Maximizing the Time of Invariance for Large Collections of Switched Systems

Petter Nilsson and Necmiye Ozay

Abstract— **Motivated by the problem of capacity allocation in aggregate systems of thermostatically controlled loads, we seek to optimize the *time of invariance*—the maximal time that a scheduler can maintain local safety constraints without exceeding an aggregate capacity constraint. Having previously characterized conditions when the time of invariance is infinite, this paper focuses on situations when no scheduler can satisfy all constraints indefinitely.**

We derive a lower bound for the optimal time of invariance, and leverage ideas from optimal control to propose control strategies with associated lower bounds on the resulting times of invariance. We also analyze the optimality gap. The results are illustrated on a 10,000-dimensional example.

I. INTRODUCTION

Thermostatically controlled loads (TCLs) include air conditioners, water heaters, refrigerators, etc., that operate around a given temperature fixed point. End users of TCLs are typically indifferent to small temperature variations around the desired fixed point. While not significant for a single system, the aggregate flexibility of large families of TCLs can be leveraged by utility companies to shape aggregate demand. The aggregate system can be thought of as a battery that can be charged and discharged [6] in order to balance the total load on the grid over time.

Efforts to control large families of TCLs frequently appear in the literature. As demonstrated in [4, 5], a universal fixed point temperature can be imposed across a family of TCLs to achieve tracking goals. However, such a universal approach may violate end user requirements that are not necessarily uniform. Another approach is to partition the temperature range into "bins" and model intra-bin transitions as a discrete-time Markov process [3]. In these and other [11, 6] methods 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.

Previous work has been demonstrated effective in simulations, but the methods do not provide formal guarantees on performance. Motivated by this shortcoming, in [9] we introduced the *mode-counting problem*: a class of problems where local safety constraints (i.e., maximal deviation from desired temperature) are in conflict with a global constraint on the number of systems that are allowed to be in a certain mode (i.e., reference tracking error). Our initial solution approach was in the form of an integer linear program from which—if feasible—a correct controller could be extracted. We further studied the TCL-inspired setting of one-dimensional subsystems with two modes (`on` and `off`) in [10], and gave a precise characterization of when the infinite-horizon mode-counting problem has a solution. However, the situations when a solution exists turned out to be rather narrow for practical purposes: we found an aggregate flexibility of around 6% (trackable signals as a fraction of the maximal power) when imposing infinite-horizon bounds, whereas applications often require a flexibility upwards of 60% but for relatively short durations [7].

In this paper we therefore further study the situation when the infinite-horizon mode-counting problem lacks a solution, and propose control strategies that satisfy the constraints *as long as possible*—we call the maximal time of satisfaction the *time of invariance*. If a low-level capacity allocator can provide guarantees on its time of invariance given a reference power signal and the current state of its system, that information can be fed to a high-level capacity allocator to make informed decisions.

The technical approach is based on a formulation of the problem as a time-optimal control problem where the objective is to *maximize* the time until a certain set is reached. We propose a heuristic value function for the optimal control problem from which a control strategy can be derived via the Hamilton-Jacobi-Bellman (HJB) equation; similar approaches have previously been considered e.g. in pursuit-evasion games [12]. By assessing how close the value function is to satisfying the HJB equation, we can also bound the optimality gap.

The paper is structured as follows. The remainder of this section introduces relevant concepts from our previous work and formalizes the problem statement. In Section II we under-approximate the time of invariance and use the lower bound to define control strategies with associated performance guarantees derived in Section III. We showcase our results on numerical examples in Section IV and conclude the paper in Section V.

A. Time to exit

In order to obtain a homogeneous measure to compare the imminence of constraint violation across heterogeneous families of TCLs, we introduced *time to exit* in [10]—the time it takes for a system to violate its safety constraints if left unattended.

This work is supported in part by NSF grants CNS-1239037, CNS-1446298 and ECCS-1553873, and DARPA grant N66001-14-1-4045.

The authors are with the Dept. of Electrical & Computer Engineering, University of Michigan, Ann Arbor, Michigan, U.S.A. E-mails: {pettni, necmiye}@umich.edu

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

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

In the following we assume that $T_f(x) < +\infty$, i.e., that the set S is transient under the flow of f. The dynamics of the time to exit are captured in the following result.

Proposition 1 ([10]): Assume that ∂S is C^1 at $U_f(x)$, where $U_f(x)$ is the *exit point*

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

Then the Lie derivative $\mathscr{L}_g T_f(x)$ of $T_f(x)$ with respect to a second vector field g is expressed as follows:

$$\mathscr{L}_g T_f(x) = -\frac{\left(\hat{n}^S_{U_f(x)}\right)^T \left.(\nabla_x \phi_f)\right|_{(x,T_f(x))} g(x)}{\left(\hat{n}^S_{U_f(x)}\right)^T f(U_f(x))}, \tag{1}$$

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

As intuition suggests, $\mathscr{L}_f T_f(x) = -1$: the time to exit with respect to f decreases with unit speed along trajectories of f itself. The expression in (1) is difficult to evaluate in general settings, but easy for the one-dimensional systems considered in this paper.

B. Problem statement

Problem 1: Let $[N]$ denote the set of integers $\{1,\ldots,N\}$. Consider a collection of N one-dimensional subsystems with states $x^i \in \mathbb{R}$, for $i \in [N]$, with dynamics

$$\frac{\mathrm{d}}{\mathrm{d}t}x^i(t) = \begin{cases} f^i_{\texttt{off}}\left(x^i(t)\right) & \text{if } \sigma^i(t) = \texttt{off}, \\ f^i_{\texttt{on}}\left(x^i(t)\right) & \text{if } \sigma^i(t) = \texttt{on}, \end{cases} \tag{2}$$

local safe sets $S^i = [\underline{a}^i, \overline{a}^i]$, and bounds $[\underline{K}_{\texttt{on}}, \overline{K}_{\texttt{on}}]$ on the number of subsystems that can simultaneously be in mode `on` (i.e. *global mode-counting bounds*). Given initial positions $\mathbf{x}_0 = [x^1_0, \ldots, x^N_0]$, compute a **time of invariance** $T_I(\mathbf{x}_0)$ and construct an aggregate switching policy $\{\sigma^i\}_{i\in[N]}$ such that for all trajectories $x^i : [0, T_I(\mathbf{x}_0)] \to \mathbb{R}$:

$$x^i(t) \in S^i, \quad \forall i \in [N], t \in [0, T_I(\mathbf{x}_0)],$$
$$\underline{K}_{\texttt{on}} \leq \sum_{i\in[N]} \mathbb{1}_{\{\texttt{on}\}}\left(\sigma^i(t)\right) \leq \overline{K}_{\texttt{on}}, \quad \forall t \in [0, T_I(\mathbf{x}_0)],$$

where $\mathbb{1}_X(x)$ is the indicator function of the set X that is equal to 1 if $x \in X$ and to 0 otherwise. □

We are interested in making $T_I(\mathbf{x}_0)$ as large as possible and assess how it compares to the optimal time of invariance $T^*_I(\mathbf{x}_0)$. We make an assumption on the signs of the vector fields, namely that the `off` field always transports the state "upwards", and that the `on` field transports the state "downwards", as depicted in Fig. 1.

Assumption 1: For $i \in [N]$ and all $x \in S^i$, $f^i_{\texttt{off}}(x) > 0$ and $f^i_{\texttt{on}}(x) < 0$.

With this assumption the systems can be thought of as air condition units operating in hot weather conditions: when they are turned `on` the temperature decreases, and, when they are turned `off` the temperature converges towards the warmer outdoor temperature.

