrix} s} ds.$$

We define

$$Q_m := Q(\bar{\tau}_m), \quad R_m := R(\bar{\tau}_m), \quad S_m := S(\bar{\tau}_m), \quad m \in \mathcal{M}.$$

2.4 Problem statement

The objective is to find the best σ_ℓ and u_ℓ for each $\ell \in \mathbb{N}$ such that (2) is minimal. Thus, we can formulate the problem as follows. Find a *switching and actuation policy* (co-design problem), i.e., a sequence $\pi = (\mu_0, \mu_1, \dots)$ of multivariate functions that provide the switching and actuation inputs at sampling times t_ℓ

$$(\sigma_\ell, u_\ell) = \mu_\ell(\mathcal{I}_\ell), \quad \ell \in \mathbb{N}, \quad (10)$$

based on the information vector available at the controller

$$\mathcal{I}_\ell := \{\mathcal{I}_{\ell-1}, \sigma_{\ell-1}, u_{\ell-1}, y_{\ell-1}, \gamma_{\ell-1}\}, \quad \ell \in \mathbb{N},$$

with initial information $\mathcal{I}_0 := \{\bar{x}_0, \Phi^{x_0}\}$. Note that at t_ℓ , the information $y_{\ell-1}$ has just arrived at the controller.

Formally, we can write the control design objective as finding a policy π to minimize the expected cost

$$J_\pi := \{J_C^a \mid (3) - (6), \pi\}. \quad (11)$$

We first observe that finding optimal switching and actuation policies, i.e. finding $\pi^* = \arg \min_\pi J_\pi$, is a hard problem in general (see, e.g., Antunes and Heemels (2014)). Here, we take a sub-optimal approach for which we will guarantee improvement over non-switching policies where σ_ℓ is constant (i.e. only one switching method is selected).

Let J_b denote J_π for a non-switching policy, i.e. where $\sigma_\ell = b \in \mathcal{M}$ for all $\ell \in \mathbb{N}$ and an optimal policy for the control input is selected. This optimal control input policy will be discussed in the sequel. In what follows, we address such policies as *base policies*. Furthermore, we denote one optimal choice of b by b^* , i.e. $J_{b^*} \leq J_b$ for all $b \in \mathcal{M}$.

Definition 1. A proposed policy π for the system (1)-(6) is said to expectedly improve over any non-switching policy if it achieves

$$J_\pi \leq J_{b^*} \leq J_b, \quad \text{for all } b \in \mathcal{M}. \quad (12)$$

To achieve our results, we need the following assumption.

Assumption 1. $\tilde{A}_b := A_b \sqrt{(1 - \bar{\gamma}_b)}$ is Schur for all $b \in \mathcal{M}$.

Assumption 1 guarantees that solutions of the base policies remain bounded, i.e. $J_b < \infty$, and is satisfied if $\bar{\gamma}_b$ satisfies the critical observation arrival probability condition in Sinopoli et al. (2003); Schenato et al. (2007) for all $b \in \mathcal{M}$.


Remark 2. Note that it is not specified which policy $b \in \mathcal{M}$ corresponds to b^* , which is different from the assumption made in van Horssen et al. (2015).

3. PROPOSED SWITCHING AND CONTROL POLICY

In this section, we describe the proposed control policy for the system leading to our main result. Details on implementation are given subsequently.

3.1 Switching policy

The switching policy we propose is based on stochastic approximate dynamic programming and is known as a *rollout method* (see Bertsekas (2005)). In the policy, a switching decision is made at t_ℓ on which algorithm to choose next, assuming that afterwards the optimal base policy b^* is always used, as is illustrated in the following diagram.



This is then repeated in a receding horizon fashion, i.e., the same procedure is repeated at times $t_{\ell+1}, t_{\ell+2}, \dots$. Our actuation policy is based on the standard optimal LQG controller design (see Åström (1970)).

Let

$$J_{C, [t, t+T]}^a := \frac{1}{T} \int_t^{t+T} \mathbb{E} [g_C(x_C(t), u_C(t))] dt. \quad (13)$$

be the cost incurred after some time t and up to time $t + T$. The switching criterion is defined as the choice of the method $m \in \mathcal{M}$ that minimizes (13) when that method is first selected assuming that afterwards methods b^* is always selected. Formally,

$$\sigma_\ell := \arg \min_{m \in \mathcal{M}} J_\ell^m \quad (14)$$

where

$$J_\ell^m := \{J_{C, [t_\ell, t_\ell + T]}^a \mid \mathcal{I}_\ell, \sigma_\ell = m \in \mathcal{M}, \\ \sigma_l = b^*, \text{ for all } l > \ell\}. \quad (15)$$

While this defines a family of policies parameterized by T , later, we will let $T \rightarrow \infty$ and consider the unique resulting policy. Note that taking the limit in (13) directly would leave (14) ill-defined. For now, we assume that T is a large number.

