Event-Driven Control With Deadline Optimization for Linear Systems With Stochastic Delays

Eelco van Horssen, Suraj Prakash, Duarte Antunes, *Member, IEEE*, and Maurice Heemels, *Fellow, IEEE*

***Abstract*—This work presents a novel control strategy for systems with actuation delays with known stochastic distribution, which improves upon previously proposed deadline-driven and event-driven strategies. In the event-driven strategy, the control input is immediately updated after the delay, whereas in the deadline-driven strategy, the actuation is updated in a periodic fashion, where the sampling period sets a deadline; if the delay is larger than this deadline, the actuation is not updated. Our method switches between these two strategies and guarantees better performance, in a linear–quadratic–Gaussian sense, than either method considered separately. An extension of the novel method with a deadline-optimization scheme is shown to improve the performance even further. Simulation results illustrate the effectiveness of the proposed methods.**

***Index Terms*—Data loss, dynamic programming, event-driven control, sampled-data control, self-triggered control, stochastic optimal control, stochastic time delay.**

I. Introduction

DELAYS ARE present in many control applications, resulting from timing effects in the loop such as control computation, communication between distributed components, or measurement acquisitions [1]. These control delays can lead to significant performance degradation in various control settings, especially in industry that requires embedded systems [2], [3]; (shared) communication networks [3]–[5]; and/or data-intensive processing [6], [7].

Many works in the literature addressing control problems with uncertain delays use a worst-case approach, taking into account the largest possible delays, and exploit robust stability analysis techniques (see, e.g., [8]–[10]). In particular, in a traditional control setting, a sufficiently large sampling time is chosen such that the worst-case delays, which may be very large, are accommodated. Naturally, this approach is conservative and can lead to poor closed-loop performance. In this paper, we take an alternative approach, exploiting knowledge of the probability distribution of the delays when selecting the sampling intervals. However, selecting sampling intervals shorter than the delay causes a data dropping effect for which some solutions have been proposed (see, e.g., [11]–[18]). The sampling interval thus plays the role of a deadline and, as such, we denote this approach *deadline driven*. This method leads to a tradeoff between "data-loss" and control rate. In [13], this tradeoff was studied in the context of reliability analysis of a networked control system with energy constraints, while in this work, we consider linear–quadratic–Gaussian (LQG)-type performance. Alternatively, some works address the stochastic nature explicitly [19]–[21], proposing the so-called *event-driven* strategies. Such event-driven strategies (differing from state-dependent event-triggered control; see, for example, [22] and [23]) consider that the control input is immediately updated after the delay and, therefore, the update intervals are equal to the stochastic delay [24]. Results for systems with stochastic parameters [25] can be used to find optimal control strategies for this case.

The current work extends the results in our preliminary work [24]. The work [24] concluded that event-driven or deadline-driven approaches do not necessarily perform better than one another (in an LQG performance sense). Here, we propose a novel switching strategy that is guaranteed to result in better performance than that of event-driven and deadline-driven approaches by switching between them. The switching strategy combines strategies that are event driven, deadline driven, and/or *event driven with a deadline*, where the control input is updated after the delay, except when the delay exceeds a deadline, in which case a data drop occurs. In addition, we show that the novel switching strategy can be combined with a deadline-optimization scheme to obtain additional performance benefits.

The new results are obtained in the setting of linear continuous-time systems with Gaussian disturbances.

The stochastic delays in the control-to-actuation link are assumed to be independent and identically distributed (i.i.d.), as is very common in the networked control systems community and justified in several contexts involving computation (see, e.g., [26]) or communication delays (see, e.g., [27] and [28]). Digital control with delayed zero-order hold inputs is used, as illustrated in Fig. 1. The closed-loop performance is evaluated by an infinite horizon average cost function as in a standard LQG framework. The performance gain of the proposed policies is

Manuscript received July 7, 2017; revised September 29, 2017; accepted October 3, 2017. Date of publication October 16, 2017; date of current version December 14, 2018. This work was supported by the Netherlands Organisation for Scientific Research (NWO-TTW) under Grant 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" (11382) awarded by NWO-TTW. Recommended by Associate Editor L. Schenato. *(Corresponding author: Eelco van Horssen.)*

The authors are with the Control Systems Technology Group, Eindhoven University of Technology, 5600 MB Eindhoven, The Netherlands (e-mail: e.p.v.horssen@tue.nl; s.prakash@tue.nl; d.antunes@tue.nl; m.heemels@tue.nl).

Digital Object Identifier 10.1109/TCNS.2017.2763749

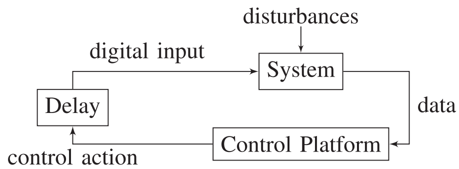

Fig. 1. Control loop with actuation delay.

illustrated by numerical examples. The proofs of the main results resort to Doob's optional sampling theorem [29].

The remainder of this paper is organized as follows. The control setup with actuation delay and the problem formulation are detailed in Section II, where the deadline-driven and event-driven strategies are also discussed, as well as the event-driven with deadline strategy. Using a discrete-time model of the system, the performance analysis of the nonswitching approaches leads to a preliminary result in Section III. Section IV provides the main results for the proposed control policies. Numerical examples in Section V illustrate the novel results and the benefits of the new method. Concluding remarks are given in Section VI. The proofs of the main results can be found in the Appendices.

II. Problem Formulation and Background

In this section, we first discuss the control setting and the formal control problem. Subsequently, we discuss several basic control strategies that serve as a benchmark for the methods proposed in this paper.

A. Problem Setting

We consider a continuous-time plant modeled by the stochastic differential equation

$$dx_c = (A_c x_c + B_c u_c)dt + B_w dw(t),\ x_c(0) = x_0,\ t \in \mathbb{R}_{\geq 0} \tag{1}$$

where $x_c(t) \in \mathbb{R}^{n_x}$ is the state and $u_c(t) \in R^{n_u}$ is the applied control input at time $t \in \mathbb{R}_{\geq 0}$, and w is an n_w-dimensional Wiener process with incremental covariance $I_{n_w} dt$ [30]. We assume that (A_c, B_c) is controllable and B_c has full rank.

As in the standard LQG framework, the average quadratic cost

$$J := \limsup_{T\to\infty} \frac{1}{T}\int_0^T \mathbb{E}[g_c(x_c(t), u_c(t))]dt \tag{2}$$

is chosen as the performance criterion, where $g_c(x,u) := x^\intercal Q_c x + u^\intercal R_c u$ with positive-definite matrix $R_c \succ 0$ and positive-semidefinite matrix $Q_c \succeq 0$ for which $(A_c, Q_c^{\frac{1}{2}})$ is observable.

We consider a setup with a simple hold device at the plant input such that the plant actuation signal is held constant between discrete *update instances* t_k, $k \in \mathbb{N}$, with $t_{k+1} > t_k$, for all $k \in \mathbb{N}$, and $t_0 = 0$. In particular, we write

$$u_c(t) = \hat{u}_k,\ \text{for all } t \in [t_k, t_{k+1})$$

where $\hat{u}_k \in \mathbb{R}^{n_u}$ is the *digital input* value held at time t_k, $k \in \mathbb{N}$. We assume that $\hat{u}_0 := u_c(0)$ is known.

We assume, for now, that the plant may be sampled at any time instance and discuss how to relax this assumption in Remark 2. In particular, we choose the sampling instances to coincide with the actuation update instances, that is, the plant is sampled at times t_k, $k \in \mathbb{N}$, with $t_{k+1} > t_k$ for all $k \in \mathbb{N}$. The time-varying "sampling" intervals can then be defined as

$$h_k := t_{k+1} - t_k, \quad k \in \mathbb{N}\,.$$

At every sampling instance t_k, we assume that the sensor provides a measurement of the full state $x_c(t_k)$ and denote this by $x_k := x_c(t_k)$, $k \in \mathbb{N}$.

The sampling intervals h_k, $k \in \mathbb{N}$, will take different values, detailed later, depending on the chosen control strategy. We assume that there exists a (possibly small) $h^{\min} \in \mathbb{R}_{>0}$ such that $h_k \geq h^{\min}$ for all $k \in \mathbb{N}$, which imposes a minimal interval.

The computation of a new *control action* $u_k \in \mathbb{R}^{n_u}$ by the controller starts immediately after a new sample is obtained, that is, at time t_k for all $k \in \mathbb{N}$ as a function of all the information in the control platform at time t_k. Due to computational delays or communication delays, the new control action can only be applied after a delay $\tau_k \in \mathbb{R}_{>0}$ for all $k \in \mathbb{N}$. The delays τ_k, $k \in \mathbb{N}$, are i.i.d. with the known delay distribution defined by the probability measure μ.

The support of μ is allowed to be unbounded, but we assume that $\mu((0,\infty]) = 1$ and $\mu(\{0\}) = 0$. The measure μ can be decomposed into continuous and discrete components as $\mu = \mu_c + \mu_d$ with $\mu_c((0,s)) = \int_0^s f^c(\tau)d\tau$, where f^c is a measurable function, and μ_d is a discrete measure that captures possible point masses at $b_i \in \mathbb{R}_{>0} \cup \{\infty\}$, $i \in \mathcal{I} \subseteq \mathbb{N}$, such that $\mu(\{b_i\}) = w_i$, $i \in \mathcal{I}$. The (Lebesgue–Stieltjes) integral of some function W with respect to the measure μ is defined as

$$\int_0^t W(s)\mu(ds) := \int_0^t W(s) f^c(s)ds + \sum_{i\in\mathcal{I}:b_i\in(0,t]} w_i W(b_i).$$

The cumulative distribution function (cdf) $F : \mathbb{R}_{>0} \cup\{\infty\} \to \mathbb{R}_{[0,1]}$ is given by $F(\tau) := \mu((0,\tau]) = \int_0^\tau f^c(s)ds + \sum_{i\in\mathcal{I}:b_i\in(0,\tau]} w_i$, for $\tau \in (0,\infty]$, which is equal to $\mathbb{P}(\tau_k \leq \tau), k \in \mathbb{N}$, where $\mathbb{P}$ denotes probability. The probability distribution function (pdf) associated with F is denoted as $f : \mathbb{R}_{>0} \cup \{\infty\} \to \mathbb{R}_{\geq 0}$. The expected value of the delay is equal for all $k \in \mathbb{N}$ and is denoted by $\bar{\tau} := \mathbb{E}[\tau_k]$.

