LUND UNIVERSITY

Sporadic Control of Scalar Systems with Delay, Jitter and Measurement Noise

Cervin, Anton; Johannesson, Erik

2008

Sporadic Control of Scalar Systems with Delay, Jitter and Measurement Noise

Anton Cervin and Erik Johannesson

Automatic Control LTH, Lund University, Sweden
(email: anton@control.lth.se, erik@control.lth.se)

Abstract: Event-triggered control is a promising alternative to time-triggered control, especially for severely resource-constrained networked embedded systems. Previous work has shown that event-triggered control can reduce both the output variance and the average control rate in scalar linear stochastic systems compared to time-triggered control. It has also been shown how a minimum inter-control interval can be imposed, hence the term "sporadic control". In this work we extend the analysis of event-triggered impulse control of first-order linear stochastic systems to handle general sampling intervals and minimum inter-control intervals, control delay and control jitter, and measurement noise. The results show that the advantage of sporadic control remains also in these cases.

1. INTRODUCTION

Most digital feedback controllers operate in a time-triggered fashion, where the sampling, control computation, and actuation are performed periodically. An alternative approach is to sample and control only when certain events occur. Such schemes may be more natural in systems that include event-based sensors (e.g. encoder wheels) or event-based dynamics (e.g. queueing systems). Event-based control could also be useful for saving CPU and network bandwidth in networked embedded control systems, computing and communicating only when something significant has occurred in the system.

This paper considers event-triggered impulse control of a scalar linear system driven by white noise. Control actions are issued only if the magnitude of the output exceeds a certain threshold. Being very simple, the system readily allows optimal controllers to be computed and compared for a range of different parameters. Higher-order systems introduce a wide range of interesting problems, which will be tackled in future papers.

1.1 Previous Work

The idea of event-based control is not new (for a survey of previous applications, see Åström [2007]), but only in the last few years have some theoretical results started to appear.

Event-based control of first-order stochastic systems was first studied by Åström and Bernhardsson [1999]. It was shown that event-based impulse control of a Wiener process requires, on average, only one third of the sampling rate of a time-triggered controller to achieve the same output variance. The work of Åström and Bernhardsson has recently been extended and elaborated upon in the PhD thesis Rabi [2006], which explores, among other things, event-based control with piecewise constant control signals and level-triggered sampling.

Event-based control of higher-order systems over networks was considered in Hristu-Varsakelis and Kumar [2002],
which explored policies for deciding what control loop should be closed at any given time. They also derived a sufficient condition for stability of multiple networked control loops.

Optimal state estimation and optimal control for first-order systems with a limited, pre-specified number of measurements or controls have been studied in Imer and Basar [2005] and Imer and Basar [2006]. Assuming a discrete-time process, optimal time-varying threshold policies were derived using dynamic programming.

The problem of optimal threshold-triggered sampling of higher-order discrete-time systems was studied in Cogill et al. [2006]. They provided an algorithm for computing a threshold guaranteed to incur a cost within a factor of six of the optimal achievable cost.

In a recent paper, Johannesson et al. [2007], we introduced the notion of *sporadic* event-based control. A sporadic controller may not issue control actions more often than every T_c seconds. Thus, such a controller is implementable and can be predictably scheduled in a real-time system. This is in contrast to Åström and Bernhardsson's *aperiodic* controller, which might generate events infinitely often.

1.2 Outline and Contributions

In the present paper, we make several extensions to the analysis of event-based control of first-order stochastic systems. First (Section 2), we generalize the notion of sporadic control to allow for different intervals for sampling and control. This is relevant for networked embedded systems, where it may be cheap to sample but expensive to communicate. Then (Section 3), we explore how network-induced delay and jitter affect the performance of sporadic control systems. Finally (Section 4) we extend the system model with measurement noise and study optimal filtering and sporadic control. Throughout, numerical results are reported for stable, marginally stable, and unstable systems.

2. SPORADIC CONTROL WITH ARBITRARY SAMPLING INTERVAL

The process to be controlled is given by the stochastic differential equation

$$dx = ax\,dt + u\,dt + dw, \quad x(0) = 0, \tag{1}$$

where x is the scalar state, a is the process pole, u is the control signal, and w is a Wiener process with unit incremental variance. The control signal is zero except at discrete sampling instants t_k, when it is allowed to be a Dirac pulse of magnitude u_k:

$$u(t) = \sum_{k=0}^{\infty} \delta(t - t_k) u_k. \tag{2}$$

If $u_k = 0$, we say that no control action is issued.

The performance of the system is measured by the stationary state cost,

$$J_x = \limsup_{t \to \infty} \frac{1}{t} \mathrm{E} \left\{ \int_0^t x^2(s)ds \right\}, \tag{3}$$

and by the average control rate (or control cost),

$$J_u = \limsup_{t \to \infty} \frac{1}{t} \mathrm{E} \left\{ N_u(0, t) \right\}, \tag{4}$$

where $N_u(0, t)$ is the number of (non-zero) control actions in the interval $(0, t)$. The total cost to be minimized is given by

$$J = J_x + \rho J_u, \tag{5}$$

where $\rho \geq 0$ is a weight. The cost function (5) reflects a trade-off between the regulation performance and the average resource (e.g. CPU or network) consumption.

