Real-time optimal feedback control of switched autonomous systems

Xu Chu Ding [*] **Axel Schild** [**] **Magnus Egerstedt** [*] **Jan Lunze** [**]

[*] *School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA 30332, USA.*
e-mail: {ding,magnus}@ece.gatech.edu
[**] *Department of Electrical Engineering and Information Sciences, Ruhr-Universitaet Bochum, 44780 Bochum, Germany.*
e-mail: {schild,lunze}@ atp.rub.de

Abstract: This paper presents a novel optimal adaptive feedback control concept for nonlinear switched autonomous systems. Our proposed controller implements a closed-loop switching law, which is represented by a parametrized switching plane in the state space. During the plant operation, the plane parameters are incrementally adapted to state measurements in real time by iteratively solving an associated nonlinear optimization problem over a finite time horizon. The combination of this control approach with ideas from model predictive control yields the so-called *crawling window optimal control* scheme that achieves the optimal stabilization of periodically operated switched systems.

Keywords: Switched Systems, Optimal Control, Receding Horizon Control

1. INTRODUCTION

This paper addresses the subject of optimal feedback control of switched autonomous systems that are subject to disturbances. A switched system refers to a continuous plant, which provides a finite set of operational modes, where each mode is associated to different continuous dynamics. The control task involves regulation of the state evolution according to a given objective function, which must be achieved solely by switching the plant among its modes in the right order at proper time instants.

Commonly, the mode sequence follows directly from the plant set-up, so that only the switching times constitute free decision variables. The off-line computation of optimal transition times for a predetermined mode sequence has been extensively studied, e.g. by Egerstedt et al. (2006); Shaikh and Caines (2007); Xu and Antsaklis (2002). In the presence of disturbances, such a feed-forward control often result in a degraded performance.

To alleviate the effects of unknown disturbances, Boccadoro et al. (2005) proposed an optimal parametrization of switching surfaces that implement an *explicit* state-feedback switching policy. As these surfaces are designed with respect to *one* specific reference trajectory, the state evolution is generally not optimal in the presence of large deviations from the predicted nominal path. One possibility to cope with large deviations from the anticipated path is provided by the alternative *implicit* feedback control scheme, which was pursued in Wardi et al. (2007). The latter approach is essentially based on a real-time optimization of the switching times during the plant operation. Its large computational complexity, however, may cause an unsatisfactory loop response under rapidly varying deviations.

Motivated by the limitations of existing approaches, this paper presents a novel *optimal adaptive control scheme* for switched systems, which merges implicit and explicit feedback control into a unified framework. The proposed cascaded control scheme enjoys the benefits of both worlds and due to its true feedback nature still produces a satisfactory performance when existing approaches fail to work in the presence of unknown disturbances. As a core feature, the state-dependent switching policy is expressed in terms of a *time-varying* switching plane, which triggers mode transitions when being intersected by the state trajectory. We explain how to successively update these switching planes in real time on the basis of state measurements to asymptotically obtain local approximations of the optimal switching surface.

By combining the adaptive control strategy with ideas from receding horizon control, we also outline a control strategy for periodically switched systems to achieve orbital stability of a predetermined limit cycle under a suboptimal infinite-horizon performance. As a key feature, this strategy allows us to trade complexity against performance. The space constraints only allow for a broad summary of the control concept and its properties. Proofs to all theorems can be found in the extended version in Schild et al. (2009b).

2. FEEDBACK CONTROL OF SWITCHED SYSTEMS

2.1 Control loop structure

This paper addresses the design of optimal switching laws

$$q(t^+) = c(x(t), q(t)), \quad q(0) = q_0 \ , \tag{1}$$

i.e. optimal feedback controller for the switched system

$$\dot{x}(t) \ = f_{q(t)}(x(t)) + w(t) \tag{2}$$

subject to unknown disturbances $w(t)$. Here, $q(t)$ denotes the operation mode of the switched system at time t and is assumed to follow a predetermined finite sequence $\overline{q} = \{q_0, q_1, \ldots, q_{N-1}\}$.

[⋆] This work is partially supported by the Deutsche Forschungsgemeinschaft (LU462/21-3), the German Academic Exchange Service (D/08/45420) and the US National Science Foundation (Grant 0509064).

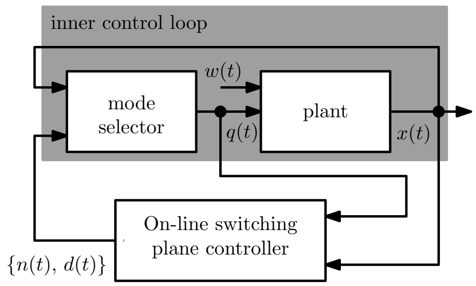

Fig. 1. Cascaded control loop structure

Our realization of the law (1) utilizes a cascaded control loop, as shown in Figure 1, which involves two basic components that operate on different time scales. The first component is a continuous-time *mode selector*. Together with the plant, it constitutes the *inner control loop* and implements a parametrized event function

$$\Phi(x(t), n(t), d(t)) = d(t) - n^{\mathrm{T}}(t)x(t) \tag{3}$$

that can be evaluated with negligible computational effort. The mode selector outputs a piecewise constant mode signal $q(t)$ which changes its value when $\Phi(x(t), n(t), d(t)) \leq 0$ and thus activates the next mode in the sequence $\overline{q}$. Adequate parameter trajectories $d(\cdot)$ and $n(\cdot)$ are provided by the second component, the *switching plane controller*

$$x_c(t_{s+1}) = f_c(x_c(t_s), x(t_s), q(t_s), r(t_s)), \quad x_c(0) = x_{c,0} \tag{4}$$

$$\left(n^{\mathrm{T}}(t_s),\ d(t_s)\right) = h_c\left(x_c(t_s), x(t_s), q(t_s)\right) \quad , \tag{5}$$

which constitutes a discrete-time dynamical system whose state $x_c(t)$ is updated every Δt time units at sampling times t_s on the basis of the hybrid state $(x(t), q(t))$. Together with the inner control loop, this component forms the outer control loop.