If the sampling interval h_k has a maximum value $D_k \in \mathbb{R}_{>0}, k \in \mathbb{N}$ imposed, then this works as a *deadline*. If the delay exceeds the deadline, that is, if $\tau_k > D_k$ for some $k \in \mathbb{N}$, then the newly computed control action u_k is dropped and the previous actuation signal $\hat{u}_k$ is held constant. We assume that there exists a (possibly large) $D^{\max} \in \mathbb{R}_{>0}$ such that $D_k \leq D^{\max}$ for all $k \in \mathbb{N}$, which imposes a maximum deadline value. The new plant input $\hat{u}_{k+1}$, after the interval h_k, becomes, for all $k \in \mathbb{N}_{\geq 0}$

$$\hat{u}_{k+1} = \begin{cases} u_k, & \text{if } \tau_k \leq D_k \\ \hat{u}_k, & \text{otherwise.} \end{cases} \tag{3}$$

We assume that only one message, that is, a control action, is allowed in the actuation channel within each sampling interval, and that a deadline is known at the actuator (if a deadline is applied).

We use a Bernoulli random variable γ_k, $k \in \mathbb{N}$, to capture the occurrence of data drops. In particular, $\gamma_k = 1$ denotes that the control input u_k has been successfully applied to the system, while $\gamma_k = 0$ denotes that u_k has been dropped. This is described by the dropping mechanism

$$\gamma_k = \begin{cases} 1, & \text{if } \tau_k \leq D_k \\ 0, & \text{if } \tau_k > D_k. \end{cases}$$

As a consequence, (3) can be rewritten as

$$\hat{u}_k = \gamma_{k-1} u_{k-1} + (1 - \gamma_{k-1}) \hat{u}_{k-1}, \quad k \in \mathbb{N}_{>0}. \tag{4}$$

Remark 1: A typical actuation channel cannot be instantaneous; therefore, a minimal delay is always present. Hence, it is easy to determine some $h^{\min} \in \mathbb{R}_{>0}$ such that $F(h^{\min}) = 0$. Otherwise, one can consider a new probability measure $\tilde{\mu}$ with the probability $F(h^{\min}) \neq 0$ accumulated at $\mu(\{h^{\min}\})$ and artificially delay the system if $\tau_k < h^{\min}$ for some $k \in \mathbb{N}$.

Remark 2: The problem setting can also capture the sampled-data scenario where the sensors can only be sampled at discrete intervals but at a fast rate, in the sense that the sampling period is much smaller than typical delay values. Delaying the actuation updates to the next sampling instance causes the delays to take values in a countable set, which can be captured by a piecewise constant cdf. Since only one message is allowed in the actuation channel, extra samples taken within the actuation update interval are discarded.

Remark 3: The setup requires that either the sensor has knowledge of the actuation update instances, which informs the controller by sending a new measurement, or that the controller has knowledge of the channel, such that it can trigger the sensor to provide a new measurement. This can be satisfied by, for example, a collocated sensor–actuator at the plant or a channel-sensing mechanism at the controller.

B. Control Problem

The control problem is the design of a methodology to obtain suitable control actions u_k, $k \in \mathbb{N}$, and delay deadlines D_k, $k \in \mathbb{N}$, such that the performance index (2) is smaller than for known existing methods.

To minimize performance index (2), that is, to solve the control problem optimally, by, for example, the use of dynamic programming [31] is not possible due to the curse of dimensionality. As such, we opt to design a suboptimal methodology that is better than current practice. In particular, our goal is to obtain a *control policy* π that provides u_k and D_k, i.e.,

$$(u_k, D_k) = \pi(\mathrm{I}_k), \quad k \in \mathbb{N}$$

as a function of the information available for control at time t_k, being

$$\mathrm{I}_k := \{x_k\} \cup \{x_l, u_l, D_l, h_l, \gamma_l | l \in \mathbb{N}_{[0,k)}\} \cup \{\hat{u}_0\}.$$

In this paper, we provide control policies that are guaranteed to improve over both optimal event-driven and deadline-driven strategies as proposed in [24], in the sense that the performance index (2) is smaller or equal. Simulation results are evidence that significant improvement can be realized.

C. Basic Control Strategies (Background)

In this paper, we consider the following basic strategies or *base policies*, which we will indicate by $d, e,$ and ed, respectively.

1) Periodic Deadline-Driven Control (d): This typical design approach sets a fixed deadline D_d^p for each interval and the control update interval coincides with this deadline, that is, $D_k = D_d^p$ and $h_k = D_d^p$ for all $k \in \mathbb{N}$. This results in dropping u_k with probability $1 - F(D_d^p)$, that is, $\mathbb{P}(\gamma_k = 0) = 1 - F(D_d^p)$. Note that imposing a deadline is a natural way to deal with large delays. In practice, however, the deadline is imposed without further analysis of the dropping effect, while this may significantly impact the stability and/or performance, as we will see.

2) Event-Driven Control (e): This aperiodic strategy updates $\hat{u}_k$ directly after the delay without considering a deadline, that is, $h_k = \tau_k$ (and $D_k = \infty$) for all $k \in \mathbb{N}$. Note that u_k is never dropped, that is, $\mathbb{P}(\gamma_k = 0) = 0$. When considering this case, we make the additional assumption $\mu(\{\infty\}) = 0$, which is necessary for stabilizability and is equivalent to all $b_i < \infty$ for all $i \in \mathcal{I}$.

3) Periodic Event-Driven Control With Deadline (ed): This aperiodic strategy updates $\hat{u}_k$ directly after the delay if the delay is less than the set deadline D_{ed}^p and at the deadline when the delay is larger, that is, $h_k = \min(\tau_k, D_k)$ and $D_k = D_{ed}^p$ for all $k \in \mathbb{N}$. This results in dropping u_k with probability $1 - F(D_{ed}^p)$, that is, $\mathbb{P}(\gamma_k = 0) = 1 - F(D_{ed}^p)$, but updating $\hat{u}_{k+1}$ earlier at time $t_k + \tau_k$ with probability $f(\tau_k)$ for each value of $\tau_k \leq D_k$. Note that assuming $\mu(\{\infty\}) = 0$ is not necessary for stabilizability in this case.

The methods d and e were previously discussed in our preliminary work [24], where it was suggested to use event-driven approaches to improve performance over conservative but easy-to-implement deadline-driven approaches that are typically adopted in practice.

Through an example, it was found that event-driven approaches can indeed give significant improvement, but this is not necessarily always the case, as was illustrated by another example that showed better performance for periodic control. This motivated the design and investigation of the ed strategy, proposed here.

Building upon the results in [25], for each of the above methods, an analytical expression for the value of performance index (2) can be obtained. A method to calculate this cost value is explained in Section III-B. The cost is given by

$$J_b := \frac{1}{\bar{h}_b} c_b, \quad c_b := \mathrm{tr}(P_b W_b) + \alpha_b, \quad b \in \{d, e, ed\} \tag{5}$$

where $\bar{h}_b$ denotes the average interval h_k, that is, $\bar{h}_b := \mathbb{E}[h_k]$, and where, as in standard LQG, P_b corresponds to the solution of a Riccati equation, as described in Appendix A, W_b is a noise term from the Wiener process w, also given in Appendix A, and

α_b is a term resulting from the intersampling behavior of the system, given in Appendix B.[1]

The values of $J_b, b \in \{d, e, ed\}$ are called the *base costs* and will serve as a reference to compare our newly proposed methods. The main results in this work derive control policies π that guarantee that $J_\pi \leq J_b$ while typically performing (significantly) better in the sense that $J_\pi < J_b$ for all $b \in \{d, e, ed\}$. For d and ed, the cost (5) depends on the chosen deadline D, and the corresponding costs can be denoted by $J_d(D)$ and $J_{ed}(D)$, respectively. Optimal deadline values that minimize the cost (5), are denoted by $D_d^\star$ and $D_{ed}^\star$, respectively. The costs (5) corresponding to the basic strategies (with optimal deadline) are denoted by $J_{d^\star} := J_d(D_d^\star)$, J_e, $J_{ed^\star} := J_{ed}(D_{ed}^\star)$, respectively. These costs correspond to parameters with the same subscripts $\bar{h}_{d^\star}$, $P_{d^\star}$, $W_{d^\star}$, $\alpha_{d^\star}$ for periodic deadline-driven and analogously e and $ed^\star$ for event-driven and periodic event-driven with deadline, respectively. Note that $\bar{h}_{d^\star} = D_d^\star$, $\bar{h}_e = \bar{\tau}$, and $\bar{h}_{ed^\star} = \mathbb{E}[\min(\tau, D_{ed}^\star)]$.

Remark 4: Although of interest, it is beyond the scope of this paper to establish a guarantee of strict performance improvement of the proposed strategies. However, we do prove $J_\pi \leq J_b$ and show the strict improvements via various numerical examples. In addition, note that we believe that strict performance improvement guarantees could be derived by following a similar approach to the one in [32], where a switched system derived in a different context was studied. Such an approach entails rather long arguments, requiring concepts such as ergodicity, and it is, therefore, not pursued here.

