Large-Scale Invariant Sets for Safe Coordination of Thermostatic Loads

Sunho Jang, Necmiye Ozay, and Johanna L. Mathieu

***Abstract*— Systems often face constraints at multiple levels. For example, in coordinating a collection of thermostatically controlled loads to provide grid services, the controller must ensure temperature constraints for each load (local constraints) and distribution network constraints (global constraints) are satisfied. In this paper, we leverage invariant sets to ensure safe coordination of systems with both local and global constraints. Specifically, we develop a method for constructing a controlled invariant set for a collection of subsystems, modeled as transition systems, to ensure they indefinitely satisfy the constraints, based on cycles in individual transition systems. Then, we develop a control algorithm that keeps the state inside the maximal controlled invariant set. We apply these algorithms to a demand response problem, specifically, the tracking of a power trajectory (e.g., a frequency regulation signal) by a population of homogeneous air conditioners. The algorithm simultaneously maintains local temperature requirements and aggregate power consumption limits, ensuring the control is nondisruptive to consumers and benign to the distribution network.**

I. Introduction

In cyber-physical systems with many subsystems, there are local constraints for the safety of the individual subsystems as well as global constraints on their collective behavior. By considering the collection of subsystems as one large monolithic system, the problem of safe control synthesis can be rendered as the problem of computation of an invariant set. For each state in an invariant set, there exists a control input that guarantees the next state remains in this invariant set. Therefore, it is possible to guarantee constraint satisfaction indefinitely [1]–[3]. There has been some recent interest in scalable algorithms for computing invariant sets (see e.g., [4], [5] for the case of linear systems). However, these results are not applicable to very large scale systems with control inputs restricted to discrete switches or modes.

In this paper, we develop an approach to compute invariant sets for a system consisting of a large number of switched subsystems subject to constraints on subsystems' states and on the number of subsystems in certain modes. A variant of this problem was recently studied in [6], where an algorithm for finding open-loop periodic switching sequences that guarantee constraint satisfaction is proposed. In contrast to open-loop control, in this paper we show how to construct *implicitly defined* invariant sets and how to incorporate them into a Model Predictive Control (MPC) algorithm allowing us to optimize additional control objectives while ensuring safety and recursive feasibility. We begin by constructing an abstracted system that is bisimilar to the original system, and we show that a controller safe for the abstracted system is also safe for the original system. Then, we construct a system consisting of a large number of homogeneous subsystems and synthesize a controlled invariant set and a safe controller.

Our formulation and approach are motivated by the problem of controlling large numbers of thermostatically controlled loads (TCLs), such as residential air conditioners and water heaters, to provide services to the electric power system. TCLs maintain a temperature within a narrow range by switching on and off. Hundreds to thousands of TCLs can be coordinated to provide power system services by collectively tracking a power consumption signal to help the grid balance supply and demand on timescales of seconds to minutes. A number of papers (e.g., [7]–[9]) use models of the dynamics of large collections of TCLs to synthesize reference tracking controllers that keep TCL temperatures within prescribed ranges; this is often referred to as non-disruptive control. However, TCL collections may face additional constraints, e.g., on the power consumption of the aggregation (or portions of the aggregation) to ensure that TCL control actions do not cause distribution network constraint violations [10]. There are a variety of approaches to include distribution network constraints when controlling distributed energy resources [11]–[17]; however, these approaches do not provide formal safety guarantees (i.e., ensure constraint satisfaction) considering TCL dynamics and constraints. Ref. [18] uses ideas from formal methods, specifically [19], to constrain TCL aggregate variability, but does not consider reference tracking or supervision.

The contributions of this paper are threefold. First, we develop a method for generating a controlled invariant set of a transition system. This method selects some cycles of a transition system and finds a set of states derived from safe circulation around the selected cycles. Second, we develop a control algorithm using the controlled invariant set. By keeping the state inside the maximal controlled invariant set, the algorithm's feasibility is recursively guaranteed. Third, we apply the control algorithm to a TCL power tracking problem with constraints on the temperature of each TCL and power consumption of the aggregation. We benchmark our approach against three other approaches.

This paper is organized as follows. The notation is given next and the problem setting is described in Section II. Section III derives the abstraction of the original subsystem and the aggregate system. In Section IV, we develop a method for construction of a controlled invariant set of a collection of subsystems and, in Section V, we synthesize a safe control algorithm. We test our approach through simulations in Section VI and conclude in Section VII.

This work was supported by US NSF Grant CNS-1837680. The authors are with the Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI 48109 USA {sunhoj,necmiye,jlmath}@umich.edu.

Notation: The set of non-negative integers is denoted by $\mathbb{N}_0 = \mathbb{N} \cup \{0\}$, $[X]$ denotes the set of the integers $\{1, \ldots, X\}$, and $[X]_0$ denotes $[X] \cup \{0\}$. The indicator function on set A is denoted by $\mathbf{1}_A$. We write the Minkowski sum as $A \oplus B = \{a + b \mid a \in A, b \in B\}$, while the subtraction $A \ominus B$ is defined as the largest solution to $X \oplus B = A$. We write the infinity norm as $\|\cdot\|$. We denote the ball with radius r centered at θ by $\mathcal{B}(\theta, r) := \{x \mid \|x - \theta\| \leq r\}$. The identity function in space $\mathbb{R}^d$ is denoted by $\mathrm{Id}_{\mathbb{R}^d}$. The least common multiplier of $l_1, \ldots, l_n$ is written as $\mathrm{lcm}(l_1, \ldots, l_n)$.

II. Problem Setting

A. Overview

In this section, we describe the application-domain problem setting. Specifically, we consider a power reference tracking problem of a collection of TCLs under both local and global constraints. The aggregate power consumption of the collection should track a continuous reference signal, such as an automatic generation control signal (normalized between -1 and 1) scaled by the power capacity of the TCLs (i.e., the range within which the reference signal may vary). However, each TCL can only be switched on/off. Further, each TCL should maintain its temperature within a prescribed range, referred to as a dead-band. In addition, we assume the collection has a global constraint that the aggregate power consumption of all TCLs is bounded by prescribed upper/lower bounds to avoid distribution network problems, such as over/undervoltages and transformer overloading. Under these constraints and given the power reference signal, the problem of interest is to control the on/off modes of the TCLs to minimize the difference between the aggregate power and the reference.

This problem corresponds to a practical setting. The amount of power capacity that can be deployed safely varies as a function of distribution network loading since high (low) loading might restrict allowable increases (decreases) in TCL demand to avoid voltage issues. The power capacity must be committed to the ancillary services market in advance, and so an estimate is made based on load and renewables forecasts. However, the actual network-safe power bounds are not known until real-time. This paper explores cases in which more capacity is committed than deliverable in real-time because of poor forecasts, specifically, cases in which the reference signal leaves the range between the upper/lower power bounds needed to protect the distribution network. While it would seem that this contradiction between the reference signal and the power bounds could be resolved by the grid, in U.S. competitive electricity markets, the Independent System Operator procures and deploys ancillary services and the utility operates the distribution network, and the two entities do not yet coordinate to resolve such issues.

Fig. 1 provides a concrete example of the problem setting. The TCL aggregation has committed in advance to provide the power capacity between the dashed black lines; however, in real time, the TCL aggregate power is constrained between the solid black lines, e.g., to mitigate unforecasted voltage issues. The reference signal (blue) should be tracked as well as possible given the power bounds, and so the new target is the truncated reference (dotted yellow line). The simplest option would be to design a controller to track the truncated reference; however, there is no guarantee that this reference can be tracked perfectly while satisfying each TCL's temperature constraints. Another option could be to develop an MPC approach that requires every state in the horizon to belong to the constraint set; however, there is no guarantee that this MPC problem will remain feasible recursively. To overcome these issues, we augment the MPC with constraints coming from an implicit controlled invariant set, which guarantees both constraint satisfaction and recursive feasibility. We compare our approach with these two simpler approaches through numerical simulations in Section VI.

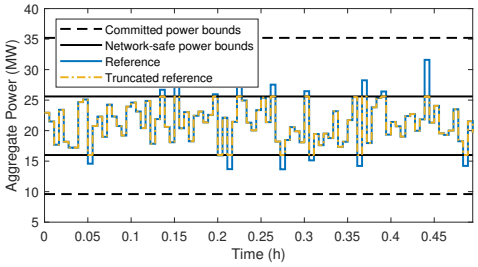

Fig. 1: Reference tracking with aggregate power constraints.