A sporadic event-based controller is defined by three parameters: the event detection threshold r, the sampling interval T_s, and the minimum inter-control interval $T_c \geq T_s$. We assume that T_c and T_s are chosen in accordance with the available computing and communication resources, while r should be chosen to optimize (5). (Previous work, Rabi [2006], has shown that a simple threshold policy is indeed optimal for event-based control.)

At time $t_0 = 0$, the controller starts in the *active state*, where it samples the process every T_s seconds, taking discrete measurements $x_k = x(t_k)$. If $|x_k| \geq r$, the controller issues the control action $u_k = -x_k$, effectively resetting the process state to zero. The controller then enters the *inactive state* where it stays for T_c seconds. Immediately after returning to the active state, the sampling starts again. The sampling instants hence progress as

$$t_{k+1} = \begin{cases} t_k + T_s, & |x_k| < r, \\ t_k + T_c, & |x_k| \geq r. \end{cases} \tag{6}$$

A typical realization of a sporadic control process with $a = 0$ (an integrator), $r = 1$, $T_s = 0.25$ and $T_c = 1$ is shown in Fig. 1. It is seen that the output is basically kept within the limits. However, the process may drift outside the detection band during the inactive intervals or between two sampling instants.

2.1 Evaluation of Cost

To evaluate the cost functions (3) and (4) for a given value of r, we compute the stationary distribution of the state in
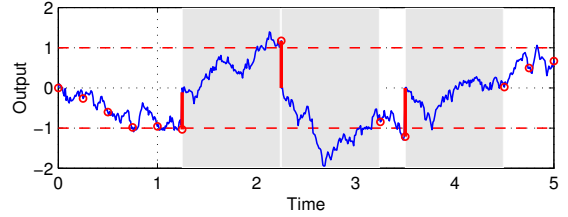

Fig. 1. Sporadic control of an integrator process with detection threshold $r = 1$, sampling interval $T_s = 0.25$ and minimum inter-control interval $T_c = 1$. The circles indicate samples, and the shaded regions show the inactive intervals.

the sampling instants t_k (before a possible control action has been issued). Note that these instants are irregularly spaced in time if $T_s \neq T_c$.

Between two sampling instants, the state evolves as

$$x_{k+1} = \begin{cases} e^{aT_s}x_k + v_k(T_s), & |x_k| < r, \\ v_k(T_c), & |x_k| \geq r, \end{cases} \tag{7}$$

where $v_k(t)$ is a Gaussian random variable with zero mean and variance

$$P(t) = \int_0^t e^{2as}ds = \begin{cases} \frac{e^{2at}-1}{2a}, & a \neq 0, \\ t, & a = 0. \end{cases} \tag{8}$$

The accumulated state cost from time t_k to time t_{k+1} is given by

$$J_{\text{active}} = Q(T_s)\,\mathrm{E}\left\{ x_k^2 \mid |x_k| < r \right\} + J_w(T_s) \tag{9}$$

if the controller stays in the active state and by

$$J_{\text{inactive}} = J_w(T_c) \tag{10}$$

if the controller enters the inactive state. Here,

$$Q(t) = \int_0^t e^{2as}ds = \begin{cases} \frac{e^{2at}-1}{2a}, & a \neq 0, \\ t, & a = 0 \end{cases} \tag{11}$$

is the state weight in the sampled cost function, while

$$J_w(t) = \int_0^t P(s)ds = \begin{cases} \frac{e^{2at}-2at-1}{4a^2}, & a \neq 0, \\ \frac{t^2}{2}, & a = 0 \end{cases} \tag{12}$$

accounts for the inter-sample noise, see e.g. Åström [1970].

Finally, assuming stationarity, the costs per time unit become

$$J_x = \frac{p_{\text{active}}J_{\text{active}} + p_{\text{inactive}}J_{\text{inactive}}}{p_{\text{active}}T_s + p_{\text{inactive}}T_c},$$
$$J_u = \frac{p_{\text{inactive}}}{p_{\text{active}}T_s + p_{\text{inactive}}T_c}, \tag{13}$$

where

$$p_{\text{active}} = \text{Prob}\left\{|x_k| < r\right\} = 1 - p_{\text{inactive}}. \tag{14}$$

The stationary distribution of x_k can be found numerically by discretizing the state space and then iterating the distribution according to (7) until convergence.