III. Preliminary Results

In this section, for reasons of completeness and self-containedness, we discuss shortly the analysis needed to obtain the results in our preliminary work [24], which will be used as a benchmark, and how this leads to an initial result for the ed policy in Lemma 1. In order to analyze the proposed strategies, it is convenient to obtain a discrete-time description of the system, which we provide next.

A. Discretization

By discretization of system (1) at times t_k, $k \in \mathbb{N}$, we obtain

$$x_{k+1} = A(h_k)x_k + B(h_k)\hat{u}_k + w_k \tag{6}$$

where $A(h) := e^{A_c h}$ and $B(h) := \int_0^h e^{A_c s} B_c ds$. The disturbance is a sequence of zero-mean independent random vectors $w_k \in \mathbb{R}^{n_w}$, $k \in \mathbb{N}$, with covariance $\mathbb{E}[w_k(w_k)^\intercal] = W(h_k)$, where $W(h) := \int_0^h e^{A_c s} B_w B_w^\intercal e^{A_c^\intercal s} ds$.

We augment the state with the current input and define $\xi_k := [x_k^\intercal \; \hat{u}_k^\intercal]^\intercal$. The state evolution of the augmented system can then be written as

$$\xi_{k+1} = \mathcal{A}_{\gamma_k}(h_k)\xi_k + \mathcal{B}_{\gamma_k} u_k + \widehat{w}_k \tag{7}$$

[1]The additional factor α_b in (5) was not considered in our preliminary work [24], where the cost due to intersampling behavior was neglected. Typically, the value of α_b is small compared to $\text{tr}(P_b W_b)$, as was the case in [24], and a good approximation of the actual cost can be obtained by neglecting α_b.

where $\mathcal{A}_\gamma(h) := \begin{bmatrix} A(h) & B(h) \\ 0 & (1-\gamma)I_{n_u} \end{bmatrix}$, $\mathcal{B}_\gamma := \begin{bmatrix} 0 \\ \gamma I_{n_u} \end{bmatrix}$, and where $\widehat{w}_k = [w_k^\intercal \; 0]^\intercal$ with covariance $\widehat{W}(h) := \begin{bmatrix} W(h) & 0 \\ 0 & 0 \end{bmatrix}$.

The average cost can be written as

$$J = \limsup_{T\to\infty} \frac{1}{T}\mathbb{E}\left[\sum_{k=0}^{N(T)-1} g(\xi_k, h_k)\right] \tag{8}$$

where $N(T) := \min\{L \in \mathbb{N}_{[1,\infty]} \mid \sum_{k=0}^{L} h_k > T\}$ and $g(\xi, h) := \xi^\intercal Q(h)\xi + \alpha(h)$, with

$$Q(h) := \int_0^h e^{\begin{bmatrix} A_c & B_c \\ 0 & 0 \end{bmatrix}^\intercal s} \begin{bmatrix} Q_c & 0 \\ 0 & R_c \end{bmatrix} e^{\begin{bmatrix} A_c & B_c \\ 0 & 0 \end{bmatrix} s} ds$$

and

$$\alpha(h) := \text{tr}\left(Q_c \int_0^h \int_0^t e^{A_c s} B_w B_w^\intercal e^{A_c^\intercal s} ds dt\right) \tag{9}$$

which is the cost associated with the intersampling behavior of (1).

The model (7) can be used to describe the behavior for all proposed strategies. Later, we sometimes use the notation $\gamma|D$ and $h|D$ to indicate that the probability distributions of those variables depend on the deadline D.

B. Performance of the Basic Control Strategies

Each cost J_b in (5) is associated with an optimal control policy (note that τ_k is not known at time t_k)

$$u_k = -K_b\xi_k, \quad b \in \{d^\star, e, ed^\star\} \tag{10}$$

where the expressions for the gains are given in Appendix A. We assume that mild conditions for mean-square stabilizability hold (see [33, Prop. 3.42] and [25, Theor. 6.1]), such that solutions (10) are well defined. Note that for the event-driven case e, $\gamma_k = 1$ for all $k \in \mathbb{N}$. While for the d and e cases, the use of the results in [25] is straightforward, for the ed case, it is possible to observe that a new probability distribution for h_k can be defined as a function of the probability distribution of τ_k, determined by the cdf

$$F^h(\tau, D) := \begin{cases} F(\tau), & \text{if } \tau < D \\ 1, & \text{if } \tau \geq D \end{cases}$$

and the results in [25] also apply.

For compactness, we introduce the following notation. For a Bernoulli variable γ and a random variable h, random matrices X and Y as in Section III-A that depend on γ and h, and some matrix P, the expected value $\mathbb{E}_{\gamma,h}[X_\gamma(h)^\intercal PY_\gamma(h)]$ is denoted as $\overline{X_\gamma(h)^\intercal PY_\gamma(h)}$, and analogously, $\mathbb{E}_{\gamma,h}[X_\gamma(h)]$ is denoted as $\overline{X_\gamma(h)}$.[2]

[2]As a special case, we have that $\overline{X_\gamma^\intercal PY_\gamma} = p(X_1^\intercal PY_1) + (1-p)(X_0^\intercal PY_0)$, where X_γ and Y_γ are random matrices depending on Bernoulli random variable γ and p is the probability of success, given by $p = \Pr[\gamma = 1]$. Additionally, we have that $\overline{X(\tau)^\intercal PY(\tau)} = \int_0^\infty [X(s)^\intercal PY(s)]dF(s)$. Finally, if a deadline D is given, for τ with cdf $F^h(\tau, D)$ and $\gamma = 1$ if $\tau \leq D$ and $\gamma = 0$ if $\tau > D$, we have that $\overline{X_\gamma(h)^\intercal PY_\gamma(h)} = \int_0^D [X_1(s)^\intercal PY_1(s)]dF(s) + (1-F(D))[X_0(D)^\intercal PY_0(D)]$.

Intuitively, the $ed^\star$ strategy seems to be better than both the $d^\star$ and e strategies. From the derivation in this section, we obtain directly the following result.

Lemma 1 ($ed^\star$ is better than e): The cost (2) of the event-driven policy with optimal deadline is not larger than that of the event-driven policy, i.e.,

$$J_{ed^\star} \leq J_e.$$

The proof follows directly from the fact that $J_{ed}(D) \to J_e$ for $D \to \infty$ and the policy e is contained in the class of policies ed parameterized by D. ■

There may exist (pathological) cases for which updating the control before the deadline has a negative effect on performance. Thus, a guarantee analogous to Lemma 1 for d does not directly exist. However, the main results of this paper, for a new strategy, do give such a guarantee.

IV. Control Policy and Main Results

In this section, the main results are presented. First, we propose the novel switching strategy that leads to a performance guarantee, which is formalized in the main theorem. Subsequently, we present a switching strategy that extends the main result with a deadline-optimization scheme.

A. Two-Policy Control (&)

We propose using a switched approach to the problem formulated in Section II-B. In particular, we allow the system to choose online which type of strategy, that is, $d/e/ed$, to use for the next update instance. Actually, the result is derived for the combination of only $d^\star$ and $ed^\star$ because Lemma 1 shows that the performance of $ed^\star$ is not larger than that of e. The proposed control policy for this case is denoted as $d^\star \& ed^\star$. The result in Theorem 1 shows that this policy leads to better performance than using either $d^\star$ or $ed^\star$ all the time.

The idea behind our policy is to choose, at each sampling instance, the control strategy to use during the next interval, denoted by $\sigma_k \in \{d^\star, ed^\star\}$, while assuming that either of the *base* policies, denoted by $b_k \in \{d^\star, ed^\star\}$, can be used all the time afterwards, such that the expected future cost is the smallest. After the next interval, the impact of disturbances is neglected in the switching condition to ensure that the cost of the lookahead predictions can be computed. Now, at each t_k, $k \in \mathbb{N}$, four switching options are available, and we establish switching conditions to determine the best option.

Let $p^\star := \arg\min_{p\in\{d^\star, ed^\star\}} J_p$ select the best periodic base policy with an optimal deadline from the possible base policies, whose costs were defined as J_b in (5). Now, we define two functions that are to be used in the switching conditions. First, we define a value function $V_p(\xi) := \xi^\intercal P_p \xi$, where $p \in \{d^\star, ed^\star\}$ and P_p is defined as in (5), that is, as solutions to the Riccati equations given in Appendix A. Second, we define a difference function

$$V^\Delta(\xi_k, m, b) := \mathbb{E}\left[V_{p^\star}(\xi_{k+1}) - V_b(\xi_{k+1}) \mid \xi_k, \begin{bmatrix}\sigma_k \\ b_k\end{bmatrix} = \begin{bmatrix} m \\ b\end{bmatrix}\right]$$

where ξ_{k+1} follows (by a prediction step) from (7). In particular, the value of ξ_{k+1} follows from (7) given that $\sigma_k = m$ and $b_k = b$, meaning that in (7), u_k is the optimal input for the next interval, which is defined below in (12), and that γ_k and h_k are random variables that depend on the deadline D_k corresponding to the choice $\sigma_k = m$, that is, $\gamma_k | D^\star_{\sigma_k}$ and $h_k | D^\star_{\sigma_k}$ [see also (13)]. Furthermore, we define a set-valued function

$$\mathcal{S}(\xi) := \{m, b \in \{d^\star, ed^\star\} \mid V^\Delta(\xi, m, b) \leq 0\}$$

mapping ξ into the choices m and b that guarantee that V^Δ is nonpositive. Note that, by definition, for any ξ, m, $V^\Delta(\xi, m, p^\star) = 0$ and, therefore, the set $\mathcal{S}(\xi)$ is nonempty.