B. Assumptions, Models, and Formal Problem Statement

We assume i) TCLs are homogeneous in thermal/electrical parameters, experience the same ambient temperature, and are not affected by noise; ii) the ambient temperature is constant over time; iii) the upper/lower power bounds are constant over time; and iv) the temperature setpoint and dead-band are constant over time. The first two assumptions can be easily relaxed to handle mild heterogeneity and small deviations in ambient temperature (as explained at the end of Section III-A), but we keep them to simplify the exposition of the main results. The last two assumptions are required for our current formulation; we will explore ways to make the setting more realistic in future work.

The system of interest is constructed as follows. Let the number of TCLs in the system be N_{TCL} and the temperature of ith TCL at time t be $T_i(t)$ whose domain is $[T_{\text{low}}, T_{\text{up}}]$. We assume that temperature evolution of each TCL follows specific discrete-time dynamics, which depend on whether it is turned on or off. We write the on/off mode of ith TCL as $\mu_i(t)$. Then, we can represent the temperature evolution dynamics using the affine model developed in [20]

$$\begin{cases} T_i(t+1) = aT_i(t) + (1-a)(T_a - R_{\text{th}} p_{\text{tr}}) & \text{if } \mu_i(t) = \text{on} \\ T_i(t+1) = aT_i(t) + (1-a)T_a & \text{if } \mu_i(t) = \text{off}, \end{cases} \tag{1}$$

where T_a is the ambient temperature. Parameter a equals $\exp(-\Delta t/(R_{\text{th}} C_{\text{th}}))$, where Δt is the sampling time, R_{th} is the thermal resistance, and C_{th} is the thermal capacitance. Parameter p_{tr} represents the thermal energy transfer rate, which is positive for a cooling TCL and negative for a heating TCL.

If we denote the temperature dead-band by $[\underline{T}, \overline{T}]$ ($\overline{T} < T_{\text{up}}$, $\underline{T} > T_{\text{low}}$), the following becomes the local temperature constraint of each TCL

$$T_i(t) \in [\underline{T}, \overline{T}] \quad \forall t \in \mathbb{N}_0, i \in [N_{\text{TCL}}]. \tag{2}$$

The power consumption of each TCL in on mode, denoted p, equals p_{tr}/ζ, where ζ is the coefficient of performance. The aggregate power consumption $P_{\text{agg}}(t)$ equals $\sum_{i=1}^{N_{\text{TCL}}} p\mathbf{1}_{\text{on}}(\mu_i(t))$ and the aggregate power constraint is

$$\underline{P}_{\text{agg}} \leq \sum_{i=1}^{N_{\text{TCL}}} p\mathbf{1}_{\text{on}}(\mu_i(t)) \leq \overline{P}_{\text{agg}} \quad \forall t \in \mathbb{N}_0, \tag{3}$$

where $\overline{P}_{\text{agg}}, \underline{P}_{\text{agg}}$ are the upper and lower aggregate power bounds. Since, for homogeneous TCLs, the aggregate power is proportional to the number of on-mode TCLs, this constraint can be converted to

$$\underline{N}_{\text{on}} \leq N_{\text{on}}(t) \leq \overline{N}_{\text{on}} \quad \forall t \in \mathbb{N}_0, \tag{4}$$

where $N_{\text{on}}(t) = \sum_{i=1}^{N_{\text{TCL}}} \mathbf{1}_{\text{on}}(\mu_i(t))$ is the number of on-mode TCLs, and $\overline{N}_{\text{on}}, \underline{N}_{\text{on}}$ are equal to $\left\lfloor \frac{\overline{P}_{\text{agg}}}{p} \right\rfloor, \left\lceil \frac{\underline{P}_{\text{agg}}}{p} \right\rceil$.

Problem 1. *Given the system described above, synthesize a controller to choose $\mu_i(t)$ for all $i = 1, \ldots, N_{TCL}$ that guarantees the satisfaction of the constraints* (2)*,* (4)*, while trying to minimize the performance measure $|P_{agg}(t) - r(t)|$, where $r(t)$ is a reference signal.*

III. Abstraction and aggregation of a collection of switched subsystems

We now consider a general setup where the goal is to coordinate a collection of switched subsystems subject to safety constraints. We first define an original subsystem with continuous state, and construct a state-abstracted subsystem by discretizing the state space. Furthermore, we construct an aggregate system from the state-abstracted subsystems; we use the aggregate system in the following sections to synthesize controllers.

A. System models & abstraction

We start by introducing the *transition systems* formalism we use to model discrete-time dynamics [21].

Definition 1. *A transition system T is a tuple $(X, U, \rightarrow, Y)$, where X is a set of states, U a set of actions, $\rightarrow \subset X \times U \times X$ a transition relation, and $Y : X \rightarrow \mathbb{R}^n$ an output function.*

We denote $(x, u, x') \in \rightarrow$ as $x \xrightarrow{u} x'$ for short.

Definition 2. *Given a transition system $T = (X, U, \rightarrow, Y)$, a safe set $X_{safe} \subset X$, and input constraints $U_{safe} \subset U$, a set X_{inv} is a controlled invariant with respect to (T, X_{safe}, U_{safe}) if $X_{inv} \subset X_{safe}$ and for all $x \in X_{inv}$ there exists $u \in U_{safe}$ such that for all x' with $x \xrightarrow{u} x'$, we have $x' \in X_{inv}$. The union of all controlled invariant sets with respect to (T, X_{safe}, U_{safe}) is called the maximal controlled invariant set for (T, X_{safe}, U_{safe}).*

Definition 3. *Given N identical copies of a transition system $T = (X, U, \rightarrow, Y)$, the product transition system is given by $T^{\times N} = (X^N, U^N, \xrightarrow[\times N]{}, Y^{\times N})$, where $(x_1, \ldots, x_N) \xrightarrow[\times N]{(u_1, \ldots, u_N)} (x'_1, \ldots, x'_N)$ if and only if $x_i \xrightarrow{u_i} x'_i$ for all $i \in [N]$; and $Y^{\times N} : (x_1, \ldots, x_N) \mapsto (Y(x_1), \ldots, Y(x_N))$.*

We employ the following relation between transition systems that gives a notion of closeness between two systems. We later use this notion to construct finite representations for discrete-time dynamics.

Definition 4. *Two transition systems $T_1 = (X_1, U, \underset{1}{\rightarrow}, Y_1)$ and $T_2 = (X_2, U, \underset{2}{\rightarrow}, Y_2)$ are ϵ-approximately bisimilar if there exists a relation $R \subset X_1 \times X_2$ such that the sets $R(x_1) = \{x_2 : (x_1, x_2) \in R\}$ and $R^{-1}(x_2) = \{x_1 : (x_1, x_2) \in R\}$ are non-empty for all x_1, x_2, and such that for all $(x_1, x_2) \in R$, all of the following are satisfied.*

- $\|Y_1(x_1) - Y_2(x_2)\| \leq \epsilon$.
- *if $x_1 \xrightarrow[1]{u} x'_1$, there exists $x_2 \xrightarrow[2]{u} x'_2$ s.t. $(x'_1, x'_2) \in R$.*
- *if $x_2 \xrightarrow[2]{u} x'_2$, there exists $x_1 \xrightarrow[1]{u} x'_1$ s.t. $(x'_1, x'_2) \in R$.*

We consider a system which includes a collection of N homogeneous switched individual subsystems with M different modes. Let $\theta_i(t) \in \Theta$ be the state of the ith subsystem at time step t, where $\Theta \subset \mathbb{R}^d$ is a compact invariant domain. The state $\theta_i(t)$ obeys the difference equation

$$S^i : \ \theta_i(t+1) = f_{\mu_i(t)}(\theta_i(t)), \quad \mu_i : \mathbb{N}_0 \rightarrow [M], \tag{5}$$

where $\mu_i(t)$ is the mode of system i at time t and is the control input. For TCLs, θ_i corresponds to T_i and f_μ corresponds to the affine dynamics in (1). For convenience, we drop the index i in the rest of this subsection.

The discrete-time dynamics in (5) can be equivalently represented as a transition system

$$S = (\Theta, [M], \rightarrow, \text{Id}_{\mathbb{R}^d}), \tag{6}$$

where $\theta \xrightarrow{\mu} \theta'$ if and only if $\theta' = f_\mu(\theta)$. We want to construct another transition system, called an *abstraction* of S, that is approximately bisimilar to S and that has finitely many states.