The resulting overall control loop operates as follows: after the activation of the k-th mode, the state $x(t)$ evolves uncontrolled according to the k-th *modal function* f_{q_k} and the disturbance $w(t)$. This evolution continues until satisfaction of the switching condition $\Phi(x(t), n(t), d(t)) \leq 0$ is detected. The exact *switching time* $\tau_{k+1} = \min_{t \geq \tau_k} t \ : \ d(t) - n^{\mathrm{T}}(t)\, x(t) = 0$ indicates the first intersection of the state trajectory $x(\cdot)$ and the time-varying switching plane

$$\mathcal{T}(t) = \{x : d(t) - n^{\mathrm{T}}(t)x = 0\} \tag{6}$$

in the state space, upon which a transition from $q(t) = q_k$ to

$$q_{k+1} = g(q_k), \quad q_0 \text{ given} \tag{7}$$

is initiated. Subsequently, the same procedure is repeated for all remaining modes. Provided a sufficiently high update frequency and under reasonable assumptions, the switching plane controller incrementally shifts and rotates $\mathcal{T}(t)$ in state space during operation of the plant.

For compactness, let $\overline{\tau} = \{\tau_k\}_{k=1}^{N}$ and $\overline{x} = \{x_k\}_{k=1}^{N}$ denote the *switching time sequence* and the *switching point sequence*, which contains all switching points $x_k = x(\tau_k)$. Moreover, we call $\delta_k = \tau_{k+1} - \tau_k$ the activation duration of mode q_k and $\overline{\delta} = \{\delta_k\}_{k=0}^{N-1}$ the corresponding *activation duration sequence*.

2.2 Finite-horizon loop specifications

In finite-horizon control, the aim is to optimally transfer the state from any x_0 inside a given initial set $\mathcal{X}_0 \subset \mathbb{R}^n$ to a desired terminal region $\mathcal{X}_T$, which leads to the following design task:

Problem 1. **Given** a plant (2) operating in a noise-free environment ($w(t) = 0$) and an initialized mode selector (3), (7) that generates a predetermined finite mode sequence $\overline{q}$, **determine** a switching plane controller (4), (5), such that the closed-loop generates sequences $\overline{\tau} = \overline{\tau}(x_0, q_0, n(\cdot), d(\cdot))$ that locally minimize

$$J_{\overline{q}}(x_0, \overline{\tau}) = \sum_{k=0}^{N-1} \left(\int_{\tau_k}^{\tau_{k+1}} L_k(x(t))\mathrm{d}t + \phi_k(x_k) \right) + \phi_N(x_N) \tag{8}$$

subject to the dynamics (2) and

$$\tau_0 \leq \tau_1, ..., \leq \tau_N \tag{9}$$

$$\psi(x_N) \leq 0 \tag{10}$$

for any initial state $x_0 \in \mathcal{X}_0$.

Here, the terminal constraint $\psi(x)$ encodes the terminal region $\mathcal{X}_T$, while the input constraints (9) arise naturally to ensure the correct mode sequence. Functions L_k encode instantaneous costs, which assess the transient behavior, ϕ_k accounts for potential switching costs and ϕ_N represents terminal costs.

The aforementioned formulation ignores disturbance effects, as the considered system class does not allow for devising a controller that rigorously ensures an optimal performance under disturbances. Our design procedure for the maps (4) and (5), however, goes beyond the minimal specifications of Problem 1 and explicitly ensures a locally optimal loop behavior under disturbances in a meaningful sense.

2.3 Infinite-horizon specifications for switched periodic systems

Infinite-horizon optimal control problems arise in the stabilization of periodically switched systems that recurrently execute a given *mode sequence*: $\overline{q}^{\Gamma} = (q_0\, q_1 \ldots q_{p-1})$.

Definition 1. A trajectory $x^{\Gamma}(\cdot)$ that starts at x_0^{Γ} and evolves according to (2) under the mode sequence $\overline{q}^{\Gamma}$ is periodic of order p, if it satisfies $x^{\Gamma}(t) = x^{\Gamma}(t+T)$, $\forall t > t_0$ with $T = \sum_{k=0}^{p-1} \delta_k^{\Gamma}(x_0^{\Gamma})$.

Definition 2. A limit cycle of a switched system (2) represents a closed orbit $\Gamma = \{x : \exists t \in [0, T], x = x^{\Gamma}(t)\}$ in the state space.

Definition 3. A limit cycle Γ is asymptotically orbitally stable, if for every $\varepsilon > 0$, $\exists \delta(\varepsilon) > 0$, such that $\forall x_0$ satisfying $\mathrm{dist}(x_0, \Gamma) \leq \delta(\varepsilon)$ the following holds:

$$\mathrm{dist}(x(t), \Gamma) \leq \varepsilon, \ \forall t > t_0, \quad \lim_{t \to \infty} \mathrm{dist}(x(t), \Gamma) = 0 \ . \tag{11}$$

To simplify the presentation, we limit our attention to the stabilization of *fundamental* limit cycles, for which $\overline{q}^{\Gamma}$ only contains distinct elements (hence $q_i \neq q_j$ if $i \neq j$).