2.2 Example

To investigate the impact of the sampling interval on performance, we fix $T_c = 1$ and vary T_s. For each parameter configuration, we compute the stationary distribution of x_k. In Figs. 2 and 3, we show the state cost J_x and the control cost J_u as functions of the threshold r for three different systems. It is seen that, when $T_s < T_c$, the state
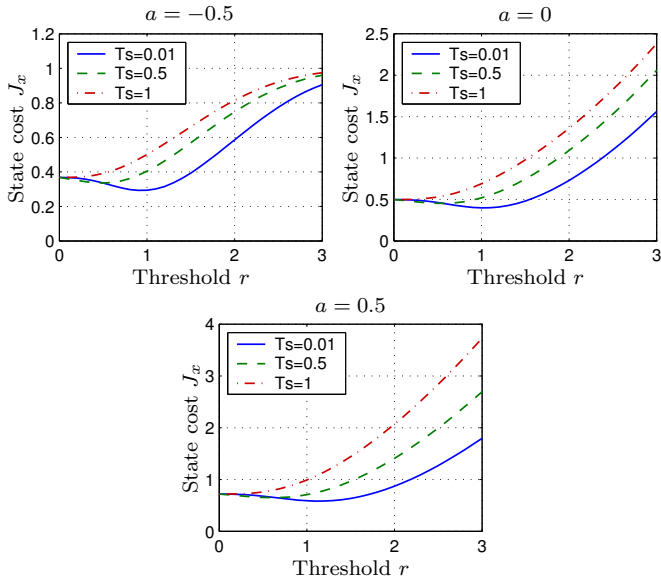
Fig. 2. State cost J_x for three systems with varying sampling interval T_s. Note the different vertical scales.
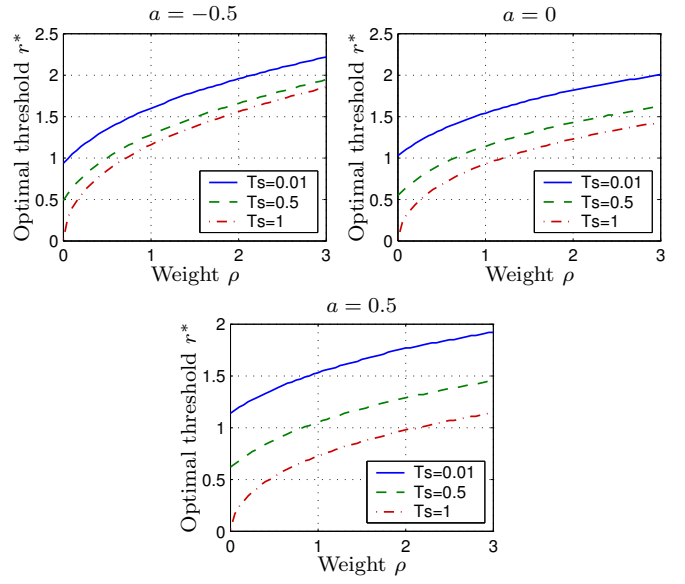

Fig. 4. Optimal threshold r^* for three systems with varying sampling interval T_s. Note the different vertical scales.
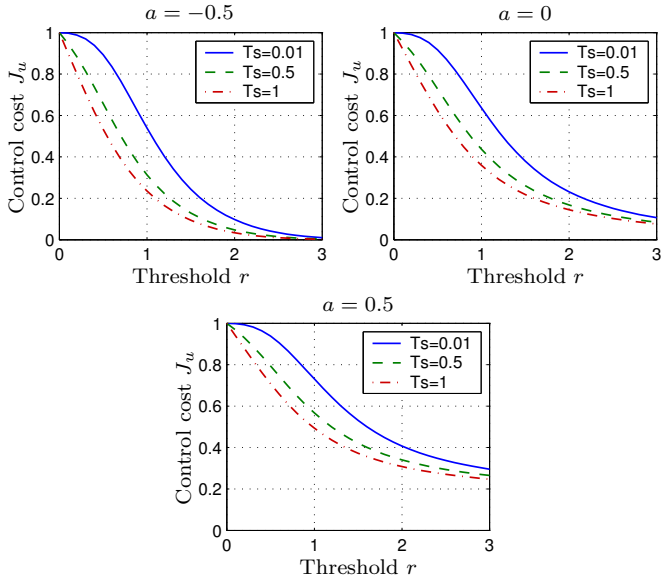

Fig. 3. Control cost J_u for three systems with varying sampling interval T_s.

Fig. 5. Minimum achievable cost J^* as a function of relative cost of control ρ, using sporadic and periodic controllers. The system is an integrator with different sampling intervals.

Fig. 6. Trade-off between state cost and control cost using sporadic and periodic controllers for the integrator.

cost initially decreases when r is increased and achieves a local minimum for some $r > 0$. At the same time, the control cost decreases monotonically. This shows that sporadic control can increase performance and decrease resource consumption at the same time. Overall, a shorter sampling interval means faster detection of deviations and hence lower state cost but also higher control cost.

The optimal threshold r^* as a function of the weight ρ is shown in Fig. 4. Here it is again seen that the optimal threshold is actually non-zero for $\rho = 0$ if $T_s < T_c$. Further, a shorter sampling interval implies that the threshold should be set higher.