To guarantee the existence of an abstraction that is approximately bisimilar to S, we make the following assumption, which also holds for the temperature evolution dynamics in (1).

Assumption 1. *For every $m \in [M]$, f_m is a local contraction, that is, there exist constants $L_m \in [0, 1)$, $c_m > 0$ such that*

$$\|f_m(\theta_1) - f_m(\theta_2)\| \leq L_m \|\theta_1 - \theta_2\| \tag{7}$$

for every $\theta_1, \theta_2 \in \Theta$ with $\|\theta_1 - \theta_2\| \leq c_m$.

For the original system S, we construct an abstraction by uniformly discretizing Θ. For a given grid size $\eta > 0$, we define an abstraction function $\gamma_\eta : \Theta \rightarrow \Theta \oplus \mathcal{B}(0, \eta/2)$ as

$$\gamma_\eta(\theta) = \eta \cdot \left\lfloor \frac{\theta}{\eta} \right\rfloor + \frac{\eta}{2}\mathbf{1}. \tag{8}$$

The inverse mapping γ_η^{inv} can also be defined for any $\Xi \subset \gamma_\eta(\Theta)$ as $\gamma_\eta^{\text{inv}}(\Xi) := \{\theta \in \Theta \mid \exists \xi \in \Xi \text{ s.t. } \gamma_\eta(\theta) = \xi\}$. Using this abstraction function γ_η, we define the state-abstracted system of S as

$$S_\eta = (\gamma_\eta(\Theta), [M], \underset{\eta}{\rightarrow}, \text{Id}_{\mathbb{R}^d}), \tag{9}$$

where $\xi \xrightarrow[\eta]{\mu} \xi'$ if and only if $\gamma_\eta(f_\mu(\xi)) = \xi'$. That is, the transition relation between the states of S_η is constructed by finding the closest grid point ξ' that will be reached at the next time step in a certain mode if the original system starts from a given grid point ξ.

The following lemma, a discrete-time variant of those in [22], [23], states a condition for S and S_η to be bisimilar.

Lemma 1. *If ϵ and η are such that $(1 - L_m)\epsilon \geq \eta/2$ for all $m \in [M]$, then S and S_η are $\epsilon-$approximately bisimilar.*

One important consequence of S and S_η being $\epsilon-$approximately bisimilar is that if S and S_η have the initial conditions $\theta(0)$ and $\xi(0) \doteq \gamma_\eta(\theta(0))$, respectively, and if they are driven by the same input sequence μ for all time, then their states θ and ξ remain ϵ close for all time, i.e., $\|\theta(t) - \xi(t)\| \leq \epsilon$ for all $t \in \mathbb{N}_0$.

A few remarks on Assumption 1 are in order. As shown in [22], for any continuous-time incrementally input-to-state stable system, there is a time discretization that ensures satisfaction of this assumption. Also note that, by Assumption 1, for any $0 < \epsilon \leq \min_{m\in[M]} c_m$, there is a small enough $\eta > 0$ that satisfies the condition of Lemma 1. Moreover, if the contraction in each f_m is strong enough, it is possible to handle disturbances (e.g., different ambient temperatures in the TCL problem) or mild parameter variations in the dynamics as in [6].

B. Collection of homogeneous subsystems with safety constraints

Now, consider a collection of homogeneous subsystems subject to both local and global constraints similar to (2) and (4). Suppose that we have N switched subsystems (which corresponds to N_{TCL} in the TCL problem), each modeled with an identical transition system S as in (6). Then the collection can be represented with the product transition system $S^{\times N}$, together with an abstract version $S_\eta^{\times N}$. We denote the state and input of $S_\eta^{\times N}$ by $\boldsymbol{\theta} := (\theta_1, \ldots, \theta_N)$ and $\boldsymbol{\mu} := (\mu_1, \ldots, \mu_N)$, where each θ_i, μ_i is the state and input of ith subsystem.

We would like to choose the inputs $\boldsymbol{\mu}(t)$ in a way that the state of each subsystem always avoids a prescribed common unsafe set $\mathcal{U}$ (which corresponds to $[T_{\text{low}}, \underline{T}) \cup (\overline{T}, T_{\text{up}}]$ in the TCL problem). That is, defining the safe set of product system $\Theta_{\text{safe}} := (\Theta \setminus \mathcal{U})^N$, we want $\boldsymbol{\theta}(t)$, to be in $\Theta_{\text{safe}} \subset \Theta^N$ for all times. In addition, we impose input constraints, namely mode-counting constraints, that restrict the number of subsystems in mode m at any given time into an interval $[\underline{N}_m, \overline{N}_m]$ (which corresponds to $[\underline{N}_{\text{on}}, \overline{N}_{\text{on}}]$ in the TCL problem). That is, the actions of $S^{\times N}$ should be limited to

$$U_{\text{safe}} := \{(\mu_1, \ldots, \mu_N) \mid \underline{N}_m \leq \sum_{i=1}^{N} \mathbf{1}_m(\mu_i) \leq \overline{N}_m, \quad \forall m \in [M]\},$$

which can be used to capture constraints of the form (4).

Now, we can establish the relationship between controlled invariant sets of $S^{\times N}$ and of $S_\eta^{\times N}$ via the following theorem.

Theorem 1. *Assume that S and S_η are ϵ-approximately bisimilar. Let Θ^{max} be the maximal controlled invariant set for $(S^{\times N}, \Theta_{safe}, U_{safe})$, $\Theta_\eta^{max-}(\delta_1)$ be the maximal controlled invariant set for $(S_\eta^{\times N}, \gamma_\eta(\Theta_{safe} \ominus \mathcal{B}(0, \delta_1)), U_{safe})$, and $\Theta_\eta^{max+}(\delta_2)$ be the maximal controlled invariant set for $(S_\eta^{\times N}, \gamma_\eta(\Theta_{safe} \oplus \mathcal{B}(0, \delta_2)), U_{safe})$. If $\delta_1 > \epsilon + \eta/2$ and $\delta_2 > \epsilon - \eta/2$, then the following statements hold.*

1) *$\gamma_\eta^{inv}(\Theta_\eta^{max-}(\delta_1)) \oplus \mathcal{B}(0, \epsilon - \eta/2) \subset \Theta^{max} \subset \gamma_\eta^{inv}(\Theta_\eta^{max+}(\delta_2)) \ominus \mathcal{B}(0, \epsilon - \eta/2)$.*
2) *If the trajectory of $S_\eta^{\times N}$ starting at $\boldsymbol{\xi} \in \Theta_\eta^{max-}(\delta_1)$ stays inside $\gamma_\eta(\Theta_{safe} \ominus \mathcal{B}(0, \delta_1))$ under an input sequence $\boldsymbol{\mu}^* : \mathbb{N}_0 \to U_{safe}$, then for every $\boldsymbol{\theta} \in \Theta^N$ satisfying $\|\boldsymbol{\theta} - \boldsymbol{\xi}\| \leq \epsilon$, the trajectory of $S^{\times N}$ starting at $\boldsymbol{\theta}$ stays inside the safe set Θ_{safe} under the same input sequences $\boldsymbol{\mu}^*$.*

The first statement of the theorem says that one can use the collection of abstracted subsystems to closely approximate the maximal invariant set of the original collection. The second statement says control inputs computed for the collection $S_\eta^{\times N}$ of abstracted subsystems can be used to guarantee the safety of the original collection. Hence, we focus on computing controlled invariant sets for $(S_\eta^{\times N}, \gamma_\eta(\Theta_{\text{safe}} \ominus \mathcal{B}(0, \delta_1)), U_{\text{safe}})$ where δ_1 is bigger than $\epsilon + \eta/2$.

C. Aggregate system

As an alternative representation of the collection $S_\eta^{\times N}$ of state-abstracted subsystems, we now construct an aggregate system following [23]. Our goal is to use the aggregate system to compute (implicitly defined) invariant sets to guarantee safety for arbitrary controllers, as opposed to the specific open-loop control approach in [23].

We first represent the state-abstracted system as a graph. Suppose that $\gamma_\eta(\Theta) = \{\widetilde{\xi}_1, \ldots, \widetilde{\xi}_K\}$, where K is the number of element of this set. Then, S_η can be represented as a graph $G = (V, E)$, each node $\nu_k \in V$ of which corresponds to a state $\widetilde{\xi}_k$, and each edge of which corresponds to potential transitions in S_η labeled with corresponding set of modes. That is, $(\nu_i, \nu_j) \in E$ if and only if there exists $\mu \in [M]$ such that $\widetilde{\xi}_i \xrightarrow[\eta]{\mu} \widetilde{\xi}_j$ and the edge (ν_i, ν_j) is labeled with all such modes μ.

