On a Class of Maximal Invariance Inducing Control Strategies for Large Collections of Switched Systems

Petter Nilsson

Dept. of Electrical Engineering and Computer Science
University of Michigan
Ann Arbor, MI
pettni@umich.edu

Necmiye Ozay

Dept. of Electrical Engineering and Computer Science
University of Michigan
Ann Arbor, MI
necmiye@umich.edu

ABSTRACT

Modern control synthesis methods that are capable of delivering safety guarantees typically rely on finding invariant sets. Computing and/or representing such sets becomes intractable for high-dimensional systems and often constitutes the main bottleneck of computational procedures. In this paper we instead analytically study a particular high-dimensional system and propose a control strategy that we prove renders a set invariant whenever it is possible to do so. The control problem—the mode-counting problem with two modes in one dimension—is inspired by scheduling of thermostatically controlled loads (TCLs) and exhibits a trade-off between local safety constraints and a global counting constraint. We improve upon a control strategy from the literature to handle heterogeneity and derive sufficient conditions for the strategy to solve the problem at hand. In addition, we show that the conditions are also necessary for the problem to have a solution, which implies a type of optimality of the proposed control strategy. We outline more general problem instances where the same control strategy can be implemented and we give sufficient (but not necessary) conditions for the closed-loop system to satisfy its specification. We illustrate our results on a TCL scheduling example.

Keywords

Control of switched systems; Energy applications

1. INTRODUCTION

Formal safety verification of control systems and control synthesis techniques that automatically generate controllers with safety guarantees provide principled alternatives to testing and simulation before system deployment. Most of these techniques rely on computation of invariant sets or controlled invariant sets in order to show that the system trajectories do not leave the part of the state-space that is deemed safe [3]. Computing and/or representing invariant sets often becomes infeasible as the geometry of such sets can be fairly

complex for high dimensional hybrid systems [16]. To partly alleviate this challenge, Hafner and Del Vecchio show that, for monotone systems, it is easier to devise an algorithm that checks whether a given state is inside, at the boundary, or outside of the maximal invariant set [6]. Instead of computing the invariant set explicitly, such an algorithm is then used within a supervisory controller in order to guarantee safety. A related recent result for cooperative systems [14] shows that if the goal is to find a (not necessarily maximal) controlled invariant set contained inside a rectangular set, there is almost no loss of generality in restricting attention to periodic input sequences, which again enables focusing on the controller instead of sets.

In this paper we seek another alternative: instead of computing the maximal invariant set, can we propose a controller and show that it enforces invariance of the maximal controlled invariant set, whenever that set is nonempty? This question is non-trivial to answer for arbitrary systems; therefore, we focus on a particular class of high dimensional switched systems, motivated by the thermostatically controlled load (TCL) coordination problem.

TCLs include air conditioners, water heaters, refrigerators, etc., that operate within a certain temperature range, called *dead band*, around a temperature set point. TCL owners are typically indifferent to small temperature perturbations around their desired set point. The idea behind TCL coordination is that an electric utility company can leverage this flexibility—which becomes meaningful for large collections of TCLs—to shape aggregate demand on the grid.

The efforts to model aggregate TCL systems can be traced back to 1985, when an aggregate model based on the Fokker-Planck partial differential equation (PDE) was proposed [10]. The paper borrowed ideas from statistical physics: just as a large collection of particles can be described by a density, a large number of TCLs can be approximately represented by a “TCL density” function whose evolution is governed by a PDE. More recent work along the same lines have focused on models which are more amenable to control so that the flexibility in TCL operating conditions can be used to help improve power grid operations. In [5] a PDE-model with exposed parameters for the temperature dead band is derived, thus enabling analysis of control of these parameters. Another area of work has been focused on developing finite-state models. These are obtained by partitioning the TCL temperature range into “bins”, and using discrete-time Markov chains to describe the evolution of a probability distribution over the resulting discrete state space [8, 15, 4]. In

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

HSCC '17, April 18–20, 2017, Pittsburgh, PA, USA.

© 2017 ACM. ISBN 978-1-4503-4590-3/17/04...\$15.00

DOI: <http://dx.doi.org/10.1145/3049797.3049810>

addition, estimation techniques for these models have been proposed to lower the amount of information required at a central control node [11].

As illustrated in [4, 5], if a universal set point temperature is imposed across a collection of TCLs, it can be treated as a control input that can be shaped in order to track reference demand. However, the tracking capability comes at the cost of modifying the set point temperature, which may result in a conflict with end user preferences. In [15, 7] a different approach is followed: the on/off-state of each individual TCL is treated as a potential control input. This additional control freedom can leverage more of the flexibility within a temperature range and can therefore potentially achieve reference tracking without violating end user constraints.

This approach is generalized in [13] into the so-called *mode-counting problem*. The fundamental challenge in such a problem is to satisfy local state constraints while simultaneously controlling the aggregate number of subsystems that are in a given mode. Such a constraint is important in TCL coordination: turning too many TCLs on or off at the same time may overload the electricity grid. The cited paper proposes a controller synthesis technique based on an approximately bisimilar abstraction and solves a discrete mode-counting problem on the abstraction via a linear program. In theory, the method can solve the mode-counting problem to an arbitrary precision, but it assumes knowledge of initial conditions.

In this paper we first focus on a special case of the mode-counting problem that is inspired by the TCL scheduling problem, where each individual subsystem has two modes and one-dimensional continuous dynamics. We show that a generalization of a control policy proposed in [7] is in a certain sense optimal for a heterogeneous family of TCLs, and derive analytic feasibility limits. The analysis shows that the control policy enforces invariance of the maximal controlled invariant set, without the need to explicitly deal with or express the controlled invariant set itself.

The problem at hand is related to schedulability theory. Previous work has identified the similar feasibility limits in the linear setting [12], but without showing that the control policy in [7] is in fact optimal. Another related work is [2] which considers schedulability of hybrid constant-rate systems, subsequently generalized to bounded-rate systems [1]. These papers present feasibility results in the form of linear programs that scale with the number of aggregate dynamical modes—in our case 2^N . By contrast, we provide closed-form formulas and an easy-to-implement control policy that achieves invariance whenever it is theoretically feasible.

This paper is structured as follows: the remainder of this section introduces relevant notation and formalizes the problem statement. In the following Section 2 we introduce time to exit—a key concept that is a novel addition to the policy in [7] and provides an abstraction that makes heterogeneity invisible to the central coordinator. Subsequently, in Section 3, we propose a control strategy to solve the problem and present our main results regarding its performance. Sufficient conditions for the strategy that apply in more general settings are presented in Section 4. The results are illustrated with simulations in Section 5 before the paper is concluded in Section 6. The proofs of the main results have been deferred to the appendix.

1.1 Preliminaries

We first introduce some notation. A set of integers from 1 to N is written as $[N] = \{1, 2, \dots, N\}$. The indicator function of a set A is denoted $\mathbf{1}_A(x)$ and is equal to 1 if $x \in A$ and to 0 otherwise. We use the Nabla operator $\nabla_x f$ to denote the Jacobian of f with respect to x .

We will work with the usual definition of the flow $\phi_f(x, t)$ of a vector field $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$. It is the solution at time t of the following differential equation:

$$\begin{cases} \frac{d}{ds}y(s) = f(y(s)), \\ y(0) = x. \end{cases}$$

Clearly, the flow operator satisfies the relation

$$\frac{d}{dt}\phi_f(x, t) = f(\phi_f(x, t)),$$

a fact that will be used extensively later in the paper.

1.2 Problem statement

In most of the literature on TCL modeling and control a one-dimensional linear switched ODE is used to model an individual TCL. A TCL has the two modes **on** and **off**, each inducing a globally stable equilibrium point. The temperature range between the two equilibria constitute the attainable temperatures. Typically, a subrange of this temperature range (the dead band) forms acceptable end user states, i.e., local safety constraints. Since the dead band is necessarily located between the equilibria, it follows that the flows corresponding to the two modes are sign-definite and have opposing signs. The following generalized problem formulation captures these characteristics.