The derivation of J_ℓ^m is detailed in the next sections, leading to the switching condition (32), provided below.

Remark 3. An important challenge introduced with respect to van Horssen et al. (2015) is that the future decision moments become decision-dependent and are not aligned in time for all $m \in \mathcal{M}$.

3.2 Optimal estimation of the current state

In this section, the optimal estimation of the current state is detailed when the Kalman filter is used.

Since the system is time-varying, the time-varying Kalman filter provides optimal state estimates in the least-squares sense. Due to the processing delay, an additional prediction step is needed to acquire an estimate of the current state. The best estimate of the current state is given by the estimator-predictor

$$\bar{x}_\ell := \mathbb{E}[x_\ell \mid \mathcal{I}_\ell] \\ = A_{\sigma_{\ell-1}} \bar{x}_{\ell-1} + B_{\sigma_{\ell-1}} u_{\ell-1} + L_{\ell-1} (y_{\ell-1} - \bar{x}_{\ell-1}). \quad (16)$$

with \bar{x}_0 as initial estimate (see Åström (1970)). The time-varying Kalman gain $L_{\ell-1}$ and the covariance $\Theta_\ell := \mathbb{E}[(x_\ell - \bar{x}_\ell)(x_\ell - \bar{x}_\ell)^\top \mid \mathcal{I}_\ell]$ are given by recursions of the Riccati equation

$$L_{\ell-1} := \gamma_{\ell-1} A_{\sigma_{\ell-1}} \quad (17)$$

$$\Theta_\ell = (A_{\sigma_{\ell-1}} - L_{\ell-1}) \Theta_{\ell-1} (A_{\sigma_{\ell-1}} - L_{\ell-1})^\top + \Phi_{\sigma_{\ell-1}}^\omega, \quad (18)$$

where $\Theta_0 = \Phi^{x_0}$.

We can distinguish two cases for (16) and (18) according to whether the last measurement was successfully processed or not. Specifically,

$$\bar{x}_\ell = \begin{cases} A_{\sigma_{\ell-1}} y_{\ell-1} + B_{\sigma_{\ell-1}} u_{\ell-1}, & \text{if } \gamma_{\ell-1} = 1, \\ A_{\sigma_{\ell-1}} \bar{x}_{\ell-1} + B_{\sigma_{\ell-1}} u_{\ell-1}, & \text{if } \gamma_{\ell-1} = 0, \end{cases} \quad (19)$$

$$\Theta_\ell = \begin{cases} \Phi_{\sigma_{\ell-1}}^\omega, & \text{if } \gamma_{\ell-1} = 1, \\ A_{\sigma_{\ell-1}} \Theta_{\ell-1} A_{\sigma_{\ell-1}}^\top + \Phi_{\sigma_{\ell-1}}^\omega, & \text{if } \gamma_{\ell-1} = 0, \end{cases} \quad (20)$$

where we see that Θ resets to Φ^ω after a measurement is correctly processed within the available time. We note here that $(\bar{x}_\ell, \Theta_\ell)$ is a sufficient statistic for x_ℓ at time t_ℓ Åström (1970). Furthermore, the conditioning of σ_ℓ on \mathcal{I}_ℓ guarantees the optimality of the Kalman filter.

3.3 Control policy

The proposed switching policy (14) compares different scheduling options where the future choices of σ are assumed to be known. Hence, for each of these options, A_l, B_l , for $l > \ell$ are assumed to be known and we denote them by A_{σ_l} and B_{σ_l} , respectively. The optimal actuation policy is then known to satisfy the separation principle Schenato et al. (2007), and is given by the time-varying LQG controller (see e.g. Åström (1970)), which consists

of a time-varying Kalman filter and a time-varying LQR controller. The time-varying Kalman filter is given by (17),(18). We detail now the time-varying LQR.