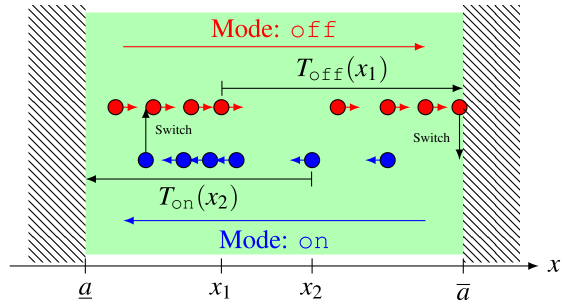

Fig. 1. Illustration of TCL system with homogeneous constraints $[\underline{a}, \overline{a}]$. The **time to `off` exit** $T_{\texttt{off}}(x)$ is the time it takes to reach the upper boundary $\overline{a}$ for a system in mode `off`, and conversely for $T_{\texttt{on}}(x)$.

In the following the notation is simplified and we write $\mathscr{L}^i_{\texttt{on}} = \mathscr{L}_{f^i_{\texttt{on}}}$, $T^i_{\texttt{on}} = T_{f^i_{\texttt{on}}}$, and analogously for the `off` mode. For the one-dimensional subsystems in Problem 1 the times to exit then satisfy [10]:

$$\mathscr{L}^i_{\texttt{off}} T^i_{\texttt{on}}(x) = -\frac{f^i_{\texttt{off}}(x)}{f^i_{\texttt{on}}(x)}, \quad \mathscr{L}^i_{\texttt{on}} T^i_{\texttt{off}}(x) = -\frac{f^i_{\texttt{on}}(x)}{f^i_{\texttt{off}}(x)}.$$

In our previous paper [10] we showed that if

$$\sum_{i\in[N]} \frac{\mathscr{L}^i_{\texttt{off}} T^i_{\texttt{on}}(\underline{a}^i)}{1 + \mathscr{L}^i_{\texttt{off}} T^i_{\texttt{on}}(\underline{a}^i)} > \underline{K}_{\texttt{on}}, \text{ and,} \tag{3a}$$

$$\sum_{i\in[N]} \frac{1}{1 + \mathscr{L}^i_{\texttt{on}} T^i_{\texttt{off}}(\overline{a}^i)} < \overline{K}_{\texttt{on}}, \tag{3b}$$

then $T^*_I(\mathbf{x}_0) = +\infty$ for any $\mathbf{x}_0$ s.t. $x^i_0 \in S^i$, and we provided a control strategy that achieves invariance indefinitely. On the other hand, we also gave a converse result based on the following assumption:

Assumption 2: The functions $f^i_{\texttt{on}}$ and $f^i_{\texttt{off}}$ are monotonically decreasing in S^i for all $i \in [N]$.

This assumption corresponds to flow velocities that decrease in the directions of the flows. Such a condition is typically true in a system that is converging towards its equilibrium. Under this assumption, if the strict version of either inequality in (3) is violated, then no control strategy can achieve indefinite invariance, so $T^*_I(\mathbf{x}_0) < +\infty$ for any $\mathbf{x}_0$. In this paper we study this latter case. If (3b) is strictly violated then (3a) holds, and vice versa. Due to symmetry we can therefore without loss of generality restrict attention to the case when (3b) is strictly violated:

Assumption 3: The following inequality holds:

$$\sum_{i\in[N]} \frac{1}{1 + \mathscr{L}^i_{\texttt{on}} T^i_{\texttt{off}}(\overline{a}^i)} > \overline{K}_{\texttt{on}}. \tag{4}$$

Going back to the air conditioner analogy, Assumption 3 represents a situation when $\sum_{i\in[N]} \frac{1}{1+\mathscr{L}^i_{\texttt{on}} T^i_{\texttt{off}}(\overline{a}^i)}$ air condition units must be allowed to turn `on` simultaneously in order to indefinitely maintain all temperatures below the upper limits. However, there is a too restrictive power consumption constraint $\overline{K}_{\texttt{on}}$ on the maximal number that can simultaneously be turned `on`, which means that it is impossible to satisfy the temperature constraints $x^i(t) \leq \overline{a}^i$ over an infinite horizon. In the rest of the paper we disregard the lower bounds $\underline{a}^i$ since they can be satisfied over an infinite horizon, and focus on satisfying $x^i \leq \overline{a}^i$ for as long as possible.

II. CONTROL STRATEGIES

We re-formulate Problem 1 as an optimal control problem, and leverage tools from optimal control theory to construct a control strategy C. The strategy is based on a lower bound of the optimal time of invariance and we show that C will always achieve a time of invariance greater than the lower bound. Moreover, we investigate how much the lower bound differs from the optimum.

A. Reformulation as an optimal control problem

We first introduce a convexified version of the dynamics (2), where the mode of a subsystem is described by a number $\alpha \in [0,1]$ rather than by the discrete action set $\{\texttt{off}, \texttt{on}\}$. The convexified dynamics are as follows:

$$\frac{\mathrm{d}}{\mathrm{d}t} x^i(t) = (1-\alpha^i(t)) f^i_{\texttt{off}}\left(x^i(t)\right) + \alpha^i(t) f^i_{\texttt{on}}\left(x^i(t)\right). \quad (5)$$

As can be seen $\alpha^i(t) = 0$ in (5) corresponds to $\sigma^i(t) = \texttt{off}$ in (2), and, conversely, $\alpha^i(t) = 1$ corresponds to $\sigma^i(t) = \texttt{on}$. Although the set of permissible control actions in (2) is a strict subset of that in (5), any trajectory of (5) can be tracked to arbitrary precision by the dynamics in (2) by sliding mode control and fast switching[1]. We let $\mathbf{x} = [x^1, \ldots, x^N]^T$ be the state of the aggregate system, and can write the aggregate dynamics as

$$\begin{aligned} &\frac{\mathrm{d}}{\mathrm{d}t}\mathbf{x}(t) = (\mathbf{1}_N - \alpha(t)) \odot \mathbf{f}_{\texttt{off}}(\mathbf{x}(t)) + \alpha(t) \odot \mathbf{f}_{\texttt{on}}(\mathbf{x}(t)), \\ &\mathbf{x}(0) = \mathbf{x}_0, \\ &\alpha : \mathbb{R}^+ \to \Delta_N\left(\overline{K}_{\texttt{on}}\right), \end{aligned} \quad (6)$$

where $\odot$ denotes element-wise product, $\mathbf{1}_N$ is the column vector of N ones, and the set of admissible controls $\Delta_N\left(\overline{K}_{\texttt{on}}\right)$ is defined as follows:

$$\Delta_N(K) = \left\{ \alpha \in \mathbb{R}^N \mid 0 \le \alpha^i \le 1, \sum_{i\in[N]} \alpha^i = K \right\}.$$

We point out that in the absence of lower state bounds, there is no benefit in selecting α such that $\sum_{i\in[N]} \alpha^i < \overline{K}_{\texttt{on}}$ whereby we may restrict control to $\Delta_N\left(\overline{K}_{\texttt{on}}\right)$. If we let $S = \Pi_{i=1}^N S^i$ be the aggregate safe set we can now associate $T_I^*(\mathbf{x}_0)$ to the value $V^*(\mathbf{x}_0)$ of an optimal control problem

$$V^*(\mathbf{x}_0) = \sup_{\alpha} \left\{ \int_0^\infty \mathbb{1}_S(\mathbf{x}(s))\, \mathrm{d}s \text{ s.t. (6) holds} \right\}.$$

It is obvious that $T_I^*(\mathbf{x}_0) \le V^*(\mathbf{x}_0)$; we conjecture that $T_I^*(\mathbf{x}_0) = V^*(\mathbf{x}_0)$ for all $\mathbf{x}_0 \in S$, but the applicability of our results is not affected by whether that conjecture is true or not. By applying the principle of optimality it follows that $V^*(\mathbf{x})$ is a viscosity solution [1, 2] of the following stationary partial differential equation (PDE):

$$\max_{\alpha\in\Delta_N\left(\overline{K}_{\texttt{on}}\right)} \sum_{i\in[N]} \frac{\partial V(\mathbf{x})}{\partial x^i} \begin{bmatrix} (1-\alpha^i) f^i_{\texttt{off}}(x^i) \\ +\alpha^i f^i_{\texttt{on}}(x^i) \end{bmatrix} = -\mathbb{1}_S(\mathbf{x}). \quad (7)$$

The left-hand side maximizes the time derivative of V which, along trajectories generated by selecting the maximizing α, should be equal to -1 inside S and to 0 otherwise.