The proposed control policy $d^\star \& ed^\star$ is the following function of the state ξ_k to be evaluated for all t_k, $k \in \mathbb{N}$:

$$\begin{bmatrix}\sigma_k \\ b_k\end{bmatrix} = \arg \min_{(m,b)^\intercal \in \mathcal{S}(\xi_k)} \xi_k^\intercal Z^\star_{m,b} \xi_k + \beta^\star_{m,b} \tag{11}$$

$$u_k = -K^\star_{\sigma_k, b_k} \, \xi_k \tag{12}$$

where the arguments are given by

$$\begin{aligned} Z^\star_{m,b} &:= \overline{\mathcal{A}_{\gamma_k}(h_k)^\intercal P_b \mathcal{A}_{\gamma_k}(h_k)} + \overline{Q(h_k)} \\ &\quad - K^\star_{m,b}{}^\intercal (\overline{\mathcal{B}^\intercal_{\gamma_k} P_b \mathcal{B}_{\gamma_k}}) K^\star_{m,b} \\ K^\star_{m,b} &:= \overline{\mathcal{B}^\intercal_{\gamma_k} P_b \mathcal{B}_{\gamma_k}}^\dagger \overline{\mathcal{B}^\intercal_{\gamma_k} P_b \mathcal{A}_{\gamma_k}(h_k)} \\ \beta^\star_{m,b} &:= \mathrm{tr}(P_b W_m) + \alpha_m - \frac{\bar{h}_m}{\bar{h}_{p^\star}} c_{p^\star} \end{aligned}$$

with $c_{p^\star}$ given in (5), and $\alpha_m = \mathbb{E}[\alpha(h_k) \mid \sigma_k = m]$ as defined in Appendix B, and where the distribution for h_k (and γ_k) depends on the value of the deadline $D^\star$ in

$$h_k = \begin{cases} D^\star_d, & \text{if } \sigma_k = d^\star \\ \min\{\tau_k, D^\star_{ed}\}, & \text{if } \sigma_k = ed^\star \end{cases} \tag{13}$$

corresponding to the value of $\sigma_k = m$. The symbol $\dagger$ denotes the pseudoinverse. The expectations can be numerically computed (using footnote 2). In the first two terms, the scalars $\beta^\star_{m,b}$ contain the cost due to noise over interval h_k, and in the third term, they contain a correction for the time difference between (the expectation of) the interval h_k and that of the best base policy. The policy selects (σ_k, b_k) as the values that minimize the right-hand side of (11) subject to the condition that $V^\Delta(\xi_k, \sigma_k, b_k) \leq 0$. The condition $V^\Delta(\xi_k, \sigma_k, b_k) \leq 0$ guarantees that switching to a different base policy while neglecting the disturbances after the next interval does not cause a performance loss.

Let the value of performance index (2) obtained for the policy (11)–(13) be denoted as $J_{d^\star \& ed^\star}$. The following result is the main result of this paper.

Theorem 1: The cost (2) of the two-policy approach given by (11)–(13) is not larger than that of both base policies ed and d in the sense that

$$J_{d^\star \& ed^\star} \leq J_{ed^\star} \qquad \text{and} \qquad J_{d^\star \& ed^\star} \leq J_{d^\star}.$$

The proof is given in Appendix C and resorts to Doob's optional sampling theorem [29]. ■

The following remark explains a relaxation of the switching condition, which will be used in the numerical example in Section V.

Remark 5: From the proof of Theorem 1, one can see that Theorem 1 also holds if the condition $V^{\Delta}(\xi_k, m, b) \leq 0$ on $\mathcal{S}(\xi)$ is relaxed to $V^{\Delta}(\xi_k, m, b) \leq \Delta(\xi_k, p^\star, p^\star, m, b)$, where Δ is defined in (18).

Remark 6: Note that $V^{\Delta}(\xi_k, m, b) \leq 0$ is directly satisfied for all ξ_k and m if $P_{p^\star} \preceq P_b$ for all b.

Remark 7: The result of Theorem 1 would directly extend to a policy $d^\star \& e$ if $D^\star_{ed} = \infty$ would be selected. The derivation of the policy is omitted for brevity. The result for this case is summarized in the following corollary.

Corollary 1: The cost (2) of the two-policy approach given by (11)–(13) when $D^\star_{ed} \to \infty$ is not larger than that of both base policies e and d in the sense that

$$J_{d^\star \& e} \leq J_e \qquad \text{and} \qquad J_{d^\star \& e} \leq J_{d^\star}.$$

The proof follows the same arguments as the ones used to prove Theorem 1. ∎

B. Online Deadline Optimization (s)

In our preliminary work [24], we proposed the idea of deadline optimization in the form of a self-triggered policy for the periodic deadline-driven controller. Here, we show that our idea of online deadline optimization can be extended to all policies that consider a deadline, including the two-policy case and the ed case.

Next, we describe extended switching conditions that include an optimization procedure for the deadline for the two-policy strategy. Analogous to $\mathcal{S}(\xi)$, we define the extended set that includes a deadline variable

$$\mathcal{S}^s(\xi) := \{m \in \{d, ed\}, D \in \mathcal{D}, \\ b \in \{d^\star, ed^\star\} \mid V^{\Delta D}(\xi, m, D, b) \leq 0\}$$

where, for mathematical and practical convenience, $D^\star_b \in \mathcal{D}$ for all $b \in \{d^\star, ed^\star\}$, where $\mathcal{D} \subset \mathbb{R}_{>0}$ is a finite but possibly arbitrarily large set of allowable deadlines, and m corresponds to a method with deadline, and

$$V^{\Delta D}(\xi_k, m, D, b) := \mathbb{E}\left[V_{p^\star}(\xi_{k+1}) - V_b(\xi_{k+1}) \mid \xi_k, \begin{bmatrix} \sigma_k \\ b_k \\ D_k \end{bmatrix} = \begin{bmatrix} m \\ b \\ D \end{bmatrix} \right]$$

now has an additional argument $D \in \mathcal{D}$ for the choice of deadline compared to $V^{\Delta}(\xi_k, m, b)$.

We propose using the following control policy:

$$\begin{bmatrix} \sigma_k \\ D_k \\ b_k \end{bmatrix} = \arg \min_{(m,D,b)^\intercal \in \mathcal{S}^s(\xi)} \xi_k^\intercal Z^s_{m,b}(D)\xi_k + \beta^s_{m,b}(D) \quad (14)$$

$$u_k = -K^s_{\sigma_k, b_k}(D_k)\xi_k \quad (15)$$

where

$$Z^s_{m,b}(D) := \overline{\mathcal{A}_{\gamma_k}(h_k)^\intercal P_b \mathcal{A}_{\gamma_k}(h_k)} + \overline{Q(h_k)} \\ - K^s_{m,b}(D)^\intercal(\mathcal{B}_1^\intercal P_b \mathcal{B}_1 F(D))K^s_{m,b}(D)$$

$$K^s_{m,b}(D) = (\mathcal{B}_1^\intercal P_b \mathcal{B}_1 F(D))^\dagger (\overline{\mathcal{B}^\intercal_{\gamma_k} P_b \mathcal{A}_{\gamma_k}(h_k)})$$

$$\beta^s_{m,b}(D) = \mathrm{tr}(P_b W_m(D)) + \alpha_m(D) - \frac{\bar{h}_{m,D}}{\bar{h}_{p^\star}} c_{p^\star}$$

and where the distribution for h_k (and γ_k) depend on the value of the deadline D in

$$h_k = \begin{cases} D_k, & \text{if } \sigma_k = d \\ \min\{\tau_k, D_k\}, & \text{if } \sigma_k = ed \end{cases} \quad (16)$$

corresponding to the value of $\sigma_k = m$. In particular, $\bar{h}_{m,D} = \mathbb{E}[h_k \mid D_k = D, \sigma_k = m]$, $W_m(D) := \mathbb{E}[\widehat{W}(h_k) \mid D_k = D, \sigma_k = m]$, and $\alpha_m(D) := \mathbb{E}[\alpha(h_k) \mid D_k = D, \sigma_k = m]$.

Let the cost of the above policy (14)–(16) be denoted as $J^s_{d\&ed}$. We obtain the following result, which can be seen as an extension of Theorem 1.

Theorem 2: The cost (2) of the online deadline-optimization policy (14)–(16) is not larger than that of the two-policy method $d\&ed$ with fixed optimal deadlines in the sense that

$$J^s_{d\&ed} \leq J_{d^\star \& ed^\star}.$$

The proof is given in Appendix D. Again, the condition $V^{\Delta D}(\xi_k, m, D, b) \leq 0$ can be relaxed as explained in Remark 5. ∎

Let the cost of the policy (14)–(16), when $b_k = ed^\star$ and $\sigma_k = ed$ for all $k \in \mathbb{N}$, be denoted as J^s_{ed}. By restricting the choice of the base policy, we obtain the following result.

Corollary 2: The performance (2) of the online deadline-optimization policy (14)–(16), when $b_k = ed^\star$ for all $k \in \mathbb{N}$, is not larger than that of the base policy $ed^\star$ in the sense that

$$J^s_{ed} \leq J_{ed^\star}.$$

In the next section, we show numerical results for the proposed policies and the performance gain that can be achieved by addressing the delay probability directly. ∎

Remark 8: Note that the above approach requires the computation of the argument in (14) for many different options. By reducing the search space for the deadline, the computational load can be easily reduced to match the available computational capacity.

