

Event-triggered Pulse Control with Model Learning (if Necessary)

Dominik Baumann^{1,2}, Friedrich Solowjow¹, Karl Henrik Johansson², and Sebastian Trimpe¹

Abstract—In networked control systems, communication is a shared and therefore scarce resource. Event-triggered control (ETC) can achieve high performance control with a significantly reduced amount of samples compared to classical, periodic control schemes. However, ETC methods usually rely on the availability of an accurate dynamics model, which is oftentimes not readily available. In this paper, we propose a novel event-triggered pulse control strategy that learns dynamics models if necessary. In addition to adapting to changing dynamics, the method also represents a suitable replacement for the integral part typically used in periodic control.

I. INTRODUCTION

In modern engineering, control systems are often connected over communication networks. Common examples for these networked control systems (NCS) include automation industry, where multiple plants have to be controlled often by a remote controller, building automation, where sensors and actuators are deployed to regulate the indoor climate, and coordinated flight of a swarm of drones. Communication then is a shared and therefore limited resource, thus, traditional, periodic control approaches are not feasible.

In this paper, we present a new architecture for event-triggered pulse control that quantifies model accuracy and, if necessary, automatically identifies system dynamics through learning. The decision, whether learning is necessary, is taken by a learning trigger. The proposed framework can cope with changing dynamics and load disturbances, that way replacing the integrator from periodic control.

A block diagram of the approach is provided in Fig. 1. We consider a plant with sensors and actuators, subject to process noise and disturbances (v and ϵ), and input saturations u_{\max} . Controller and actuator are connected over a communication network. Since communication is a scarce resource, periodic communication is not desirable and, therefore, we employ an event-triggered design (block ‘State Trigger’). In case of an event, we apply a pulse with length t_{inp} to reset the system to its equilibrium state. The pulse length naturally depends on the system dynamics. To obtain an accurate model of the system dynamics, we leverage system identification techniques to learn the model from data. As learning may be expensive, *e.g.*, due to the involved computations, we only learn a new model if necessary, for instance, in case of a poor initial model or if the dynamics have changed.

¹Intelligent Control Systems Group, Max Planck Institute for Intelligent Systems, Stuttgart/Tübingen, Germany. Email: dbaumann@tuebingen.mpg.de, fsolowjow@is.mpg.de, trimpe@is.mpg.de

²Division of Decision and Control Systems, KTH Royal Institute of Technology, Stockholm, Sweden. Email: kallej@kth.se

This work was supported in part by the German Research Foundation (DFG) within the priority program SPP 1914 (grant TR 1433/1-1), the Cyber Valley Initiative, the IMPRS-IS, and the Max Planck Society.

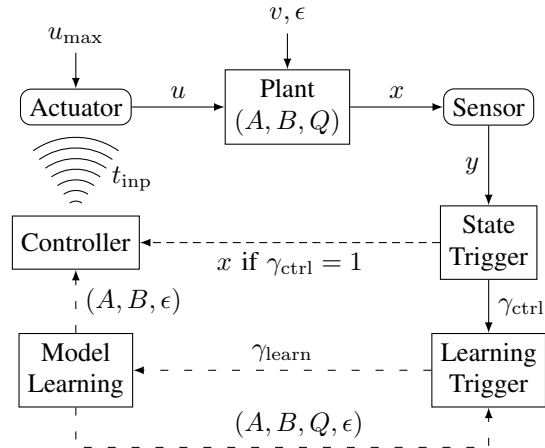


Fig. 1. Block diagram of the proposed control design. Dashed lines represent connections that are only active in case of an event.

This decision process is made by the ‘Learning Trigger’. Based on a statistical analysis of the time between events, the learning trigger decides whether the model of the system dynamics is accurate enough. If not, learning of a new model is triggered. We thus have two different triggers: the state trigger (γ_{ctrl}), which triggers communication of control commands if necessary, and a learning trigger (γ_{learn}), which triggers learning in case of bad performance.

Contributions: We make the following contributions:

- A new architecture for event-triggered pulse control that automatically learns dynamics model to cope with changing dynamics;
- Development of a learning trigger for ETC, which allows to automatically identify system dynamics if necessary;
- Handling load disturbances by learning and compensating for them, thus replacing the integrator typically used in periodic control in a way suitable for ETC.