Problem 2. **Given** a plant (2) operating in a noise-free environment ($w(t) = 0$), an initialized mode selector (3), (7) that generates a predetermined periodic mode sequence $\overline{q} = (\overline{q}^{\Gamma}\overline{q}^{\Gamma} \ldots)$ and a desired limit cycle Γ, **determine** a controller (4), (5), such that the closed-loop generates sequences $\overline{\tau} = \overline{\tau}(x_0, q_0, n(\cdot), d(\cdot))$ that orbitally stabilize Γ and locally minimize the objective

$$J_{\infty}(x_0, \overline{\tau}, \Gamma) = \sum_{k=0}^{\infty} \left(\int_{\tau_k}^{\tau_{k+1}} L_k(x(t))\mathrm{d}t + \phi_k(x_k) \right) \tag{12}$$

subject to dynamics (2) and

$$\tau_k \leq \tau_{k+1}, \quad \forall k \geq 0 \tag{13}$$

for any initial state x_0 from inside a feasible region $\mathcal{X}_0$.

Problem 2 is meaningful, only if the costs $J_\infty(x_0, \overline{\tau}, \Gamma)$ stay bounded. Hence, L_k and ϕ_k cannot be arbitrary, but must suitably encode the orbital stabilization goal (11).

3. OPTIMAL ADAPTIVE FEEDBACK CONTROL OF SWITCHED SYSTEMS OVER A FINITE HORIZON

In this section we describe the algorithm of a switching plane controller (4), (5) and show that the closed-loop operation meets the specifications of Problem 1, if sufficient time for convergence is available. To simplify the presentation, we make the following assumptions:

(1) Problem 1 is feasible.
(2) The input constraints (9) are ineffective during operation.
(3) All functions f_k, L_k, ϕ_k, ϕ_N and ψ are sufficiently smooth.

3.1 Definition of predicted signal trajectories

In order to explain the working principle of the switching plane controller, we need to introduce the predicted loop behavior with respect to a *relative time frame*.

At a time $t \geq \tau_k$ after the k-th mode transition, the system's state is $x(t)$ and the active mode is $q(t) = q_k$. Denote the remaining mode sequence of length $N(t) = N - k$ as

$$\overline{\theta}(t) = \{\theta_0(t), ..., \theta_{N(t)-1}(t)\} = \{q_k, ..., q_{N-1}\} \quad . \tag{14}$$

In connection with (14), denote $\overline{\chi}(t) = \{\chi_0(t), ..., \chi_{N(t)-1}(t)\}$ as the *predicted activation duration sequence* and $\overline{\rho}(t) = \{\rho_1(t), ..., \rho_{N(t)}(t)\}$ as the *predicted switching time sequence* for the remaining $N(t)$ mode transitions in the simulated time frame. Note, that the time dependence of all predicted sequences actually implies the dependence on the hybrid state $(x(t),\ q(t))$. For brevity of notation, we use $\chi^\star(t) = \chi^\star(x(t), q(t))$ to refer to a locally optimal predicted activation duration sequence at time t and use $\overline{\chi}^\star(\cdot)$ to refer to a trajectory of sequences related to state trajectories $x(\cdot)$ and $q(\cdot)$.

Similarly, let $\xi(\cdot|t)$ refer to the *predicted state trajectory* over the relative time interval $[0, \rho_{N(t)}(t)]$, which is obtained by forward integration of the undisturbed plant model

$$\dot{\xi}(r|t) = f_{\theta_k(t)}(\xi(r|t)), \quad r \in [\rho_k(t), \rho_{k+1}(t)) \tag{15}$$

with the initial condition $\xi(0|t) = x(t)$. The *predicted switching points* of this simulated trajectory are consequently denoted as $\xi_k(t) = \xi(\rho_k(t)|t)$. An example of an elapsed trajectory $x(\cdot)$ and the attached prediction $\xi(\cdot|t_s)$ is shown in Figure 2.

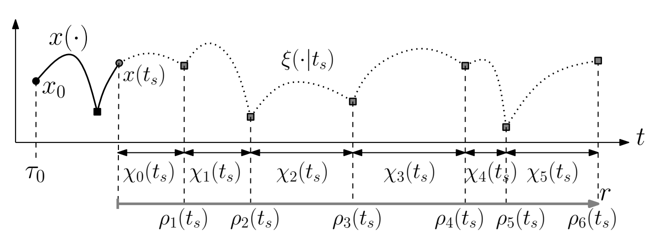

Fig. 2. Example of an elapsed state trajectory and its attached prediction (dotted line) at time t.

3.2 Optimal on-line adaptation of switching planes

With the previous definitions, we are in the position to summarize our procedure for incrementally adapting the parametrized switching plane $\mathcal{T}(t)$. The foundation for the determination of suitable values $n(t_s)$ and $d(t_s)$ at time t_s is the knowledge of the triple $\left(x(t_s),\ \overline{\rho}(t_s),\ \overline{\theta}(t_s)\right)$. In particular, we will see that optimal values $n^\star(t_s)$ and $d^\star(t_s)$ are obtained, if and only if $\overline{\rho}(t_s) = \overline{\rho}^\star(x(t_s), \overline{\theta}(t_s))$ constitutes a local minimizer of the cost-to-go problem

$$\min_{\overline{\rho}(t)} J_{\overline{\theta}(t_s)}(x(t_s), \overline{\rho}(t_s)) = \sum_{j=0}^{N(t_s)-1} \left(\int_{\rho_k(t)}^{\rho_{k+1}(t)} L_k(\xi(r|t_s))\mathrm{d}r + \phi_k(\xi_k(t_s)) \right)$$
$$+ \phi_N(\xi_{N(t_s)}(t_s)) \ \text{ subject to the dynamics (15) and}$$
$$\psi(\xi_{N(t_s)}(t_s)) \leq 0 \ ,$$

where L_k, ϕ_k, ϕ_N and ψ are all identical to the functions of the original Prob. 1. As explained in Wardi et al. (2007), the above problem can be solved iteratively in real-time by determining the Newton direction $\Delta\overline{\rho}(t_s)$ at each sample time t_s. This descent direction allows for improving the current guess $\overline{\rho}(t_s) = \overline{\rho}(t_{s-1}) + \Delta\overline{\rho}(t_s)$ and with it allows for updating the plane offset $d(t_s)$ as the optimal switching plane $\mathcal{T}^\star(x(t_s))$ must inevitable pass through the optimal switch point $\xi_1(t_s)$.