We define the state of the aggregate system $x \in \mathbb{N}_0^{MK}$ to be the number of subsystems in each node of the graph G in a specific mode. Here, a subsystem being in a node ν_k means, the state of that subsystem has the value $\widetilde{\xi}_k$. The $(m-1)K+k$th element of x, denoted by $x_{m,k}$, corresponds to the number of subsystems in node ν_k with mode m. Also, we let the input $u \in \mathbb{N}_0^{M(M-1)K}$ be the number of subsystems changing from a mode to another. The element $u_{m_1,m_2,k}$

of input represents the number of subsystems in node ν_k changing the mode from m_1 to m_2. Then, it is shown in [23] that x evolves according to the linear dynamics

$$\Gamma_\eta : x(t+1) = Ax(t) + Bu(t), \tag{10}$$

where A, B are based on the incidence matrices of the graph G at each mode and they describe how the subsystems move around the nodes of the graph depending on the chosen mode. The state space X and the admissible input space $U(x)$ of this system are obtained as follows.

$$\begin{aligned} X &= \left\{ x \in \mathbb{N}_0^{MK} : \sum_{m=1}^{M} \sum_{k=1}^{K} x_{m,k} = N \right\} \\ U(x) &= \left\{ u \in \mathbb{N}_0^{M(M-1)K} : 0 \leq \sum_{m_2} u_{m_1,m_2,k} \leq x_{m_1,k} \right\} \end{aligned} \tag{11}$$

The local and global constraints can also be imposed on the aggregate system as state constraints

$$\begin{aligned} X_{\text{safe}} := \{x \in X : \ \forall m \in [M],\ \forall k \in \tilde{\mathcal{I}},\ x_{m,k} = 0, \\ \underline{N}_m \leq \sum_{k=1}^{K} x_{m,k} \leq \overline{N}_m \Big\}, \end{aligned} \tag{12}$$

where $\tilde{\mathcal{I}} := \{k \mid \widetilde{\xi}_k \in (\gamma_\eta(\Theta) \setminus \gamma_\eta((\Theta \setminus \mathcal{U}) \ominus \mathcal{B}(0, \delta_1)))\}$ denotes the indices of unsafe state values in $\gamma_\eta(\Theta)$.

A mapping from the states of the aggregate system Γ_η to the states of the state-abstraction system S_η is

$$\begin{aligned} \Theta_{\text{map}}(x) := \Bigg\{ \boldsymbol{\xi} \in \gamma_\eta(\Theta^N) : \sum_{m=1}^{M} x_{m,k} = \sum_{i=1}^{N} \mathbf{1}_{\widetilde{\xi}_k}(\xi_i) \\ \forall k \in [K] \Big\}, \end{aligned}$$

The next theorem shows that inclusion in a controlled invariant set is preserved by this mapping.

Theorem 2. *For any controlled invariant set X' of Γ_η in X_{safe}, the set $\Theta_{map}(X')$ is a controlled invariant set with respect to $(S_\eta^{\times N}, \gamma_\eta(\Theta_{safe} \ominus \mathcal{B}(0, \delta_1)), U_{safe})$.*

The aggregate system Γ_η is a linear system with integer-valued states and inputs. Its state space, control constraints and the safe set are all described by linear inequalities (11),(12). Crucially, the dimension of the state Γ_η does not depend on the number of subsystems. If we can compute an invariant set for this aggregate system, then by Theorems 1 and 2, we can use this invariant set to compute a safe policy for $S^{\times N}$. However, there are two challenges: (i) the integrality constraints, (ii) even if we relax the integrality constraints, the dimension of Γ_η is still out of reach of the existing invariant set computation tools for linear systems [4], [6]. In the next section, we show how we can utilize the structure in (10) to compute invariant sets for this system.

IV. Controlled invariant set generation using cycles

In this section, we introduce a method which constructs a controlled invariant set for the aggregate system using cycles of the graph G. The invariant set generated by this method includes states that have a periodic input sequence leading to periodic trajectories inside X_{safe}.

Suppose that C is a cycle in the graph G whose transitions between nodes are labeled with their corresponding modes. We denote the nodes of C by $(\tilde{\nu}_{C,1}, \ldots, \tilde{\nu}_{C,|C|})$ and its labels by $\tilde{\mu}_C = (\tilde{\mu}_{C,1}, \ldots, \tilde{\mu}_{C,|C|}) \in [M]^{|C|}$; if the state of S_η is in $\tilde{\nu}_{C,l}$ and the input is chosen to be $\tilde{\mu}_{C,l}$, the state moves to the next node $\tilde{\nu}_{C,l+1}$ (or, $\tilde{\nu}_{C,1}$ if $l = |C|$).

The first step of constructing an invariant set is to choose some labeled cycles $C_1, \ldots, C_n$ from graph G which are composed of only the nodes corresponding to the safe states; every node $\tilde{\nu}_{C_j,l}$ of each cycle C_j satisfies $\tilde{\nu}_{C_j,l} \neq \nu_k$ for all $k \in \tilde{\mathcal{I}}$. Hence, every subsystem circulating around one of these cycles always satisfies the safe constraint on the state.

Now, we show how to assign all the subsystems over the nodes of $C_1, \ldots, C_n$ at initial time step in a way that the mode-counting constraints would never be violated if every subsystem circulates around its assigned cycle from that initial assignment. Let l_j be the length of cycle C_j and let an initial assignment to the cycle C_j be a vector $\beta_j = (\beta_j(1), \ldots, \beta_j(l_j))$, where $\beta_j(l)$ represents the number of subsystems assigned to $\tilde{\nu}_{C_j,l}$. Since each subsystem should be assigned to one and only one cycle at anytime, $\sum_{j=1}^{n} \sum_{l=1}^{l_j} \beta_j(l) = N$ should be satisfied. Also, under the assumption that every subsystem circulates around its assigned cycle, the number of subsystems in mode m after q steps is given by

$$\sum_{j=1}^{n} \sum_{\substack{l \text{ s.t.} \\ \tilde{\mu}_{C_j,l}=m}} \beta_j((-q+l \mod l_j)+1). \tag{13}$$

Note that, due to periodicity, after $\text{lcm}(l_1, \ldots, l_n)$ steps, all subsystems will come to their initial assignments. Therefore, the mode-counting constraints $\{\underline{N}_m, \overline{N}_m\}_{m \in [M]}$ hold if the number of subsystems in mode m is between $\underline{N}_m$ and $\overline{N}_m$ for all $q = 1, \ldots, \text{lcm}(l_1, \ldots, l_n)$, which is equivalent to the following condition

$$\begin{aligned} \underline{N}_m \leq \sum_{j=1}^{n} \sum_{\substack{l \text{ s.t.} \\ \tilde{\mu}_{C_j,l}=m}} \beta_j((-q+l \mod l_j)+1) \leq \overline{N}_m \\ \forall q \in [\text{lcm}(l_1, \ldots, l_n)],\ \forall m \in [M]. \end{aligned} \tag{14}$$

Concepts related to cycles are illustrated in Fig. 2.