Related work: Reducing communication is one key aspect of NCS and has been addressed by introducing event-triggered methods [1], [2]. While most of these approaches are based on a dynamical model of the system, the model is typically assumed to be given and not learned from data as proposed herein.

The ‘Learning Trigger’ in Fig. 1 is based on a framework developed in recent work [3]. There, ETL was introduced to trigger learning experiments in event-triggered state estimation (ETSE). Here, we extend this framework to control. We look at a straightforward type of ETC, namely event-triggered pulse control. In contrast to [3], which considered

a linear system perturbed only by Gaussian noise, we also consider load disturbances herein.

Using Dirac inputs for ETC, as we do for developing the learning trigger, has also been investigated in other works on ETC such as [4]–[6]. However, the approximation we propose to take into account input saturations has not been discussed therein. Moreover, none of these references considers learning approaches to cope with changing system dynamics or disturbances.

The problem of finding a replacement for continuous or periodic integral control that is suitable for ETC (*e.g.*, to deal with load disturbances) has for instance been addressed in [7], where a disturbance observer is used. Instead of introducing a disturbance observer, we directly include the load disturbances in the learning framework. As PID-controllers are the most common controllers used in industry, event-triggered PID-control has also been investigated starting from [8]. A particular problem here is the replacement of the integral part of the PID-controller [9]. Mostly, a network between sensor and controller is considered, thus the main problem for the integral part is the non-constant sampling time of the event-triggered mechanism. In [10], this is dealt with by explicitly taking into account the actual sampling time instead of assuming a nominal, constant sampling time. A different approach is presented in [11], where the event detector is connected to the sensor. Instead of looking at the absolute value of the integrator, the difference between the current value and the value at the last triggering instant is used to trigger communication, as a constant value of the integrator indicates a control error of zero. Replacing the integrator through model learning has not been proposed yet.

Event-triggered controllers can also be learned from data without learning a model. Such approaches are proposed, for example, in [12]–[15]. In contrast to those approaches, we use a specific control design and use learning to obtain accurate dynamic models.

Outline. In the following section, we formulate the problem setting. After that, we introduce the approach for event-triggered pulse control and discuss the concrete implementation of the learning trigger in Sec. III. In Sec. IV, we will present a numerical study and conclude with a discussion in Sec. V.

II. PROBLEM FORMULATION

We consider linear, time-invariant systems of the form

$$dx(t) = Ax(t) dt + Bu(t) dt + \epsilon dt + Q dW(t), \quad (1)$$

with the state $x(t) \in \mathbb{R}^n$, the control input $u(t) \in \mathbb{R}^m$, the constant load disturbance $\epsilon \in \mathbb{R}^n$, and $W(t) \in \mathbb{R}^n$ a multidimensional Wiener process representing process noise. We assume that we can measure the full state, thus, $y = x$ in Fig. 1.

As depicted in Fig. 1, control commands have to be transmitted over a communication network. We thus employ an event-triggered design, with the block ‘State Trigger’ implemented by

$$\gamma_{\text{ctrl}} = 1 \iff \|x(t)\|_2 \geq \delta, \quad (2)$$

where δ is a user-defined threshold and is essentially the deviation from the equilibrium that we are willing to tolerate. In case of an event, *i.e.*, $\gamma_{\text{ctrl}} = 1$, we apply a pulse to reset the system to its equilibrium,

$$u(t) = \begin{cases} 0 & \text{if } t - t_k > t_{\text{inp}} \vee t_k = 0 \\ \phi_{u_{\text{max}}}(\hat{A}, \hat{B}, \hat{\epsilon}, x_{t_k}) & \text{if } t - t_k \leq t_{\text{inp}} \wedge t_k \neq 0, \end{cases} \quad (3)$$

where $\phi_{u_{\text{max}}}$ is the pulse generating policy, $(\hat{A}, \hat{B}, \hat{\epsilon})$ captures the model of the system dynamics, x_{t_k} is the state of the system at the triggering instant t_k , t_{inp} is the pulse length (see Fig. 1), and u_{max} the maximum input the actuator can apply. By applying a pulse with appropriate length, we can reset the system to its equilibrium state. This, however, requires that we have a model that accurately describes the true system dynamics. We obtain this model, and adapt it in case the dynamics change, via model learning techniques. Model learning may be expensive due to the involved computations or required exploration, therefore, we only want to learn in case the estimated dynamics $(\hat{A}, \hat{B}, \hat{\epsilon})$ deviate too much from the true dynamics (A, B, ϵ) . Since the true dynamics are unknown, the decision needs to be based on some implicit feature, which will be the communication signal. Developing such a learning scheme for ETC is the main objective of this paper.