The key for obtaining a suitable switching plane orientation $n(t_s)$, on the other hand, is to transform the above optimization problem into an equivalent well-known discrete-time optimal control problem

$$\min_{\overline{\chi}(t)} \tilde{J}_{\overline{\theta}(t_s)}(x(t_s), \overline{\chi}(t_s)) = \sum_{k=0}^{N(t_s)} g_k(\xi_k(t_s), \chi_k(t_s)) \quad \text{subject to} \tag{16}$$

$$\xi_{k+1}(t_s) = h_k(\xi_k(t_s), \chi_k(t_s)), \quad \xi_0(t_s) = x(t_s) \tag{17}$$

$$\psi_N(\xi_{N(t_s)}(t_s)) \leq 0 \tag{18}$$

with discretized dynamics and costs

$$h_k(\xi_k(t_s), \chi_k(t_s)) = \xi_k + \int_0^{\chi_k(t_s)} f_{\theta_k(t_s)}(\xi(r|t_s))\mathrm{d}r \tag{19}$$

$$g_k(\xi_k(t_s), \chi_k(t_s)) = \int_0^{\chi_k(t_s)} L_{\theta_k(t_s)}(\xi(r|t_s))\mathrm{d}r + \phi_{\theta_k(t_s)}(\xi_k(t_s)) \tag{20}$$

$$g_{N(t_s)}(\xi_k(t_s), \chi_k(t_s)) = \phi_N(\xi_{N(t_s)}(t_s)). \tag{21}$$

This exact discretization comprises two central benefits: First, due to Dunn and Bertsekas (1989); Dohrmann and Robinett (1999) quadratic programming (QP) provides the means for obtaining the equivalent Newton direction $\Delta\overline{\chi}(t_s)$ that is dual to $\Delta\overline{\rho}(t_s)$, namely by solving a constrained discrete-time LQR problem. For brevity of notation, we refer to this problem as LQR$(x, \overline{\chi})$. Second, this standard optimal control problem yields a well-defined notion of a neighboring extremal solution (see Ghaemi et al. (2008)), which can be exploited for the determination of a suitable plane normal $n(t_s)$. For investigating both aspects, let us introduce the Hamiltonian

$$H_k(\xi_k(t), \chi_k(t), p_{k+1}(t)) = p_{k+1}^{\mathrm{T}}(t) h_k(\xi_k(t), \chi_k(t)) + g_k(\xi_k(t), \chi_k(t))$$

and the costate $p_k(t)$ at time t determined via backward iteration

$$p_k^{\mathrm{T}}(t) = p_{k+1}^{\mathrm{T}}(t) A_k(t) + a_k(t), \quad p_N(t) = \frac{\partial \psi}{\partial x}(\xi_{N(t)}(t)) + \nu(t)$$

with $\nu(t)$ being a Lagrange multiplier that enforces the terminal constraint (18). The invertible matrix $A_k(t)$ and $a_k(t)$ denote

$$A_k(t) = \frac{\partial h_k}{\partial \xi}(\xi_k(t), \chi_k(t)), \quad a_k(t) = \frac{\partial g_k}{\partial \xi}(\xi_k(t), \chi_k(t)) \ .$$

With the above, the augmented performance index

$$\tilde{J}'_{\overline{\theta}(t_s)}(x(t), \overline{\chi}(t)) = \phi_{N(t)}(\xi_{N(t)}(t)) - p_{N(t)}^{\mathrm{T}}(t)\xi_{N(t)} + H_0(\xi_0(t), \chi_0(t), p_1(t))$$
$$+ \sum_{k=1}^{N(t)-1} \left(H_k(\xi_k(t), \chi_k(t), p_{k+1}(t)) - p_k^{\mathrm{T}}(t)\xi_k(t) \right) \tag{22}$$

is obtained by adjoining the constraints (17), (18) to the original index (16). It transforms the constrained optimization problem into an equivalent unconstrained one. Let $\Delta x(t_{s+1}) = x(t_{s+1}) - x(t_s)$. Assuming the Hessian

$$H(t_s) = \frac{\partial^2 \tilde{J}'_{\overline{\theta}(t_s)}}{\partial \overline{\chi}^2}(x(t_s), \chi(t_s))$$

to be positive definite, the Newton direction $\Delta\overline{\chi}^{\star}(t_{s+1})$ at time t_{s+1} minimizes the second order Taylor expansion of (22) at time t_{s+1} around the previous values $x(t_s)$, $\overline{\chi}(t_s)$, and is given by

$$\Delta\overline{\chi}(t_{s+1}) = \overbrace{-H^{-1}(t_s)\frac{\partial \tilde{J}'_{\overline{\theta}(t_s)}}{\partial \overline{\chi}}(x(t_s), \overline{\chi}(t_s))}^{\Delta\overline{\chi}'(t_{s+1})} - H^{-1}(t_s)\frac{\partial^2 \tilde{J}'_{\overline{\theta}(t_s)}}{\partial x \partial \overline{\chi}}(x(t_s), \overline{\chi}(t_s))\Delta x(t_{s+1}) \ . \quad (23)$$