Then, we define the set of "good" aggregate cycle assignments, which satisfy the conditions described above, as follows

$$\begin{aligned} \Omega := \{(\beta_1, \ldots, \beta_n) \in \mathbb{N}_0^{\sum_j l_j} : \\ \sum_{j=1}^{n} \sum_{l=1}^{l_j} \beta_j(l) = N \text{ and (14) holds} \Big\}. \end{aligned} \tag{15}$$

Also, we define the mapping from the space of assignments of a cycle into the state space of Γ_η, which computes the number of subsystems in each mode at each node of G corresponding to a cycle assignment vector. For a cycle C

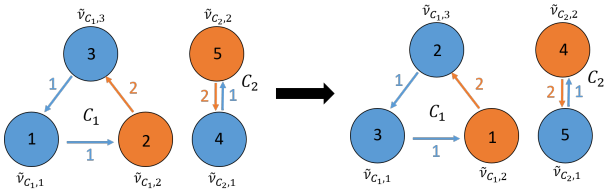

Fig. 2: Cycle C_1 has nodes $(\tilde{\nu}_{C_1,1}, \tilde{\nu}_{C_1,2}, \tilde{\nu}_{C_1,3})$ and labels $\tilde{\mu}_{C_1} = (1,2,1)$ and cycle C_2 has nodes $(\tilde{\nu}_{C_2,1}, \tilde{\nu}_{C_2,2})$ and labels $\tilde{\mu}_{C_2} = (1,2)$. On the left, the cycle assignments (shown inside the nodes) are $\beta_1 = (1,2,3)$ and $\beta_2 = (4,5)$. At step $q = 1$ (left), the 1-mode count, given by Eq. (13), is $1+3+4 = 8$ and the 2-mode count is $2+5$. After applying $(\tilde{\mu}_{C_1}, \tilde{\mu}_{C_2})$, at step $q = 2$ (right), the 1-mode and 2-mode counts are 10 and 5, respectively. The mode-counting conditions on the assignments β_1 and β_2 should be checked for $\text{lcm}(3,2) = 6$ steps.

with labels $\tilde{\mu}_C$, this linear mapping $\Phi_{C,\tilde{\mu}_C} : \mathbb{N}_0^l \to \mathbb{N}_0^{MK}$ is defined as

$$\left(\Phi_{C,\tilde{\mu}_C}(\beta)\right)_{m,k} = \sum_{\substack{l:\tilde{\nu}_{C,l}=\nu_k \\ \tilde{\mu}_{C,l}=m}} \beta(l), \tag{16}$$

where $\left(\Phi_{C,\tilde{\mu}_C}(\beta)\right)_{m,k}$, the $(m-1)K+k$th element of $\Phi_{C,\tilde{\mu}_C}(\beta)$, corresponds to the number of subsystems in mode m with state at node ν_k assigned to cycle C.

Let X_{cyc} be the set obtained by applying the above mapping to Ω with respect to the cycles $C_1, \ldots, C_n$, specifically,

$$X_{\text{cyc}} = \{x \in \mathbb{N}_0^{MK} : \exists(\beta_1, \ldots, \beta_n) \in \Omega, \text{ s.t. } x = \sum_{j=1}^{n} \Phi_{C_j,\tilde{\mu}_{C_j}}(\beta_j)\}. \tag{17}$$

Observe that X_{cyc} corresponds to a set of states of Γ_η, every element of which has a periodic input sequence which rotates all the subsystems around $C_1, \ldots, C_n$ they are assigned to, while ensuring the mode-counting constraint indefinitely. From this fact, we obtain the following theorem.

Theorem 3. *X_{cyc} is a controlled invariant set of the system Γ_η under the constraint set X_{safe}.*

An explicit representation of X_{cyc} can be obtained by projection of the integral-valued set defined by linear equations and inequalities in the variable $(x,\beta) \in \mathbb{N}_0^{\sum_{i=1}^n l_i + MK}$ to the state space $X \subset \mathbb{N}_0^{MK}$. However, projection is intractable [24] when the dimension of the state and the cycle assignments are large. On the other hand, for the purposes of checking whether a state x is in X_{cyc} or computing a $u \in U(x)$ that guarantees invariance, we do not need an explicit representation; we can just use the linear representation in (x,β)-space, which renders both operations integer linear programs. Hence, we work with the implicit representation of X_{cyc} and incorporate it into the safe control algorithm introduced in the next section. Before we move on, the relationship between the main results regarding invariance are summarized in Fig. 3.

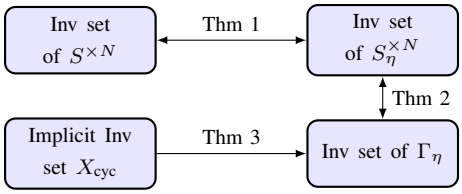

Fig. 3: Relation between invariant sets of the systems.

V. Invariant-set-driven MPC

In this section, we propose an MPC-based control algorithm, referred to as Invariant-set-driven MPC. The algorithm keeps the state of the aggregate system inside its maximal controlled invariant set using X_{cyc} as a terminal condition. This algorithm solves the following program at every time step $t \in \mathbb{N}_0$ for a given horizon length h

$$\begin{aligned}
\min \ & \sum_{\tau=0}^{h} c(x^{t+\tau|t}) \\
\text{s.t. } & x^{t+\tau+1|t} = Ax^{t+\tau|t} + Bu^{t+\tau|t} && \forall \tau \in [h-1]_0 && \text{(18a)} \\
& u^{t+\tau|t} \in U(x^{t+\tau|t}) && \forall \tau \in [h-1]_0 && \text{(18b)} \\
& x^{t|t} = x(t) && && \text{(18c)} \\
& x^{t+h|t} = \sum_{j=1}^{n} \Phi_{C_j,\tilde{\mu}_{C_j}}(\beta_j) && && \text{(18d)} \\
& (\beta_1, \ldots, \beta_n) \in \Omega, && && \text{(18e)}
\end{aligned}$$

whose variables are $x^{t+\tau|t} \in \mathbb{N}_0^{MK}$ $(\forall \tau \in [h]_0)$, $u^{t+\tau|t} \in \mathbb{N}_0^{M(M-1)K}$ $(\forall \tau \in [h-1]_0)$, $(\beta_1, \ldots, \beta_n) \in \mathbb{N}_0^{\sum_j l_j}$, and c is the cost function. We select the optimal $u^{t|t}$ as the input $u(t)$ which drives the state to $x(t+1) = x^{t+1|t}$. The formulation forces the last state in the horizon $x^{t+h|t}$ to be included in the controlled invariant set X_{cyc}. Terminal constraints that guarantee recursive feasibility are common in the MPC literature [25]. Different from explicitly-defined invariant sets, we exploit the fact that every element of X_{cyc} is mapped from an element of Ω (from (17)) and use this implicit representation in cycle assignments space. Use of the terminal constraint leads to the following result.

Theorem 4. *Suppose that the initial state $x(0)$ belongs to X_{cyc}. Then, the program (18) has a feasible solution at any time step $t \in \mathbb{N}_0$. In addition, the trajectory $x(t)$ generated by above algorithm always belongs to the maximal controlled invariant set X^{max} of the aggregate dynamics Γ_η for every $t \in \mathbb{N}_0$.*

The performance of the algorithm improves with more cycles $C_1, \ldots, C_n$ and/or larger horizons h. However, this also increases the number of variables and inequalities, which increases the online computational burden. This trade-off needs to be taken into account when these choices are made.

VI. Simulation Results

In this section, we compare the performance of the Invariant-set-driven MPC to several benchmark algorithms. We first describe the benchmark algorithms. Then, we run

several numerical experiments that show how the benchmarks can fail while the Invariant-set-driven MPC is able to ensure safety.

In all experiments, we use a cost function of the form

$$c(x) = \left| \sum_{k=1}^{K} p x_{\text{on},k} - r \right| = |P_{\text{agg}} - r|, \tag{19}$$

which penalizes the difference between the aggregate power consumption and a reference r. We control a collection of air conditioners and set $T_a = 32$°C, $C_{\text{th}} = 1.8$ kWh/°C, $R_{\text{th}} = 1.8$ °C/kW; $p = 6.4$ kW, $p_{\text{tr}} = 16$ kW, $[T_{\text{low}}, T_{\text{up}}] = [3.2, 32]$ °C, $[\underline{T}, \overline{T}] = [22, 23]$ °C, $\mathcal{U} = [3.2, 22) \cup (23, 32]$ °C, $N_{\text{TCL}} = 10000$, $[\underline{P}_{\text{agg}}, \overline{P}_{\text{agg}}] = [16, 25.6]$ MW, $\overline{N}_{\text{on}} = 4000$, $\underline{N}_{\text{on}} = 2500$, $\overline{N}_0 = N$, $\underline{N}_0 = 0$, $\overline{N}_1 = \overline{N}_{\text{on}}$, and $\underline{N}_1 = \underline{N}_{\text{on}}$. The TCL parameters are informed by [7].