III. EVENT-TRIGGERED PULSE CONTROL WITH MODEL LEARNING

In this section, we present the event-triggered control and learning framework. We start with a derivation of the learning trigger assuming Dirac impulse control, then show how we learn the system dynamics, and finally detail practical pulse control (with bounded pulses) for first-order systems.

A. Event-triggered Learning for Control

The learning trigger is based on the framework presented in [3] for ETSE. Here, we extend this framework to ETC. For the theoretical analysis, we assume a control strategy based on Dirac impulses, *i.e.*, here, we will ignore the assumption of having an input saturation at the actuator. Using Dirac impulses, we can, similarly as in ETSE [3], reset the error to zero at communication times. We then have a control law of the form

$$u(t) = F \delta_{t_k}(t), \quad (4)$$

where F is the control gain and δ_{t_k} the Dirac impulse. In particular, the control input is zero apart from the triggering times t_k , where t_k corresponds to $\gamma_{\text{ctrl}} = 1$ in (2). To further

analyze this scheme, we write (1) in integrated form,

$$\begin{aligned}
x(t_k) &= \int_{t_{k-1}}^{t_k} e^{A(t_k-t)} B u(t) dt \\
&+ \underbrace{\int_{t_{k-1}}^{t_k} e^{A(t_k-t)} \epsilon dt + \int_{t_{k-1}}^{t_k} e^{A(t_k-t)} Q dW(t)}_{:=N(t_k)} \quad (5) \\
&= e^{A(t_k-t_k)} B F + N(t_k) \\
&= B F + N(t_k),
\end{aligned}$$

where we assume that the process starts in $x(t_{k-1}) = 0$ and $N(t_k)$ is the measurement we get before applying the impulse. By setting (5) to zero we can (assuming B invertible) show that $F = -B^{-1}N(t_k)$ resets the system to zero. Implementing such a control law then also fulfills the prior assumption of $x(t_{k-1}) = 0$, as the system starts in zero after every triggering instant. In Sec. III-C, we will drop the assumption of being able to apply Dirac impulses as inputs and instead apply pulses with the maximum input u_{\max} for a given time.

Considering a control law as proposed in (4), we thus have a random process that always starts in zero. This is only true, if the input matrix B is known exactly. In that case, the sole cause of an error would be propagated noise and the load disturbance ϵ . Therefore, in case of no communication, we obtain

$$x(t) = \int_0^t e^{A(t-s)} \epsilon ds + \int_0^t e^{A(t-s)} Q dW(s). \quad (6)$$

We can now define a stopping time τ as the first moment the state crosses the threshold δ , which resets the error to zero,

$$\tau := \inf\{t : \|x(t)\|_2 \geq \delta\}. \quad (7)$$

The stopping times defined in (7) coincide with the time between communication, hence, ‘stopping times’ and ‘inter-communication times’ will be used synonymously hereafter. We can now further define the expected value of these stopping times, $\mathbb{E}[\tau|x(0) = 0]$, which is the average inter-communication time of the system. This expected value can be obtained via Monte Carlo simulations (for a more detailed discussion, see [3]).

If we had a perfect model of the system dynamics, the average inter-communication times that we observe in the system should approach the expected value of the stopping time. If both values deviate by too much, we have evidence that the model is inaccurate and can trigger learning of a new model. Precisely, we define the learning trigger in Fig. 1 as

$$\gamma_{\text{learn}} = 1 \iff \left| \frac{1}{N} \sum_{i=1}^N \tau_i - \mathbb{E}[\tau] \right| \geq \kappa. \quad (8)$$

In this equation, $\gamma_{\text{learn}} = 1$ indicates that a new model shall be learned, $\mathbb{E}[\tau]$ is approximated using Monte Carlo simulations, *i.e.*, $\mathbb{E}[\tau] \approx \frac{1}{M} \sum_{i=1}^M \tau_i^{\text{sim}}$, and $\tau_1, \tau_2, \dots, \tau_N$ define the

last N empirically observed inter-communication times. Due to the randomness of the process, it can still happen that we trigger learning despite the model being perfect. Assuming that the stopping times are bounded by τ_{\max} , the confidence level can be quantified using Hoeffding’s inequality [16] and influenced through the design parameter η .