For details of this derivation, please see Schild et al. (2009b). It is important to note that, except for $\Delta x(t_{s+1})$, the expression (23) only involves information, which is already available at the previous sample time t_s. This fact is crucial concerning a real-time controller implementation, as it allows for a one-step-ahead preparation of information that needs a considerable amount of time for being computed. The approach of Wardi et al. (2007) does *not* possess this remarkable property. Moreover, the second term on the RHS of (23) represents a so called *perturbation feedback controller*, which according to Ghaemi et al. (2008) is of linear form

$$u_k^{\mathrm{T}} H^{-1}(t_s)\frac{\partial^2 \tilde{J}'_{\overline{\theta}(t_s)}}{\partial \xi \partial \overline{\chi}}(\xi(t_s), \overline{\chi}(t_s))\Delta x(t_{s+1}) = k_k^{\mathrm{T}}(t_s)\Delta\xi_k, \quad (24)$$

with $\Delta\xi_{k+1} = (A_k(t_s) - f_{\theta_k(t_s)}(\xi_{k+1}(t_s))k_k^{\mathrm{T}}(t_s))\Delta\xi_k$, $\Delta\xi_0 = \Delta x(t_{s+1})$ and u_k denoting the k-th unit vector. The state-feedback gains $k_k^{\mathrm{T}}(t_s)$ are obtained as a by-product of the QP solution to LQR$(x(t_s), \overline{\chi}(t_s))$ (see Dunn and Bertsekas (1989)). A fundamental result, which was shown in Schild et al. (2009a), states that parameter sequences

$$n_k^{\mathrm{T}}(t_{s+1}) = k_k^{\mathrm{T}}(t_s)A_k^{-1}(t_s), \quad d_k(t_{s+1}) = n_k^{\mathrm{T}}(t_s)\xi_{k+1}(t_s) \quad (25)$$

allow for evaluating the perturbation control law (24) at time t_{s+1} simply by means of an event-driven forward-in-time simulation of the plant dynamics with initial state $x(t_{s+1})$. Thus, compared to Wardi et al. (2007), the determination of switching plane parameters and the evaluation of the perturbation controller (24) only causes marginal additional effort, whereas the potential for improving the loop performance is significant.

With the above, let us summarize the overall real-time algorithm, which solves Problem 1 under reasonable assumptions.

Algorithm 1: Optimal adaptive feedback control of switched systems.

Initialize: Set $t = 0$, $\overline{\theta}(0) = \overline{q}$, specify tolerance $\varepsilon > 0$ and step size Δt. Determine a feasible initial sequence $\overline{\chi}(0)$ for given x_0.

Iterate: until mode q_{N-1} gets deactivated at $t = \tau_N > T$

1. At current time t_s, obtain the new measurement $x(t_s)$.
2. Use the results of the previous iteration to obtain new parameter sequences $\overline{n}(t_s)$, $\overline{d}(t_s)$. Substitute $n(t_s) = n_0(t_s)$ and $d(t_s) = d_0(t_s)$ into (3).

If $\chi_0(t) > \Delta t$ **then**

3.1 Execute an event-triggered forward integration to obtain the intermediate sequence $\overline{\chi}'(t_s)$. Update the controller state $x_c(t_s) = \overline{\chi}'(t_s) + \Delta\chi'(t_s)$.

3.2 Solve LQR$(x(t_s), x_c(t_s))$ to obtain the intermediate update $\Delta\overline{\chi}'(t_{s+1})$, which does not yet include the corrective effects of the perturbation control law (24).

3.3 While executing steps 3.1-2, check for collision of $x(\cdot)$ with $\mathcal{T}(t_s)$, i.e. for $d(t_s) - n^{\mathrm{T}}(t_s)x(t) < -\varepsilon$. **If** so, interrupt all computations, set $\Delta\chi'(t_{s+1}) = 0$ and jump to (4.2).

3.4 Otherwise wait until Δt time units have elapsed and return.

else

4.1 Wait until the collision of $x(\cdot)$ with $\mathcal{T}(t_s)$ is detected at time $\tau_{q(t_s)+1} \geq t_s$. Set $t_{s+1} = \tau_{q(t)+1}$ and $\Delta\chi'(t_{s+1}) = \underline{\Delta}\chi'(t_s)$.

4.2 Set $N(t_{s+1}) = N(t_s) - 1$, shrink mode sequence $\overline{\theta}(t_{s+1}) = \{\theta_k(t_s)\}_{k=1}^{N(t_s)}$, controller state $x_{c,(k-1)}(t_{s+1}) = x_{c,k}(t_s)$, $k = 1..N(t_s)$ and return.

Result: Control loop (1), (2), (4), (5) that generates closed loop trajectories $x(\cdot)$ that are optimal with respect to (8), provided that Δt is sufficiently small and that $\overline{\chi}(0)$ is sufficiently good.

Figure 3 illustrates the algorithm's working principle. After completing the computation of step 2 at time t_s, an updated switching plane $\mathcal{T}(t_s)$ is available, which passes through the predicted switching point $\xi_1(t_{s-1})$ determined during the previous iteration in step 3.2. According to step 3.1, we use this plane (and the plane parameters of all remaining modes $\overline{\theta}(t_s)$) in an event-driven simulation to obtain the intermediate switching point $\xi_1'(t_s)$ and the associated activation duration sequence $\overline{\chi}'(t_s)$. The latter is shifted by $\Delta\chi'(t_s)$ and then stored as the current controller state. Step 3.2 involves the most expensive operations as it solves for the prediction $\Delta\chi'(t_{s+1})$ as well as for all information, which is needed to obtain $\mathcal{T}(t_{s+1})$ at the beginning of the next iteration. In particular, this requires the evaluation of second order partial derivatives of (19)-(20).