PROBLEM 1. *Given N subsystems with states $\{x^i\}_{i \in [N]}$ s.t. $x^i \in \mathbb{R}$, obeying the dynamics*

$$\frac{d}{dt}x^i(t) = \begin{cases} f_{\text{off}}^i(x^i(t)) & \text{if } \sigma^i(t) = \text{off}, \\ f_{\text{on}}^i(x^i(t)) & \text{if } \sigma^i(t) = \text{on}, \end{cases} \quad (1)$$

local safe sets $S^i = [\underline{a}^i, \bar{a}^i]$, and global mode-counting bounds $[\underline{K}_{\text{on}}, \bar{K}_{\text{on}}]$, construct an aggregate switching policy $\{\sigma^i\}_{i \in [N]}$ such that for all trajectories $x^i : \mathbb{R}^+ \rightarrow \mathbb{R}$:

$$x^i(t) \in S^i, \quad \forall i \in [N], t \in \mathbb{R}^+, \quad (2)$$

$$\underline{K}_{\text{on}} \leq \sum_{i=1}^N \mathbf{1}_{\{\text{on}\}}(\sigma^i(t)) \leq \bar{K}_{\text{on}}, \quad \forall t \in \mathbb{R}^+. \quad (3)$$

In accordance with the TCL setting, we assume that f_{off}^i is strictly positive on S and that f_{on}^i is strictly negative. With these assumptions, the **on** mode will always transport the state “downwards”, while the **off** mode transports “upwards”, as illustrated in Figure 1. In order for the problem to have a solution the initial conditions of all x^i ’s must evidently be inside their safe sets S^i .

Eq. (2) represents a local safety constraint while (3) is a global constraint that requires coordination between the different subsystems. Both constraints are trivial to satisfy by themselves, but when they have to be taken into account simultaneously a conflict arises. In this paper we propose a control strategy that resolves this conflict and prove that under certain assumptions—that apply for TCL dynamics—it is the best possible strategy.

Figure 1: Illustration of the TCL scheduling problem with homogeneous local safety constraints; $\underline{a}^i = \underline{a}$ and $\bar{a}^i = \bar{a}$ for all $i \in [N]$. Each subsystem is illustrated with a colored circle; those that are in mode on are in blue and those that are in mode off are in red. The local safety constraints stipulate that each subsystem must remain in the interval $[\underline{a}, \bar{a}]$, while the global mode-counting constraint restricts the number of subsystems that can be in mode on at a given time. When a subsystem switches mode the direction of movement is reversed due to the assumption on signs of f_{on} and f_{off} .

The aggregate dynamics can be seen as an N -dimensional switched system with 2^N aggregate modes. For N 's in the tens or hundreds of thousands, conventional computer-assisted analysis methods that do not exploit the symmetries of the problem become intractable.

2. TIME TO EXIT

In order to treat a heterogeneous family of subsystems in a cohesive manner we introduce *time to exit* as a way to assess the urgency to switch a given subsystem. As a one-dimensional quantity that abstracts away heterogeneity, the concept is well suited for general problems with constraints on mode counts. We therefore present formulas for arbitrary n -dimensional systems, although the main focus in this paper is the special setting with one-dimensional subsystems. The time to exit with respect to a set $S \subset \mathbb{R}^n$ is the minimal time it takes for the flow to reach the boundary of S .

DEFINITION 1. Given a set $S \subset \mathbb{R}^n$, for $x \in S$ and a vector field $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$, the **time to exit** $T_f(x)$ is the time it takes for the flow of f starting in x to reach ∂S :

$$T_f(x) = \inf \{ \tau \geq 0 : \phi_f(x, \tau) \in \partial S \}.$$

If the set S is not transient, $T_f(x)$ might be equal to $+\infty$. In the following we assume that the time to exit is finite, which is the case in Problem 1. We next define an operator that maps a given point x to the point on ∂S where the flow of f starting in x exits S .

DEFINITION 2. For $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ and $x \in S$ with $T_f(x) < +\infty$, the **exit point** $U_f(x)$ is

$$U_f(x) = \phi_f(x, T_f(x)).$$

We now analyze the dynamic evolution of the time to exit T_f . It is clear that T_f decreases with unit speed along a trajectory of f itself. Using the Lie derivative, this fact is expressed as $\mathcal{L}_f T_f(x) = -1$. However, we are also interested in how the time to exit with respect to f , T_f , varies

along trajectories of a second vector field g . The infinitesimal change is given in terms of the Lie derivative $\mathcal{L}_g T_f$.

PROPOSITION 1. Assume that ∂S is C^1 at $U_f(x)$. Then the Lie derivative of $T_f(x)$ with respect to g is expressed as follows:

$$\mathcal{L}_g T_f(x) = - \frac{\left(\hat{n}_{U_f(x)}^S \right)^T (\nabla_x \phi_f)|_{(x, T_f(x))} g(x)}{\left(\hat{n}_{U_f(x)}^S \right)^T f(U_f(x))}, \quad (4)$$

where $\hat{n}_x^S$ is the outward-pointing unit normal of ∂S at x .

PROOF. By definition of the Lie derivative:

$$\begin{aligned} \mathcal{L}_g T_f(x) &= \frac{d}{ds} \Big|_{s=0} T_f(\phi_g(x, s)) \\ &= \frac{d}{ds} \Big|_{s=0} \inf \{ \tau : \phi_f(x + g(x)s + \mathcal{O}(s^2), \tau) \in \partial S \} \\ &= \frac{d}{ds} \Big|_{s=0} \inf \left\{ \tau : \begin{aligned} &\phi_f(x, \tau) + (\nabla_x \phi_f)|_{(x, \tau)} g(x)s \\ &+ \mathcal{O}(s^2) \in \partial S \end{aligned} \right\} \\ &= \frac{d}{ds} \Big|_{s=0} T_f(x) + \inf \left\{ \tau : \begin{aligned} &\phi_f(x, T_f(x) + \tau) \\ &+ (\nabla_x \phi_f)|_{(x, T_f(x) + \tau)} g(x)s \\ &+ \mathcal{O}(s^2) \in \partial S \end{aligned} \right\}. \end{aligned}$$

For a small τ , $\phi_f(x, T_f(x) + \tau) = U_f(x) + f(U_f(x))\tau + \mathcal{O}(\tau^2)$. Since $U_f(x) \in \partial S$, it follows that the τ that achieves the infimum must counteract the effect of the s -term along the normal of ∂S , as displayed in Figure 2. Ignoring second-order terms, this is expressed by the relation

$$0 = \left\langle \hat{n}_{U_f(x)}^S, f(U_f(x))\tau + (\nabla_x \phi_f)|_{(x, T_f(x) + \tau)} g(x)s \right\rangle,$$

which, given that $\left\langle \hat{n}_{U_f(x)}^S, f(U_f(x)) \right\rangle \neq 0$, by the implicit function theorem describes an implicit mapping $s \mapsto \tau$ defined around the origin, such that $0 \mapsto 0$. By writing

$$(\nabla_x \phi_f)|_{(x, T_f(x) + \tau)} = (\nabla_x \phi_f)|_{(x, T_f(x))} + \mathcal{O}(\tau);$$

differentiating the implicit relation with respect to s ; and letting $s \rightarrow 0$; the result is obtained. $\square$

The identity in Proposition 1 depends in an intricate way on the geometry of S and the vector field flow lines. We remark that the derivative of the flow operator with respect to the initial condition x can be written as the solution $\eta(t)$ of a matrix-valued ODE obtained by linearizing f along the flow [17, p. 154]:

$$\eta(0) = I, \quad \frac{d}{ds} \eta(s) = (\nabla_x f)|_{\phi_f(x, s)} \eta(s). \quad (5)$$

3. A SWITCHING STRATEGY FOR PROBLEM 1

We now apply the theory developed in the previous section to the study of Problem 1—the mode-counting problem for a collection of one-dimensional two-mode switched systems. In the interest of keeping notation simple, we omit the f 's in subscripts and write ϕ_{on}^i instead of $\phi_{f_{\text{on}}^i}$, and similarly for T , U , and $\mathcal{L}$.

First we adapt Proposition 1 for the simplified geometry of Problem 1; one-dimensional sign-definite vector fields allow us to evaluate the exit point mapping without solving an ODE.

Figure 2: Illustration of the effect an infinitesimal movement along the vector field $g(x)$ has on the exit time $T_f(x)$ with respect to f . A movement $g(x)s$ at x is propagated along the flow line of f and results in the infinitesimal movement $\nabla_x \phi_f(x, T_f(x))g(x)s$ at $U_f(x)$. To compensate for this movement, τ must be chosen so that $f(U_f(x))\tau$ counteracts the effect in the normal direction.