Theorem 1: Let the parameters $\eta, N, M > N$, and τ_{\max} be given, τ_1, \dots, τ_N and $\tau_1^{\text{sim}}, \dots, \tau_M^{\text{sim}}$ independent and identically distributed, and assume a perfect model. For

$$\kappa = \tau_{\max} \sqrt{-\frac{2}{N} \ln \frac{\eta}{4}} \quad (9)$$

we obtain

$$\mathbb{P} \left[\left| \frac{1}{N} \sum_{i=1}^N \tau_i - \frac{1}{M} \sum_{i=1}^M \tau_i^{\text{sim}} \right| \geq \kappa \right] < \eta. \quad (10)$$

Proof: We compare stopping times obtained via Monte Carlo simulations with stopping times observed from the real process. In both cases, we have a random process that always starts in zero. This is the same setting as investigated in [3], thus, the theorem can be proven as shown therein. ■

Boundedness of the stopping times can easily be ensured in practice by applying a control input the latest when τ_{\max} is reached. The confidence parameter η then basically defines the tradeoff between accepting an inaccurate model or triggering learning despite the model being perfect. Intuitively, η defines the probability that the error κ is observed, while empirical and expected stopping times are drawn from the same distribution (*i.e.*, we have a perfect model). If this probability is below a predefined threshold, we trigger learning of a new model.

B. Model Learning

For the derivation of the stopping times as well as for the final controller design, we need knowledge of the full system dynamics (matrices A and B) and the load disturbance ϵ . To calculate the stopping times, *e.g.*, via Monte Carlo simulations, we additionally need knowledge of the process noise variance Q . To estimate the model, we rewrite system (1) in discrete time,

$$\begin{aligned}
x(k+1) &= A_d x(k) + B_d u(k) + \epsilon + v(k) \\
&= (A_d \quad B_d \quad \epsilon) \begin{pmatrix} x(k) \\ u(k) \\ 1 \end{pmatrix} + v(k), \quad (11)
\end{aligned}$$

with A_d, B_d the discrete-time system and input matrix, respectively, and $v(k)$ the discrete-time process noise. That way, we can learn the system dynamics with standard least-squares techniques, as we will demonstrate in Sec. IV.

Having knowledge of the load disturbances, we can incorporate them in the control design in Sec. III-C. This represents a suitable solution to replace the integral part of standard, periodic controllers.

Remark 1: Another problem that may be considered with this approach is the knowledge of the zero-level of the system. We are considering an equilibrium at $x(t) = 0$, but the measurements are actually voltage signals from a sensor

and what zero means for that system is not clear from the beginning. We can model this as a sensor bias, *i.e.*, we would have the following system dynamics

$$x(k+1) = A_d x(k) + B_d u(k) + v(k) \quad (12a)$$

$$y(k) = x(k) + \xi, \quad (12b)$$

where ξ is the sensor bias. Rewriting this yields

$$\begin{aligned} y(k+1) &= x(k+1) + \xi \\ &= A_d(y(k) - \xi) + B_d u(k) + v(k) \\ &= \begin{pmatrix} A_d & B_d & (I - A_d)\xi \end{pmatrix} \begin{pmatrix} y(k) \\ u(k) \\ 1 \end{pmatrix} + v(k). \end{aligned} \quad (13)$$

Thus, we can identify the system dynamics and the sensor bias via least-squares techniques. Estimating both, a sensor bias and a load disturbance, is, however, not possible, as they are not distinguishable given the output data.

C. Implementation of Event-triggered Pulse Control with Input Saturation

If the state of the system exceeds the threshold δ , we want to quickly reset it to zero. Application of Dirac impulses is not compatible with our assumption of an input saturation at the actuator.