A. Benchmark algorithms

1) Benchmark 1: We use a basic MPC algorithm and require every state in horizon to be within X_{safe}

$$\begin{aligned} \min \ & \sum_{\tau=0}^{h} \left| \sum_{k=1}^{K} p x_{\text{on},k}^{t+\tau|t} - r(t+\tau) \right| \\ \text{s.t. } & \text{(18a), (18b), (18c) hold} \\ & x^{t+\tau|t} \in X_{\text{safe}} \qquad \forall \tau \in [h]_0 \end{aligned}$$

2) Benchmark 2: We truncate the reference signal as shown by the dotted yellow line in Fig. 1 and use MPC to track the truncated signal $\hat{r}(t)$, computed as

$$\hat{r}(t) = \begin{cases} r(t) & \text{if } \underline{P}_{\text{agg}} \le r(t) \le \overline{P}_{\text{agg}} \\ \underline{P}_{\text{agg}} & \text{if } r(t) < \underline{P}_{\text{agg}} \\ \overline{P}_{\text{agg}} & \text{if } r(t) > \overline{P}_{\text{agg}}. \end{cases}$$

The MPC problem is

$$\begin{aligned} \min \ & \sum_{\tau=0}^{h} \left| \sum_{k=1}^{K} p x_{\text{on},k}^{t+\tau|t} - \hat{r}(t+\tau) \right| \\ \text{s.t. } & \text{(18a), (18b), (18c) hold,} \\ & x_{m,k}^{t+\tau|t} = 0 \qquad \forall m \in \{\text{on}, \text{off}\}, k \in \tilde{\mathcal{I}}, \tau \in [h]_0. \end{aligned} \tag{20}$$

where (20) imposes the local temperature constraints on TCLs.

3) Benchmark 3: Again using MPC, this algorithm requires the state at next time step to belong to the invariant set X_{cyc}

$$\begin{aligned} \min \ & \sum_{\tau=0}^{h} \left| \sum_{k=1}^{K} p x_{\text{on},k}^{t+\tau|t} - r(t+\tau) \right| \\ \text{s.t. } & \text{(18a), (18b), (18c) hold} \\ & x^{t+1|t} = \sum_{j=1}^{n} \Phi_{C_j, \tilde{\mu}_{C_j}}(\beta_j) \\ & (\beta_1, \ldots, \beta_n) \in \Omega \end{aligned}$$

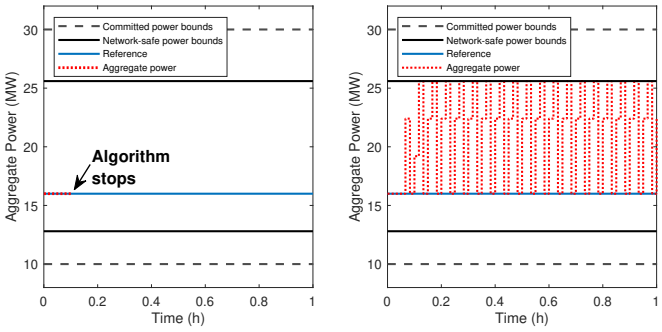

(a) Benchmark 1 (b) Invariant-set-driven MPC

Fig. 4: Benchmark 1 can fail because the approach does not guarantee recursive feasibility (left), while invariant-set-driven MPC is recursively feasible and ensures safety (right).

B. Numerical Experiments

1) Experiment 1: In this experiment, we demonstrate a case where Benchmark 1 fails but Invariant-set-driven MPC ensures constraint satisfaction though tracking performance suffers. We choose a constant reference signal that is safe (i.e., within the network-safe power bounds) but untrackable because every TCL is initialized near the boundary of its dead-band. For Invariant-set-driven MPC, the set of chosen cycles include every elementary cycle. We set $\Delta t = 60$ s, $\epsilon = 0.3$, $\eta = 0.003$, $\delta_1 = 3.003$, and $h = 1$.

Fig. 4 shows the tracking performance of both controllers. Benchmark 1 (left) stops when the trajectory goes outside of the maximal invariant set because the program does not have feasible solution. This occurs when a large number of off-mode TCLs are next to the upper bounds of their temperature dead-bands, and they are forced to be turned on to stay inside the dead-band. In contrast, Invariant-set-driven MPC (right) turns on the TCLs in advance to prevent too many TCLs being simultaneously turned on at the boundary of the dead-band, which makes the state stay inside the maximal controlled invariant set.

2) Experiment 2: In this experiment, we demonstrate the drawbacks of Benchmarks 2 and 3 and the superior performance of Invariant-set-driven MPC by using an unsafe reference signal. Parameters and selected cycles for Benchmark 3 and Invariant-set-driven MPC are the same as in the first experiment, except $T = 180$.

Fig. 5 shows the tracking performance of each controller and Table I compares them. The reference signal is safe at first, but then goes outside of the network-safe power bounds. Benchmark 2 (left), which uses the truncated reference signal, violates the power bounds; it is unable to track the truncated reference signal and, in trying to, TCLs synchronize and the aggregate power starts to oscillate. Both Benchmark 3 and Invariant-set-driven MPC stay within the power bounds but Benchmark 3 has a larger tracking error since it is more conservative; notice the tracking errors even during the first part of the simulation when the reference signal is safe and trackable. Therefore, we can conclude that Invariant-set-driven MPC has the best overall performance in that it has lower tracking error than Benchmark 3 and ensures all constraints are satisfied unlike Benchmark 2. All

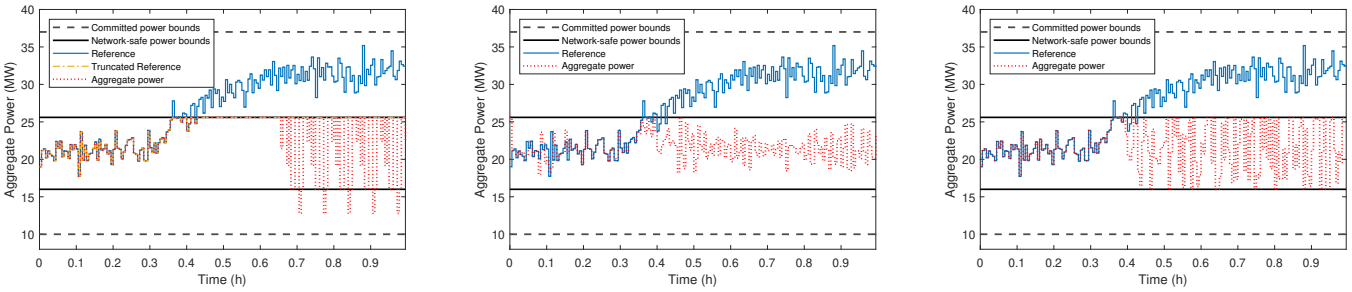

(a) Benchmark 2 (BM 2) (b) Benchmark 3 (BM 3) (c) Invariant-set-driven MPC (Invset MPC)

Fig. 5: Tracking of unsafe reference signal. BM 2 violates network-safe power bounds. BM 3 is conservative compared to Invset MPC.

TABLE I: Performance comparison

	BM 2	BM 3	Invset MPC
Tracking Error	4.56×10^3	5.75×10^3	5.66×10^3
Constraint Violations	Yes	None	None
Average Time (s)	0.44	4.48	2.62

computation times are shorter than Δt so all methods could be employed in practice.

VII. Conclusions

This paper proposed a method to construct controlled invariant sets by finding states that have a periodic trajectory inside a safe set. It also proposed an MPC-based safe control algorithm for systems composed of homogeneous switched subsystems with local and global constraints. The control algorithm uses an implicit representation of the invariant set, which mitigates the computational burden. It also requires the state to be within the invariant set only in the last step of the MPC horizon, which is shown to be less conservative than a benchmark approach that requires the state to always remain within the invariant set. The approach is further benchmarked against two naive strategies that are shown to break down under certain reference signals. In contrast, the proposed approach always satisfies the safety constraints.

Future work will explore way to relax some of the assumptions used to derive the results in an effort to make the approach more applicable to the TCL power tracking problem that motivated this work. In particular, we will explore ways of handling time-varying ambient temperature, power bounds, and temperature setpoints, along with compressor lock-out constraints.

References

[1] D. Bertsekas, "Infinite time reachability of state-space regions by using feedback control," *IEEE Transactions on Automatic Control*, vol. 17, no. 5, pp. 604–613, 1972.

[2] J.-P. Aubin, A. Bayen, and P. Saint-Pierre, *Viability theory: new directions*. Springer Science & Business Media, 2011.

[3] F. Blanchini, "Set invariance in control," *Automatica*, vol. 35, no. 11, pp. 1747–1767, 1999.

[4] T. Anevlavis and P. Tabuada, "Computing controlled invariant sets in two moves," in *IEEE CDC*, 2019, pp. 6248–6254.

[5] M. Fiacchini and M. Alamir, "Computing control invariant sets in high dimension is easy," *arXiv preprint arXiv:1810.10372*, 2018.