Finally, we compute the optimal achievable cost J^* for the integrator for different weights ρ and compare sporadic control to standard periodic control in Fig. 5. It is seen that the sporadic controller outperforms the periodic one, and
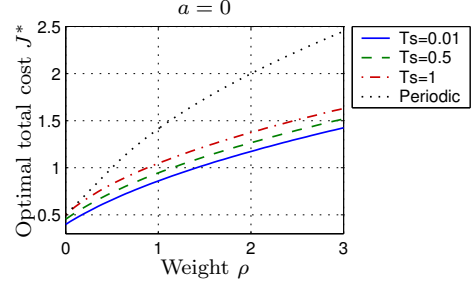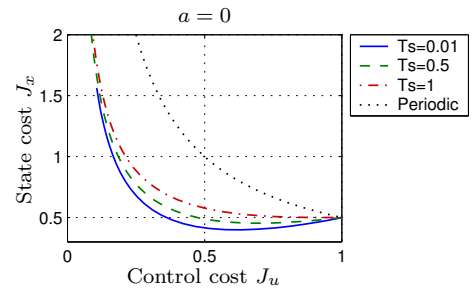
that a shorter sampling interval means lower total cost. The plot also indicates that there is a smooth transition from the case $T_s = 0$ (continuous sampling) to $T_s = T_c$. The trade-off between state cost and control cost for the integrator is made explicit in Fig. 6. It is seen that the average control rate can be substantially decreased before the state cost starts to increase if $T_s < T_c$.

3. DELAY AND JITTER

One of the motivations for event-based control is to save communication bandwidth in networked control. However, control over networks always means that delays are in-

duced in the feedback loop. If many applications share the same network, there may also be delay jitter due to unpredictable medium access times.

3.1 Constant Delay

We first consider the case of a constant control delay $\tau \leq T_c$, meaning that the control action computed at time t_k will not be applied to the process until τ seconds later. The control signal is now given by

$$u(t) = \sum_{k=0}^{\infty} \delta(t - t_k - \tau)u_k. \qquad (15)$$

To evaluate the costs, it is necessary to reconsider the inter-sample behavior in the inactive state. From time t_k up to $t_k + \tau$ (just before the control action has been applied), the process evolves as

$$x(t_k + \tau) = e^{a\tau}x_k + v_k(\tau). \qquad (16)$$

The extra state cost during this interval becomes

$$Q(\tau)\mathrm{E}\left\{x_k^2 \mid |x_k| \geq r\right\}. \qquad (17)$$

We assume that the delay is known at design time, so that the controller can compensate for it by issuing the modified control action

$$u_k = -e^{a\tau}x_k. \qquad (18)$$

This will bring the expected value of the state to zero for the remainder of the interval, eventually leading to

$$x_{k+1} = v_k(T_c). \qquad (19)$$

Note that this expression is identical to the update equation (7) for the case $|x_{k+1}| \geq r$. Hence, viewing the system only at the sampling instants, the stationary distribution of x_k will remain the same as without the delay. This implies that J_{active} and J_u will not be affected. J_x will however increase, the new expression for J_{inactive} in (13) being

$$J_{\text{inactive}} = Q(\tau)\mathrm{E}\left\{x_k^2 \mid |x_k| \geq r\right\} + J_w(T_c). \qquad (20)$$

3.2 Delay Jitter

In the case of variable delay, or *delay jitter*, it is not realistic to assume that the actual value of the delay is known when the control action is calculated. Rather, we suppose that the delay compensation is based on the estimated average delay, $\bar{\tau}$. The real delay in each sample is assumed to be given by a random variable τ_k with a given probability density function.

The control signal is now given by

$$u(t) = \sum_{k=0}^{\infty} \delta(t - t_k - \tau_k)u_k,$$

where $u_k = -e^{a\bar{\tau}}x_k$. Similar to before, from time t_k up to $t_k + \tau_k$ in the inactive state, the process will evolve as

$$x(t_k + \tau_k) = e^{a\tau_k}x_k + v_k(\tau_k). \qquad (21)$$

The extra state cost during this interval is given by

$$Q(\tau_k)\mathrm{E}\left\{x_k^2 \mid |x_k| \geq r\right\}.$$

Because of the inexact delay compensation, the state just after the control action will no longer have expected value zero, but is given by

$$x(t_k + \tau_k^+) = \left(e^{a\tau_k} - e^{a\bar{\tau}}\right)x_k + v_k(\tau_k). \qquad (22)$$

The extra state cost during the final part of the inactive interval becomes

$$Q(T_c - \tau_k)\left(e^{a\tau_k} - e^{a\bar{\tau}}\right)^2 \mathrm{E}\left\{x_k^2 \mid |x_k| \geq r\right\}.$$

Finally, the new state update equation becomes

$$x_{k+1} = e^{aT_c}\left(1 - e^{a(\bar{\tau}-\tau_k)}\right)x_k + v_k(T_c). \qquad (23)$$

In general, $\bar{\tau} \neq \tau_k$, which means that the stationary distribution of x_k will be different from the jitter-free case. This implies that both J_x and J_u are affected by the jitter. Taking the expected value over the different possible τ_k, the new expression for the cost during an inactive interval becomes

$$J_{\text{inactive}} = \mathop{\mathrm{E}}_{\tau_k}\left\{\left(Q(\tau_k) + Q(T_c - \tau_k)\left(e^{a\tau_k} - e^{a\bar{\tau}}\right)^2\right)\right.$$
$$\left. \times \mathrm{E}\left\{x_k^2 \mid |x_k| \geq r\right\}\right\} + J_w(T_c). \qquad (24)$$

3.3 Example

To investigate the impact of delay and jitter on sporadic control we fix $T_s = T_c = 1$ and then compare three cases:

- No delay.
- Constant delay $\tau = 0.5$.
- Random delay $\tau_k \in \{0, 1\}$, where both outcomes have equal probability. The average delay is $\bar{\tau} = 0.5$.