PROPOSITION 2. For the dynamics (1) and the safe set $S^i = [\underline{a}^i, \bar{a}^i]$, the expressions in Proposition 1 simplify to:

$$\mathcal{L}_{off}^i T_{on}^i(x) = -\frac{f_{off}^i(x)}{f_{on}^i(x)}, \quad \mathcal{L}_{on}^i T_{off}^i(x) = -\frac{f_{on}^i(x)}{f_{off}^i(x)}. \quad (6)$$

PROOF. In the interest of brevity we do not solve (5) but opt for a simpler argument. For a one-dimensional vector field f we can for a small h write

$$\begin{aligned} \phi_f(x+h, t) &= \phi_f(x, t+h/f(x) + \mathcal{O}(h^2)) \\ &= \phi_f(x, t) + h/f(x) + \mathcal{O}(h^2) \\ &= \phi_f(x, t) + hf(\phi_f(x, t))/f(x) + \mathcal{O}(h^2). \end{aligned}$$

Thus we get,

$$\nabla_x \phi_f(x, t) = \lim_{h \rightarrow 0} \frac{h \frac{f(\phi_f(x, t))}{f(x)} + \mathcal{O}(h^2)}{h} = \frac{f(\phi_f(x, t))}{f(x)}.$$

Turning to (4), the simple geometry yields $U_{on}^i(x) = \underline{a}^i$ and $\hat{n}_{U_{on}^i(x)} = -1$. Thus,

$$f_{on}^i(\phi_{on}^i(x, T_{on}^i(x))) = f_{on}^i(U_{on}^i(x)) = f_{on}^i(\underline{a}^i),$$

which gives

$$\mathcal{L}_{off}^i T_{on}^i(x) = -\frac{f_{on}^i(\underline{a}^i)f_{off}^i(x)}{f_{on}^i(x)} / f_{on}^i(\underline{a}^i) = -\frac{f_{off}^i(x)}{f_{on}^i(x)}.$$

The expression for $\mathcal{L}_{on}^i T_{off}^i(x)$ is derived in an analogous fashion. $\square$

Now, motivated by the control strategy proposed in [7], we propose a strategy for Problem 1 that only switches a subsystem if either the local state constraint or the global mode-counting bounds are about to be violated. In the referenced paper, the strategy was found to have good tracking performance¹ and robustness in a mildly heterogeneous setting, but no formal analysis was carried out. A key novelty in our work is that we use time to exit—rather than state value—to select which subsystem to switch to preserve mode-counting bounds. This modification allows us to work with truly heterogeneous collections of subsystems.

¹Using the terminology of this paper, tracking translates to time-varying mode-counting bounds.

STRATEGY 1. Switch subsystem i at a time instant t if one of the following conditions occur:

1. If $T_{off}^i(x^i(t)) = 0$, switch subsystem i to *on*,
2. If $T_{on}^i(x^i(t)) = 0$, switch subsystem i to *off*,
3. If $\sum_{j \in [N]} \mathbf{1}_{\{on\}}(\sigma^j(t^+)) < \underline{K}_{on}$ for $t^+ > t$, select the subsystem i in mode *off* with the largest time to on-exit, i.e.

$$i = \arg \max_{j \in [N] : \sigma^j(t) = off} T_{on}^j(x^j(t)),$$

and switch it to *on*. If the bound is still violated at t^+ , repeat step 3.

4. If $\sum_{j \in [N]} \mathbf{1}_{\{on\}}(\sigma^j(t^+)) > \overline{K}_{on}$ for $t^+ > t$, select the subsystem i in mode *on* with the largest time to off-exit, i.e.

$$i = \arg \max_{j \in [N] : \sigma^j(t) = on} T_{off}^j(x^j(t)),$$

and switch it to *off*. If the bound is still violated at t^+ , repeat step 4.

REMARK 1. In the event that more than $\overline{K}_{on}$ subsystems simultaneously satisfy the condition in 1., or analogously for $\underline{K}_{on}$ and condition 2., the strategy above is not well defined with respect to the dynamics. This is however a measure-zero event in the space of initial conditions. Let's call aggregate states where several subsystems have identical and non-zero times to exit **degenerate**. Since a mode perturbation (i.e., switching off some systems in the *on* mode while simultaneously switching on some systems in the *off* mode) can be applied whenever the system is in a degenerate state, we can without loss of generality disregard degenerate states provided that the initial condition itself is not degenerate.

Alternatively, degenerate cases can be handled by adding a time margin $\tau > 0$ to the switching condition: if the times to exit of several subsystems cross into the margin region $[0, \tau]$ simultaneously, switch one subsystem i_1 when it crosses the margin ($T^{i_1}(x^{i_1}) = \tau$) and space out the remaining switches linearly such that the last switch for subsystem i_k occurs at the boundary when $T^{i_k}(x^{i_k}) = 0$.

Our results regarding the performance of this strategy rely on two assumptions. The first is a mild assumption that allows us to analyze the flow in a vicinity of the end points.

ASSUMPTION 1. The functions f_{on}^i and f_{off}^i are continuous at $\underline{a}^i$ and at $\bar{a}^i$ for all $i \in [N]$.

Secondly, we state an assumption that is crucial for proving “optimality” of Strategy 1.

ASSUMPTION 2. The functions f_{on}^i and f_{off}^i are monotonically decreasing in S^i for all $i \in [N]$.

Crucially, this assumption holds for typical TCL models, since the (absolute) flow velocity of a stable one-dimensional linear system is necessarily monotonically decreasing towards its equilibrium point. We are now in a position to state the main results of this paper.

THEOREM 1. Assume that Assumption 1 holds and that the initial condition is not degenerate in the sense of Remark 1. Then, if

$$\sum_{i \in [N]} \frac{\mathcal{L}_{\text{off}}^i T_{\text{on}}^i(\underline{a}^i)}{1 + \mathcal{L}_{\text{off}}^i T_{\text{on}}^i(\underline{a}^i)} > \underline{K}_{\text{on}}, \text{ and,} \quad (7a)$$

$$\sum_{i \in [N]} \frac{\mathcal{L}_{\text{on}}^i T_{\text{off}}^i(\bar{a}^i)}{1 + \mathcal{L}_{\text{on}}^i T_{\text{off}}^i(\bar{a}^i)} > N - \bar{K}_{\text{on}}, \quad (7b)$$

then Strategy 1 solves Problem 1.

Strategy 1 can also be used to satisfy time-varying mode-counting bounds.

COROLLARY 1. Consider a generalization of Problem 1 with piecewise constant time-varying mode-counting bounds $\underline{K}_{\text{on}}(t)$ and $\bar{K}_{\text{on}}(t)$. If (7) holds for all $t \in \mathbb{R}^+$, then Strategy 1 enforces (2)-(3) for all $t \in \mathbb{R}^+$.

The next result shows that under the additional monotonicity assumption the inequalities (7) are also (almost) necessary conditions for Problem 1 to have a solution.

THEOREM 2. Assume that Assumption 1 and 2 hold. If the strict version of (7) is violated, i.e.,

$$\sum_{i \in [N]} \frac{\mathcal{L}_{\text{off}}^i T_{\text{on}}^i(\underline{a}^i)}{1 + \mathcal{L}_{\text{off}}^i T_{\text{on}}^i(\underline{a}^i)} < \underline{K}_{\text{on}}, \text{ or,} \quad (8a)$$

$$\sum_{i \in [N]} \frac{\mathcal{L}_{\text{on}}^i T_{\text{off}}^i(\bar{a}^i)}{1 + \mathcal{L}_{\text{on}}^i T_{\text{off}}^i(\bar{a}^i)} < N - \bar{K}_{\text{on}}, \quad (8b)$$

then Problem 1 has no solution.

The proofs of these two results are presented in the appendix. Together, Theorem 1 and 2 allow us to make a conclusion about the maximal controlled invariant² set contained in the global safety set

$$\mathfrak{S} = \left\{ \prod_{i \in [N]} \{(x^i, \sigma^i)\} : \begin{array}{l} \underline{K}_{\text{on}} \leq \sum_{i \in [N]} \mathbb{1}_{\{\text{on}\}}(\sigma^i) \leq \bar{K}_{\text{on}}, \\ x^i \in S^i, \quad \forall i \in [N] \end{array} \right\}. \quad (9)$$

COROLLARY 2. Assume that Assumptions 1 and 2 hold. Then, if (8) holds, the maximal controlled invariant set contained in $\mathfrak{S}$ is empty. Otherwise, the maximal controlled invariant set contained in $\mathfrak{S}$ is equal, up to closure, to $\mathfrak{S}$ itself.