Instead of applying Dirac impulses, we propose to apply the maximum input and vary the *duration* of the pulse. This has two main benefits: 1) The input cannot exceed the saturation and, thus, will drive the system state to its desired value, as it will not be limited by the saturation; 2) The system will be driven to zero as fast as possible, that way coming as close to the idealized Dirac input as possible. This represents a straightforward approach to lift the idea of ETL to the ETC setting. For the derivation of the length of the pulse, we will restrict ourselves to first-order systems (*i.e.*, $n = 1$ in (1) with scalar variables a and b) and later comment on extensions to higher-order systems.

To derive the length of an impulse, we look at the system equation in integrated form. If no event is triggered, *i.e.*, if the system is close to its desired state, we have $u(t) = 0$. At triggering times t_k we apply the maximum input,

$$x(t) = e^{at} x(t_k) + \int_{t_k}^t e^{a(t-s)} (bu_{\max} + \epsilon) ds. \quad (14)$$

The input shall be applied for long enough such that the state becomes zero. We thus set (14) to zero and solve for t (setting $t_k = 0$),

$$0 \stackrel{!}{=} e^{at} x(0) + \int_0^t e^{a(t-s)} (bu_{\max} + \epsilon) ds,$$

which leads to

$$t_{\text{inp}} = \frac{1}{a} \ln \left(\frac{bu_{\max} + \epsilon}{ax(0) + bu_{\max} + \epsilon} \right). \quad (15)$$

In (14) and (15), we assumed the noise to be zero during the application of the pulse. For the Dirac impulse, this

holds, as the time of the application tends to zero. Here, we explicitly derive how long the input will be applied, hence, the system will during this time also be excited by noise and we will not be able to exactly drive it to zero. We take this into account when computing the stopping times. Instead of starting at zero, we consider a process that starts in $x(0) \sim \mathcal{N}(0, \Sigma_0)$, with variance Σ_0 .

Remark 2: For a system (1) with state dimension $n > 1$, a single pulse will generally not be sufficient to drive the state to zero. Instead, we would have to change between maximum and minimum input, which leads to a bang-bang controller [17]. Then, t_{inp} in Fig. 1 would be a vector containing the switching points. Different algorithms have been proposed in literature to find these switching points [18], [19]. This paper is a first approach towards combining the ideas of ETL with ETC and the extension to bang-bang type of controllers is beyond the scope of this work. Generalizations to such or other control structures are, however, certainly important questions for future work.

IV. NUMERICAL STUDY

For the numerical study, we will consider a collection of first-order processes. For each system, we assume a remote controller that is colocated with the sensor, but needs to transmit its actuation commands over a communication network, where all controllers share the same network. For all examples, we assume a threshold of $\delta = 0.02$. As parameters of the learning trigger, we choose a confidence level $\eta = 0.05$, $N = 2000$, $M = 10\,000$, and $\tau_{\max} = 1$ s. According to Theorem 1, we then obtain $\kappa \approx 0.066$.

In [20, p. 227], a batch of process models that are common in process industry is collected. Among others, typical parameters for first-order systems with time delay are provided. We will start the numerical investigation of the proposed framework in Sec. IV-A by looking at these models, but we neglect the time delay. The models have stable dynamics as is a common property in process industry. To showcase the capability of the framework to also deal with unstable systems, we will consider such examples in Sec. IV-B. For all investigated systems, we will consider additional process noise and load disturbances. We model the load disturbance to enter with the input, similar as for instance done in [21, p. 54].

As in (1), we assume continuous-time systems

$$dx(t) = ax(t) dt + b(u(t) + \epsilon) dt + Q dW(t), \quad (16)$$

which we discretize with a sample time of 1 ms. The sample time is not equal to the update interval of the communication system and is only limited by the maximum frequency of the timers in the processors used for controller and actuator. A fine discretization is necessary, as we will derive a continuous pulse length. The finer the discretization is, the more accurate is the application of the pulse (and the earlier we notice if the system is outside the tolerable range).