[1] Sliding modes require a weak notion of solutions (e.g. Carathéodory) but in this paper we abstain from precise technical discussions.

Unfortunately this PDE is difficult to solve for the type of high-dimensional systems we are interested in. However, from (7) we know that the optimal control input α^* should be selected so as to achieve the maximum in (7). By disregarding the term that is constant in α, the optimal input can be written as the feedback

$$\alpha^*(\mathbf{x}) \in \underset{\alpha\in\Delta_N\left(\overline{K}_{\texttt{on}}\right)}{\arg\max} \sum_{i\in[N]} \alpha^i \frac{\partial V^*(\mathbf{x})}{\partial x^i} (f^i_{\texttt{on}}(x^i) - f^i_{\texttt{off}}(x^i)). \quad (8)$$

Motivated by this way to extract an input, we pursue a three-stage approach.

1) Construct an analytical under-approximation $T_I^-(\mathbf{x})$ of the optimal time of invariance $T_I^*(\mathbf{x})$.
2) Substitute $T_I^-(\mathbf{x})$ for $V^*(\mathbf{x})$ in (8) to construct a strategy C which achieves an unknown time of invariance $T_I^C(\mathbf{x}) > T_I^-(\mathbf{x})$.
3) Analytically bound the optimality gap between $V^*(\mathbf{x})$ and $T_I^-(\mathbf{x})$.

B. Under-approximation of the time of invariance

We derive a lower bound of the time of invariance for linear systems, the reason being that in order to accurately assess the time of invariance some future information must be taken into account. If a nonlinear system can be under-approximated by a linear system, in the sense that $T^i_{\texttt{off}}(x^i)$ is always smaller in the approximation, the results still hold. We also give a looser lower bound in the end of this section that holds without further assumptions on the dynamics. For now, we assume that

$$f^i_{\texttt{off}}(x^i) = \gamma^i\left(-x^i + \overline{b}^i\right), \quad f^i_{\texttt{on}}(x^i) = \gamma^i\left(-x^i + \underline{b}^i\right). \quad (9)$$

Here, $\underline{b}^i$ and $\overline{b}^i$ are the mode on and off equilibrium points of subsystem i, respectively, and Assumption 1 is equivalent to the condition $\underline{b}^i < \overline{a}^i < \overline{b}^i$, which we assume to be fulfilled in the following. In addition, Assumption 2 is fulfilled due to linearity.

For linear systems, the time to off exit given a constant input $\alpha^i \in [0,1]$ (the time it takes to reach $\overline{a}^i$, c.f. Fig. 1) is

$$\begin{aligned} T^i_{\texttt{off}}(x^i_0, \alpha^i) &= \frac{1}{\gamma^i} \log\left(\frac{\tilde{b}^i(\alpha^i) - x^i_0}{\tilde{b}^i(\alpha^i) - \overline{a}^i}\right) \\ &= \frac{1}{\gamma^i} \log\left(1 + \frac{\overline{a}^i - x^i_0}{\tilde{b}^i(\alpha^i) - \overline{a}^i}\right), \end{aligned} \quad (10)$$

where $\tilde{b}^i(\alpha^i) = \underline{b}^j\alpha^j + (1-\alpha^j)\overline{b}^j$ is the equilibrium point of (9) resulting from the input α^i. Using the inequality $\log(1+x) \ge 2x/(2+x)$ for $x \ge 0$, it follows that $T^i_{\texttt{off}}(x^i_0, \alpha^i)$ can be lower bounded by $T^{i-}_{\texttt{off}}(x^i_0, \alpha^i)$ defined as follows:

$$T^{i-}_{\texttt{off}}(x^i_0, \alpha^i) = \frac{1}{\gamma^i}\left(\frac{\overline{a}^i - x^i_0}{\overline{b}^i - \alpha^i(\overline{b}^i - \underline{b}^i) - \dfrac{\overline{a}^i + x^i_0}{2}}\right). \quad (11)$$

Our strategy to lower bound the time of invariance amounts to selecting a weight vector $\alpha \in \Delta_N\left(\overline{K}_{\texttt{on}}\right)$ that maximizes the minimal lower bound across all N subsystems. Then the strategy of applying α as a constant input until the first subsystem violates its constraint will achieve a time of invariance of at least $\min_{i\in[N]} T^{i-}_{\texttt{off}}(x^i_0, \alpha^i)$.

A reasonable approach would be to select a weight vector that results in lower bounds that are uniform across subsystems. However, if there are large disparities in the initial conditions such a selection is not possible with $\alpha \in \Delta_N\left(\overline{K}_{\mathtt{on}}\right)$—$\alpha^i$ can become negative for some subsystems. To arrive at a feasible α we use Algorithm 1 to iteratively sort out such subsystems before assigning weights to remaining subsystems that result in uniform lower bounds.

Algorithm 1: Selection of subsystem weights.

Data: N subsystem positions $x_0^i \in S^i$ for $i \in [N]$
Result: Weights $\alpha \in \Delta_N\left(\overline{K}_{\mathtt{on}}\right)$, index set $J^* \subset [N]$

1 Set $J = [N]$, $\alpha^i = -1 \; \forall \, i \in J$
2 **while** $\min_i \alpha^i < 0$ **do**
3   **forall** $i \in J$ **do**
4 $$\alpha^i = \left(\overline{K}_{\mathtt{on}} + \left(\frac{\overline{b}^i - \frac{1}{2}\left(\overline{a}^i + x_0^i\right)}{\overline{a}^i - x_0^i}\right) \sum_{j\in J} \frac{\gamma^i}{\gamma^j} \frac{\overline{a}^j - x_0^j}{\overline{b}^j - \underline{b}^j} - \sum_{j\in J} \frac{\overline{b}^j - \frac{1}{2}\left(\overline{a}^j + x_0^j\right)}{\overline{b}^j - \underline{b}^j}\right) \Bigg/ \left(\frac{\overline{b}^i - \underline{b}^i}{\overline{a}^i - x_0^i} \sum_{j\in J} \frac{\gamma^i}{\gamma^j} \frac{\overline{a}^j - x_0^j}{\overline{b}^j - \underline{b}^j}\right)$$
5     $J = J \setminus \{i \in J : \alpha^i < 0\}$
6   **forall** $i \in [N] \setminus J$ **do**
7         $\alpha^i = 0$
8 **return** $\alpha, J^* = J$

Since the index set J is strictly decreasing over algorithm iterations, the algorithm will terminate. Below we will prove that $T_I^-(\mathbf{x}_0, J^*)$ is a lower bound on the time of invariance resulting from the weight selection, where

$$T_I^-(\mathbf{x}_0, J) = \frac{\displaystyle\sum_{i\in J} \frac{\overline{a}^i - x_0^i}{\gamma^i\left(\overline{b}^i - \underline{b}^i\right)}}{\displaystyle\sum_{i\in J} \frac{\overline{b}^i - \frac{1}{2}\left(\overline{a}^i + x_0^i\right)}{\overline{b}^i - \underline{b}^i} - \overline{K}_{\mathtt{on}}}. \tag{12}$$

Proposition 2: In the inner **forall** loop α^is are assigned in a way such that $T^{i-}_{\mathtt{off}}(x_0^i, \alpha^i) = T_I^-(\mathbf{x}_0, J)$ for all $i \in J$.