PROOF. Only the case when neither (7) nor (8) is true remains. In this case, due to monotonicity we can pick an $\epsilon > 0$ such that (7) is satisfied for the modified bounds $(\underline{a}^i)' = \underline{a}^i + \epsilon$, $(\bar{a}^i)' = \bar{a}^i - \epsilon$. Hence, the set

$$\mathfrak{S}' = \left\{ \prod_{i \in [N]} \{(x^i, \sigma^i)\} : \begin{array}{l} \underline{K}_{\text{on}} \leq \sum_{i \in [N]} \mathbb{1}_{\{\text{on}\}}(\sigma^i) \leq \bar{K}_{\text{on}} \\ \underline{a}^i + \epsilon \leq x^i \leq \bar{a}^i - \epsilon \quad \forall i \in [N] \end{array} \right\}.$$

is a controlled invariant set contained in $\mathfrak{S}$. Letting $\epsilon \rightarrow 0$ gives the result. $\square$

²A set S is controlled invariant if for all initial states $x(0) \in S$, there is a switching policy such that $x(t) \in S$ for all $t \geq 0$.

REMARK 2. If the monotonicity assumption Assumption 2 does not hold, a variation of Strategy 1 may still be optimal. Assume that two points $(\underline{a}^i)'$ and $(\bar{a}^i)'$ can be found such that $(\underline{a}^i)' < (\bar{a}^i)'$, and such that

$$(\underline{a}^i)' \in \arg \max_{x \in S^i} \mathcal{L}_{\text{off}}^i T_{\text{on}}^i(x), \quad (\bar{a}^i)' \in \arg \max_{x \in S^i} \mathcal{L}_{\text{on}}^i T_{\text{off}}^i(x).$$

Then, if Strategy 1 is redefined to force switches at $(\underline{a}^i)'$ and $(\bar{a}^i)'$, the results from Theorem 1 and 2 still apply provided that initial conditions are within the range $[(\underline{a}^i)', (\bar{a}^i)']$.

4. CONTROL STRATEGY GENERALIZATIONS

In the previous section we proposed Strategy 1 and proved that it enforces invariance whenever it is possible to do so in the setting with one-dimensional subsystems. Here we outline more general problem instances where the strategy can be applied, and give conditions analogous to (7) that guarantee that the strategy solves Problem 1. However, the additional generality comes at the cost of potential conservativeness: there are no results analogous to Theorem 2 in these cases. The proofs of these generalizations follow that of Theorem 1.

4.1 Uncertainty in vector fields

Suppose that there is some bounded parametric uncertainty $d^i \in D^i$ present in the dynamics, i.e.,

$$\frac{d}{dt} x^i(t) = f_{\sigma^i(t)}^i(x^i(t), d^i(t)), \quad d^i(t) \in D^i, \quad (10)$$

but that $f_{\text{on}}^i(x, d) < 0$ for all $(x, d) \in S^i \times D^i$ and that $f_{\text{off}}^i(x, d) > 0$ for all $(x, d) \in S^i \times D^i$.

Time to exit is no longer defined for such a system, but if the “undisturbed” time to exit corresponding to $d = 0$ is used, Strategy 1 can still be implemented. If the worst-case disturbance is taken into account at the boundaries of the S^i ’s, the following inequalities are obtained, that—if fulfilled—guarantee that the strategy solves Problem 1.

$$\sum_{i \in [N]} \min_{d^i \in D^i} \frac{f_{\text{off}}^i(\underline{a}^i, d^i)}{-f_{\text{on}}^i(\underline{a}^i, d^i) + f_{\text{off}}^i(\underline{a}^i, d^i)} > \underline{K}_{\text{on}}, \text{ and,} \quad (11a)$$

$$\sum_{i \in [N]} \min_{d^i \in D^i} \frac{f_{\text{on}}^i(\bar{a}^i, d^i)}{-f_{\text{on}}^i(\bar{a}^i, d^i) + f_{\text{off}}^i(\bar{a}^i, d^i)} > N - \bar{K}_{\text{on}}. \quad (11b)$$

4.2 Higher-dimensional systems

The concept of time to exit is defined for systems of arbitrary dimension, so Strategy 1 can be implemented also for arbitrary mode-counting problems provided that the time to exit can be computed. To arrive at sufficient conditions for correctness, we divide the boundary ∂S^i of the local safety set S^i into parts ∂S_{on}^i and $\partial S_{\text{off}}^i$ where the **on** and **off** modes may exit, respectively:

$$\partial S_{\text{off}}^i = \left\{ x \in \partial S^i : \langle \hat{n}_x^{S^i}, f_{\text{off}}^i(x) \rangle \geq 0 \right\}, \quad (12a)$$

$$\partial S_{\text{on}}^i = \left\{ x \in \partial S^i : \langle \hat{n}_x^{S^i}, f_{\text{on}}^i(x) \rangle \geq 0 \right\}. \quad (12b)$$

We will propose sufficient conditions to guarantee that the maximal invariant set contained in $\mathfrak{S}$ is equal to $\mathfrak{S}$ itself, which is sufficient for Strategy 1 to solve the generalized version of Problem 1. First, we require that the distance between the sets $\partial S_{\text{off}}^i$ and ∂S_{on}^i is lower bounded by some

$\epsilon > 0$ for all $i \in [N]$. If, furthermore, the following generalizations of (7) hold:

$$\sum_{i \in [N]} \min_{\underline{a}^i \in \partial S_{\text{off}}^i} \frac{\mathcal{L}_{\text{off}}^i T_{\text{on}}^i(\underline{a}^i)}{1 + \mathcal{L}_{\text{off}}^i T_{\text{on}}^i(\underline{a}^i)} > \underline{K}_{\text{on}}, \text{ and,} \quad (13a)$$

$$\sum_{i \in [N]} \min_{\bar{a}^i \in \partial S_{\text{on}}^i} \frac{\mathcal{L}_{\text{on}}^i T_{\text{off}}^i(\bar{a}^i)}{1 + \mathcal{L}_{\text{on}}^i T_{\text{off}}^i(\bar{a}^i)} > N - \bar{K}_{\text{on}}, \quad (13b)$$

then $\mathfrak{S}$ is controlled invariant.

For higher-dimensional systems the maximal controlled invariant set contained in $\mathfrak{S}$ may exhibit a complicated geometry. In the one-dimensional case under the monotonicity assumption, it is, as stated in Corollary 2, either empty or equal (up to closure) to $\mathfrak{S}$ itself. In higher dimensions there is no corresponding result. Therefore, even if (13) does not hold, Strategy 1 may be able to enforce invariance of a smaller set.

5. EXAMPLE: TCL SCHEDULING

We illustrate the results with simulations of an aggregate TCL system. Following [7], the dynamics of an individual TCL can be modeled as follows:

$$\frac{d}{dt} \theta_i(t) = -a(\theta_i(t) - \theta_a) - bP_m \times \mathbb{1}_{\{\text{on}\}}(\sigma_i(t)). \quad (14)$$

We generated a heterogeneous collection of 1000 TCL's by sampling the parameter values a, θ_a, b and P_m around nominal values³, and sampled individual temperature dead band parameters $\underline{a}^i$ and $\bar{a}^i$ uniformly from [19, 20.9] and [21, 23], respectively. In the sampling we made sure that fundamental problem assumptions were satisfied; in particular, that $f_{\text{on}}^i < 0$ and $f_{\text{off}}^i > 0$ on $[\underline{a}^i, \bar{a}^i]$. When this is true, both f_{on}^i and f_{off}^i are also monotonically decreasing on the safe set. Therefore Assumptions 1 and 2 hold and therefore condition (7) is tight.

Evaluating (7) showed that the largest possible $\underline{K}_{\text{on}}$ is 323, and the smallest possible $\bar{K}_{\text{on}}$ is 250 for the generated collection. In other words, any mode-counting bounds $[\underline{K}_{\text{on}}, \bar{K}_{\text{on}}]$ with $\underline{K}_{\text{on}} \leq 323$ and $\bar{K}_{\text{on}} \geq 250$ can be satisfied indefinitely while also respecting local safety constraints. Figure 3 displays the resulting mode-on-count during a simulation of Strategy 1 where the imposed mode-counting bounds are varying within the feasible values discussed above. Information about the times to exit is shown in Figure 4 and temperature traces of three individual TCL's are depicted in Figure 5.