C. Computational Complexity

Here, we consider the computational complexity of the online optimization methods considered. For each argument for the minimization (11) and (14), the matrices Z and the scalars β can be computed offline *a priori* and, for example, stored in memory. The same holds for the gains in (12) and (15). Therefore, to compute the optimal arguments, it is required to compute and compare the terms $\xi^\intercal Z \xi + \beta$ for $\#deadlines \times \#switchingpolicies \times \#basepolicies$ switching options in the sets $\mathcal{S}$ and $\mathcal{S}^s$. Note that the size of the sets $\mathcal{S}$ and $\mathcal{S}^s$ depends on the current state and is upper bounded by the total number of switching options. Knowledge of sets $\mathcal{S}$ and $\mathcal{S}^s$ may be used to reduce computations by limiting online the number of options in (11) and (14), but in some implementations, it may be more convenient to compute all possible options. The computation of V^{Δ} for the sets $\mathcal{S}$ and $\mathcal{S}^s$ requires two computations of the form $\xi^\intercal X\xi + Y$, where X and Y are a matrix and a scalar, which can be computed offline for each switching option, taking the forms $(A - BK)^\intercal P(A - BK)$ and $\mathrm{tr}(PW)$, respectively. The number of computations in the terms $\xi^\intercal Z\xi + \beta$ scales quadratically with the state dimension and linearly with the number of switching parameters, but they can be computed in parallel for all elements in $\mathcal{S}$ or $\mathcal{S}^s$ (or all switching options). For the control inputs, the multiplication Kx scales linearly with the state.

From this analysis, we conclude that, typically, the time required to compute (11) and (14) is small when compared to the communication or data-processing computation delay modeled by F. However, in cases in which these computation times are non-negligible (because the initial computation/communication delay modeled by F can be small), it can be incorporated in a new probability distribution modeling the sum of delays, say $\bar{F}$. Therefore, the methods in this paper can be used to analyze this case as well. Note, however, that in the latter case, one should compare the simpler $d^\star$ and e methods considering the initial distribution F with the $ed^\star$ method considering $\bar{F}$ and, therefore, the method $ed^\star$ does not guarantee better performance than e *a priori*.

V. Numerical Results

In this section, we compare the performance of the proposed strategies on a second-order system taking the form (1) with

$$A_c = \begin{bmatrix} 0 & 1 \\ \frac{g}{l} & -\frac{d}{ml} \end{bmatrix}, \quad B_c = \begin{bmatrix} 0 \\ 1 \end{bmatrix}, \quad B_w = \begin{bmatrix} 0.05 & 0 \\ 0 & 0.05 \end{bmatrix}$$

which represents a linearized inverted pendulum system with force input, gravitational acceleration $g = 10\ \mathrm{ms}^{-1}$, mass of pendulum $m = 0.25$ kg, length $l = 0.5$ m, and damping coefficient $d = 1\ \mathrm{N} \cdot m/rad \cdot \mathrm{s}^{-1}$.

The cost function matrices in (2) are taken as

$$Q_c = \begin{bmatrix} 20 & 1 \\ 1 & 20 \end{bmatrix}, \qquad R_c = \begin{bmatrix} 3 \end{bmatrix}.$$

We consider for the delay both a Gamma distribution f_1 with shape and scale parameters $k = 3$ and $\theta = 4/100$, respectively, and the piecewise constant two-block distribution

$$f_2(\tau) = \begin{cases} 4.5, & \text{if } \tau \in [0.05, 0.25) \\ 0.1, & \text{if } \tau \in [0.50, 0.60) \\ 0, & \text{otherwise.} \end{cases}$$

Note that although f_1 does not satisfy the condition $F_1(\epsilon) = 0$, it can easily be adapted (see Remark 1) to meet such an assumption, with a small ϵ, without impacting on the results.

First, for f_1, the optimal solutions for the base policies are computed. For D in the interval $[10^{-3}, 1]$, the stochastic Riccati equations for $P_d(D)$, P_e, and $P_{ed}(D)$ (see Appendix A) are solved iteratively, with initial value $P = 10^{-4} I_{n_x}$, up to accuracy 10^{-4} of the mean square error. All cost values are depicted as a function of the deadline in Fig. 2. Subsequently, the optimal values $J_{d^\star}$ and $J_{ed^\star}$ of $J_d(D)$ and $J_{ed}(D)$, respectively, are found for their respective optimal deadlines $D_d^\star$ and $D_{ed}^\star$. The optimal deadline values, are found to be $D_d^\star = 0.1508$ and $D_{ed}^\star = 0.3307$. Furthermore, $\bar{\tau} = 0.1200$ and $\bar{h}_{ed^\star} = 0.1194$. The optimal costs $J_{d^\star}$, J_e, and $J_{ed^\star}$ for f_1 are given in Table I. Although the contribution of the intersampling behavior is very small, since $\alpha_b \approxeq 0.0014$ for all the base policies, it has been taken into account in the calculation of the costs and our simulations.

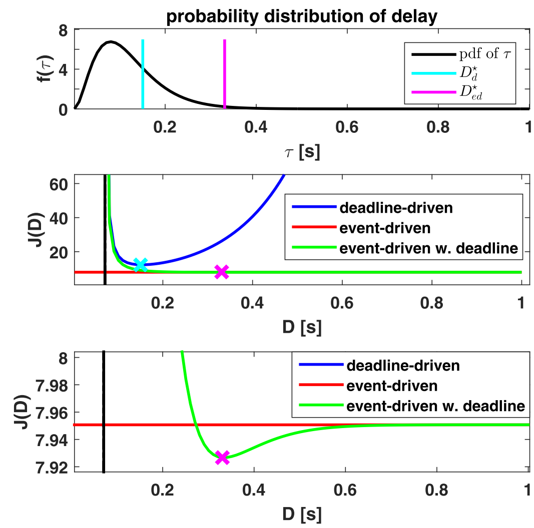

Fig. 2. Comparison of the performance of the three base policies varying with the deadline for $f = f_1$. The optimal deadlines are marked (x) and shown in the pdf (top figure). The bottom figure is a zoom of the middle figure. The black vertical line indicates the stability limit for the deadlines.

TABLE I
Values of Performance Index (2) for $f = f_1$

Method	Base cost J_b (analytical)	Base cost J_b (simulation)	Cost with deadline optimization J^s (simul.)
d	12.3283	12.3305 ($=J_{d^\star}$)	11.3743 ($=J_d^s$)
e	7.9507	7.9591 ($=J_e$)	n.a.
ed	7.9267	7.9307 ($=J_{ed^\star}$)	7.2421 ($=J_{ed}^s$)
$d\&ed$	n.a.	7.7993 ($=J_{d^\star\&ed^\star}$)	7.3047 ($=J_{d\&ed}^s$)
$d\&e$	n.a.	7.7850 ($=J_{d^\star\&e}$)	7.3514 ($=J_{d\&e}^s$)

For this example, observe that the green line in Fig. 2 is below the blue line everywhere, indicating that even for a suboptimal deadline, the ed approach outperforms the deadline-driven approach. As expected from Lemma 1, $J_{ed}(D)$ performance approximates J_e for large D.

For each delay value and each switching option, all variables are computed *a priori* to speed up computation for Monte Carlo (MC) simulations. For $\xi_0^\intercal = [0, 0, 0]^\intercal$, we run 40 "long" MC simulations for $t \in [0, 24000]$ s such that the average costs have approximately converged for each simulation. Then, the costs are averaged over the MC simulations, and the values are given in Table I. Due to the limited simulation time and the limited

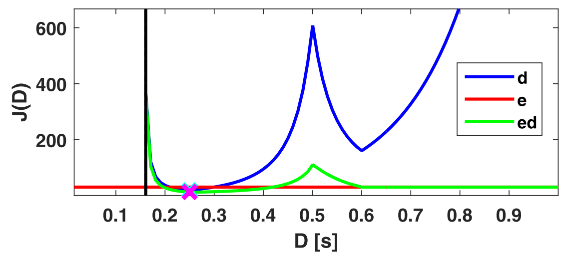

Fig. 3. Comparison of the performance of the three base policies varying with the deadline for $f = f_2$.

TABLE II
VALUES OF PERFORMANCE INDEX (2) FOR $f = f_2$

Method	Base cost J_b (analytical)	Base cost J_b (simulation)	Cost with deadline optimization J^s (simul.)
d	19.5069	19.5230 ($=J_{d^\star}$)	18.0640 ($=J^s_d$)
e	30.6565	36.9563 ($=J_e$)	n.a.
ed	11.9640	12.0625 ($=J_{ed^\star}$)	11.2330 ($=J^s_{ed}$)
$d\&ed$	n.a.	12.0023 ($=J_{d^\star\&ed^\star}$)	11.2774 ($=J^s_{d\&ed}$)

number of MC simulations, the cost is not completely averaged over the probability space, leaving a small error.

The cost differences that support our theorems are indeed visible, thereby underlining the results. Note that for the two-policy approaches, the relaxed switching conditions (see Remark 5) are used, giving a small additional performance gain of 1–2%. It is notable that, while adding an optimal deadline only gives small improvement over the event-driven case ($J_{ed^\star} \approx J_e$), the fact that deadline optimization can be enabled brings significant advantage of 6–9% compared to the nonswitching case (e.g., $J^s_{ed} < J_e$ and $J^s_{d\&ed} < J_{d^\star\&ed^\star}$). Furthermore, the strategy with deadline optimization on ed, which builds upon our previously proposed policies, performs better than the two-policy approach in the sense that $J^s_{ed} < J^s_{d\&ed}$. However, such a performance improvement is not guaranteed formally, and the converse, i.e., $J^s_{ed} > J^s_{d\&ed}$, may also occur for different examples. Furthermore, a quantification of the performance difference for the deadline-optimization approaches may only be obtained through simulation or experiments.

For f_2, the results are given in Fig. 3 and Table II. Similar cost benefits as for f_1 are observed. For this example, it is found that the optimal deadlines are the same ($D^\star_d = D^\star_{ed} = 0.2508$), but again the ed policy performs better. Furthermore, $\bar{\tau} = 0.1900$ and $\bar{h}_{ed^\star} = 0.1592$. Both d and ed are better than e for this case, showing that pure event-driven control is not necessarily better than periodic control with a deadline. The large deviation of e from the analytical values is due to the fact that the cost has not yet converged.