In Figs. 7 and 8, we show the state cost J_x and the control cost J_u as functions of the threshold r for three different systems. It is seen that both delay and jitter increase the state cost, which is expected. Jitter is better than a constant delay for the stable system, while the opposite is true for the unstable system. As predicted, the control cost is unaffected by the constant delay, while the change is very small in the jitter case (indicating that the distribution does not change much). In the integrator case, it can be noted that the extra state cost due to delay is always equal to τ. (Actually, this is true regardless of the sampling scheme used: event-based, periodic, random, etc.)
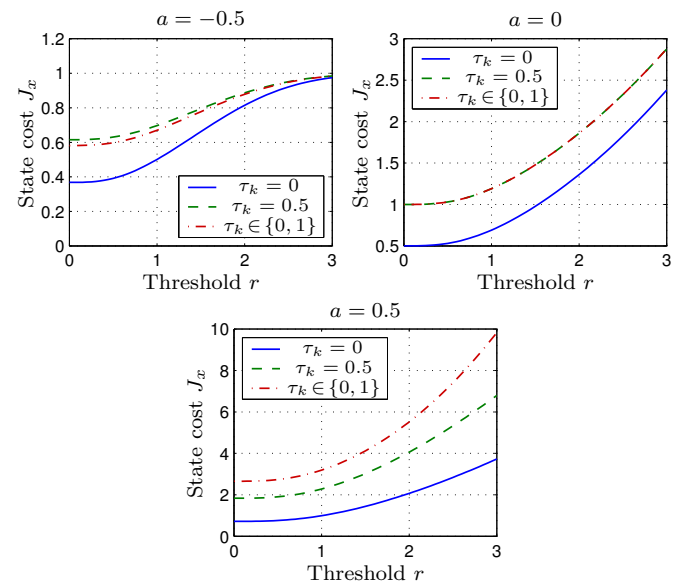
Fig. 7. State cost J_x for three systems with constant delay or jitter. Note the different vertical scales.
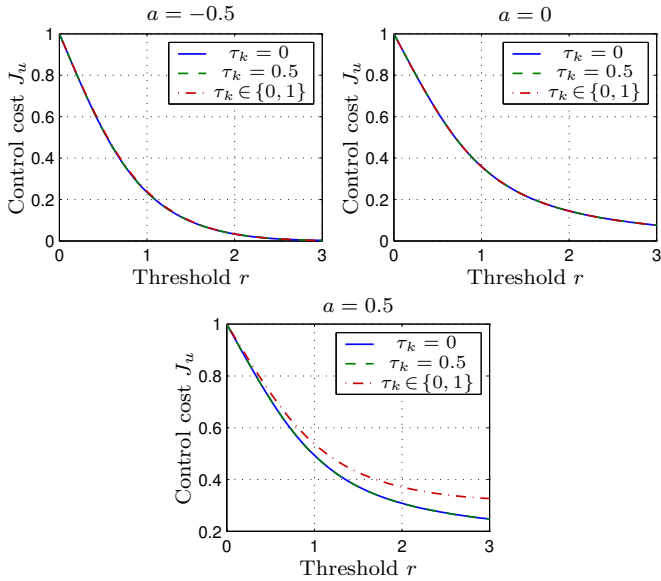
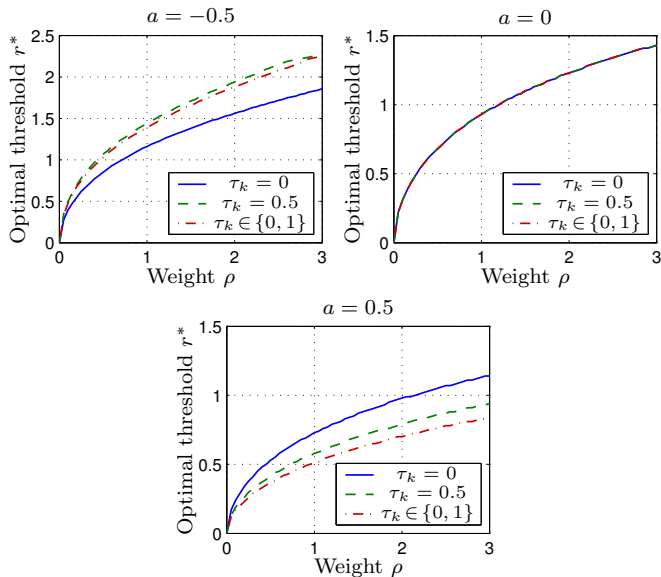Fig. 8. Control cost J_u for three systems with constant delay or jitter.

Fig. 9. Optimal threshold r^* for three systems with constant delay or jitter. Note the different vertical scales.

Fig. 10. Minimum achievable cost J^* as a function of relative cost of control ρ, using sporadic and periodic controllers. The system is an integrator with and without jitter.

The optimal threshold r^* as a function of the weight ρ is shown in Fig. 9. We observe that a delay increases the optimal threshold for stable systems but decreases it for unstable systems. For stable systems, the controller need not care about small deviations, since it is likely that the state will decay before the control action takes effect. The opposite reasoning holds for unstable systems.