REMARK 3. *In the simulation a sample time 0.01 h was used to implement Strategy 1. Rather than switching a TCL at the instant it reaches the boundary of S^i , it is switched at the sample instant if it exited S^i during the last sample period. This implementation is practical due to its simplicity, but will lead to brief violations of local safety constraints. However, it follows from the monotonicity property and the reasoning in the appendix that no single state constraint will be permanently violated, and there is an upper bound on the magnitude of the violation that goes to 0 as the sample time decreases. The violations could also be corrected for by shrinking the local safety constraints by an appropriate margin.*

³The nominal values were chosen as $a = 0.25$, $b = 1.25$, $\theta_a = 28.6$, $P_m = 5.6$.

Figure 3: Actual mode-on-count (solid blue) and imposed mode on counting bounds (dashed green) in a Simulation of Strategy 1 for an example with 1000 heterogeneous TCL's. As can be seen, the imposed bounds are satisfied throughout the simulation.

If Strategy 1 is implemented in continuous time, it is shown in the appendix that it will not result in Zeno behavior as long as (7) holds.

6. CONCLUSIONS

In this paper we studied the mode-counting problem with two modes in one dimension, and proposed a control strategy that we proved to be optimal in a certain sense. Our analysis is grounded on careful consideration of the hybrid nature of the closed loop dynamics, where the control strategy imposes switching based on local end user constraints and global mode-counting constraints. The results imply tight bounds on the performance that can be achieved in steady state in TCL coordination applications, as demonstrated in Section 5.

Future work will be focused on two objectives. Firstly, we are interested in further investigating the short-horizon control properties of TCL scheduling. It can be shown that Strategy 1 is not the best possible strategy if (8) holds, in the sense that there is a different strategy that maintains $\mathfrak{S}$ invariant for a longer time (although no strategy can achieve invariance indefinitely). By finding strategies that maximize the time until local safety constraint violation we hope to put forward control schemes with short-term optimality conditions and simultaneously derive analytical bounds on tracking performance for larger classes of time-varying counting constraints. Similar control schemes will also be relevant to additional control objectives like minimizing the number of switches uniformly or on average while maintaining local safety and counting constraints.

Secondly, we hope to discover additional settings where the same approach can be applied: to propose a strategy and prove that it enforces invariance under some assumptions. When this is achievable, the control strategy becomes an implicit representation of a controlled invariant set. In this work the maximal controlled invariant set turned out to be either empty, or equal (up to closure) to the safe set itself. This is typically not the case; for more complex situations an implicit definition of a controlled invariant set through a control strategy may be the only practical option.

Figure 4: Time to exit for the simulation in Figure 3. The plots show maximal, average, and minimal times to exit for the two modes. As can be seen, the minimal times to exit remain above 0 which implies that local safety constraints are satisfied. In the latter part of the simulation, the lower mode-on-bound $\underline{K}_{on}$ is equal to 320 which is close to the largest feasible value 323. As a result, the subsystems congregate at the lower boundary which is illustrated by the fact that the times to on exit approach 0. This confirms the tightness of condition (7).

Figure 5: Illustration of temperature movement of three individual TCL's depicted in red, blue, and green. The dashed lines indicate heterogeneous local state constraints. As can be seen, Strategy 1 switches an individual TCL whenever those local safety constraints are about to be violated. Other switches are the result of necessity to enforce global mode-counting constraints. In the latter part of the simulation the imposed mode-on-count is high, as shown in Figure 4. As a consequence, the individual TCL's group in the lower parts of their temperature spectra, since the on fields f_{on}^i are relatively weaker there, thus enabling staying in mode on during a larger fraction of the time. The reverse holds around $t = 5$ h where the imposed mode-on-count is low.

APPENDIX

The proofs of Theorem 1 and 2 rely on a series of lemmas that are presented first. By without loss of generality excluding degenerate cases (c.f. Remark 1), subsystems can be assumed to arrive at boundaries $\underline{a}^i$ and $\bar{a}^i$ one by one. In this case, Strategy 1 enforces the constraints (2)-(3) for all times that the closed loop trajectory is defined. The challenge is to show that the trajectory is defined for all times which amounts to ruling out *Zeno behavior* [9]—accumulation of switching instances at a time t_Z beyond which the solution is undefined.

LEMMA 1. *Zeno behavior can only occur for Strategy 1 if a group $I \subset [N]$ of subsystems all congregate at either their lower or upper boundaries, but not both. That is, either $x^i(t) \rightarrow \underline{a}^i$ for all $i \in I$, or, $x^i(t) \rightarrow \bar{a}^i$ for all $i \in I$.*

PROOF. By the definition of Zeno behavior, an infinite number of switches must take place over a finite time interval $[0, t_Z)$. These switches must necessarily be undertaken by subsystems in some subset $I \subset [N]$, and we may assume that all subsystems in I switch an infinite amount of times by excluding those that do not. Consider a single subsystem $i \in I$ and some interval $[t_0, t_Z)$ so that the duration between switches of i is less than δ , where δ can be chosen arbitrarily small by adjusting t_0 . By inspection of Strategy 1, it follows that either $T_{on}^i(x^i(t)) \leq \delta$ or $T_{off}^i(x^i(t)) \leq \delta$ for all $t \in [t_0, t_Z)$. Hence $x^i(t)$ approaches either $\underline{a}^i$ or $\bar{a}^i$.

We show that all systems in I must congregate at either the lower boundaries $\underline{a}^i$ or the upper boundaries $\bar{a}^i$, but not both. This trivially holds when I is a singleton. Now, assume for contradiction that I can be partitioned into non-empty sets $\underline{I}$ and $\bar{I}$ such that $x^i(t) \rightarrow \underline{a}^i$ for $i \in \underline{I}$ and $x^j(t) \rightarrow \bar{a}^j$ for $j \in \bar{I}$. By the same construction as above, we pick a $\delta > 0$ and an associated time interval $[t_0, t_Z)$ such that $T_{on}^i(x^i(t)) \leq \delta$ and $T_{off}^j(x^j(t)) \leq \delta$ for all $i \in \underline{I}$, all $j \in \bar{I}$, and all $t \in [t_0, t_Z)$. This is a contradiction of the behavior of Strategy 1. Indeed, for δ small enough $T_{on}^j(x^j) > T_{on}^i(x^i)$ for all $i \in \underline{I}, j \in \bar{I}$, in which case 1) x^i would never be switched from off to on unless all systems in $\bar{I}$ are already in on, and 2) x^j would never be switched from on to off unless all systems in $\underline{I}$ are already in off. But in such a situation no switches occur at all. Therefore either $\underline{I} = I$ or $\bar{I} = I$, which shows that either all subsystems congregate at $\underline{a}^i$, or all subsystems congregate at $\bar{a}^i$. $\square$

LEMMA 2. *If Zeno behavior occurs by a group subsystems $I \subset [N]$ congregating at their lower boundaries $\underline{a}^i$, then*

$$\lim_{t \rightarrow t_Z} \sum_{i \in I} T_{on}^i(x^i(t)) = 0,$$

where t_Z is the Zeno time.

PROOF. It suffices to remark that the time between switches for a subsystem $i \in I$ goes to 0 as $t \rightarrow t_Z$, and that $\sup_{t' > t} T_{on}^i(x^i(t'))$ is upper bounded by the maximal time between switches on the interval $[t, t_Z)$. $\square$

LEMMA 3. *Zeno behavior at the lower boundaries $\underline{a}^i$ for subsystems $I \subset [N]$ implies that exactly*

$$\max(0, \underline{K}_{on} - N + |I|)$$

of the subsystems in I are in mode on for times close to the Zeno time.

PROOF. We first show that all subsystems that are not in I must be in mode **on**. Assume that Zeno behavior occurs for a subset $I \subset [N]$ at the lower boundary, and assume for contradiction that there is a subsystem $j \in [N] \setminus I$ that is in mode **off** for all times $t \in [t_0, t_Z]$ for some t_0 . By Lemma 2, $T_{\text{on}}^i(x^i(t)) \rightarrow 0$ for all $i \in I$ as $t \rightarrow t_Z$. It follows that $T_{\text{on}}^j(x^j(t)) > T_{\text{on}}^i(x^i(t))$ for t close enough to t_Z , so Strategy 1 switches subsystem j to **on** before switching subsystem i to **on**, for all $i \in I$. This is a contradiction.