Proof: Follows by inserting the expression for α^i on Line 4 into (11). ■

We also need to establish that the α^i's are feasible.

Proposition 3: Algorithm 1 terminates with a non-empty index set J^* and $\alpha \in \Delta_N\left(\overline{K}_{\mathtt{on}}\right)$.

Proof: It is easy to verify that $\sum_{i\in J} \alpha^i = \overline{K}_{\mathtt{on}}$ after line 4 in Algorithm 1, which implies that at least one α^i must be positive and hence J^* is non-empty at termination. It is also clear that $\alpha^i \geq 0$ for all $i \in [N]$ at termination. What remains to show is that $\alpha^i \leq 1$; we do this by establishing that

$$\overline{K}_{\mathtt{on}} < \sum_{i\in J} \frac{\overline{b}^i - \frac{1}{2}(\overline{a}^i + x_0^i)}{\overline{b}^i - \underline{b}^i} \tag{13}$$

is true at the beginning of each iteration, from which it follows that $\alpha^i \leq 1$ for all i by inspection of Line 4.

We now establish (13). First remark that, when $J = [N]$, (13) is a result of Assumption 3 and the fact that $x_0^i \leq \overline{a}^i < \overline{b}^i$. Assume for induction that (13) holds at the beginning of the **while** loop, we prove that it holds also for $J \setminus J'$ where $J' \subset J$ is the set of indices excluded on Line 5. For all $i \in J'$, $\alpha^i < 0$ and therefore by Proposition 2

$$T^{i-}_{\mathtt{off}}(x_0^i, 0) > T^{i-}_{\mathtt{off}}(x_0^i, \alpha^i) = T_I^-(\mathbf{x}_0, J), \tag{14}$$

from which the following can be derived by summing over $i \in J'$:

$$\frac{\displaystyle\sum_{i\in J'} \frac{\overline{a}^i - x_0^i}{\gamma^i\left(\overline{b}^i - \underline{b}^i\right)}}{\displaystyle\sum_{i\in J'} \frac{\overline{b}^i - \frac{1}{2}\left(\overline{a}^i + x_0^i\right)}{\overline{b}^i - \underline{b}^i}} > \frac{\displaystyle\sum_{j\in J} \frac{\overline{a}^j - x_0^j}{\gamma^j\left(\overline{b}^j - \underline{b}^j\right)}}{\displaystyle\sum_{j\in J} \frac{\overline{b}^j - \frac{1}{2}\left(\overline{a}^j + x_0^j\right)}{\overline{b}^j - \underline{b}^j} - \overline{K}_{\mathtt{on}}}. \tag{15}$$

This expression can be used to derive

$$\sum_{i\in J\setminus J'} \frac{\overline{b}^i - \frac{1}{2}\left(\overline{a}^i + x_0^i\right)}{\overline{b}^i - \underline{b}^i} > \overline{K}_{\mathtt{on}} + \sum_{i\in J'} \frac{\overline{b}^i - \frac{1}{2}\left(\overline{a}^i + x_0^i\right)}{\overline{b}^i - \underline{b}^i} \times \left[\frac{\sum_{i\in J} \frac{1}{\gamma^i}\left(\overline{a}^i - x_0^i\right)\left(\overline{b}^i - \underline{b}^i\right)}{\sum_{i\in J'} \frac{1}{\gamma^i}\left(\overline{a}^i - x_0^i\right)\left(\overline{b}^i - \underline{b}^i\right)} - 1\right],$$

from which (13) follows also for J replaced by $J \setminus J'$, by positivity of the term in brackets. This completes the induction argument. ■

Finally, we conclude that $T_I^-(\mathbf{x}_0, J^*)$ is a lower bound on the time of invariance achieved by the selection in Algorithm 1.

Theorem 1: Given a point $\mathbf{x}_0 \in S$, the strategy of assigning the constant α obtained in Algorithm 1 results in a time of invariance of at least $T_I^-(\mathbf{x}_0, J^*)$.

Proof: We have already established that $T_I^-(\mathbf{x}_0, J^*) = T^{i-}_{\mathtt{off}}(x_0^i, \alpha^i)$ for all $i \in J^*$. All subsystems $i \in [N] \setminus J^*$ have been assigned $\alpha^i = 0$, and we have that

$$T^{i-}_{\mathtt{off}}(x_0^i, 0) > T_I^-(\mathbf{x}_0, J)$$

for some $J \supset J^*$. We show that $T_I^-(\mathbf{x}_0, J)$ is decreasing over the **while** iterations, from which it follows that $T^{i-}_{\mathtt{off}}(x_0^i, 0) > T_I^-(\mathbf{x}_0, J^*)$ as required. Remark that the following two algebraic conditions are equivalent for $c_i \geq 0$ and $c_5 < c_3$:

$$\frac{c_2}{c_4} > \frac{c_1 + c_2}{c_3 + c_4 - c_5} \iff \frac{c_1 + c_2}{c_3 + c_4 - c_5} > \frac{c_1}{c_3 - c_5}.$$

This equivalence can be used to transform (15) (which is on the left-hand side form) into the inequality $T_I^-(\mathbf{x}_0, J) > T_I^-(\mathbf{x}_0, J \setminus J')$ (which can be written on the right-hand side form), which completes the proof. ■

We point out that the reverse of inequality (14) necessarily holds for all $i \in J^*$, that is,

$$i \in J^* \iff T^{i-}_{\mathtt{off}}(x_0^i, 0) \leq T_I^-(\mathbf{x}_0, J^*). \tag{16}$$

In the following we just write $T_I^-(\mathbf{x}_0) = T_I^-(\mathbf{x}_0, J^*)$ since J^* depends on $\mathbf{x}_0$ via Algorithm 1, and we also define $T_I^-(\mathbf{x}_0)$ to be zero outside of S:

$$T_I^-(\mathbf{x}_0) = 0 \quad \text{for } \mathbf{x}_0 \notin S.$$

C. An easy-to-implement strategy

We now use the under-approximation $T_I^-(\mathbf{x})$ to propose a strategy C that achieves a time of invariance of *at least* $T_I^-(\mathbf{x})$. Motivated by the selection of controls in (8) we propose to dynamically select feedback weights $\alpha^C(\mathbf{x}(t))$ such that

$$\alpha^C(\mathbf{x}(t)) \in \underset{\alpha\in\Delta_N\left(\overline{K}_{\mathtt{on}}\right)}{\arg\max} \sum_{i\in J^*} \frac{\partial T_I^-(\mathbf{x}(t))}{\partial x^i}\gamma^i \begin{bmatrix} -x^i(t)+\overline{b}^i \\ -\alpha^i(\overline{b}^i-\underline{b}^i) \end{bmatrix}. \quad (17)$$

Two subtleties occur here: a subsystem i only contributes to $T_I^-(\mathbf{x}(t))$ if $i \in J^*$, where J^* is determined by Algorithm 1 in a way such that (16) holds for $\mathbf{x}(t)$. Furthermore, $T_I^-(\mathbf{x}(t)) = 0$ if $\mathbf{x}(t) \notin S$. By incorporating these side conditions in the partial derivative[2], the following two rules must be taken into account when selecting the maximizing α^C in (17):

1) If $x^i(t) = \overline{a}^i$, then α^i must be such that
$$(1-\alpha^i)\overline{b}^i + \alpha^i\underline{b}^i \le \overline{a}^i \quad (18)$$
in order to prevent x^i from exiting S^i.
2) If $T^{i-}_{\mathtt{off}}(x^i(t),0) = T_I^-(\mathbf{x})$, then α^i must be such that
$$\frac{\mathrm{d}}{\mathrm{d}t}T^{i-}_{\mathtt{off}}(x^i(t),0) \le \frac{\mathrm{d}}{\mathrm{d}t}T_I^-(\mathbf{x}(t)) \quad (19)$$
in order to maintain i in the index set J^*.