Finally, we compare the achievable optimal performance J^* of sporadic control to periodic control for the integrator process in Fig. 10. The sporadic controller outperforms the periodic controller (except for $\rho = 0$ where the controllers coincide, since $T_c = T_s$). In both cases, a jitter $\tau_k \in \{0, 1\}$ with mean value 0.5 increases the total cost by the same amount.
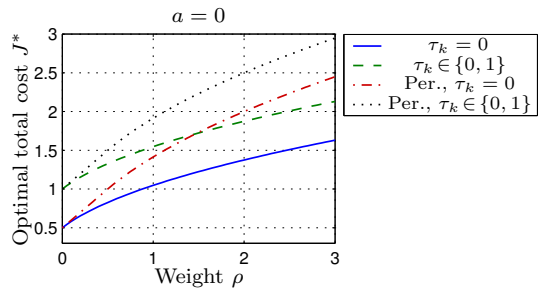
4. MEASUREMENT NOISE

So far, we have assumed that the controller has access to perfect measurements of the state variable. We now investigate the notion of sporadic control in the case when the measurements are distorted by noise. For simplicity, it is assumed that $T_s = T_c = T$.

The sampled system is

$$x_{k+1} = \Phi x_k + u_k + v_k,$$
$$y_k = x_k + e_k,$$

where $\Phi = e^{aT}$. The noise v_k and e_k are (mutually) independent zero-mean Gaussian random variables with $\mathrm{V}\{v_k\} = P(T)$ and $\mathrm{V}\{e_k\} = \sigma_e^2$.

4.1 Controller Structure

Since the system is linear with Gaussian noise, the optimal state estimate is given by the standard Kalman filter,

$$\hat{x}_k = \Phi \hat{x}_{k-1} + u_{k-1} + K(y_k - \Phi \hat{x}_{k-1} - u_{k-1})$$
$$= (1 - K)\left(\Phi \hat{x}_{k-1} + u_{k-1}\right) + K(x_k + e_k),$$

where K is the Kalman filter gain.

Given $\hat{x}_k$, then x_{k+1} is Gaussian and

$$\mathrm{E}\left\{x_{k+1}^2 | \hat{x}_k, u_k\right\} = \mathrm{V}\left\{x_{k+1} | \hat{x}_k, u_k\right\} + \mathrm{E}\left\{x_{k+1} | \hat{x}_k, u_k\right\}^2$$
$$= \mathrm{V}\left\{x_{k+1} | \hat{x}_k\right\} + \left(\Phi \hat{x}_k + u_k\right)^2.$$

The conditional variance of x_{k+1} is not affected by u_k, so the expected contribution to the cost at time $k + 1$, given $\hat{x}_k$, is minimized by $u_k = -\Phi \hat{x}_k$. The same holds for succeeding terms. Thus, the optimal control law is

$$u_k = \begin{cases} 0, & |\hat{x}_k| < r, \\ -\Phi \hat{x}_k, & |\hat{x}_k| \geq r, \end{cases} \tag{25}$$

where the threshold r remains to be determined.

4.2 Evaluation of Cost

The cost for a given threshold can be evaluated through numerical computation of the stationary distribution of the system and filter states. The closed-loop system may be regarded as a two-dimensional stochastic process. Let $f_{x,\hat{x}}(v_1, v_2)$ be the probability density of the stationary distribution just prior to any control event. The costs can then be expressed as

$$J_x^{d-} = \iint v_1^2 f_{x,\hat{x}}(v_1, v_2) \, dv_1 dv_2$$
$$J_u = \iint_{|v_2| \geq r} f_{x,\hat{x}}(v_1, v_2) \, dv_1 dv_2. \tag{26}$$

Here, the discrete-time state cost $J_x^{d-} = \mathrm{E}\left\{x_k^2\right\}$ is evaluated at sampling instants before any control action. The variance after control action can be found by going one sample backwards in time:

$$J_x^{d-} = \Phi^2 J_x^{d+} + \mathrm{V}\left\{v_k(T)\right\}.$$

The continuous-time state cost is then given by

$$J_x = \frac{Q(T)}{T} J_x^{d+} + J_w(T).$$

The stationary distribution is found numerically by propagating the distribution of the two-dimensional state on a discretized state-space, until convergence. In order to do this, we need to know the joint distribution of $(x_{k+1}, \hat{x}_{k+1})$ conditioned on $Z = (x_k, \hat{x}_k, u_k)$. Some simple calculations give that

$$\mathrm{E}\left\{x_{k+1}|Z\right\} = \Phi x_k + u_k,$$
$$\mathrm{V}\left\{x_{k+1}|Z\right\} = \sigma_w^2,$$
$$\mathrm{E}\left\{\hat{x}_{k+1}|x_{k+1}, Z\right\} = (1 - K)(\Phi\hat{x}_k + u_k) + Kx_{k+1},$$
$$\mathrm{V}\left\{\hat{x}_{k+1}|x_{k+1}, Z\right\} = K^2\sigma_e^2.$$