Now, a subsystem in I can only be switched to mode **on** when condition 3. of Strategy 1 occurs. It follows that there are exactly $\underline{K}_{\text{on}}$ subsystems in mode **on** for all times $t \in [t_0, t_Z]$, and out of those,

$$\underline{K}_{\text{on}} - \min(\underline{K}_{\text{on}}, N - |I|) = \max(0, \underline{K}_{\text{on}} - N + |I|)$$

are members of I . $\square$

LEMMA 4. If (7) holds, then

$$-\max(\underline{K}_{\text{on}} + |I| - N, 0) + \sum_{i \in I} \frac{\mathcal{L}_{\text{off}}^i T_{\text{on}}^i(\underline{a}^i)}{1 + \mathcal{L}_{\text{off}}^i T_{\text{on}}^i(\underline{a}^i)} > 0$$

for all sets $I \subset [N]$.

PROOF. If the max evaluates to 0, the inequality is evidently valid. If not, it suffices to note that

$$\mathcal{L}_{\text{off}}^i T_{\text{on}}^i(\underline{a}^i) / (1 + \mathcal{L}_{\text{off}}^i T_{\text{on}}^i(\underline{a}^i)) < 1 \text{ for all } i \in [N]$$

to obtain

$$\begin{aligned} -\underline{K}_{\text{on}} + (N - |I|) + \sum_{i \in I} \frac{\mathcal{L}_{\text{off}}^i T_{\text{on}}^i(\underline{a}^i)}{1 + \mathcal{L}_{\text{off}}^i T_{\text{on}}^i(\underline{a}^i)} \\ \geq -\underline{K}_{\text{on}} + \sum_{i \in [N]} \frac{\mathcal{L}_{\text{off}}^i T_{\text{on}}^i(\underline{a}^i)}{1 + \mathcal{L}_{\text{off}}^i T_{\text{on}}^i(\underline{a}^i)} > 0. \end{aligned}$$

$\square$

LEMMA 5. Under Assumption 2, $\mathcal{L}_{\text{off}}^i T_{\text{on}}^i(x)$ is monotonically decreasing in S^i .

PROOF. By assumption, both $f_{\text{on}}^i < 0$ and $f_{\text{off}}^i > 0$ are monotonically decreasing in S^i . From Proposition 2:

$$\mathcal{L}_{\text{off}}^i T_{\text{on}}^i(x) = f_{\text{off}}^i(x) / (-f_{\text{on}}^i(x)).$$

As a quotient of a positive monotonically decreasing function and a positive monotonically increasing function, it is monotonically decreasing. $\square$

The following proofs use the concept of a *cycle* in the state space which we introduce next.

DEFINITION 3. A **cycle** for subsystem i is a continuous trajectory segment $x^i(t)$ on the bounded interval $[0, \tau_{\text{on}} + \tau_{\text{off}}]$ with exactly one switch such that $x^i(0) = x^i(\tau_{\text{on}} + \tau_{\text{off}})$.

In the following we use a wider notion of a cycle as a collection of trajectory segments that—when patched together—satisfy Definition 3. The subsequent proofs rely on partitioning a trajectory into such generalized cycles, as illustrated in Figure 6. Any trajectory can be divided into cycles with the possible exception of bounded time intervals. We state this precisely in the next lemma.

Figure 6: Illustration of cycle partitioning of a trajectory $x^i(t)$. The segment plotted in solid blue is a true cycle in the sense of Definition 3: it starts in mode **off** at $\underline{b}_1$, reaches $\bar{b}_1$, switches to mode **on**, and returns to $\underline{b}_1$. Therefore $\bar{b}_1 = \phi_{\text{off}}^i(\tau_{\text{off}}, \underline{b}_1)$ and $\underline{b}_1 = \phi_{\text{on}}^i(\tau_{\text{on}}, \bar{b}_1)$, where τ_{off} and τ_{on} are the times spent in mode **off** and **on**, respectively. The segments marked in red can be patched together to form a cycle.

LEMMA 6. A trajectory $x^i(t)$ defined on an interval $[t_0, t_1]$ that satisfies the state constraints (2) can be divided into cycles in a way such that the duration of time not captured by cycles is at most

$$\max_{t \in [t_0, t_1]} (T_{\text{on}}^i(x^i(t)), T_{\text{off}}^i(x^i(t))). \quad (15)$$

PROOF. Any segment $[t'_0, t'_1]$ such that $x^i(t'_0) = x^i(t'_1)$ can be recursively partitioned into cycles, as illustrated in Figure 6. The bound (15) captures the maximal possible time duration during which $x^i(t) \neq x^i(t')$ for all $t \neq t'$. $\square$

From the definition of a cycle it follows that the time interval for a cycle starting in mode **off** can be divided into a mode-**off** interval $[0, \tau_{\text{off}}]$ and a mode-**on** interval $[\tau_{\text{off}}, \tau_{\text{off}} + \tau_{\text{on}}]$. The next lemma relates these times to the relative strengths of the vector fields, which is captured by $\mathcal{L}_{\text{off}}^i T_{\text{on}}^i$. This result also applies for cycles that are not connected in time, such as the red cycle in Figure 6.

LEMMA 7. Consider a cycle for subsystem i with time τ_{on} in mode **on** and time τ_{off} in mode **off**. Let $\underline{b}$ and $\bar{b}$ be the extremal values of $x^i(t)$ during the cycle. Then,

$$\frac{\int_{s=0}^{\tau_{\text{off}}} \mathcal{L}_{\text{off}}^i T_{\text{on}}^i(\phi_{\text{off}}^i(s, \underline{b})) \, ds}{\tau_{\text{on}} + \tau_{\text{off}}} \leq \frac{\mathcal{L}_{\text{off}}^i T_{\text{on}}^i(\bar{b})}{1 + \mathcal{L}_{\text{off}}^i T_{\text{on}}^i(\bar{b})}. \quad (16)$$

Furthermore, if f_{off}^i and f_{on}^i are continuous at $\underline{b}$:

$$\lim_{\tau_{\text{off}} \rightarrow 0} \frac{\int_{s=0}^{\tau_{\text{off}}} \mathcal{L}_{\text{off}}^i T_{\text{on}}^i(\phi_{\text{off}}^i(s, \underline{b})) \, ds}{\tau_{\text{on}} + \tau_{\text{off}}} = \frac{\mathcal{L}_{\text{off}}^i T_{\text{on}}^i(\underline{b})}{1 + \mathcal{L}_{\text{off}}^i T_{\text{on}}^i(\underline{b})}.$$

PROOF. Let $\bar{b} = \phi_{\text{off}}^i(\tau_{\text{off}}, \underline{b})$ be the maximal value of $x^i(t)$ during the cycle. By twice re-parameterizing an integral we

can write τ_{on} as

$$\begin{aligned}\tau_{\text{on}} &= \int_{s=0}^{\tau_{\text{on}}} ds = \int_{x=\bar{b}}^{\bar{b}} \frac{dx}{f_{\text{on}}^i(x)} \left(\begin{array}{l} \text{by ch. of vars:} \\ x = \phi_{\text{on}}^i(s, \bar{b}) \\ dx = f_{\text{on}}^i(x) ds \end{array} \right) = \int_{x=\bar{b}}^{\bar{b}} -\frac{dx}{f_{\text{on}}^i(x)} \\ &= \int_{s=0}^{\tau_{\text{off}}} -\frac{f_{\text{off}}^i(\phi_{\text{off}}^i(s, \underline{b}))}{f_{\text{on}}^i(\phi_{\text{off}}^i(s, \underline{b}))} ds \left(\begin{array}{l} \text{by ch. of vars:} \\ x = \phi_{\text{on}}^i(s, \underline{b}) \\ dx = f_{\text{off}}^i(\phi_{\text{off}}^i(s, \underline{b})) ds \end{array} \right) \\ &= \int_{s=0}^{\tau_{\text{off}}} \mathcal{L}_{\text{off}}^i T_{\text{on}}^i(\phi_{\text{off}}^i(s, \underline{b})) ds.\end{aligned}$$

It follows that the left-hand side of (16) is actually an expression for $\tau_{\text{on}}/(\tau_{\text{on}} + \tau_{\text{off}})$ —the fraction of time spent in mode **on** during the cycle. Furthermore, by denoting $\tau_{\text{on}}(\tau_{\text{off}}) = \int_{s=0}^{\tau_{\text{off}}} \mathcal{L}_{\text{off}}^i T_{\text{on}}^i(\phi_{\text{off}}^i(s, \underline{b})) ds$,

$$\frac{\int_{s=0}^{\tau_{\text{off}}} \mathcal{L}_{\text{off}}^i T_{\text{on}}^i(\phi_{\text{off}}^i(s, \underline{b})) ds}{\tau_{\text{on}} + \tau_{\text{off}}} = \frac{\frac{1}{\tau_{\text{off}}} \tau_{\text{on}}(\tau_{\text{off}})}{1 + \frac{1}{\tau_{\text{off}}} \tau_{\text{on}}(\tau_{\text{off}})}. \quad (17)$$