Both of these conditions are linear in α and are therefore easy to write as a linear program. In addition, the α returned by Algorithm 1 is a feasible solution with objective value equal to -1. The feasibility is thus guaranteed, but it may be possible to find an α that achieves a larger objective value than -1, which may result in a time of invariance larger than $T_I^-(\mathbf{x}_0)$.

D. Relaxed "lazy" strategy

The strategy C will in general introduce sliding surfaces in the state space, which means that it generates trajectories with non-integer switching signals. Since a physical subsystem can only be in mode `on` or `off` this is impossible in practice. Although non-integer trajectories can be approximated by fast switching (chattering in the limit), such fast switching is not user-friendly and may cause premature equipment failure. For this reason it is appropriate to give up some time of invariance in favor of reducing the switching frequency.

In our previous work [10] (see also [8, 6]) we proposed a "lazy" switching strategy that only applies integer inputs 0 and 1 (corresponding to the physical modes `off` and `on`). Inspired by this we propose an event-based relaxed strategy L that always maintains $\overline{K}_{\mathtt{on}}$ system in mode `on`, and that only acts if either a safety constraint $x^i \in S^i$ is about to be violated, or if a subsystem i is about to exit the index set J^*.

"Lazy" strategy L: Initialize by switching `on` any $\overline{K}_{\mathtt{on}}$ subsystems in J^*. At time t only perform a switch if one of the following conditions occur:

1) If $x^i(t) \ge \overline{a}^i$, switch i to `on` while switching `off` subsystem j^* for
$$j^* \in \underset{j\in J_{\mathtt{on}}}{\arg\max}\, T^j_{\mathtt{off}}(x^j(t)),$$
where $J_{\mathtt{on}}$ is the set of subsystems in J^* that are in mode `on`.
2) If $T^{i-}_{\mathtt{off}}(x^i(t),0) \ge T_I^-(\mathbf{x}(t))$, switch i to `off` while switching `on` subsystem j^* for
$$j^* \in \underset{j\in J_{\mathtt{off}}}{\arg\min}\, T^j_{\mathtt{off}}(x^j(t)),$$
where $J_{\mathtt{off}}$ is the set of subsystems in J^* that are in mode `off`.

E. Myopic strategy for nonlinear systems

We finally revisit the nonlinear setting and propose a general strategy based on the lower bound

$$T_I^{\ominus}(\mathbf{x}) = \frac{\displaystyle\sum_{i\in J^*} \frac{T^i_{\mathtt{off}}(x^i)}{1+\mathscr{L}^i_{\mathtt{on}}T^i_{\mathtt{off}}(x^i)}}{\displaystyle\sum_{i\in J^*} \frac{1}{1+\mathscr{L}^i_{\mathtt{on}}T^i_{\mathtt{off}}(x^i)} - \overline{K}_{\mathtt{on}}}. \quad (20)$$

Just as in Algorithm 1, the set J^* should be determined recursively in a way such that $T_I^{\ominus}(\mathbf{x}) \ge T^i_{\mathtt{off}}(x^i)$ for all $i \in J^*$. The lower bound (20) is myopic in the sense that it only considers the current times to `off` exit $T^i_{\mathtt{off}}(x^i)$ and their instantaneous derivatives $\mathscr{L}^i_{\mathtt{on}}T^i_{\mathtt{off}}(x^i)$. However, as demonstrated later it can still achieve good results when used as a foundation in the switching rule (8). Also this strategy can be turned into a lazy strategy by disregarding optimization and just ensuring that the index set J^* remains the same.

III. Comparisons of times of invariance

In the following we study the value of $\frac{\mathrm{d}}{\mathrm{d}t}T_I^-(\mathbf{x}(t))$ along different trajectories by further investigating the left-hand side of (7). Roughly speaking, if there is no way to achieve a time derivative much *larger* than -1, then $T_I^-(\mathbf{x})$ should be close to the optimal value $V^*(\mathbf{x})$, and hence also to the optimal time of invariance $T_I^*(\mathbf{x})$. Conversely, if even a non-optimizing strategy such as the lazy strategy does not result in a time derivative much *smaller* than -1, then that strategy should achieve a time of invariance close to $T_I^-(\mathbf{x})$. By expanding the partial derivative, as shown in the Appendix, the following expression is obtained for $\alpha \in \Delta_N\left(\overline{K}_{\mathtt{on}}\right)$ in the linear case.

$$\sum_{i\in J^*}\frac{\partial T_I^-(\mathbf{x})}{\partial x^i}\gamma^i\left[-x^i+\alpha^i\underline{b}^i+(1-\alpha^i)\overline{b}^i\right] = -1 + \frac{\displaystyle\sum_{i\in J^*}\frac{T_I^-(\mathbf{x})\gamma^i\left(\overline{b}^i-x^i\right)-\left(\overline{a}^i-x^i\right)}{2(\overline{b}^i-\underline{b}^i)} - \frac{T_I^-(\mathbf{x})\gamma^i\alpha^i}{2}}{\displaystyle\sum_{j\in J^*}\frac{\overline{b}^j-\frac{1}{2}(\overline{a}^j+x^j)}{\overline{b}^j-\underline{b}^j} - \overline{K}_{\mathtt{on}}}. \quad (21)$$

As a side note, we can here identify a simple switching rule for the strategy C, which simply is to maximize α^i for the subsystems with the smallest γ^i's while respecting (18)-(19).

The main theoretical results of this paper are based on upper and lower bounds on the expression in (21): an upper bounds gives information about how close $T_I^-(\mathbf{x})$ is to the optimal value function of (7), while a lower bound allows us to give a guaranteed performance bound for a "lazy" strategy that does not attempt to optimize.

[2] The boundary analysis can be done rigorously with Heaviside and Dirac delta functions, but we opted for a (hopefully) more readable presentation.

Proposition 4: Provided that $\gamma^i T_I^-(\mathbf{x}_0) \le 2$ for all i, then (21) is bounded above by $-1+\overline{\varepsilon} T_I^-(\mathbf{x})$ along all trajectories starting at $\mathbf{x}_0$ regardless of α, where $\overline{\varepsilon}$ is as follows:

$$\overline{\varepsilon} = \frac{\frac{1}{2}\left(\sum_{i\in J^*}\gamma^i \frac{\overline{b}^i - \overline{a}^i}{\overline{b}^i - \underline{b}^i}\frac{2+\gamma^i T_I^-(\mathbf{x}_0)}{2-\gamma^i T_I^-(\mathbf{x}_0)} - \min_{\alpha\in\Delta_N(\overline{K}_{\mathrm{on}})}\sum_{i\in J^*}\alpha^i\gamma^i\right)}{\sum_{i\in J^*}\frac{\overline{b}^i - \overline{a}^i}{\overline{b}^i - \underline{b}^i} - \overline{K}_{\mathrm{on}}}.$$

The denominator in this bound is large when there is a large gap in (4). A larger gap implies a smaller time of invariance, meaning that the bound is tighter when the time of invariance is small.