By the definition of conditional density, we have that

$$f_{x_{k+1}, \hat{x}_{k+1}|Z}(v_1, v_2|z) = f_{\hat{x}_{k+1}|x_{k+1}, Z}(v_2|v_1, z) \cdot f_{x_{k+1}|Z}(v_1|z).$$

It can be shown that the joint distribution is Gaussian,

$$(x_{k+1}, \hat{x}_{k+1})^T|Z \in N\left(\mu(z), R\right),$$

with parameters given by

$$\mu(z) = \begin{bmatrix} \Phi x_k + u_k \\ (1 - K)(\Phi\hat{x}_k + u_k) + K(\Phi x_k + u_k) \end{bmatrix},$$
$$R = \begin{bmatrix} \sigma_w^2 & K\sigma_w^2 \\ K\sigma_w^2 & K^2(\sigma_w^2 + \sigma_e^2) \end{bmatrix}.$$

The computations required for a good approximation are quite demanding since the computation time scales roughly with the number of grid points to the power of 4.

4.3 Example

To investigate the impact of measurement noise on sporadic control, we evaluate the performance for different values of σ_e. The state cost J_x is plotted in Fig. 11 as a function of the threshold r. Not surprisingly, it is seen that measurement noise yields additional state variance. When $a < 0$, J_x approaches a constant as the threshold increases, for any σ_e. This is because the open-loop variance is finite ($= -1/2a$) and provides an upper bound. For $a \geq 0$, the extra cost seems to be fairly constant for moderate thresholds.

The corresponding control cost J_u is shown in Fig. 12. It appears that the control rate changes with the level of measurement noise when $a \neq 0$. The reason is that when $\sigma_e = 0$ it is both necessary and sufficient for $|x|$ to exceed r at one sample in order to issue a control action. When there is more measurement noise, the Kalman filter puts more emphasis on the model and less on the measurements, which are averaged out. In the stable case this makes it less likely for $|\hat{x}|$ to exceed r than what it is for $|x|$, and thus the control rate decreases. The situation is reversed in the unstable case. For the integrator process, these effects balance out, with the interesting result that the average control rate is invariant to σ_e.
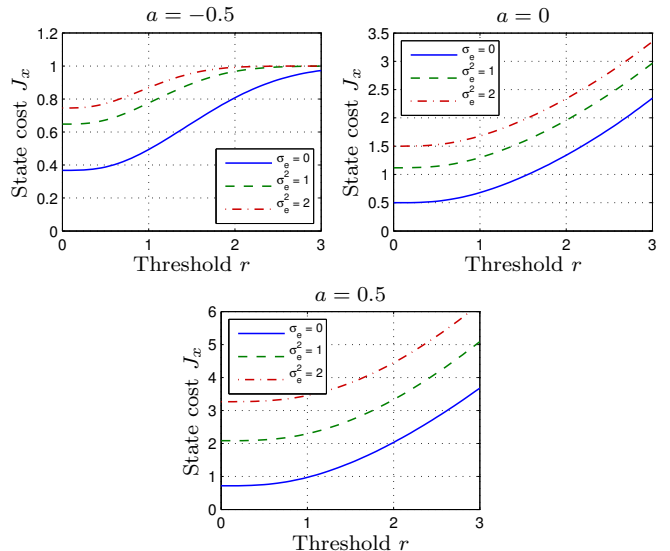

Fig. 11. State cost J_x for three systems with varying measurement noise. Note the different vertical scales.
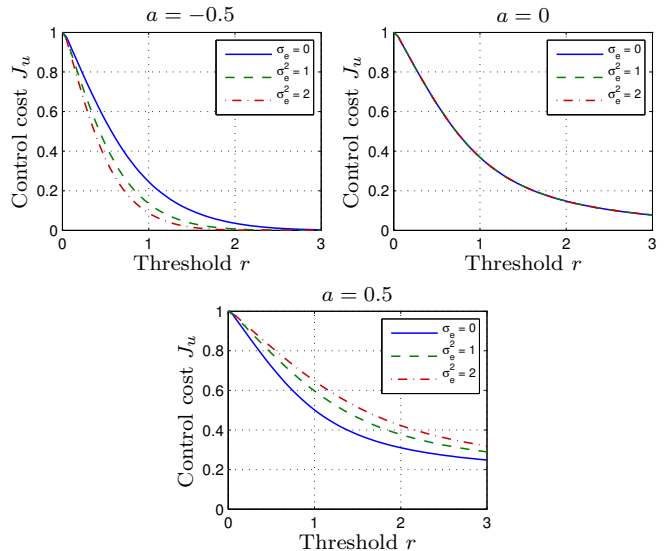

Fig. 12. Control cost J_u for three systems with varying measurement noise.

It is even more interesting to see how measurement noise affects the optimal threshold r^*, which is plotted in Fig. 13 as a function of ρ. It turns out that for the integrator process, the optimal threshold is independent of the measurement noise. For other systems, the difference is small. The conclusion is that the certainty equivalence principle does not hold for sporadic control with measurement noise, although there seems to be *almost* separation between estimation and control for moderate values of σ_e and ρ.