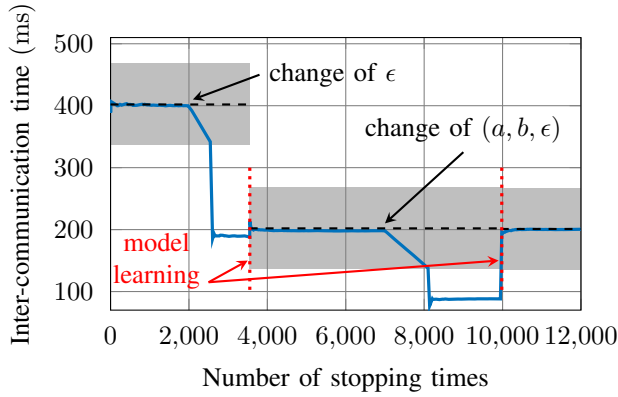


Fig. 2. Average inter-communication times during one simulation for Example 1. The solid line shows the empirically observed inter-communication times computed as a moving average over 2000 stopping times. The moving average is reset in case learning is triggered and when a new model has been learned. The dashed line indicates the expected inter-communication times with confidence interval $\pm\kappa$ highlighted in gray. The dynamics change after 2000 and after 7000 stopping times. The points in time where new system matrices are learned are marked with vertical, red, dotted lines. In both cases, this causes a decrease in the inter-communication time (*i.e.*, more communication). The inter-communication time increases again after learning new matrices.

A. Stable Dynamics

As a proof of concept, we first consider one specific from [20] (note that the first-order systems provided therein always lead to $a = b$):

Example 1: System (16) with $a = b = -0.01$, process noise $Q = 10^{-4}$, load disturbance $\epsilon = 5$, and maximum input $u_{\max} = 100$. In case learning is triggered, we always collect data for 200 s and then use these data points to identify the system dynamics.

In Fig. 2, the average and expected inter-communication times are shown. The average inter-communication times are computed with a moving average over 2000 stopping times, which we reset in case learning is triggered and after deriving new system matrices. In the beginning, we assume that we have an accurate model, hence, the observed inter-communication times approach the expected ones. After 2000 stopping times, we set the load disturbance to $\epsilon = 10$. As expected, the inter-communication times decrease and learning is triggered. After learning (first vertical line in Fig. 2), the empirical inter-communication times again approach the expected ones and we reduce communication. In a second change, after 7000 stopping times, we have $\epsilon = 20$ and $a = b = -0.05$. Similar as before, this leads to a decrease of the average inter-communication time and learning is triggered. Having learned new dynamics (second vertical line in Fig. 2), the empirical inter-communication times again approach the expected ones, *i.e.*, average communication is reduced through learning.

We now consider the case where the initial matrices that we have of the process are wrong. Both, nominal and true dynamics of the systems are again taken from [20]:

Example 2: Two systems of the form (16) with nominal dynamics $a = b \in \{-0.1, -1\}$, load disturbances $\epsilon \in$

TABLE I

COMPARISON OF THE AVERAGE INTER-COMMUNICATION TIMES FOR EXAMPLE 2 BEFORE AND AFTER LEARNING. SYSTEM 1 IS SHOWN IN THE LEFT COLUMN, SYSTEM 2 IN THE RIGHT.

a, b	Before	After	a, b	Before	After
-10	1 ms	21 ms	-0.25	1 ms	80 ms
-4/3	1 ms	39 ms	-0.05	69 ms	132 ms
-0.5	114 ms	215 ms	-0.02	85 ms	431 ms
-0.25	111 ms	428 ms	-0.01	62 ms	384 ms
-1/6	101 ms	321 ms	-0.005	115 ms	850 ms

TABLE II

COMPARISON OF THE AVERAGE INTER-COMMUNICATION TIMES FOR EXAMPLE 3 BEFORE AND AFTER LEARNING.

System	Before	After	System	Before	After
1	44 ms	239 ms	6	128 ms	271 ms
2	136 ms	349 ms	7	56 ms	240 ms
3	107 ms	283 ms	8	55 ms	219 ms
4	135 ms	256 ms	9	89 ms	279 ms
5	163 ms	363 ms	10	34 ms	273 ms

$\{0.1, 1\}$, $Q = 10^{-4}$, and $u_{\max} \in \{1, 100\}$. For the true dynamics, a and b are given in Table I, the noise variance and load disturbance are sampled from uniform distributions over the intervals $Q \in [10^{-4}, 10^{-3}]$, $\epsilon \in [0.1, 0.2]$ (for system 1), and $\epsilon \in [1, 5]$ (for system 2). In case learning is triggered, we use all data collected so far to identify the system dynamics.