At times t_ℓ , actuation updates take place according to

$$u_\ell = -K_{\sigma_\ell} \bar{x}_\ell, \quad (21)$$

where the control gains \check{K}_m , $m \in \mathcal{M}$, are given by

$$\check{G}_m = R_m + B_m^\top \bar{P}_{b^*} B_m \quad (22)$$

$$\check{K}_m = (\check{G}_m)^{-1} (B_m^\top \bar{P}_{b^*} A_m + (S_m)^\top) \quad (23)$$

$$\check{P}_m = A_m^\top \bar{P}_{b^*} A_m + Q_m - \check{K}_m^\top \check{G}_m \check{K}_m, \quad (24)$$

solved backward, where $\bar{P}_{b^*} = \check{P}_{b^*}$ is the stationary solution to the infinite horizon LQR problem when using only base policy b^* . Additionally, we define $\check{G}_{b^*} = \check{G}_{b^*}$ and $\check{K}_{b^*} = \check{K}_{b^*}$.

3.4 Expected cost

To establish an estimate for the cost incurred after the decision moment, predictions of the future state and covariance are needed.

Let

$$\check{L}_{\ell, h} := \mathbb{E}[L_{\ell+h} \mid \mathcal{I}_\ell] = \bar{\gamma}_{\sigma_{\ell+h}} A_{\sigma_{\ell+h}} \quad (25)$$

We can make predictions of the future state estimates and their variance according to

$$\check{x}_{\ell, h+1} := \mathbb{E}[\bar{x}_{\ell+h+1} \mid \mathcal{I}_\ell] \\ = A_{\sigma_{\ell+h}} \check{x}_{\ell, h} + B_{\sigma_{\ell+h}} u_{\ell+h} \quad (26)$$

with initial condition $\check{x}_{\ell, 0} = \bar{x}_\ell$, and

$$\check{\Theta}_{\ell, h+1} := \mathbb{E}[(x_{\ell+h+1} - \bar{x}_{\ell, h+1})(x_{\ell+h+1} - \bar{x}_{\ell, h+1})^\top \mid \mathcal{I}_\ell] \\ = \Phi_{\sigma_{\ell+h}}^\omega + \left[A_{\sigma_{\ell+h}} \check{\Theta}_{\ell, h} A_{\sigma_{\ell+h}}^\top \right] \cdot (1 - \bar{\gamma}_{\sigma_{\ell+h}}) \quad (27)$$

with initial condition $\check{\Theta}_{\ell, 0} = \Theta_\ell$. Note that (27) follows directly when taking the expected value over (20).

We do not have measurement noise, hence the minimal and maximal bounds of the predictions (as e.g. detailed in Schenato et al. (2007)) coincide. Assumption 1 guarantees that for the base policies the expected covariance converges to some $\bar{\Theta}_b$ that satisfies

$$\bar{\Theta}_b = \Phi_b^\omega + [A_b \bar{\Theta}_b A_b^\top] \cdot (1 - \bar{\gamma}_b) \quad (28)$$

for any initial value of $\check{\Theta}_{\ell, 0}$.

Using (Åström (1970); Schenato et al. (2007)) and in accordance with our switching policy, we can establish for each $m \in \mathcal{M}$ the expected future cost (15) up to time $t_\ell + T$, for $T = N\bar{\tau}_{b^*} + \bar{\tau}_m$, $N \in \mathbb{N}$, as

$$J_\ell^m = \limsup_{N \rightarrow \infty} \frac{1}{N\bar{\tau}_{b^*} + \bar{\tau}_m} [\bar{x}_\ell^\top \check{P}_m \bar{x}_\ell + \chi_m(\Theta_\ell, N)], \quad (29)$$

for with

$$\chi_m(\Theta_\ell, N) := \text{tr}(\check{\Theta}_{\ell, 0} \check{P}_m + \check{\Theta}_{\ell, 0} \check{K}_m^\top \check{G}_m \check{K}_m) \\ + \text{tr}(\Phi_m^\omega \bar{P}_{b^*} + \check{\Theta}_{\ell, 1} \bar{K}_{b^*}^\top \bar{G}_{b^*} \bar{K}_{b^*}) \\ + \sum_{l=2}^N \text{tr}(\Phi_{b^*}^\omega \bar{P}_{b^*} + \check{\Theta}_{\ell, l} \bar{K}_{b^*}^\top \bar{G}_{b^*} \bar{K}_{b^*}) \quad (30)$$

where $\check{\Theta}_{\ell, l}$ are forward predictions of the variance according to (27), starting with Θ_ℓ , while assuming that $\sigma_\ell = m$, and $\sigma_l = b^*$ for all $l > \ell$.