Furthermore, it is noted that when σ_e increases, the optimal threshold r^* changes in different directions depending on the sign of a. This is a consequence of the shape of J_u. Consider e.g. the case when $a = -0.5$: with measurement noise, the control actions are imperfect since they are based on an uncertain estimate. It is therefore beneficial to decrease the control rate somewhat and let the stable dynamics take care of things. Intuitively, this should lead to a higher threshold. However, as seen in Fig. 12, the
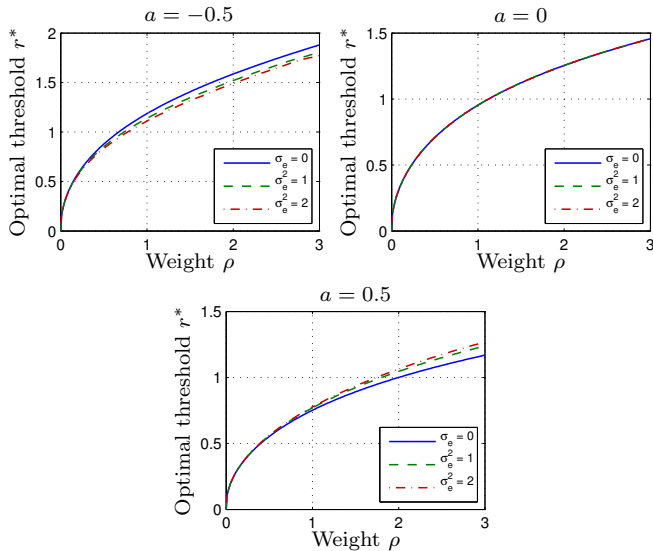
Fig. 13. Optimal threshold r^* as a function of relative cost of control ρ for three systems with varying measurement noise. Note the different vertical scales.
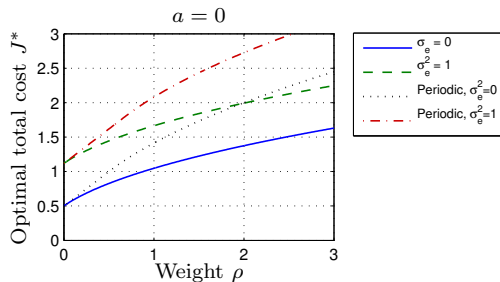

Fig. 14. Minimum achievable cost J^* as a function of relative cost of control ρ, using sporadic and periodic controllers. The system is an integrator with and without measurement noise.

control rate is already smaller for a given r. In fact, it is so much smaller that the optimal threshold becomes somewhat *lower*. It could be said that the filtering itself decreases the control rate just a little more than needed, making r^* insensitive to σ_e.

The optimal performance of the sporadic controller is compared to the standard periodic LQG controller, for an integrator process with and without measurement noise, in Fig. 14. It is clear that sporadic control delivers superior performance in terms of the specified cost function. Moreover, using sporadic control, the performance loss due to measurement noise is constant when ρ changes. For periodic control however, there is a slight increase in the performance loss. Accordingly, the gain of using sporadic control is actually somewhat larger if there is measurement noise.

5. CONCLUSIONS

This paper has provided several extensions to the concept of sporadic impulse control of first-order linear stochastic systems. Some of the results of this analysis deserve special attention: a generalization of the two sampling patterns analyzed in Johannesson et al. [2007] is provided, exhibiting a smooth transition between the two special cases $T_s = 0$ and $T_s = T_c$. It can be noted that the advantage of sporadic control over periodic control remains even if there is delay, jitter, or measurement noise. Surprisingly, the detection threshold should be adjusted in opposite directions if the delay increases and if the measurement noise increases. We also note that there is approximate separation between control and estimation for moderate parameter values.

REFERENCES

Karl Johan Åström. Event based control. In A. Astolfi and L. Marconi, editors, *Analysis and Design of Nonlinear Control Systems*. Springer Verlag, 2007.

Karl Johan Åström. *Introduction to Stochastic Control Theory*. Academic Press, New York, January 1970.

Karl Johan Åström and Bo Bernhardsson. Comparison of periodic and event based sampling for first-order stochastic systems. In *Preprints 14th World Congress of IFAC*, Beijing, P.R. China, July 1999.

Randy Cogill, Sanjay Lall, and Joao P. Hespanha. A constant factor approximation algorithm for event-based sampling. In *Proc. Allerton Conference on Communication, Control, And Computing*, 2006.

D. Hristu-Varsakelis and P. R. Kumar. Interrupt-based control over a shared communication medium. In *Proc. IEEE Conference on Decision and Control*, 2002.

Ohran C. Imer and Tamer Basar. Optimal estimation with limited measurements. In *Proc. IEEE Conference on Decision and Control*, 2005.

Ohran C. Imer and Tamer Basar. Optimal control with limited controls. In *Proc. American Control Conference*, 2006.

Erik Johannesson, Toivo Henningsson, and Anton Cervin. Sporadic control of first-order linear stochastic systems. In *Proc. 10th International Conference on Hybrid Systems: Computation and Control*, Pisa, Italy, April 2007.

Maben Rabi. *Packet based Inference and Control*. PhD thesis, Institute for Systems Research, University of Maryland, 2006. URL: http://hdl.handle.net/1903/6612.