The inter-communication times before and after learning the dynamics are given in Table I. It can clearly be seen that learning helps in lowering communication.

B. Unstable Dynamics

For the investigation of systems with unstable dynamics, we look at the following system:

Example 3: System (16) with nominal dynamics $a = 5$, $b = 3$, $\epsilon = 0.01$, $Q = 10^{-4}$, and maximum input $u_{\max} = 1$. The true dynamics are sampled from random distributions over the intervals $a \in [1, 10]$, $b \in [1, 2]$, $\epsilon \in [0.01, 0.02]$, and $Q \in [10^{-4}, 10^{-3}]$. As in Example 2, we use all data we have collected so far to identify a new model in case learning is triggered.

In Table II, we compare the average inter-communication times of all systems before and after deriving new system matrices. For all of them, we observe a significant increase in the inter-communication times after learning, *i.e.*, communication is reduced. This demonstrates that the approach is also suitable for unstable systems.

In Fig. 3, one specific system from Example 3 is shown before (Fig. 3a) and after (Fig. 3b) learning new system matrices. Due to the error in the initial matrices, the system is not reset to zero with the pulses before learning and, thus, new control inputs have to be generated very frequently. After learning, the pulse length is such that the system is successfully reset, which also results in increased inter-communication times, and, therefore, less communication. This is especially emphasized as in Fig. 3a, before learning, we show only 1 s, while in Fig. 3b we show 2 s and still observe far less pulses.

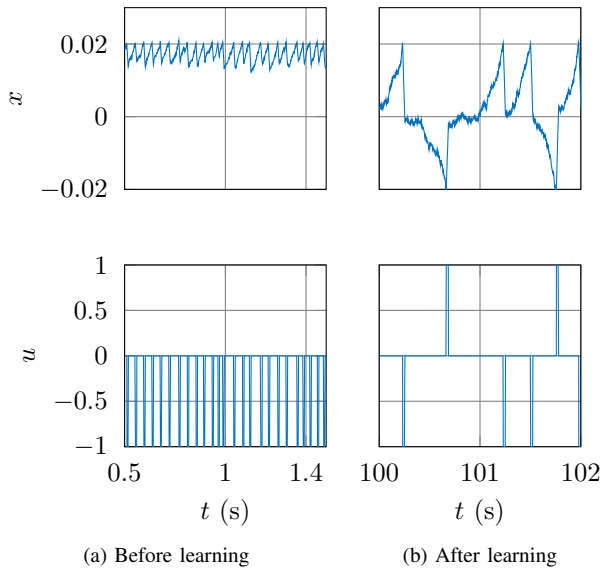


Fig. 3. Performance of one specific system of Example 3 before (left) and after (right) learning. It can be seen that before learning, the pulses are too short and the system is not reset to zero, while after learning the pulse length is appropriate. Further, communication is significantly reduced through learning.

The study reveals that the proposed architecture enables us to increase inter-communication times through learning. We are able to learn system dynamics and subsequently reset the state of the system to zero in case it leaves its tolerable range. Through learning load disturbances, the architecture is a suitable replacement for integral control in event-triggered settings.

V. CONCLUSION

In NCS, communication is a scarce and limited resource. In this work, we presented a framework for event-triggered pulse control for NCS. Most common event-triggered control approaches rely on the availability of an accurate dynamics model. Contrary to that, the proposed framework does not rely on this assumption, but uses model learning instead. As learning is expensive (*e.g.*, due to the involved computations), we only learn if necessary using the ETL framework. By observing the communication behavior, we quantify the accuracy of the model and trigger learning of a new model only in case the accuracy is not sufficient. The presented control design respects input saturations and can also handle load disturbances, thus essentially replacing the integral part of common periodic controllers.

A numerical study demonstrates the applicability of the approach and the benefit of learning the system dynamics. After learning, we observe a significant increase in the inter-communication time. However, the presented examples are first-order systems. While we have outlined how higher-order systems could be treated (Remark 2), the actual extension to such systems is an interesting topic for future work. Moreover, we assumed that we are able to perfectly measure the full state of the system. Incorporating Gaussian measurement noise is already possible with the presented approach. How

to extend the ETL framework to partial state measurements is subject to ongoing research.