The mapping $x \mapsto x/(1+x)$ is strictly increasing for $x \geq 0$, therefore (17) attains its maximal value for the largest possible $\frac{1}{\tau_{\text{off}}} \tau_{\text{on}}(\tau_{\text{off}})$. By Lemma 5;

$$\begin{aligned}\frac{1}{\tau_{\text{off}}} \tau_{\text{on}}(\tau_{\text{off}}) &= \frac{1}{\tau_{\text{off}}} \int_{s=0}^{\tau_{\text{off}}} \mathcal{L}_{\text{off}}^i T_{\text{on}}^i(\phi_{\text{off}}^i(s, \underline{b})) ds \\ &\leq \frac{1}{\tau_{\text{off}}} \sup_{u \in [0, \tau_{\text{off}}]} \mathcal{L}_{\text{off}}^i T_{\text{on}}^i(\phi_{\text{off}}^i(u, \underline{b})) = \mathcal{L}_{\text{off}}^i T_{\text{on}}^i(\underline{b}).\end{aligned}$$

Plugging this into (17) proves the first claim. By taking the limit of the integral, it follows given the continuity assumption that the maximum is approached as $\tau_{\text{off}} \rightarrow 0$. $\square$

PROOF OF THEOREM 1. We need to prove that the strategy does not induce Zeno behavior. Assume for contradiction that Zeno behavior occurs at $t = t_Z$ but that (7) holds. Then by Lemma 1 we can select a subset of subsystems $I \subset [N]$ that congregate at the (w.l.o.g.) lower boundaries $\underline{a}^i$.

Consider a time $t < t_Z$ such that all switches that occur after t are switches of subsystems in I . We can pick a $\tau > 0$ such that the trajectory is defined on the whole time interval $[t, t + \tau]$. Define $\tau_{\text{on}}^* = \max_{i \in I} \sup_{s \geq t} T_{\text{on}}^i(s)$. By Lemma 2, $\tau_{\text{on}}^* \rightarrow 0$ as $t \rightarrow t_Z$. Consider the evolution of $\sum_{i \in I} T_{\text{on}}^i(x^i(t))$ on the interval $[t, t + \tau]$:

$$\begin{aligned}\sum_{i \in I} T_{\text{on}}^i(x^i(t + \tau)) &= \sum_{i \in I} T_{\text{on}}^i(x^i(t)) \\ &\quad + \int_{s=t}^{t+\tau} \frac{d}{ds} \sum_{i \in I} T_{\text{on}}^i(x^i(s)) ds.\end{aligned} \quad (18)$$

By Lemma 3, out of the $|I|$ subsystems in I , $\max(\underline{K}_{\text{on}} + |I| - N, 0)$ are in mode **on**. We recall that

$$\frac{d}{ds} T_{\text{on}}^i(x^i(s)) = \begin{cases} -1 & \text{if } \sigma^i(s) = \text{on}, \\ \mathcal{L}_{\text{off}}^i T_{\text{on}}^i(x^i(s)) & \text{if } \sigma^i(s) = \text{off}. \end{cases}$$

By splitting the integral we can therefore write

$$\begin{aligned}\int_{s=t}^{t+\tau} \sum_{i \in I} \frac{d}{ds} T_{\text{on}}^i(x^i(s)) ds &= \sum_{i \in I} \int_{s=t}^{t+\tau} \frac{d}{ds} T_{\text{on}}^i(x^i(s)) ds \quad (19) \\ &= \sum_{i \in I} \left(\int_{\substack{s \in [t, t+\tau] \\ \sigma^i(s) = \text{on}}} \frac{d}{ds} T_{\text{on}}^i(x^i(s)) ds + \int_{\substack{s \in [t, t+\tau] \\ \sigma^i(s) = \text{off}}} \frac{d}{ds} T_{\text{on}}^i(x^i(s)) ds \right) \\ &= -\max(\underline{K}_{\text{on}} + |I| - N, 0) \tau + \sum_{i \in I} \int_{\substack{s \in [t, t+\tau] \\ \sigma^i(s) = \text{off}}} \mathcal{L}_{\text{off}}^i T_{\text{on}}^i(x^i(s)) ds.\end{aligned}$$

We consider the remaining integral. Each subsystem i in I will reach $\underline{a}^i$, switch to **off**, travel “upwards” until $T_{\text{on}}^i(x^i(t)) = \tau_{\text{on}}^{i,c}$ for some $\tau_{\text{on}}^{i,c} < \tau_{\text{on}}^*$, switch to **on**, and finally reach $\underline{a}^i$ again. We can thus partition the trajectory into a set of cycles Cyc_i in the sense of Definition 3. We split the integral above accordingly and obtain the following inequality due to possibly omitting times at the beginning and the end of $[t, t + \tau]$ that are not part of complete cycles (c.f. Lemma 6).

$$\int_{\substack{s \in [t, t+\tau] \\ \sigma^i(s) = \text{off}}} \mathcal{L}_{\text{off}}^i T_{\text{on}}^i(x^i(s)) ds \geq \sum_{c \in Cyc_i} \int_{s=0}^{\tau_{\text{off}}^{c,i}} \mathcal{L}_{\text{off}}^i T_{\text{on}}^i(\phi_{\text{off}}^i(\underline{a}^i, s)) ds,$$

where $\tau_{\text{off}}^{c,i}$ is the time spent in mode **off** by subsystem i during cycle c . By Lemma 7 we can bound this from below as

$$\begin{aligned}\sum_{c \in Cyc_i} \int_{s=0}^{\tau_{\text{off}}^{c,i}} \mathcal{L}_{\text{off}}^i T_{\text{on}}^i(\phi_{\text{off}}^i(\underline{a}^i, s)) ds \\ = \sum_{c \in Cyc_i} \left(\tau_{\text{off}}^{c,i} + \tau_{\text{on}}^{c,i} \right) \left(\frac{\mathcal{L}_{\text{off}}^i T_{\text{on}}^i(\underline{a}^i)}{1 + \mathcal{L}_{\text{off}}^i T_{\text{on}}^i(\underline{a}^i)} - \delta(\tau_{\text{off}}^{c,i}) \right) \\ \geq (\tau - \delta_\tau) \left(\frac{\mathcal{L}_{\text{off}}^i T_{\text{on}}^i(\underline{a}^i)}{1 + \mathcal{L}_{\text{off}}^i T_{\text{on}}^i(\underline{a}^i)} - \delta \right),\end{aligned}$$

where $\delta = \max_i \max_c \delta(\tau_{\text{off}}^{c,i}) \rightarrow 0$ as $\tau_{\text{on}}^* \rightarrow 0$ and $\delta_\tau < \tau$ has been subtracted to compensate for time before and after cycles. We thus obtain a lower bound for (19) as

$$\begin{aligned}\int_{s=t}^{t+\tau} \sum_{i \in I} \frac{d}{ds} T_{\text{on}}^i(x^i(s)) ds \\ \geq \left[-\max(\underline{K}_{\text{on}} + |I| - N, 0) + \sum_{i \in I} \frac{\mathcal{L}_{\text{off}}^i T_{\text{on}}^i(\underline{a}^i)}{1 + \mathcal{L}_{\text{off}}^i T_{\text{on}}^i(\underline{a}^i)} - \delta |I| \right] (\tau - \delta_\tau).\end{aligned}$$

Now, by letting $\tau \rightarrow \tau_Z = t_Z - t$ in (18);

$$0 \geq \underbrace{\sum_{i \in I} T_{\text{on}}^i(x(t))}_{\geq 0} + \underbrace{\left[-\max(\underline{K}_{\text{on}} + |I| - N, 0) + \sum_{i \in I} \frac{\mathcal{L}_{\text{off}}^i T_{\text{on}}^i(\underline{a}^i)}{1 + \mathcal{L}_{\text{off}}^i T_{\text{on}}^i(\underline{a}^i)} - \delta |I| \right]}_A \underbrace{(\tau_Z - \delta_{\tau_Z})}_B.$$