Proof: First we remark that $-(a^i - x^i) \le 0$, so that term in (21) can be disregarded. Furthermore, from (16) we get for $\gamma^i T_I^-(\mathbf{x}) \le 2$

$$\overline{b}^i - x^i \le (\overline{b}^i - \overline{a}^i)\left(\frac{2+\gamma^i T_I^-(\mathbf{x})}{2-\gamma^i T_I^-(\mathbf{x})}\right). \tag{22}$$

Since the mapping $x \mapsto (2+x)/(2-x)$ is increasing with x on $[0,2)$, and we know that $T_I^-(\mathbf{x})$ decreases with time, we can replace $T_I^-(\mathbf{x})$ with $T_I^-(\mathbf{x}_0)$. Inserting these inequalities into (21) and applying (13) concludes. ■

Secondly, we give a bound from below which corresponds to an actor making the "worst possible" choices.

Proposition 5: Provided that $\gamma^i T_I^-(\mathbf{x}_0) \le 2$ for all i, then (21) is bounded below by $-1-|\underline{\varepsilon}| T_I^-(\mathbf{x})$ along all trajectories starting at $\mathbf{x}_0$ regardless of α, where $\underline{\varepsilon}$ is as follows:

$$\underline{\varepsilon} = \frac{\frac{1}{2}\left(\sum_{i\in J^*}\gamma^i \frac{\overline{b}^i - \overline{a}^i}{\overline{b}^i - \underline{b}^i}\frac{-\gamma^i T_i^-(\mathbf{x}_0)}{2-\gamma^i T_I^-(\mathbf{x}_0)} - \max_{\alpha\in\Delta_N(\overline{K}_{\mathrm{on}})}\sum_{i\in J^*}\alpha^i\gamma^i\right)}{\sum_{i\in J^*}\frac{\overline{b}^i - \overline{a}^i}{\overline{b}^i - \underline{b}^i} - \overline{K}_{\mathrm{on}}}.$$

Proof: Follows like above by noting that $\overline{b}^i - \overline{x}^i \ge \overline{b}^i - \overline{a}^i$ and deriving the following inequality from (16):

$$-(\overline{a}^i - x^i) \ge -(\overline{b}^i - \overline{a}^i)\frac{\gamma^i T_I^-(\mathbf{x})}{1-\gamma^i T_I^-(\mathbf{x})/2}. \tag{23}$$

■

The reason that these results do not hold for $\gamma^i T_I^-(\mathbf{x}_0) > 2$ can be traced back to the approximation $\log(1+x) \ge 2x/(2+x)$ that we used to lower bound the times to exit. When $x\to\infty$, $2x/(2+x)\to 2$ so the approximation becomes "non-invertible" for $\log(1+x)\ge 2$.

Next we prove that the strategy C achieves a time of invariance of at least $T_I^-(\mathbf{x})$. The strategy C *maximizes* the expression in (21), and per the discussion in Section II-C an α will always exists such that $\frac{\mathrm{d}}{\mathrm{d}t}T_I^-(\mathbf{x}(t)) \ge -1$ along trajectories of C.

Proposition 6: The time of invariance achieved by C, $T_I^C(\mathbf{x})$, satisfies

$$T_I^C(\mathbf{x}) \ge T_I^-(\mathbf{x}).$$

Proof: Consider a trajectory $\mathbf{x}^C(t)$ starting at $\mathbf{x}_0$ generated by C, we know that $\frac{\mathrm{d}}{\mathrm{d}t}T_I^-(\mathbf{x}^C(t)) \ge -1$. Integrating up to a time τ gives $T_I^-(\mathbf{x}^C(\tau)) - T_I^-(\mathbf{x}_0) \ge -\tau$. Set $\tau = T_I^C(\mathbf{x}_0)$, then $T_I^-(\mathbf{x}^C(\tau)) = 0$ and the result follows. ■

Finally, we prove a general two-sided bound that allows us to bound T_I^L and T_I^* in terms of T_I^-.

Proposition 7: Consider a a strategy D such that along trajectories $\mathbf{x}^D(t)$ generated by D,

$$\frac{\mathrm{d}}{\mathrm{d}t}T_I^-(\mathbf{x}^D(t)) \lesseqgtr -1 + \varepsilon T_I^-(\mathbf{x}^D(t)), \tag{24}$$

where $\lesseqgtr$ is either $\le$ or $\ge$. Then the time of invariance achieved by D, $T_I^D(\mathbf{x})$, satisfies

$$T_I^-(\mathbf{x}) \gtreqless \frac{1-e^{-\varepsilon T_I^D(\mathbf{x})}}{\varepsilon}.$$

Proof: Integrating (24) up to time τ along a trajectory $\mathbf{x}^D(t)$ starting at $\mathbf{x}_0$, we get

$$-e^{-\varepsilon\tau}T_I^-(\mathbf{x}^D(\tau)) - T_I^-(\mathbf{x}_0) \lesseqgtr \frac{e^{-\varepsilon\tau}-1}{\varepsilon^C}.$$

Setting $\tau = T_I^D(\mathbf{x}_0)$ gives $T_I^-(\mathbf{x}^D(\tau)) = 0$ and thus

$$T_I^-(\mathbf{x}_0) \gtreqless \frac{1-e^{-\varepsilon T_I^D(\mathbf{x}_0)}}{\varepsilon}.$$

■

Considering Proposition 4 and 5 the following two corollaries now immediately follow:

Corollary 1: Under the condition in Proposition 4, the lower approximation of time of invariance $T_I^-(\mathbf{x})$ satisfies

$$T_I^-(\mathbf{x}) \ge \frac{1-e^{-\overline{\varepsilon}T_I^*(\mathbf{x})}}{\overline{\varepsilon}}.$$

Corollary 2: Under the condition in Proposition 5, the lazy strategy L achieves a time of invariance $T_I^L(\mathbf{x})$ of at least

$$T_I^L(\mathbf{x}) \ge \frac{1}{|\underline{\varepsilon}|}\log(1+|\underline{\varepsilon}|T_I^-(\mathbf{x})).$$

To summarize we have shown that

$$\frac{1-e^{-\overline{\varepsilon}T_I^*(\mathbf{x}_0)}}{\overline{\varepsilon}} \le T_I^-(\mathbf{x}) \le T_I^C(\mathbf{x}) \le T_I^*(\mathbf{x}),$$

and, in addition, that the lazy strategy is guaranteed to achieve a time of invariance expressed in terms of $T_I^-(\mathbf{x})$.

Note that the results in this section are independent of the particular choice of $T_I^-(\mathbf{x})$ in Section II-B: any lower bound $\hat{T}_I(\mathbf{x})$ of the time of invariance defines a strategy through (17), and we can bound how far $\hat{T}_I(\mathbf{x})$ is from the optimal by assessing how close it is to satisfying the Hamilton-Jacobi-Bellman equation (7). An advantage of the under-approximation $T_I^-(\mathbf{x})$ used in this paper is that it is cheap to compute and results in a simple decision rule, which facilitates implementation.