In this work, we proposed for the first time to trigger model learning experiments in ETC to adapt to changes in the dynamics. The learning trigger compares the expected and the observed time between communication. While this is an intuitive approach, in some cases this trigger does not detect disturbed models. A more robust behavior can be achieved by triggering on the full distribution, *e.g.*, via a Kolmogorov-Smirnoff test, which is subject of current research [22].

REFERENCES

- [1] W. P. M. H. Heemels, K. H. Johansson, and P. Tabuada, "An introduction to event-triggered and self-triggered control," in *51st IEEE Conference on Decision and Control*, Dec 2012.
- [2] M. Miskowicz, *Event-Based Control and Signal Processing*. CRC Press, 2016.
- [3] F. Solowjow, D. Baumann, J. Garcke, and S. Trimpe, "Event-triggered learning for resource-efficient networked control," in *American Control Conference*, 2018.
- [4] K. J. Åström, "Event based control," in *Analysis and design of nonlinear control systems*. Springer, 2008.
- [5] T. Henningsson, E. Johannesson, and A. Cervin, "Sporadic event-based control of first-order linear stochastic systems," *Automatica*, vol. 44, no. 11, 2008.
- [6] X. Meng and T. Chen, "Optimal sampling and performance comparison of periodic and event based impulse control," *IEEE Trans. Automat. Contr.*, vol. 57, no. 12, 2012.
- [7] A. Cervin and K. J. Åström, "On limit cycles in event-based control systems," in *46th IEEE Conference on Decision and Control*, 2007.
- [8] K.-E. Årzén, "A simple event-based PID controller," *IFAC Proceedings Volumes*, vol. 32, no. 2, 1999.
- [9] J. Sánchez, A. Visioli, and S. Dormido, "Event-based PID control," in *PID Control in the Third Millennium*. Springer, 2012.
- [10] S. Durand and N. Marchand, "Further results on event-based PID controller," in *European Control Conference*, 2009.
- [11] M. Rabi and K. H. Johansson, "Event-triggered strategies for industrial control over wireless networks," in *4th Annual International Conference on Wireless Internet*, 2008.
- [12] K. G. Vamvoudakis and H. Ferraz, "Model-free event-triggered control algorithm for continuous-time linear systems with optimal performance," *Automatica*, vol. 87, 2018.
- [13] X. Zhong, Z. Ni, H. He, X. Xu, and D. Zhao, "Event-triggered reinforcement learning approach for unknown nonlinear continuous-time system," in *2014 International Joint Conference on Neural Networks*, July 2014.
- [14] X. Yang, H. He, and D. Liu, "Event-triggered optimal neuro-controller design with reinforcement learning for unknown nonlinear systems," *IEEE Trans. Syst., Man, Cybern., Syst.*, 2018.
- [15] D. Baumann, J.-J. Zhu, G. Martius, and S. Trimpe, "Deep reinforcement learning for event-triggered control," in *57th IEEE Conference on Decision and Control*, Dec. 2018.
- [16] W. Hoeffding, "Probability inequalities for sums of bounded random variables," *Journal of the American Statistical Association*, vol. 58, no. 301, 1963.
- [17] R. Bellman, I. Glicksberg, and O. Gross, "On the bang-bang control problem," *Quarterly of Applied Mathematics*, vol. 14, no. 1, 1956.
- [18] C. Y. Kaya and J. L. Noakes, "Computations and time-optimal controls," *Optimal Control Applications and Methods*, vol. 17, no. 3, 1996.
- [19] J. Wen and A. Desrochers, "An algorithm for obtaining bang-bang control laws," *Journal of Dynamic Systems, Measurement, and Control*, vol. 109, no. 2, 1987.
- [20] K. J. Åström and T. Häggglund, *Advanced PID control*. ISA-The Instrumentation, Systems, and Automation Society Research Triangle, 2006, vol. 461.
- [21] —, *PID controllers: theory, design, and tuning*. Instrument society of America Research Triangle Park, NC, 1995, vol. 2.
- [22] F. Solowjow and S. Trimpe, "Event-triggered learning," under review, available at <https://is.mpg.de/publications/solowjowetl19>.