It is known from optimal control theory (Åström (1970); Schenato et al. (2007)) that the cost for a non-switching policy, in which always method b is chosen, is given by

$$J_b := J_C^a|_{\sigma=b} = \limsup_{\ell \rightarrow \infty} \mathbb{E} \left[\frac{1}{\ell} \sum_{l=0}^{\ell-1} g(x_l, u_l, \bar{\tau}_b) \right] \\ = \frac{1}{\bar{\tau}_b} \text{tr}(\Phi_b^\omega \bar{P}_b + \bar{\Theta}_b \bar{K}_b^\top \bar{G}_b \bar{K}_b) \quad (31)$$

where the matrices follow from the stationary solutions to the Riccati equations in Section 3.3.

Observe that $\limsup_{\ell \rightarrow \infty} \bar{\Theta}_{\ell,l} = \bar{\Theta}_{b^*}$ by the proposed switching policy and Assumption 1. Hence, we can deduce that, for some large N , the matrix norm $\|\bar{\Theta}_{\ell,N} - \bar{\Theta}_{b^*}\| < \epsilon$ for some arbitrarily small ϵ . Thus, the term in the summation in (30) converges to

$$c_{b^*} := \text{tr}(\Phi_{b^*}^\omega \bar{P}_{b^*} + \bar{\Theta}_{b^*} \bar{K}_{b^*}^\top \bar{G}_{b^*} \bar{K}_{b^*})$$

which is equal to the numerator part of the cost (31) when b^* is used every instance.

We can now establish a computable switching condition.

3.5 Switching condition and main result

The next proposition provides the proposed policy (14) with (29) when $T \rightarrow \infty$, i.e. when $N \rightarrow \infty$.

Proposition 1. For $T \rightarrow \infty$, the switching condition (14) is equivalent to

$$\sigma_\ell := \arg \min_{m \in \mathcal{M}} \bar{x}_\ell^\top \check{P}_m \bar{x}_\ell + \tilde{\eta}_m(\Theta_\ell) + (1 - \frac{\bar{\tau}_m}{\bar{\tau}_{b^*}}) c_{b^*} \quad (32)$$

where

$$\tilde{\eta}_m(\Theta_\ell) := \text{tr}(\Theta_\ell \check{P}_m + \Theta_\ell \check{K}_m^\top \check{G}_m \check{K}_m) \\ + \text{tr}(\Phi_m^\omega \bar{P}_{b^*}) + \text{tr}(\mathfrak{F}_m(\Theta_\ell) \bar{K}_{b^*}^\top \bar{G}_{b^*} \bar{K}_{b^*}), \quad (33)$$

where $\mathfrak{F}_m(\Theta_\ell)$ is the solution Y to the Lyapunov equation

$$\bar{A}_{b^*} Y \bar{A}_{b^*}^\top - Y = -[\Phi_m^\omega + A_m \Theta_\ell A_m^\top \cdot (1 - \bar{\gamma}_m)], \quad (34)$$

which always has a solution by Assumption 1. \square

Proof. For brevity, only a sketch of the proof is provided. In the summation in (30), $\text{tr}(\Phi_{b^*}^\omega \bar{P}_{b^*})$ and parts of $\bar{\Theta}_{\ell,l}$, for $l \geq 1$, are not dependent on m . The remainder of $\bar{\Theta}_{\ell,l}$ can be replaced by $\bar{A}_{b^*}^{l-1} \bar{\Theta}_{\ell,1} (\bar{A}_{b^*}^{l-1})^\top$, leading to the solution of the Lyapunov equation.

Here, $\tilde{\eta}_m(\Theta_\ell)$ is defined as the method-dependent part (the terms that depend on a particular choice of m) of $\chi_m(\Theta_\ell, N)$ for $N \rightarrow \infty$. The last term in (32) corrects for the misalignment of the expected sampling times such that the term $\frac{1}{N\bar{\tau}_{b^*} + \bar{\tau}_m}$ in (29) equals to the same value of T in (15) for each $m \in \mathcal{M}$. Hence, the divisor can be eliminated from the arg min.

We can compute the right-hand side of (32) explicitly for each $m \in \mathcal{M}$ and thus our switching condition has become explicit. We can now state our main result.