As a final note, we observed a significant amount of switching occurrences without any recognizable pattern, as was expected since the switching depends on realizations of the random disturbances. While one might expect that only the best base policy, denoted by $p^\star$, is chosen in the two-policy case, we observed that this is not necessarily the case. While the number of occurrences is small, the selection $\sigma_k \neq p^\star$ is recurring, indicating that for some parts of the state space switching to a different base policy is the best strategy.

VI. Conclusion

This paper presents novel control policies for linear systems subject to actuation delays with a known probability distribution. From optimal control policies for deadline-driven, event-driven, and event-driven control with a deadline, analytic solutions for optimal performance/cost and deadline values were deduced. The proposed "switched" policy can combine the different benefits of nonswitching policies to improve closed-loop performance. The performance of this policy is proven to be better than that of any of the nonswitching policies. Furthermore, the idea of deadline optimization that was presented in preliminary work is extended to both the proposed "event-driven with deadline" policy as well as the "switched" policy. This allows for additional guaranteed performance improvement, which was not attainable in previous event-driven or deadline-driven approaches. Numerical examples illustrate the results and give insight in the tradeoffs in systems with delay, showing that gains of 6–9% can easily be obtained with the proposed policies. Performance relations between switched policies with deadline optimization are still subject of study. Moreover, future work also includes the output-feedback counterpart, which adds an estimation problem influenced by stochastic delays, and studies of robustness with respect to model uncertainty.

Appendix

A. Riccati Equations for the Base Policies

To compute the cost for event-driven control with deadline, it is required to solve, for $P_{ed}(D) \succ 0$, the generalized Riccati equation

$$P_{ed}(D) = \overline{\mathcal{A}_{\gamma|D}(h|D)^\intercal P_{ed}(D)\mathcal{A}_{\gamma|D}(h|D)} + \overline{Q(h|D)} - K_{ed}(D)^\intercal G_{ed}(D)K_{ed}(D) \quad (17)$$

$$G_{ed}(D) := \overline{\mathcal{B}^\intercal_{\gamma|D}P_{ed}(D)\mathcal{B}_{\gamma|D}} = F(D)\mathcal{B}^\intercal_1 P_{ed}(D)\mathcal{B}_1$$

$$K_{ed}(D) := G_{ed}(D)^\dagger(\mathcal{B}^\intercal_1 P_{ed}(D)\int_0^D \mathcal{A}_1(s)\mu(ds))$$

and to compute $W_{ed}(D) := \overline{\widehat{W}(h|D)} = \int_0^D \widehat{W}(s)dF(s) + (1-F(D))\widehat{W}(D)$; see also footnote 2 on page 11. The solution to the Riccati equation can be found by, e.g., the iteration $P^{k+1}_{ed}(D) = Ric(P^k_{ed}(D))$ for $k \geq 0$ where $Ric(\cdot)$ is the function of the right-hand side of (17). One can recover the Riccati equations for d and e, respectively, since considering any new probability measure with $\mu((0,D)) = 0$ and $\mu(\{D\}) = F(D)$ gives

$$P_d(D) = \overline{\mathcal{A}_\gamma(D)^\intercal P_d(D)\mathcal{A}_\gamma(D)} + Q(D) - K_d(D)^\intercal G_d(D)K_d(D)$$

$$G_d(D) := \overline{\mathcal{B}^\intercal_\gamma P_d \mathcal{B}_\gamma} = F(D)\mathcal{B}^\intercal_1 P_d(D)\mathcal{B}_1$$

$$K_d(D) := G_d(D)^\dagger(\overline{\mathcal{B}^\intercal_\gamma P_d(D)\mathcal{A}_\gamma(D)})$$

and, alternatively, by letting $D \to \infty$, we have

$$P_e = \overline{\mathcal{A}_1(\tau)^\intercal P_e \mathcal{A}_1(\tau)} + \overline{Q(\tau)} - K_e{}^\intercal G_e K_e$$
$$G_e := \mathcal{B}_1^\intercal P_e \mathcal{B}_1$$
$$K_e := G_e{}^\dagger(\mathcal{B}_1^\intercal P_e \overline{\mathcal{A}_1(\tau)}).$$

Furthermore, $W_d(D) = \widehat{W}(D)$ for a given value of D, and $W_e = \overline{\widehat{W}(\tau)}$.

B. Cost Due to Intersampling Behavior

For the base policies $b \in \{d, e, ed\}$, where $h_k, k \in \mathbb{N}$, are i.i.d., in (5), the contribution $\alpha_b, b \in \{d, e, ed\}$, of the intersampling behavior of the Wiener process is given by the average value $\mathbb{E}[\alpha(h_k)]$, where $\alpha(h)$ is given in (9). This follows from $\limsup_{T\to\infty} \frac{1}{T}\mathbb{E}[\sum_{k=0}^{N(T)-1} \alpha(h_k)] = \frac{1}{\mathbb{E}[h_k]}\,\mathbb{E}[\alpha(h_k)] = \frac{1}{\bar{h}_b}\alpha_b$, which can be concluded from [34, Prop. 3.4.1]. Specifically, for a given deadline $D \in \mathbb{R}$ and policy $b \in \{d, e, ed\}$, $\alpha_b(D) := \mathbb{E}[\alpha(h_k) \mid D_k = D, \sigma_k = b]$. Then, in (5), we have $\alpha_d(D) := \alpha(D)$, $\alpha_e := \overline{\alpha(\tau)}$, $\alpha_{ed}(D) := \overline{\alpha(h|D)}$; see also footnote 2 on page 11.

C. Proof of Theorem 1

We drop the superscript $\star$ for Z and β for brevity, and let d and ed be represented by m or b or p. Note that for each option, $\xi_k^\intercal Z_{m,b}\xi_k + tr(P_b W(h_k|D, m)) + \alpha(h_k|D, m)$ is the minimal value of the optimization

$$\min_{u_k} \mathbb{E}\left[\int_{t_k}^{t_k+h_k|D} x(t)^\top Q_c x(t) + u(t)^\top R_c u(t)dt + \xi(t_k + h_k|D)^\top P_b \xi(t_k + h_k|D)\right].$$

This follows from standard LQG arguments (see [30]).

Define the difference in arguments in (11) by

$$\Delta(\xi, m_1, b_1, m_2, b_2) := \left[\xi^\intercal Z_{m_1,b_1}\xi + \text{tr}(P_{b_1}W_{m_1}) + \alpha_{m_1} - \frac{\bar{h}_{m_1}}{\bar{h}_p}c_p\right] - \left[\xi^\intercal Z_{m_2,b_2}\xi + \text{tr}(P_{b_2}W_{m_2}) + \alpha_{m_2} - \frac{\bar{h}_{m_2}}{\bar{h}_p}c_p\right] \quad (18)$$

where c_p is given in (5), such that, for $p \in \{d^\star, ed^\star\}$, we have

$$\Delta(\xi, p, p, p, p) := 0$$
$$\Delta(\xi, p, p, m_2, b_2) := [\xi^\intercal P_p \xi + c_p] - \left[\xi^\intercal Z_{m_2,b_2}\xi + \text{tr}(P_{b_2}W_{m_2}) + \alpha_{m_2} - \left(\frac{\bar{h}_{m_2} - \bar{h}_p}{\bar{h}_p}\right)c_p\right]$$

where $\bar{h}_m = \mathbb{E}[h_k \mid \sigma_k = m]$, such that, e.g., $\bar{h}_m = D$ if $m = d$, and we use the fact that $Z_{p,p} = P_p$. Observe that the switching condition (11) aims to maximize the value of $\Delta(\xi_k, p, p, \sigma_k, b_k)$ for $p = p^\star$. Note that it is always allowed to choose the optimal base policy $\sigma_k = p, b_k = p$ since it directly satisfies the condition on V^Δ. It corresponds to a value $\Delta(\xi_k, p, p, p, p) = 0$. Hence, $\Delta(\xi_k, p, p, \sigma_k, b_k)$ is non-negative since any choice of σ_k, b_k following from (11) satisfies $\Delta(\xi_k, p, p, \sigma_k, b_k) \geq \Delta(\xi_k, p, p, p, p) = 0$.

Moreover, we consider

$$\mathbb{E}\left[g(\xi_k, h_k) + V_b(\xi_{k+1}) \mid \xi_k, \begin{bmatrix}\sigma_k\\ b_k\end{bmatrix} = \begin{bmatrix}m\\ b\end{bmatrix}\right]$$
$$= \xi_k^\intercal Z_{m,b}\xi_k + \text{tr}(P_b W_m) + \alpha_m$$
$$= \xi_k^\intercal Z_{m,b}\xi_k + \text{tr}(P_b W_m) - \xi_k^\intercal P_p \xi_k + \xi_k^\intercal P_p \xi_k + \left(\frac{\bar{h}_m - \bar{h}_p}{\bar{h}_p} - \frac{\bar{h}_m - \bar{h}_p}{\bar{h}_p}\right)c_p + \alpha_m$$
$$= \frac{\bar{h}_m}{\bar{h}_p}c_p - c_p - \xi_k^\intercal P_p \xi_k + \xi_k^\intercal P_p \xi_k + \xi_k^\intercal Z_{m,b}\xi_k + \text{tr}(P_b W_m) - \left(\frac{\bar{h}_m - \bar{h}_p}{\bar{h}_p}\right)c_p + \alpha_m$$
$$= \frac{\bar{h}_m}{\bar{h}_p}c_p - \Delta(\xi_k, p, p, m, b) + V_p(\xi_k).$$