IV. EXAMPLES

We now illustrate the results developed in the previous sections on collections of TCLs. We employ the model proposed in e.g. [6] where the dynamics of an individual TCL with state θ^i are

$$\frac{\mathrm{d}}{\mathrm{d}t}\theta^i(t) = -\gamma^i(\theta^i(t)-\theta_a^i) - b^i P_m^i \times \mathbb{1}_{\{\mathtt{on}\}}\left(\sigma^i(t)\right).$$

Equivalently, the system can be written on the previous form as

$$\frac{\mathrm{d}}{\mathrm{d}t}\theta^i(t) = \begin{cases}\gamma^i\left(-\theta^i(t)+\overline{b}^i\right) & \text{if } \sigma^i(t) = \mathtt{off},\\ \gamma^i\left(-\theta^i(t)+\underline{b}^i\right) & \text{if } \sigma^i(t) = \mathtt{on},\end{cases}$$

for $\overline{b}^i = \theta_a^i$ and $\underline{b}^i = \theta_a^i - b^i P_m^i$.

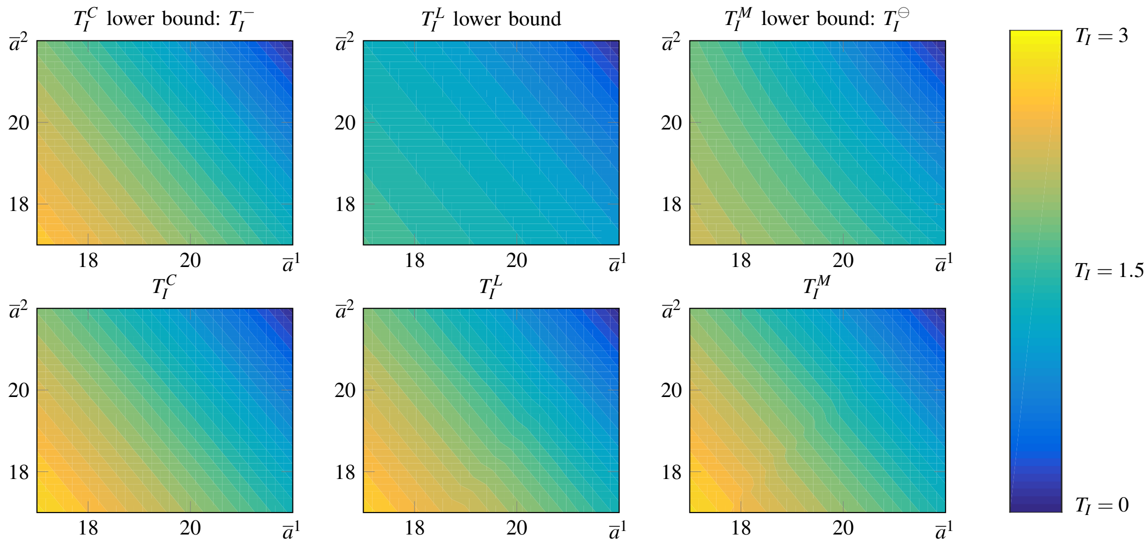

Fig. 2. Illustration of how the different times of invariance compare on an example with two subsystems. The data for the upper plots is analytic while the achieved times of invariance in the lower plots come from simulations of the strategies C, L (lazy) and M (myopic), respectively. As can be seen, all strategies achieved times of invariance above the guaranteed lower bounds, and the bounds are tighter for low times of invariance.

A. Illustrations of times of invariance in two dimensions

Fig. 2 illustrates how the different times of invariance vary over the state space for an example with two subsystems[3]. The plots show the analytic times of invariance $T_I^-(\mathbf{x})$ and $T_I^\ominus(\mathbf{x})$ that serve as guaranteed lower bounds for strategies C and the myopic strategy M based on the alternative lower bound (20), respectively. Moreover, from Corollary 2 we can obtain a guaranteed lower bound on the time of invariance achieved by the lazy strategy L. The guarantee is relatively strong for times of invariance less than 1, but for longer times of invariance it is quite conservative.

B. High-dimensional example

Next we consider a collection of 10,000 air condition units with randomized parameter values[4]. For this particular configuration, our previous results [10] imply that at least 6,391 subsystems must be allowed to be turned on simultaneously in order to keep all temperatures θ^i below the upper bounds $\overline{a}^i$ indefinitely; by contrast we imposed $\overline{K}_{\mathrm{on}} = 2000$ which implies $T_I^*(\theta_0) < +\infty$ for any initial condition θ_0. Computing the actual value of $T_I^*(\theta_0)$ requires solving a constrained 10,000-dimensional optimal control problem for each given initial condition. In contrast, our method gives a closed-form control strategy that, although not optimal, has guaranteed performance bounds which makes it useful for planning.

For this initial condition the under-approximation of the time of invariance is $T_I^-(\theta_0) = 0.95h$. Using Corollary 2 the time of invariance achieved by the "lazy" strategy is bounded below by $0.86h$, but the actual time of invariance achieved was significantly better at $T_I^L(\theta_0) = 0.96h$, as shown in Figs. 3 and 4. In addition, Corollary 1 guarantees that the optimal time of invariance is bounded as $T_I^*(\theta_0) \leq 1.1h$. The average number of switches per subsystem was 3.6 with no system switching more than 10 times, which showcases the practicality of the "lazy" strategy.

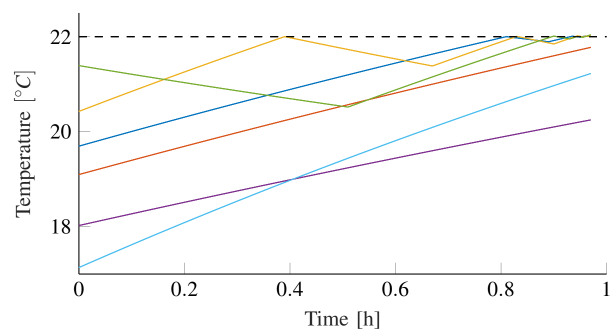

Fig. 3. Selected trajectories of six subsystems during a simulation of the lazy strategy; the states of all subsystems remain below the upper bound $\overline{a}^i = 22°C$ up until the achieved time of invariance $T_I^L(\theta_0) = 0.96h$.

For comparison, naïvely using the strategy in e.g. [6] achieved a time of invariance of 0.67h, or 30% shorter, for this example. The difference between the strategy in this paper and our earlier "lazy" strategy in [10] that accounts for heterogeneity is more subtle: that strategy achieves the same time of invariance on this example but may perform worse for certain initial conditions that lead to violations of 2) in Section II-D.

V. CONCLUSIONS

In this paper we proposed a method to approximately solve a very high-dimensional but structured optimal control problem. Our proposed method is on closed form and has guaranteed performance bounds, which makes it useful for planning; in addition we bounded the worst-case optimality gap. For practical purposes we also suggested a lazy version of the control strategy, which allows to trade off performance against a low switching frequency.

Current work is focused on tightening the theoretical guarantees, as well as incorporating the method in a higher-level planning framework that uses the theoretical guarantees

[3]The parameters in the example are $\gamma^1 = 0.25°C/h$, $\overline{a}^1 = 22°C$, $\overline{b}^1 = 35°C$, $\underline{b}^1 = 17°C$, $\gamma^2 = 0.3°C/h$, $\overline{a}^2 = 22°C$, $\overline{b}^2 = 38°C$, $\underline{b}^2 = 16°C$.

[4]Parameters were sampled uniformly: $\gamma^i \in [0.125°C/h, 0.375°C/h]$, $\overline{a}^i = 22°C$, $\theta_0^i \in [\overline{a}^i - 5°C, \overline{a}^i]$, $\overline{b}^i \in [32°C, 35°C]$, and $\underline{b}^i \in [14°C, 17°C]$.