Theorem 1. Under Assumption 1, the following holds

$$J_\pi \leq J_{b^*} \leq J_b, \text{ for all } b \in \mathcal{M}, \quad (35)$$

if the control policy functions (10) are given by (14) and (21). \square

Proof. For brevity, only a sketch of the proof is provided. Each decision incurs a cost increment of $\frac{\bar{\tau}_m}{\bar{\tau}_{b^*}} c_{b^*} -$




Fig. 2. Running average cost, averaged over all Monte Carlo simulations. Comparison to using only M_2 .

$f_m(\bar{x}_\ell, \Theta_\ell)$, where $f_m(x, \Theta) := x^\top (\bar{P}_{b^*} - \check{P}_m)x + \tilde{\eta}_{b^*}(\Theta) - \tilde{\eta}_m(\Theta) - (1 - \frac{\bar{\tau}_m}{\bar{\tau}_{b^*}}) c_{b^*}$. If m is chosen according to (32), then $f_m(\bar{x}_\ell, \Theta_\ell) \geq 0$, leading to the guarantee that

$$J_\pi = \frac{1}{\bar{\tau}_{b^*}} c_{b^*} - \limsup_{L \rightarrow \infty} \frac{1}{TL} \mathbb{E} \left[\sum_{\ell=0}^{L-1} f_{\sigma_\ell}(\bar{x}_\ell, \Theta_\ell) \mid \mathcal{I}_\ell \right] \leq J_{b^*}.$$

Theorem 1 means that our proposed policy achieves the improvement conditions in Definition 1. The effectiveness of the method is illustrated in the next section.

The following remark is relevant for implementation of the policy for resource-constrained systems such as embedded systems.

Remark 4. Since Θ_ℓ will reset regularly, it can be beneficial to precompute values of $\mathfrak{F}_m(\Theta_\ell)$ for evolutions of Θ_ℓ according to (20) with initial condition Φ_m^ω . Using the precomputed values may provide a significant reduction in on-line computational effort, at the cost of memory to store the values.

4. SIMULATION

We simulate our method for a mass-spring-damper system with negative damping

$$A_C = \begin{bmatrix} 0 & 1 \\ -1 & 0.8 \end{bmatrix}, B_C = \begin{bmatrix} 0 \\ 1 \end{bmatrix}, B_\omega = \begin{bmatrix} 0.8 \\ 0.1 \end{bmatrix}, \quad (36)$$

and a cost

$$Q_C = \begin{bmatrix} 100 & 0 \\ 0 & 100 \end{bmatrix}, R_C = 100. \quad (37)$$

We take two processing methods, given by $M_1(\bar{\tau}_1 = 1s, \bar{\gamma}_1 = 0.9)$ and $M_2(\bar{\tau}_2 = 0.62s, \bar{\gamma}_2 = 0.55)$. The expected average costs (31) of the fixed policies are found to be $J_1 = 1156.6$ and $J_2 = 1140.9$, i.e. M_1 performs 1.3% worse and thus $b^* = 2$.

We run 100 Monte Carlo simulations of the system. We use a small initial condition near the origin $x_0 \approx [0, 0]^\top$, and small initial covariance $\Phi^{x_0} \ll \Phi_m^\omega$, for all $m \in \mathcal{M}$. We take $t \in \mathbb{R}_{[0, 24000]}$ such that the cost has approximately converged. We observed that near the origin the slow method was often selected and further away the fast method was selected increasingly often.

The running average cost when averaged over all Monte Carlo simulations is depicted in Fig. 2. It also shows the cost of the base policy $b^* = 2$ and its theoretical value. We see that our switched policy (red) performs better, i.e. it has lower average cost, than using M_{b^*} all the time (blue).

	Cost	Theoretical value
Only M1	1156.1	1156.6
Only M2	1139.9	1140.9
Switched	1026.5	n.a.

Table 1. Average cost

The approximately stabilized values are given in Table 1. When the number of Monte Carlo simulations is increased, the averages converge to the theoretical values. The cost difference with the theoretical value J_{b^*} is ≈ 114.3 , which amounts to an expected cost improvement of $\approx 10.03\%$ when using our switched policy.

We observed also that at the end of the simulations the variance of the final costs over the Monte Carlo simulations was small for M_1 , a bit larger for the switched approach, and much larger (one order of magnitude) for the system when only M_2 was used. This can be explained by the fact that, within the same time span, the fast method has more instances where a dropout can occur. Increasing the number of Monte Carlo simulations or the simulation time reduces the variance, as is expected.