As defined in Section IV-A, $p^\star = \arg\min_{p\in\{d^\star, ed^\star\}} J_p$. The one-step cost of the proposed control policy is then given by the difference

$$\mathbb{E}\left[g(\xi_k, h_k) \mid \xi_k, \begin{bmatrix}\sigma_k\\ b_k\end{bmatrix} = \begin{bmatrix}m\\ b\end{bmatrix}\right] = \frac{\bar{h}_m}{\bar{h}_{p^\star}}c_{p^\star} - \Delta(\xi_k, p^\star, p^\star, m, b) + V_{p^\star}(\xi_k) - \mathbb{E}\left[V_b(\xi_{k+1}) \mid \xi_k, \begin{bmatrix}\sigma_k\\ b_k\end{bmatrix} = \begin{bmatrix}m\\ b\end{bmatrix}\right]. \quad (19)$$

Note that the condition $V^\Delta(\xi_k, m, b) \leq 0$, guarantees that

$$\mathbb{E}\left[V_{p^\star}(\xi_{k+1}) - V_b(\xi_{k+1}) \mid \xi_k, \begin{bmatrix}\sigma_k\\ b_k\end{bmatrix} = \begin{bmatrix}m\\ b\end{bmatrix}\right] \leq 0.$$

Define

$$G_N := \sum_{k=0}^{N} g(\xi_k, h_k), \quad E_N := \sum_{k=0}^{N} \mathbb{E}\left[g(\xi_k, h_k) \mid \mathrm{I}_k\right].$$

Note that $g(\xi_k, h_k)$ given I_k is a random variable since I_k includes ξ_k but h_l only for $l < k$. We have that the process $X := (X_k)_{k\in\mathbb{N}}$, with $X_k := G_k - E_k$, is a martingale with respect to the filtration associated with I_{k+1}, since

$$\mathbb{E}\left[X_{k+1} \mid \mathrm{I}_{k+1}\right] = \mathbb{E}\left[X_k + g(\xi_{k+1}, h_{k+1}) - \mathbb{E}\left[g(\xi_{k+1}, h_{k+1}) \mid \mathrm{I}_{k+1}\right] \mid \mathrm{I}_{k+1}\right] = \mathbb{E}\left[X_k \mid \mathrm{I}_{k+1}\right] = X_k.$$

Note that $N(T)$ is a stopping time w.r.t. (X_k, I_{k+1}), which has finite expectation for given T, i.e., $\mathbb{E}[N(T)] < \infty$ since $h_k \geq h^{\min} > 0$ for all $k \in \mathbb{N}$, and that $N(T) \to \infty$ as $T \to \infty$. Provided that we prove that there exists some constant $c \in \mathbb{R}$ such that $\mathbb{E}\left[|X_{k+1} - X_k| \mid \mathrm{I}_{k+1}\right] \leq c$ for all $k < N(T)$ for $k \in \mathbb{N}$, which we will do in the following, we can apply Doob's optional sampling theorem (see, e.g., [35, Th. 9, Sec. 12.5] or [29, Th. 2.2, Ch. VII]) and have

$$\mathbb{E}\left[X_{N(T)}\right] = \mathbb{E}\left[X_0\right] = 0.$$

where we use the fact that

$$\mathbb{E}\left[X_0\right] = \mathbb{E}\left[g(\xi_0, h_0) - \mathbb{E}\left[g(\xi_0, h_0) \mid \mathrm{I}_0\right]\right] = 0.$$

This implies that $\mathbb{E}[G_{N(T)}] = \mathbb{E}[E_{N(T)}]$. Furthermore, since by the law of total expectation (or tower rule), $\mathbb{E}\left[g(\xi_{N(T)}, h_{N(T)}) - \mathbb{E}\left[g(\xi_{N(T)}, h_{N(T)}) \mid \mathrm{I}_{N(T)}\right]\right] = 0$, we have that

$$\mathbb{E}\left[\sum_{k=0}^{N(T)-1} g(\xi_k, h_k)\right] = \mathbb{E}\left[\sum_{k=0}^{N(T)-1} \mathbb{E}\left[g(\xi_k, h_k) \mid \mathrm{I}_k\right]\right]. \tag{20}$$

We have that

$$\begin{aligned}
&\mathbb{E}\left[|X_{k+1} - X_k| \mid \mathrm{I}_{k+1}\right] = \\
&= \mathbb{E}\left[\left|g(\xi_{k+1}, h_{k+1}) - \mathbb{E}\left[g(\xi_{k+1}, h_{k+1}) \mid \mathrm{I}_{k+1}\right]\right| \mid \mathrm{I}_{k+1}\right] \\
&\leq \mathbb{E}\left[\left|g(\xi_{k+1}, h_{k+1})\right| + \left|\mathbb{E}\left[g(\xi_{k+1}, h_{k+1}) \mid \mathrm{I}_{k+1}\right]\right| \mid \mathrm{I}_{k+1}\right] \\
&= 2\,\mathbb{E}\left[g(\xi_{k+1}, h_{k+1}) \mid \mathrm{I}_{k+1}\right] \\
&\leq 2\left(\frac{\bar{h}_{\sigma_{k+1}}}{\bar{h}_{p^\star}} c_{p^\star} + V_{p^\star}(\xi_{k+1})\right)
\end{aligned}$$

where the last inequality follows from (19). The first term is bounded since $\bar{h}_{\sigma_{k+1}} \leq \bar{\tau}$ for all $k \in \mathbb{N}$ for strategies e and ed, and $\bar{h}_{\sigma_{k+1}} \leq D^{\max}$ for all $k \in \mathbb{N}$ for strategy d. The fact that $\mathbb{E}\left[|X_{k+1} - X_k| \mid \mathrm{I}_{k+1}\right] \leq c$ follows then from mean-square stability of ξ_k, which is proven by boundedness of $\mathbb{E}[V_{p^\star}(\xi_k)]$ for all $k \in \mathbb{N}$ as $k \to \infty$, which follows similar arguments as a similar proof in [32, Th. 4].

Summing (19) for $k \in \{0, \ldots, N(T) - 1\}$, we have

$$\begin{aligned}
\sum_{k=0}^{N(T)-1} \mathbb{E}[g(\xi_k, h_k) \mid \mathrm{I}_k] = &\sum_{k=0}^{N(T)-1} \bar{h}_{\sigma_k} \frac{1}{\bar{h}_{p^\star}} c_{p^\star} - \delta_k + \nu_k \\
&+ V_{p^\star}(\xi_0) - V_{p^\star}(\xi_{N(T)})
\end{aligned} \tag{21}$$

where

$$\delta_k := \Delta(\xi_k, p^\star, p^\star, \sigma_k, b_k) - V^\Delta(\xi_k, \sigma_k, b_k)$$

and

$$\nu_k := V_{p^\star}(\xi_{k+1}) - \mathbb{E}[V_{p^\star}(\xi_{k+1}) | \mathrm{I}_k].$$

Next, we substitute (20) and (21) into (8). Taking the expectation, we have that $\mathbb{E}[\nu_k] = 0$ for all $k \in \mathbb{N}$. Then, when taking the limit $T \to \infty$ in (8), the last two terms in (21) vanish, since $\mathbb{E}[V_{p^\star}(\xi_{N(T)})]$ is bounded as $T \to \infty$, as explained before, and $\mathbb{E}[V_{p^\star}(\xi_0)]$ is bounded by the initial condition. Furthermore, the first term becomes equal to $J_{p^\star}$ since

$$\limsup_{T \to \infty} \frac{1}{T} \mathbb{E}\left[\sum_{k=0}^{N(T)-1} \bar{h}_{\sigma_k}\right] = 1. \tag{22}$$

This holds by the fact that $(\tilde{H}_k)_{k \in \mathbb{N}}$, with $\tilde{H}_N := \sum_{k=0}^{N} h_k - \bar{h}_{\sigma_k}$, is a martingale with respect to the filtration associated with I_{k+1} and again the fact that $\bar{h}_{\sigma_k}$ is bounded by $\bar{\tau}$ or $D^{\max}$. These conditions, again by Doob's optional sampling theorem [29], [35], imply that

$$\mathbb{E}\left[\sum_{k=0}^{N(T)-1} \bar{h}_{\sigma_k}\right] = \mathbb{E}\left[\sum_{k=0}^{N(T)-1} h_k\right].$$

Then, (22) holds by the fact that $\lim_{T \to \infty} \frac{1}{T} \mathbb{E} [\sum_{k=0}^{N(T)-1} h_k] = 1$, since the discretization error vanishes in the limit.

As a result, we get, for $\pi = d^\star \& ed^\star$, that

$$\begin{aligned}
J_{d^\star \& ed^\star} &= \frac{1}{\bar{h}_{p^\star}} c_{p^\star} - \lim_{T \to \infty} \frac{1}{T} \mathbb{E}\left[\sum_{k=0}^{N(T)-1} \delta_k\right] \\
&\leq \frac{1}{\bar{h}_{p^\star}} c_{p^\star} = J_{p^\star} = \min\{J_{d^\star}, J_{ed^\star}\}
\end{aligned} \tag{23}$$

with $\delta_k \geq 0$ by the fact that Δ is non-negative and V^Δ is non-positive by definition of (11). Note that $\delta_k = 0$ for the choice $(\sigma_k, b_k) = (p^\star, p^\star)$. This proves the theorem.

D. Proof of Theorem 2 and Corollary 2

Consider all allowable choices of combinations (m, D), where D is in the finite set $\mathcal{D}$, as new methods $\tilde{m}$ such that $(\tilde{m}, b)^\mathsf{T} \in \mathcal{S}^s(\xi)$. Each method $\tilde{m}$ has the particular value of D as its optimal choice of deadline $D^\star$. By reformulation, the switching condition (14) then takes the same form as (11) and the proof of Theorem 1 applies. For Corollary 2, the switching options are limited to $b_k = ed^\star$ and $\sigma_k = ed^\star$ for all $k \in \mathbb{N}$, hence always $V^{\Delta D} \leq 0$, and the deadline D is the only switching parameter.

References

[1] B. Wittenmark, J. Nilsson, and M. Torngren, "Timing problems in real-time control systems," in *Proc. Amer. Control Conf.*, Jun. 1995, vol. 3, pp. 2000–2004.