Note that the predicted activation duration sequence $\overline{\chi}(t_s)$ constitutes the internal state of the switching plane controller, which is never used to directly control the plant. Thus, step 2, the update of $n(t)$ and $d(t)$, and step 3.3, the detection of the next switching event, can be executed in parallel. This is vital for a real-time implementation, as the control loop is never broken.

If a mode transition is about to occur immediately, steps 3.x must all be skipped, since there is not enough time to complete the associated computations. Instead, the switching plane $\mathcal{T}(t_s)$ is used to detect the exact transition time $\tau_{\overline{q}+1}$. This skipping results in prolonging the sampling period to at most $2\Delta t$.

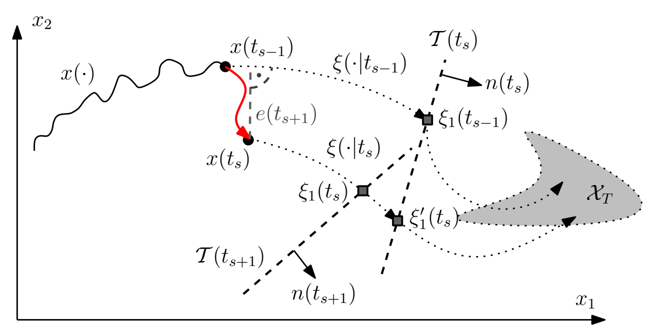

Fig. 3. On-line update of the switching plane.

Remark 1. Unlike variable structure control, we intend to avoid sliding along a switching surface. Whether a chattering behavior occurs or not, depends on the chosen cost $J_{\bar{q}}$.

3.3 Properties of the proposed control concept

An important aspect of the adaptive control scheme concerns the convergence of plane parameter trajectories $n(\cdot)$ and $d(\cdot)$ to locally optimal trajectories $n^\star(\cdot)$ and $d^\star(\cdot)$. As a preliminary for this, we can bound the rate with which the controller state $x_c(t)$ converges towards a locally optimal trajectory $x_c^\star(\cdot)$.

Theorem 1. Let no mode transition occur between t_s and t_{s+1}, and suppose that the gradient $\frac{\partial \bar{J}_{\bar{\theta}(t_s)}}{\partial \chi}(x, \bar{\chi})$ is differentiable at points $(x(t_s),\, x_c(t_s))$ and $(\xi(\Delta t|t_s),\, x_c(t_{s+1}))$. If the Hessian matrix $H(t_s)$ is positive definite, then there exists a bounded constant $K > 0$, such that

$$||x_c(t_{s+1}) - x_c^\star(t_{s+1})|| \le K \left(||x_c(t_s) - x_c^\star(t_s)|| + ||e(t_{s+1})||\right)^2 \quad (26)$$

with $e(t_{s+1}) = \mathrm{dist}(x(t_{s+1}), \xi(\cdot|t_s))$.

The proof of Theorem 1 follows arguments, which are similar to the convergence proof presented in Wardi et al. (2007).

Theo. 1 states that for a sufficiently high update rate $\Delta t \ll 1$, a small one-step deviation $||e(t)|| \ll 1,\ \forall t$, a good initial guess $\bar{\chi}(0)$ and a sufficiently large first activation duration $\delta_0^\star$, practical convergence of the controller state $x_c(t)$ to a neighborhood of a locally optimal trajectory $x_c^\star(\cdot) = \bar{\chi}^\star(\cdot)$ is ensured before $t = \tau_1^\star$. Due to non-convexity of the optimization problem, the attracting trajectory $x_c^\star(\cdot)$ depends on the initial sequence $\bar{\chi}(0)$.

Corollary 1. In a disturbance-free environment ($w(\cdot) = 0 \Rightarrow e(\cdot) = 0$) and for negligible Δt, the controller state $x_c(t)$ converges quadratically towards a locally optimal trajectory $x_c^\star(\cdot)$.

Consequently, if Δt is sufficiently small, $||\bar{\chi}(0) - \bar{\chi}^\star(0)|| \ll 1$ and $\delta_0^\star$ sufficiently large, the optimal adaptive control scheme solves Prob. 1. Moreover, the previous result implies convergence of $\mathcal{T}(t)$ towards an optimal configuration.

Theorem 2. Convergence of $x_c(t)$ to an optimal trajectory $x_c^\star(\cdot)$ implies convergence of the switching plane $\mathcal{T}(t)$ to a optimal configuration $\mathcal{T}^\star(x(t_s))$, i.e. the differences $||n(t_s) - n^\star(t_s)||$ and $||d(t_s) - d^\star(t_s)||$ vanish with increasing s.

However, in Sect. 2 we claimed to accomplish more than specified in Prob. 1, namely that our control scheme ensures an optimal behavior even in the presence of reasonable disturbances. To verify the latter, we first observe that utilizing the almost-for-free switching plane information in step 3.1 of Algorithm 1 enforces a quadratic dependence of (26) on $e(t)$, whereas the approach of Wardi et al. (2007) one exhibits a worse linear dependence. Consequently, if $||e(t)|| \ll 1$ it takes less update steps for $x_c(t)$ to settle to an optimal trajectory and the disturbance effects are less severe. Moreover, a practically converged switching plane $\mathcal{T}(t)$ guarantees an optimal compensation of small amplitude perturbations from the nominal path by the inner loop alone.

Theorem 3. If at time t_s, the controller state $x_c(t_s) = x_c^\star(t_s)$ is locally optimal with respect to (16), then the corresponding plane $\mathcal{T}(t_s) = \mathcal{T}^\star(t_s)$ gives rise to neighboring extremal solutions with respect to the nominal path $\xi^\star(\cdot|t_s)$.