As $t \rightarrow t_Z$, $B \rightarrow 0$ from above, and

$$A \rightarrow -\max(\underline{K}_{\text{on}} + |I| - N, 0) + \sum_{i \in I} \frac{\mathcal{L}_{\text{off}}^i T_{\text{on}}^i(\underline{a}^i)}{1 + \mathcal{L}_{\text{off}}^i T_{\text{on}}^i(\underline{a}^i)},$$

which is strictly positive by Lemma 4. This is a contradiction of the last inequality, and thus of the Zeno behavior. $\square$

PROOF OF THEOREM 2. We assume that the first inequality in (8) holds; the other case can be treated symmetrically. Assume for contradiction that a switching strategy $\{\sigma^i(t)\}_{i \in [N]}$ that generates infinite-time trajectories satisfying the constraints of Problem 1 exists. Then,

$$\frac{d}{dt} \sum_{i \in [N]} T_{\text{on}}^i(x^i(t)) \leq -K_{\text{on}} + \sum_{\substack{i \in [N] \\ \sigma^i(t) = \text{off}}} \mathcal{L}_{\text{off}}^i T_{\text{on}}^i(x^i(t)). \quad (20)$$

We integrate the right-hand sum over an interval $[0, t_f]$,

$$\int_{s=0}^{t_f} \sum_{\substack{i \in [N] \\ \sigma^i(t) = \text{off}}} \mathcal{L}_{\text{off}}^i T_{\text{on}}^i(x^i(s)) ds = \sum_{i \in [N]} \int_{s \in [0, t_f] \\ \sigma^i(s) = \text{off}} \mathcal{L}_{\text{off}}^i T_{\text{on}}^i(x^i(s)) ds,$$

and seek to bound the integral for each $i \in [N]$. By Lemma 6 the trajectory can be partitioned into complete cycles $c \in \text{Cyc}_i$ with extremal points $\underline{b}_c$ and $\bar{b}_c$ so that the whole interval $[0, t_f]$ is covered except possibly for a total duration $\Delta_{t_f}^i$ at the beginning and the end of the interval $[0, t_f]$. We divide the integral into contributions from these cycles and use Lemma 7,

$$\begin{aligned} \int_{\substack{s \in [0, t_f] \\ \sigma^i(s) = \text{off}}} \mathcal{L}_{\text{off}}^i T_{\text{on}}^i(x^i(s)) ds &= \sum_{c \in \text{Cyc}_i} \int_{s=0}^{\tau_{\text{off}}^{c,i}} \mathcal{L}_{\text{off}}^i T_{\text{on}}^i(\phi_{\text{off}}^i(s, \underline{b}_c)) ds \\ &\leq \sum_{c \in \text{Cyc}_i} \left(\tau_{\text{off}}^{c,i} + \tau_{\text{on}}^{c,i} \right) \frac{\mathcal{L}_{\text{off}}^i T_{\text{on}}^i(\underline{b}_c)}{1 + \mathcal{L}_{\text{off}}^i T_{\text{on}}^i(\underline{b}_c)} \\ &\leq (t_f - \Delta_{t_f}^i) \frac{\mathcal{L}_{\text{off}}^i T_{\text{on}}^i(\underline{a}^i)}{1 + \mathcal{L}_{\text{off}}^i T_{\text{on}}^i(\underline{a}^i)}. \end{aligned}$$

By Lemma 6 the non-cycle time $\Delta_{t_f}^i$ is bounded independently of t_f as $\Delta_{t_f}^i \leq \max(T_{\text{on}}^i(\bar{a}), T_{\text{off}}^i(\underline{a}^i))$.

Using this upper bound and (8), integrating (20) yields

$$\begin{aligned} \sum_{i \in [N]} T_{\text{on}}^i(x^i(t_f)) - \sum_{i \in [N]} T_{\text{on}}^i(x^i(0)) &\leq -K_{\text{on}} t_f + \sum_{i \in [N]} \frac{\mathcal{L}_{\text{off}}^i T_{\text{on}}^i(\underline{a}^i)}{1 + \mathcal{L}_{\text{off}}^i T_{\text{on}}^i(\underline{a}^i)} (t_f - \bar{\Delta}_{t_f}^i) \\ &\leq \left(-K_{\text{on}} + \sum_{i \in [N]} \frac{\mathcal{L}_{\text{off}}^i T_{\text{on}}^i(\underline{a}^i)}{1 + \mathcal{L}_{\text{off}}^i T_{\text{on}}^i(\underline{a}^i)} \right) t_f - \bar{\Delta}_{t_f} \\ &\leq -\epsilon t_f - \bar{\Delta}_{t_f} \end{aligned}$$

for some $\epsilon > 0$ and where $\bar{\Delta}_{t_f}$ is bounded independently of t_f and i since each $\Delta_{t_f}^i$ is uniformly bounded. Letting $t_f \rightarrow \infty$ results in a violation of state constraints, which is a contradiction of the correctness of the switching strategy. $\square$

Acknowledgments

The authors thank Johanna Mathieu for discussions motivating this work. Petter Nilsson is supported by NSF grant CNS-1239037. Necmiye Ozay is supported in part by NSF grants CNS-1446298 and ECCS-1553873, and DARPA grant N66001-14-1-4045.

A. REFERENCES

- [1] R. Alur, V. Forejt, S. Moarref, and A. Trivedi. Safe schedulability of bounded-rate multi-mode systems. In *Proc. of HSCC*, pages 243–252, 2013.
- [2] R. Alur, A. Trivedi, and D. Wojtczak. Optimal scheduling for constant-rate multi-mode systems. In *Proc. of HSCC*, pages 75–84, 2012.
- [3] F. Blanchini. Survey paper: Set invariance in control. *Automatica*, 35(11):1747–1767, 1999.
- [4] S. Esmail Zadeh Soudjani and A. Abate. Aggregation and control of populations of thermostatically controlled loads by formal abstractions. *IEEE Trans. on Control Systems Technology*, 23(3):975–990, 2015.
- [5] A. Ghaffari, S. Moura, and M. Krstić. Modeling, control, and stability analysis of heterogeneous thermostatically controlled load populations using partial differential equations. *ASME Journal of Dynamic Systems, Measurement, and Control*, 137(10), 2015.
- [6] M. R. Hafner and D. Del Vecchio. Computational tools for the safety control of a class of piecewise continuous systems with imperfect information on a partial order. *SIAM Journal on Control and Optimization*, 49(6):2463–2493, 2011.
- [7] H. Hao, B. M. Sanandaji, K. Poolla, and T. L. Vincent. Aggregate Flexibility of Thermostatically Controlled Loads. *IEEE Trans. on Power Systems*, 30(1):189–198, 2015.
- [8] S. Koch, J. L. Mathieu, and D. S. Callaway. Modeling and control of aggregated heterogeneous thermostatically controlled loads for ancillary services. In *Proc. of PSCC*, pages 1–7, 2011.
- [9] D. Liberzon. *Switching in Systems and Control*. Springer, 2003.
- [10] R. Malhame and C.-Y. Chong. Electric load model synthesis by diffusion approximation of a high-order hybrid-state stochastic system. *IEEE Trans. on Automatic Control*, 30(9):854–860, 1985.
- [11] J. L. Mathieu, S. Koch, and D. S. Callaway. State estimation and control of electric loads to manage real-time energy imbalance. *IEEE Trans. on Power Systems*, 28(1):430–440, 2013.
- [12] T. X. Nghiem, M. Behl, and R. Mangharam. Green scheduling of control systems for peak demand reduction. In *Proc. of IEEE CDC*, pages 5131–5136, 2011.
- [13] P. Nilsson and N. Ozay. Control synthesis for large collections of systems with mode-counting constraints. In *Proc. of HSCC*, pages 205–214, 2016.
- [14] S. Sadraddini and C. Belta. Safety control of monotone systems with bounded uncertainties. In *Proc. of IEEE CDC*, pages 4874–4879, 2016.
- [15] B. M. Sanandaji, H. Hao, and K. Poolla. Fast regulation service provision via aggregation of thermostatically controlled loads. In *Proc. of HICSS*, pages 2388–2397, 2014.
- [16] C. J. Tomlin, I. Mitchell, A. M. Bayen, and M. Oishi. Computational techniques for the verification of hybrid systems. *Proceedings of the IEEE*, 91(7):986–1001, 2003.
- [17] W. Walter. *Ordinary Differential Equations*. Springer, 1998.