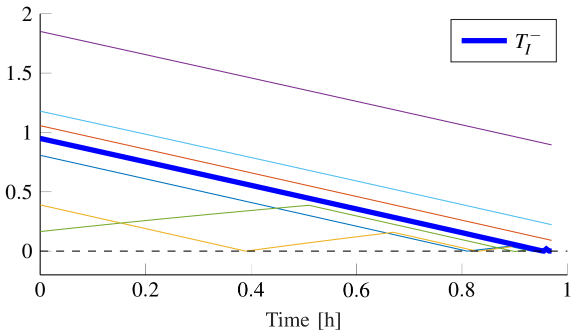

Fig. 4. Illustration of how the estimated times to exit vary during the simulation. The lower bound (14) of time to `off` exit is plotted for six different subsystems, along with the lower bound on time of invariance $T_I^-(\theta(t))$. The trajectories above $T_I^-(\theta(t))$ belong to subsystems that are not part of the index set J^*: even when turned `off` they will satisfy their safety constraint for at least a duration $T_I^-(\theta(t))$.

to distribute loads in a clever fashion. We are also interested in enhancements to provide a minimal dwell time guarantee, and in a better understanding of the trade off between minimal dwell time and time of invariance. Finally, we believe that similar ideas can be relevant in other applications as a strategy to counteract large but short-lived disturbance signals.

References

[1] M Bardi and I Capuzzo-Dolcetta. *Optimal Control and Viscosity Solutions of Hamilton-Jacobi-Bellman Equations*. Springer, 1997. DOI: `10.1007/978-0-8176-4755-1`.

[2] M G Crandall and P-L Lions. "Viscosity solutions of Hamilton-Jacobi equations". In: *Trans. American Mathematical Society* 277.1 (1983), pp. 1–42. DOI: `10.1090/S0002-9947-1983-0690039-8`.

[3] S J Crocker and J L Mathieu. "Adaptive state estimation and control of thermostatic loads for real-time energy balancing". In: *Proc. ACC*. 2016, pp. 3557–3563. DOI: `10.1109/ACC.2016.7525465`.

[4] S Esmaeil Zadeh Soudjani and A Abate. "Aggregation and control of populations of thermostatically controlled loads by formal abstractions". In: *IEEE Trans. Control Systems Technology* 23.3 (2015), pp. 975–990. DOI: `10.1109/TCST.2014.2358844`.

[5] A Ghaffari, S Moura, and M Krstić. "Modeling, Control, and Stability Analysis of Heterogeneous Thermostatically Controlled Load Populations Using Partial Differential Equations". In: *ASME. J. Dyn. Sys., Meas., Control.* 137.10 (2015). DOI: `10.1115/1.4030817`.

[6] H Hao et al. "Aggregate Flexibility of Thermostatically Controlled Loads". In: *IEEE Trans. Power Systems* 30.1 (2015), pp. 189–198. DOI: `10.1109/TPWRS.2014.2328865`.

[7] J L Mathieu, S Koch, and D S Callaway. "State estimation and control of electric loads to manage real-time energy imbalance". In: *IEEE Trans. on Power Systems* 28.1 (2013), pp. 430–440. DOI: `10.1109/TPWRS.2012.2204074`.

[8] T X Nghiem et al. "Green scheduling of control systems for peak demand reduction". In: *Proc. IEEE CDC*. 2011, pp. 5131–5136. DOI: `10.1109/CDC.2011.6161164`.

[9] P Nilsson and N Ozay. "Control Synthesis for Large Collections of Systems with Mode-Counting Constraints". In: *Proc. HSCC*. 2016, pp. 205–214. DOI: `10.1145/2883817.2883831`.

[10] P Nilsson and N Ozay. "On a class of maximal invariance inducing control strategies for large collections of switched systems". In: *Proc. HSCC*. 2017, pp. 187–196. DOI: `10.1145/3049797.3049810`.

[11] B M Sanandaji, H Hao, and K Poolla. "Fast regulation service provision via aggregation of thermostatically controlled loads". In: *Proc. HICSS*. 2014, pp. 2388–2397. DOI: `10.1109/HICSS.2014.300`.

[12] D M Stipanović, C J Tomlin, and G Leitmann. "Monotone Approximations of Minimum and Maximum Functions and Multi-objective Problems". In: *Applied Mathematics & Optimization* 66.3 (2012), pp. 455–473. DOI: `10.1007/s00245-012-9179-8`.

Appendix

Here we derive the expression for $\frac{\mathrm{d}}{\mathrm{d}t}T_I^-(\mathbf{x})$ in Section II-C. First remark that $T_I^-(\mathbf{x})$ only depends on subsystems in the index set J^* this $\frac{\partial T_I^-(\mathbf{x})}{\partial x^i} = 0$ for $i \notin J^*$. By differentiating (12) we obtain for $i \in J^*$:

$$\frac{\partial T_I^-(\mathbf{x})}{\partial x^i} = \frac{-\dfrac{1}{\gamma^i\left(\overline{b}^i - \underline{b}^i\right)} + \dfrac{T_I^-(\mathbf{x})}{2\left(\overline{b}^i - \underline{b}^j\right)}}{\displaystyle\sum_{j\in J^*} \frac{\overline{b}^j - \frac{1}{2}\left(\overline{a}^j + x^j\right)}{\overline{b}^j - \underline{b}^j} - \overline{K}_{\mathrm{on}}}.$$

Thus,

$$\begin{aligned}
\frac{\mathrm{d}}{\mathrm{d}t}T_I^-(\mathbf{x}) &= \sum_{i\in J^*} \frac{\partial T_I^-(\mathbf{x})}{\partial x^i}\gamma^i\left[-x^i + \alpha^i\underline{b}^i + (1-\alpha^i)\overline{b}^i\right] \\
&= \frac{\displaystyle\sum_{i\in J^*}\left(-\frac{1}{\overline{b}^i - \underline{b}^i} + \frac{\gamma^i T_I^-(\mathbf{x})}{2(\overline{b}^i - \underline{b}^i)}\right)\left[-x^i + \overline{b}^i - \alpha^i(\overline{b}^i - \underline{b}^i)\right]}{\displaystyle\sum_{j\in J^*}\frac{\overline{b}^j - \frac{1}{2}\left(\overline{a}^j + x^j\right)}{\overline{b}^j - \underline{b}^j} - \overline{K}_{\mathrm{on}}} \\
&= \left(\sum_{i\in J^*} - \left(\frac{\overline{b}^i - \frac{1}{2}\left(\overline{a}^i + x^i\right)}{\overline{b}^i - \underline{b}^i} - \alpha^i\right) + \frac{x^i - \frac{1}{2}\left(\overline{a}^i + x^i\right)}{\overline{b}^i - \underline{b}^i}\right. \\
&\quad \left. + \frac{\gamma^i T_I^-(\mathbf{x})}{2(\overline{b}^i - \underline{b}^i)}\left[-x^i + \overline{b}^i\right] - \frac{\gamma^i T_I^-(\mathbf{x})\alpha^i}{2}\right) \\
&\quad \Bigg/ \left(\sum_{j\in J^*}\frac{\overline{b}^j - \frac{1}{2}\left(\overline{a}^j + x^j\right)}{\overline{b}^j - \underline{b}^j} - \overline{K}_{\mathrm{on}}\right) \\
&= -1 + \frac{\displaystyle\sum_{i\in J^*}\frac{T_I^-(\mathbf{x})\gamma^i\left(\overline{b}^i - x^i\right) - \left(\overline{a}^i - x^i\right)}{2(\overline{b}^i - \underline{b}^i)} - \frac{T_I^-(\mathbf{x})\gamma^i\alpha^i}{2}}{\displaystyle\sum_{j\in J^*}\frac{\overline{b}^j - \frac{1}{2}(\overline{a}^j + x^j)}{\overline{b}^j - \underline{b}^j} - \overline{K}_{\mathrm{on}}}.
\end{aligned}$$