Indeed, the inner loop is able to conserve local optimality of the switching times up to first order, even if disturbances affect the plant in between two consecutive update times t_s and t_{s+1}. This is not the case for the control approach of Wardi et al. (2007), as the control loop is broken in between two updates.

4. CRAWLING WINDOW ADAPTIVE FEEDBACK CONTROL OF PERIODICALLY SWITCHED SYSTEMS

The infinite length of $\bar{q}$ in the infinite-horizon setting clearly prohibits the application of Algorithm 1. However, an approximate suboptimal solution to Prob. 2 can be determined by combining this adaptive control scheme with ideas from classical model predictive control. In this extension, we essentially propose to iteratively solve the cost-to-go problem (16)-(18) at each sampling time t_s over a finite prediction window that is shifted along the time axis as the operation progresses. In contrast to classical model predictive control, the length of the prediction horizon is characterized by a finite number of $N(t_s)$ switchings, not by a constant duration T. the stabilization of a limit cycle Γ and the periodicity of the mode sequence $\bar{q}$ require a variation of $N(t_s)$ during the controller operation. Metaphorically speaking, the prediction window executes a crawling motion, which provides the source for the name of the approach: *crawling window adaptive feedback control*.

To explain the latter, let us subdivide the infinite horizon into distinct *stages* and associate each stage with the execution of a complete mode cycle $\bar{q}^\Gamma$. A transition from the $(c-1)$-st to the c-th stage occurs at the switching time τ_{cp}, when q_0 gets reactivated for the $(c+1)$-st time. At the beginning of the c-th stage the prediction window length is reset to $N(\tau_{cp}) = lp$, so that the initial c-th *stage mode sequence* $\bar{\theta}^c(\tau_{cp}) = (\bar{q}^\Gamma \ldots \bar{q}^\Gamma)$ results from concatenating the subsequence $\bar{q}^\Gamma$ of the cycle Γ exactly l times. During the execution of each stage Algorithm 1 is applied. In particular, this implies iteratively shrinking the mode sequence $\bar{\theta}^c(\tau_k^+) = \{\theta_j^c(\tau_k)\}_{j=1}^{n(\tau_k)-1}$ as well as the predicted activation duration sequence $\bar{\chi}^c(\tau_k^+) = \{\chi_j^c(\tau_k)\}_{j=1}^{n(\tau_k)-1}$ at inter-cycle switching times τ_k, $k \neq lp$. At the transition to the subsequent $(c+1)$-st stage, however, the predicted mode sequence

$$\bar{\theta}^{c+1}(\tau_{(c+1)p}) = (\bar{\theta}^c(\tau_{(c+1)p})\, \bar{q}^\Gamma) \quad (27)$$

is appended by a complete cycle $\bar{q}^\Gamma$. Likewise, the predicted activation duration sequence

$$\bar{\chi}^{c+1}(\tau_{(c+1)p}) = \bar{\chi}^c(\tau_{(c+1)p})\, \check{\chi}(\tau_{(c+1)p}) \quad (28)$$

must be reinitialized with an appropriate appendix $\check{\chi}(\tau_{(c+1)p})$. The overall control scheme can be summarized as follows:

Algorithm 2: Crawling window optimal adaptive feedback control of switched autonomous systems.

Initialize: Set $\bar{\theta}^0(0) = (\bar{q}^\Gamma \ldots \bar{q}^\Gamma)$, $c = 0$, specify a tolerance $\varepsilon > 0$ and a sampling rate Δt. Determine a feasible initial sequence $x_c^0(0) = \bar{\chi}^0(0)$ for a given x_0.

Iterate:

1.1 At the beginning of the c-th stage, execute Algorithm 1 with initial values $x(\tau_{cp})$, $\theta^c(\tau_{cp})$, $x_c^c(\tau_{cp})$, ε and $N(\tau_{cp}) = lp$ to accomplish a real-time update of the switching plane $\mathcal{T}(t)$ until the completion of *one* cycle q^Γ.

1.2 Increment counter $c = c+1$ and append sequences $\bar{\theta}^c(\tau_{(c+1)p})$ and $x_c^c(\tau_{(c+1)p})$ according to (27) and (28).

Result: Control loop (1), (2), (4), (5) that orbitally stabilizes Γ and exhibits a suboptimal performance with respect to (12).

As known from classical NMPC, it is crucial to reinitialize the state $x_c^c(\tau_{(c+1)p})$ at the transition to the $(c+1)$-st stage by an admissible appendix $\check{\chi}(\tau_{(c+1)p})$ that preserves feasibility of the problem (16)-(18) for the next stage. Also, it is known that optimality of the sequence $\overline{\chi}^{c\star}(\tau_{(c+1)p})$ is generally not preserved at any stage transition. This fact highlights the importance of the quadratic convergence property (Theorem 1) of Algorithm 1.

Assuming that $x_c^c(t)$ converges to a locally optimal solution $x_c^{c\star}(\cdot)$ before the termination of the c-th stage at $\tau^\star_{(c+1)p}$, orbital stability can be shown by extending the quasi-infinite horizon concept of NMPC (see Findeisen et al. (2003); Mayne et al. (2000)). This extension, yields the following:

Theorem 4. Suppose that

1.1 the plant (2) is orbitally stabilizable at Γ,
1.2 the closed terminal region $\mathcal{X}_T$ contains the switch point $x_0^\Gamma \in \Gamma$ associated to the activation of q_0,
1.3 the terminal cost term $\phi_N(x)$ is positive semi-definite,
1.4 the region $\mathcal{X}_T$ and $\phi_N(x)$ are chosen such that for all $x_0 \in \mathcal{X}_T$, there exists an (admissible) switching time sequence $\check{\chi}$ of length p, which when applied to (17) produces a sequence ξ for which $\xi_p \in \mathcal{X}_T$ and

$$\phi_N\left(\xi_p\right) - \phi_N(x_0) + \sum_{k=0}^{p-1} g_k(\xi_k, \check{\chi}_k) < 0 \ \text{ and}$$

1.5 the finite horizon Problem (16)-(18) is initially feasible for $t = 0$.