[6] P. Nilsson and N. Ozay, "Control synthesis for permutation-symmetric high-dimensional systems with counting constraints," *IEEE Transactions on Automatic Control*, vol. 65, no. 2, pp. 461–476, 2019.

[7] J. Mathieu, S. Koch, and D. Callaway, "State estimation and control of electric loads to manage real-time energy imbalance," *IEEE Transactions on Power Systems*, vol. 28, no. 1, pp. 430–440, 2012.

[8] W. Zhang, J. Lian, C.-Y. Chang, and K. Kalsi, "Aggregated modeling and control of air conditioning loads for demand response," *IEEE Transactions on Power Systems*, vol. 28, no. 4, pp. 4655–4664, 2013.

[9] S. Tindemans, V. Trovato, and G. Strbac, "Decentralized control of thermostatic loads for flexible demand response," *IEEE Transactions on Control Systems Technology*, vol. 23, no. 5, pp. 1685–1700, 2015.

[10] S. Ross, G. Vuylsteke, and J. Mathieu, "Effects of load-based frequency regulation on distribution network operation," *IEEE Transactions on Power Systems*, vol. 34, no. 2, pp. 1569–1578, 2019.

[11] E. Vrettos and G. Andersson, "Combined load frequency control and active distribution network management with thermostatically controlled loads," in *IEEE SmartGridComm*, 2013, pp. 247–252.

[12] E. Dall'Anese, S. Guggilam, A. Simonetto, Y. C. Chen, and S. Dhople, "Optimal regulation of virtual power plants," *IEEE Transactions on Power Systems*, vol. 33, no. 2, pp. 1868–1881, 2017.

[13] A. Bernstein and E. Dall'Anese, "Real-time feedback-based optimization of distribution grids: A unified approach," *IEEE Transactions on Control of Network Systems*, vol. 6, no. 3, pp. 1197–1209, 2019.

[14] N. Nazir and M. Almassalkhi, "Convex inner approximation of the feeder hosting capacity limits on dispatchable demand," in *IEEE CDC*, 2019, pp. 4858–4864.

[15] S. Ross, N. Ozay, and J. Mathieu, "Coordination between an aggregator and distribution operator to achieve network-aware load control," in *IEEE PowerTech*, 2019.

[16] D. Molzahn and L. Roald, "Grid-aware versus grid-agnostic distribution system control: A method for certifying engineering constraint satisfaction," in *HICSS*, 2019.

[17] S. Ross and J. Mathieu, "A method for ensuring a load aggregator's power deviations are safe for distribution networks," in *PSCC*, 2020.

[18] S. Ross, P. Nilsson, N. Ozay, and J. Mathieu, "Managing voltage excursions on the distribution network by limiting the aggregate variability of thermostatic loads," in *ACC*, 2019, pp. 4260–4267.

[19] P. Nilsson and N. Ozay, "On a class of maximal invariance inducing control strategies for large collections of switched systems," in *HSCC*, 2017, pp. 187–196.

[20] R. C. Sonderegger, "Dynamic models of house heating based on equivalent thermal parameters," *PhD Thesis*, 1978.

[21] P. Tabuada, *Verification and control of hybrid systems: a symbolic approach*. Springer Science & Business Media, 2009.

[22] G. Pola and P. Tabuada, "Symbolic models for nonlinear control systems: Alternating approximate bisimulations," *SIAM Journal on Control and Optimization*, vol. 48, no. 2, pp. 719–733, 2009.

[23] P. Nilsson and N. Ozay, "Control synthesis for large collections of systems with mode-counting constraints," in *HSCC*, 2016, pp. 205–214.

[24] H. Tiwary, "On the Hardness of Computing Intersection, Union and Minkowski Sum of Polytopes," *Discrete & Computational Geometry*, vol. 40, no. 3, pp. 469–479, 2008.

[25] K. Wabersich and M. Zeilinger, "Linear model predictive safety certification for learning-based control," in *IEEE CDC*, 2018, pp. 7130–7135.

APPENDIX

Proofs for Section III:

Proof of Lemma 1. Follows by from [23] where $\alpha_m(x) = L_m x$ acts as the class $\mathcal{K}$-functions in their proof for a fixed sampling time. □

To prove the theorems in section III, the following lemmas are used.

Lemma 2. *Abstraction function has the following properties:*

P1. *For any set $\Theta' \subset \Theta$:*

$$\gamma_\eta(\Theta') \subset (\Theta' \oplus \mathcal{B}(0, \eta/2)) \tag{21}$$

$$\Theta' \subset \gamma_\eta(\Theta' \oplus \mathcal{B}(0, \eta/2)) \tag{22}$$

P2. *If $x_d \in \gamma_\eta(\Theta)$ and $x_d \notin \gamma_\eta(X)$ for a set $X \subset \Theta$, then $\|x_d - x\| \geq \eta/2$ for any $x \in X$.*

P3. *Let γ_η^{inv} be the inverse mapping of γ_η defined for any $\Xi \subset \gamma_\eta(\Theta)$ as $\gamma_\eta^{inv}(\Xi) := \{\theta \in \Theta \mid \exists \xi \in \Xi \text{ s.t. } \gamma_\eta(\theta) = \xi\}$, then*

$$\gamma_\eta^{inv}(\Xi) \subset \Xi \oplus \mathcal{B}(0, \eta/2). \tag{23}$$

Lemma 3. *If $\boldsymbol{\xi} \in \gamma_\eta(\Theta^N)$ and $\delta > \epsilon + \eta/2$, $\gamma_\eta(\Theta_{safe} \ominus \mathcal{B}(0, \delta))$ is a subset of $\Theta_{safe} \ominus \mathcal{B}(0, \epsilon)$.*

Proof. Suppose $\boldsymbol{\xi}$ is an element of $\gamma_\eta(\Theta_{\text{safe}} \ominus \mathcal{B}(0, \delta))$ and $\delta > \epsilon + \eta/2$. Then, by Lemma2, P1, we have

$$\forall \boldsymbol{u}' \in \mathcal{U}^N, \ \ \|\boldsymbol{\xi} - \boldsymbol{u}'\| \geq \delta - \frac{\eta}{2} > \epsilon, \tag{24}$$

Thus, we can obtain the following

$$\boldsymbol{\xi} \in (\Theta \setminus (\mathcal{U} \oplus \mathcal{B}(0, \epsilon)))^N \subset \Theta_{\text{safe}} \ominus \mathcal{B}(0, \epsilon).$$

Therefore, $\gamma_\eta(\Theta_{\text{safe}} \ominus \mathcal{B}(0, \delta))$ is a subset of $\Theta_{\text{safe}} \ominus \mathcal{B}(0, \epsilon)$. □

Proof of Theorem 1. The proof is separated into two parts which the first part shows $\gamma_\eta^{\text{inv}}(\Theta_\eta^{\text{max}-}(\delta_1)) \oplus \mathcal{B}(0, \epsilon - \eta/2) \subset \Theta^{\text{max}}$ and the second statement, and the second part shows $\Theta^{\text{max}} \subset \gamma_\eta^{\text{inv}}(\Theta_\eta^{\text{max}+}(\delta_2)) \ominus \mathcal{B}(0, \epsilon - \eta/2)$.