[2] G. C. Buttazzo, *Hard Real-Time Computing Systems: Predictable Scheduling Algorithms and Applications*. New York, NY, USA: Springer, 2011.

[3] J. Nilsson, "Real-time control systems with delays," Ph.D. dissertation, Dept. Autom. Control, Lund Inst. Technol., Lund, Sweden, 1998.

[4] J. P. Hespanha, P. Naghshtabrizi, and Y. Xu, "A survey of recent results in networked control systems," *Proc. IEEE*, vol. 95, no. 1, pp. 138–162, Jan. 2007.

[5] J. Araújo *et al.*, "Self-triggered control over wireless sensor and actuator networks," in *Proc. Int. Conf. Distrib. Comput. Sens. Syst. Workshops*, Jun. 2011, pp. 1–9.

[6] E. P. van Horssen, D. Antunes, and W. P. M. H. Heemels, "Switching data-processing methods in a control loop: Trade-off between delay and probability of data acquisition," in *Proc. 6th IFAC Workshop Distrib. Estimation Control Netw. Syst.*, Sep. 2016, pp. 274–279.

[7] E. P. van Horssen, D. Antunes, and W. P. M. H. Heemels, "Switching data-processing methods for feedback control: Breaking the speed versus accuracy trade-off," in *Proc. IEEE 54th Annu. Conf. Decision Control*, Dec. 2015, pp. 2313–2318.

[8] H. Pfifer and P. Seiler, "An overview of integral quadratic constraints for delayed nonlinear and parameter-varying systems," 2015. [Online]. Available: http://arxiv.org/abs/1504.02502

[9] E. Fridman, "Tutorial on Lyapunov-based methods for time-delay systems," *Eur. J. Control*, vol. 20, no. 6, pp. 271–283, 2014.

[10] W. Michiels and S.-I. Niculescu, *Stability and Stabilization of Time-Delay Systems (Advances in Design and Control)*. Philadelphia, PA, USA: SIAM, 2007.

[11] L. Schenato, B. Sinopoli, M. Franceschetti, K. Poolla, and S. Sastry, "Foundations of control and estimation over lossy networks," *Proc. IEEE*, vol. 95, no. 1, pp. 163–187, Jan. 2007.

[12] B. Sinopoli, L. Schenato, M. Franceschetti, K. Poolla, M. I. Jordan, and S. S. Sastry, "Kalman filtering with intermittent observations," *IEEE Trans. Autom. Control*, vol. 49, no. 9, pp. 1453–1464, Sep. 2004.

[13] B. Demirel, Z. Zou, P. Soldati, and M. Johansson, "Modular co-design of controllers and transmission schedules in wirelessHART," in *Proc. 50th IEEE Conf. Decision Control Eur. Control Conf.*, 2011, pp. 5951–5958.

[14] B. Demirel, "Architectures and performance analysis of wireless control systems," Ph.D. dissertation, KTH Royal Inst. Technol., Stockholm, Sweden, 2015.

[15] D. Antunes and W. Heemels, "Frequency-domain analysis of control loops with intermittent data losses," *IEEE Trans. Autom. Control*, vol. 61, no. 8, pp. 2295–2300, Aug. 2016.

[16] D. E. Quevedo, E. I. Silva, and G. C. Goodwin, "Control over unreliable networks affected by packet erasures and variable transmission delays," *IEEE J. Sel. Areas Commun.*, vol. 26, no. 4, pp. 672–685, May 2008.

[17] V. Gupta, D. Spanos, B. Hassibi, and R. M. Murray, "On LQG control across a stochastic packet-dropping link," in *Proc. Amer. Control Conf.*, Jun. 2005, pp. 360–365.

[18] T. M. P. Gommans, W. P. M. H. Heemels, N. Bauer, and N. van de Wouw, "Compensation-based control for lossy communication networks," *Int. J. Control*, vol. 86, pp. 1880–1897, Jun. 2013.

[19] J. Nilsson, B. Bernhardsson, and B. Wittenmark, "Stochastic analysis and control of real-time systems with random time delays," *Automatica*, vol. 34, no. 1, pp. 57–64, 1998.

[20] B. Lincoln and B. Bernhardsson, *Optimal Control over Networks with Long Random Delays*. Perpignan, France: Lab. Théorie des Syst., Univ. Perpignan, 2000.

[21] L. A. Montestruque and P. Antsaklis, "Stability of model-based networked control systems with time-varying transmission times," *IEEE Trans. Autom. Control*, vol. 49, no. 9, pp. 1562–1572, Sep. 2004.

[22] W. P. M. H. Heemels, K. H. Johansson, and P. Tabuada, "Event-triggered and self-triggered control," in *Encyclopedia of Systems and Control*. London, U.K.: Springer, 2015, pp. 384–391.

[23] C. G. Cassandras, "Event-driven control, communication, and optimization," in *Proc. 32nd Chinese Control Conf.*, Jul. 2013, pp. 1–5.

[24] S. Prakash, E. P. van Horssen, D. Antunes, and W. P. M. H. Heemels, "Self-triggered and event-driven control for linear systems with stochastic delays," in *Proc. Amer. Control Conf.*, May 2017, pp. 3023–3028.

[25] W. L. de Koning, "Infinite horizon optimal control of linear discrete time systems with stochastic parameters," *Automatica*, vol. 18, no. 4, pp. 443–453, Jul. 1982.

[26] S. Adyanthaya, Z. Zhang, M. Geilen, J. Voeten, T. Basten, and R. Schiffelers, "Robustness analysis of multiprocessor schedules," in *Proc. Int. Conf. Embedded Comput. Syst.: Archit., Model. Simul.*, Jul. 2014, pp. 9–17.

[27] R. Rom and M. Sidi, *Multiple Access Protocols: Performance and Analysis*. New York, NY, USA: Springer, 1990.

[28] C. Ramesh, H. Sandberg, and K. H. Johansson, "Performance analysis of a network of event-based systems," *IEEE Trans. Autom. Control*, vol. 61, no. 11, pp. 3568–3573, Nov. 2016.

[29] J. Doob, *Stochastic Processes*. New York, NY, USA: Wiley, 1953.

[30] K. J. Åström, *Introduction to Stochastic Control Theory*. New York, NY, USA: Elsevier, 1970.

[31] D. P. Bertsekas, *Dynamic Programming and Optimal Control*, vol. I. Belmont, MA, USA: Athena Scientific, 2005.

[32] D. Antunes and W. Heemels, "Rollout event-triggered control: Beyond periodic control performance," *IEEE Trans. Autom. Control*, vol. 59, no. 12, pp. 3296–3311, Aug. 2014.

[33] O. L. V. Costa, M. D. Fragoso, and R. P. Marques, *Discrete-Time Markov Jump Linear Systems*. New York, NY, USA: Springer, 2006.

[34] S. I. Resnick, *Adventures in Stochastic Processes*. Basel, Switzerland: Birkhauser, 1992.

[35] G. Grimmett and D. Stirzaker, *Probability and Random Processes*, 3rd ed. Oxford, U.K.: Oxford Univ. Press, 2001.

Eelco van Horssen received the M.Sc. degree (*cum laude*) in electrical engineering in 2013 from Eindhoven University of Technology, Eindhoven, The Netherlands, where he is currently working toward the Ph.D. degree in mechanical engineering.

His research interests include general theory of systems and control, distributed control, robust stability and performance, switched systems, networked and embedded control, dynamic programming, optimal control, and vision-based systems.

Suraj Prakash received the B.E. degree in electronics and communication from Visvesvaraya Technological University, Belagavi, India, in 2009, and the M.Sc. degree (with great appreciation) in systems and control in 2016 from the Eindhoven University of Technology, Eindhoven, The Netherlands, where he has been working toward the Ph.D. degree in mechanical engineering.

From 2009 to 2012, he was a Software Engineer with Delphi Automotive Systems, Bengaluru, India. His current research interests include optimal design and control of continuously variable transmission-based powertrains.

Duarte Antunes (M'13) was born in Viseu, Portugal, in 1982. He received the Licenciatura in electrical and computer engineering from Instituto Superior Técnico (IST), Lisbon, Portugal, in 2005, and the Ph.D. degree in automatic control from the Institute for Systems and Robotics, IST, in 2011.

Throughout his Ph.D. research, he visited the University of California, Santa Barbara, Santa Barbara, CA, USA, several times. From 2011 to 2013, he was a Postdoctoral Researcher with Eindhoven University of Techonology, Eindhoven, The Netherlands, where he is currently an Assistant Professor with the Department of Mechanical Engineering. His research interests include networked control systems, stochastic control, multirate and sampled-data systems, optimization, and dynamic programming.

Maurice Heemels (F'16) received the M.Sc. degree in mathematics and the Ph.D. degree in control theory (both *summa cum laude*) from Eindhoven University of Technology (TU/e), Eindhoven, The Netherlands, in 1995 and 1999, respectively.

From 2000 to 2004, he was with the Department of Electrical Engineering, TU/e, where he was from 2004 to 2006, with the Embedded Systems Institute. Since 2006, he has been with the Department of Mechanical Engineering, TU/e, where he is currently a Full Professor. He held visiting professor positions at the Swiss Federal Institute of Technology (ETH), Zurich, Switzerland, in 2001, and at the University of California at Santa Barbara, Santa Barbara, CA, USA, in 2008. In 2004, he was with the company Océ, Venlo, The Netherlands. His current research interests include hybrid and cyberphysical systems, networked and event-triggered control systems, and constrained systems including model-predictive control.

Dr. Heemels served/serves on the Editorial Boards of *Automatica*, *Nonlinear Analysis: Hybrid Systems*, *Annual Reviews in Control*, and the IEEE TRANSACTIONS ON AUTOMATIC CONTROL. He received a personal VICI grant awarded by STW (Dutch Technology Foundation).