5. CONCLUSION

In this paper, we proposed a switching and control policy that provides a solution to the design trade-off encountered in applications having multiple processing methods that provide information with different probability, depending on the elapsed processing time. Our policy achieves an improvement in expected cost by switching on-line. In simulations, a simple second-order motion-control example with only two processing methods showed a significant 10% performance gain. It is expected that the methods are even more beneficial for more complex systems. Connections with our previous work van Horssen et al. (2015), i.e. the interplay between (stochastic) delay, accuracy and loss probability are currently under investigation.

ACKNOWLEDGEMENTS

This work is supported by the Dutch Technology Foundation (STW) 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).

REFERENCES

Antunes, D. and Heemels, W.P.M.H. (2014). Rollout event-triggered control: Beyond periodic control performance. *Automatic Control, IEEE Transactions on*, 59(12), 3296–3311.

Antunes, D.J., Heemels, W.P.M.H., Hespanha, J.P., and Silvestre, C.J. (2012). Scheduling measurements and controls over networks - part i: Rollout strategies for protocol design. In *American Control Conference (ACC), 2012*, 2036–2041.

Araújo, J., Anta, A., Mazo, M., Faria, J., Hernandez, A., Tabuada, P., and Johansson, K.H. (2011). Self-triggered control over wireless sensor and actuator networks. In *Distributed Computing in Sensor Systems and Workshops (DCOSS), 2011 International Conf. on*, 1–9.

Åström, K.J. (1970). *Introduction to stochastic control theory*. Elsevier.

Åström, K.J. and Wittenmark, B. (2013). *Computer-controlled systems: theory and design*. Courier Corporation.

Bengea, S.C. and DeCarlo, R.A. (2005). Optimal control of switching systems. *Automatica*, 41(1), 11 – 27.

Bertsekas, D.P. (2005). *Dynamic Programming and Optimal Control*. Athena Scientific. 3rd Edit. Vol. 1 and 2.

Demirel, B., Aytakin, A., Quevedo, D.E., and Johansson, M. (2015). To wait or to drop: On the optimal number of retransmissions in wireless control. In *Control Conference (ECC), 2015 European*, 962–968.

Gommans, T.M.P., Antunes, D.J., Donkers, M.C.F., Tabuada, P., and Heemels, W.P.M.H. (2014). Self-triggered linear quadratic control. *Automatica*, 50(4), 1279 – 1287.

Hero, A.O. and Cochran, D. (2011). Sensor management: Past, present, and future. *IEEE Sensors Journal*, 11(12), 3064–3075.

Kouchiyama, Y. and Ohmori, H. (2010). Switching control algorithm for sensor scheduling. In *SICE Annual Conference 2010, Proceedings of*, 652–656.

Leong, A.S., Dey, S., and Quevedo, D.E. (2015). On the optimality of threshold policies in event triggered estimation with packet drops. In *Control Conference (ECC), 2015 European*, 927–933.

Molin, A. and Hirche, S. (2009). 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*, 5832–5838.

Molin, A. and Hirche, S. (2010). Optimal event-triggered control under costly observations. *Proceedings of the 19th international symposium on mathematical theory of networks and systems*.

Rabi, M., Johansson, K.H., and Johansson, M. (2008). Optimal stopping for event-triggered sensing and actuation. In *Decision and Control, 2008. CDC 2008. 47th IEEE Conference on*, 3607–3612.

Schenato, L., Sinopoli, B., Franceschetti, M., Poolla, K., and Sastry, S.S. (2007). Foundations of control and estimation over lossy networks. *Proceedings of the IEEE*, 95(1), 163–187.

Sinopoli, B., Schenato, L., Franceschetti, M., Poolla, K., Jordan, M.I., and Sastry, S.S. (2003). Kalman filtering with intermittent observations. In *Decision and Control, 2003. Proceedings. 42nd IEEE Conference on*, volume 1, 701–708 Vol.1.

van Horssen, E.P., Antunes, D., and Heemels, W.P.M.H. (2015). Switching data-processing methods for feedback control: Breaking the speed versus accuracy trade-off. In *Decision and Control (CDC), 2015 IEEE 54th Annual Conference on*, 2313–2318.

Vasudevan, R., Gonzalez, H., Bajcsy, R., and Sastry, S.S. (2013). Consistent approximations for the optimal control of constrained switched systems. *SIAM Journal on Control and Optimization*, 51(6), 4463–4483.

Wu, J., Jia, Q.S., Johansson, K., and Shi, L. (2013). Event-based sensor data scheduling: Trade-off between communication rate and estimation quality. *Automatic Control, IEEE Transactions on*, 58(4), 1041–1046.