In this case, the crawling window adaptive optimal control strategy of Algorithm 2 ensures asymptotic orbital stability of the orbit Γ as $t \to \infty$ and the region of attraction $\mathcal{X}_0$ consists of all states, for which a feasible mode signal $q(t)$ exists.

Theo. 4 only ensures stability, but not optimality of the resulting evolution $x(\cdot)$. To verify that Prob. 2 is solved by application of the crawling window control approach, we need to show the convergence of the discounted MPC-costs towards the desired infinite horizon costs J_∞ as the horizon length l goes to infinity. Let us introduce the infinite horizon value function $V_\infty(x_0, \Gamma) = \min_{\overline{\tau}} J_\infty(x_0, \overline{\tau}, \Gamma)$, which returns the minimum infinite-horizon performance for a trajectory starting at the initial state x_0. Moreover, let the MPC-like feedback law

$$\chi_k^\star = \kappa_{l,k}^\star(x) = \arg\min_{\chi_k}\left(V_{lp-k}(h_k(x, \chi_k)) + g_k(x, \chi_k)\right) \qquad (29)$$

return the predicted optimal activation duration $\chi_k^\star$ for the k-th mode q_k from the mode sequence $\overline{q}^\Gamma$ of the limit cycle Γ. With this law, let us introduce the MPC value function $V_{\text{MPC}}(x_0, l) = \sum_{c=0}^{\infty}\left(\sum_{k=0}^{p-1} g_k(x_{(cp)+k}, \kappa_{l,k}^\star(x_{(cp)+k}))\right)$, which returns the discounted infinite horizon MPC-costs for a trajectory starting at the initial state x_0 under horizon length l.

Theorem 5. The discounted MPC-costs $V_{\text{MPC}}(x_0, l)$ asymptotically converge to the desired infinite horizon performance $V(x_0)$ as l approaches infinity.

By Theo. 5, the application of crawling window control results in a suboptimal solution of Prob. 2, where complexity can be traded against performance. However, the value $V_{\text{MPC}}(x_0, l)$ may not be monotonic in l.

5. CONCLUSION

The central contribution of this paper is a novel optimal adaptive control scheme for switched nonlinear systems that achieves a satisfactory loop performance in the presence of disturbances. The core of the proposed strategy is a cascaded control loop with an inner loop that implements an explicit switching law in the form of a switching plane and an outer loop, which iteratively adjusts this plane in such a way that it converges to a local approximation of the corresponding optimal switching surface under reasonable assumptions. The perturbation feedback feature, which provides a major distinction from the approach pursued in Wardi et al. (2007), makes the overall control scheme well-suited for switched systems that operate in a high frequency noisy environment.

REFERENCES

Boccadoro, M., Wardi, Y., Egerstedt, M., and Verriest, E. (2005). Optimal control of switching surfaces in hybrid systems. *Discrete Event Dynamic Systems*, 15(4), 433–448.

Dohrmann, C.R. and Robinett, R.D. (1999). Dynamic programming method for constrained discrete-time optimal control. *Journal of Optimization Theory and Applications*, 101(2), 259–283.

Dunn, J. and Bertsekas, D. (1989). Efficient dynamic programming implementations of newton's method for unconstrained optimal control problems. *Journal of Optimization Theory and Applications*, 63(1), 23–38.

Egerstedt, M., Wardi, Y., and Axelsson, H. (2006). Transition-time optimization for switched-mode dynamical systems. *IEEE Transaction on Automatic Control*, 51, 110–115.

Findeisen, R., Imsland, L., Allgöwer, F., and Foss, B.A. (2003). State and output feedback nonlinear model predictive control: An overview. *European Journal of Control*, 9, 190–206.

Ghaemi, R., Sun, J., and Kolmanovsky, I. (2008). Neighboring extremal solution for discrete-time optimal control problems with state inequality constraints. In *Proc. American Control Conference*, 3823–3828. doi:10.1109/ACC.2008.4587089.

Mayne, D., Rawlings, J., Rao, C., and Scokaert, P. (2000). Constrained model predictive control: Stability and optimality. *Automatica*, 36, 789–814.

Schild, A., Ding, X., Egerstedt, M., and Lunze, J. (2009a). Design of optimal switching surfaces for switched autonomous systems. In *Proc. of 42nd IEEE CDC (submitted)*.

Schild, A., Ding, X., Egerstedt, M., and Lunze, J. (2009b). Real-time adaptive optimal feedback control of switched autonomous systems. *IEEE Transaction on Automatic Control*, (submitted).

Shaikh, M.S. and Caines, P.E. (2007). On the hybrid optimal control problem: Theory and algorithms. 52(9), 1587–1603. doi:10.1109/TAC.2007.904451.

Wardi, Y., Ding, X., Egerstedt, M., and Azuma, S.I. (2007). On-line optimization of switched-mode systems: Algorithms and convergence properties. In *Proc. on 46th IEEE Conference on Decision and Control*.

Xu, X. and Antsaklis, P. (2002). Optimal control of switched systems via non-linear optimization based on direct differentiations of value functions. *International Journal of Control*, 75, 1406–1426.