1) Let δ_1 be larger than $\epsilon - \eta/2$ and $\boldsymbol{\xi}$ be an element of $\Theta_\eta^{\text{max}-}(\delta_1)$. Then, an aggregated input sequences $\boldsymbol{\mu}^{\text{tr}}(t)$ and an aggregated trajectories $\boldsymbol{\xi}^{\text{tr}}(t)$ exist which satisfy

$$\begin{aligned}
&\boldsymbol{\xi}^{\text{tr}}(t) \xrightarrow[\eta, \times N]{\boldsymbol{\mu}^{\text{tr}}(t)} \boldsymbol{\xi}^{\text{tr}}(t+1) && \forall t \in \mathbb{N}_0 \\
&\boldsymbol{\xi}^{\text{tr}}(t) \in \gamma_\eta(\Theta_{\text{safe}} \ominus \mathcal{B}(0, \delta_1)) && \forall t \in \mathbb{N}_0 \quad (25) \\
&\boldsymbol{\mu}^{\text{tr}}(t) \in U_{\text{safe}} && \forall t \in \mathbb{N}_0 \quad (26) \\
&\boldsymbol{\xi} = \boldsymbol{\xi}^{\text{tr}}(0).
\end{aligned}$$

For any element $\boldsymbol{\theta}$ which satisfies $\|\boldsymbol{\theta} - \boldsymbol{\xi}\| \leq \epsilon$, let $\boldsymbol{\theta}^{\text{tr}}(t)$ be the aggregated trajectories started from $\boldsymbol{\theta}$ under the aggregated input sequences $\boldsymbol{\mu}^{\text{tr}}(t)$ in $S^{\times N}$ (i.e. $\boldsymbol{\theta}^{\text{tr}}(0) = \boldsymbol{\theta}$, $\boldsymbol{\theta}^{\text{tr}}(t) \xrightarrow[\times N]{\boldsymbol{\mu}^{\text{tr}}(t)} \boldsymbol{\theta}^{\text{tr}}(t+1)$). By the assumption that S and S_η is ϵ-bisimilar, $\|\boldsymbol{\xi}^{\text{tr}}(t) - \boldsymbol{\theta}^{\text{tr}}(t)\|$ is smaller than or equal to ϵ. Moreover, from Lemma 3 and (25), we have $\boldsymbol{\xi}^{\text{tr}}(t) \in \Theta_{\text{safe}} \ominus \mathcal{B}(0, \epsilon)$. Therefore, the following statement holds,

$$\boldsymbol{\theta}^{\text{tr}}(t) \in (\Theta_{\text{safe}} \ominus \mathcal{B}(0, \epsilon)) \oplus \mathcal{B}(0, \epsilon) = \Theta_{\text{safe}}. \tag{27}$$

which validates the second statement.

In addition, we can conclude that $\boldsymbol{\theta}$ is included in Θ^{max} from (26), (27). Since this holds for any element $\boldsymbol{\xi}$ in $\Theta_\eta^{\text{max}-}(\delta_1)$ and any $\boldsymbol{\theta}$ satisfying $\|\boldsymbol{\theta} - \boldsymbol{\xi}\| \leq \epsilon$, we can conclude that every element $\boldsymbol{\theta}$ in $\Theta_\eta^{\text{max}-}(\delta_1) \oplus \mathcal{B}(0, \epsilon)$ also belongs to Θ^{max}. Therefore, $\Theta_\eta^{\text{max}-}(\delta_1) \oplus \mathcal{B}(0, \epsilon) \subset \Theta^{\text{max}}$. Using the property of γ_η^{inv} in (23), the following is obtained,

$$\begin{aligned}
\gamma_\eta^{\text{inv}}(\Theta_\eta^{\text{max}-}(\beta)) \oplus \mathcal{B}(0, \epsilon - \eta/2) &\subset \Theta_\eta^{\text{max}-}(\beta) \oplus \mathcal{B}(0, \epsilon) \\
&\subset \Theta_{\text{max}}
\end{aligned} \tag{28}$$

2) For any element $\boldsymbol{\theta}$ of Θ^{max}, an aggregated input sequences $\boldsymbol{\mu}^{\text{tr}}(t)$ and an aggregated trajectories $\boldsymbol{\theta}^{\text{tr}}(t)$ exist which satisfy the following.

$$\begin{aligned}
&\boldsymbol{\theta}^{\text{tr}}(t) \xrightarrow[\times N]{\boldsymbol{\mu}^{\text{tr}}(t)} \boldsymbol{\theta}^{\text{tr}}(t+1) && \forall t \in \mathbb{N}_0 \\
&\boldsymbol{\theta}^{\text{tr}}(t) \in \Theta_{\text{safe}} && \forall t \in \mathbb{N}_0 \quad (29) \\
&\boldsymbol{\mu}^{\text{tr}}(t) \in U_{\text{safe}} && \forall t \in \mathbb{N}_0 \quad (30) \\
&\boldsymbol{\theta} = \boldsymbol{\theta}^{\text{tr}}(0).
\end{aligned}$$

For any $\boldsymbol{\xi} \in \gamma_\eta(\Theta^N)$ satisfying $\|\boldsymbol{\xi} - \boldsymbol{\theta}\| \leq \epsilon$, let $\boldsymbol{\xi}^{\text{tr}}(t)$ be the aggregated trajectories started from $\boldsymbol{\xi}$ under the aggregated input sequences $\boldsymbol{\mu}^{\text{tr}}(t)$ in $S_\eta^{\times N}$ (i.e. $\boldsymbol{\xi}(0) = \boldsymbol{\xi}$, $\boldsymbol{\xi}^{\text{tr}}(t) \xrightarrow[\eta, \times N]{\boldsymbol{\mu}^{\text{tr}}(t)} \boldsymbol{\xi}^{\text{tr}}(t+1)$). By the assumption of ϵ-bisimilarity between S and S_η, norm $\|\boldsymbol{\theta}^{\text{tr}}(t) - \boldsymbol{\xi}^{\text{tr}}(t)\|$ is equal to or smaller than ϵ. Using this and (29), the following can be shown.

$$\boldsymbol{\xi}^{\text{tr}}(t) \in \Theta_{\text{safe}} \oplus \mathcal{B}(0, \epsilon). \tag{31}$$

From (22), (31), we can show the following inclusion relationship for every $\delta_2 \geq \epsilon - \eta/2$.

$$\begin{aligned}
\boldsymbol{\xi}^{\text{tr}}(t) &\in \gamma_\eta((\Theta_{\text{safe}} \oplus \mathcal{B}(0, \epsilon)) \oplus \mathcal{B}(0, \eta/2)) \\
&\subset \gamma_\eta(\Theta_{\text{safe}} \oplus \mathcal{B}(0, \delta_2)).
\end{aligned} \tag{32}$$

From (30), (32), we can conclude that $\boldsymbol{\xi}$ is included in $\Theta_\eta^{\text{max}+}(\delta_2)$. Since this holds for any element $\boldsymbol{\theta}$ in Θ^{max}, and any $\boldsymbol{\xi} \in \gamma_\eta(\Theta^N)$ satisfying $\|\boldsymbol{\xi} - \boldsymbol{\theta}\| \leq \epsilon$, we can conclude that every element $\boldsymbol{\xi}$ in $(\Theta^{\text{max}} \oplus \mathcal{B}(0, \epsilon)) \cap \gamma_\eta(\Theta^N)$ also belongs to $\Theta_\eta^{\text{max}+}(\delta_2)$.

For any element $\boldsymbol{\theta}^+$ of $\Theta^{\text{max}} \oplus \mathcal{B}(0, \epsilon - \eta/2)$, $\boldsymbol{\xi}' = \gamma_\eta(\boldsymbol{\theta}^+)$ is an element of $\gamma_\eta(\Theta^{\text{max}} \oplus \mathcal{B}(0, \epsilon - \eta/2))$. From (21), the following is obtained.

$$\boldsymbol{\xi'} \in \gamma_\eta(\Theta^{\text{max}} \oplus \mathcal{B}(0, \epsilon - \eta/2)) \subset \Theta^{\text{max}} \oplus \mathcal{B}(0, \epsilon). \tag{33}$$

By applying first statement of the theorem, we have

$$\boldsymbol{\xi'} \in \Theta_\eta^{\text{max}+}(\delta_2),$$

which also means that $\boldsymbol{\theta}^+$ belongs to $\gamma_\eta^{\text{inv}}(\Theta_\eta^{\text{max}+}(\delta_2))$. Therefore, $\Theta_{\text{max}} \oplus \mathcal{B}(0, \epsilon - \eta/2)$ is a subset of $\gamma_\eta^{\text{inv}}(\Theta_\eta^{\text{max}+}(\delta_2))$ and therefore,

$$\Theta_{\text{max}} \subset \gamma_\eta^{\text{inv}}(\Theta_\eta^{\text{max}+}(\delta_2)) \ominus \mathcal{B}(0, \epsilon - \eta/2)$$

By 1) and 2), the statement holds. □

Proof of Theorem 2:

Proof. Suppose $\boldsymbol{\xi} = (\xi_1, \dots, \xi_n)$ is an element of $\Theta_{\text{map}}(X')$. Then, $x \in X'$ exists which satisfies $\boldsymbol{\xi} = \Theta_{\text{map}}(x)$. Since x belongs to X_{safe}, the following is obtained from (12)

$$\sum_{i=1}^{N} \mathbf{1}_{\tilde{\xi}_k}(\xi_i) = \sum_{m=1}^{M} x_{m,k} = 0 \qquad \forall k \in \tilde{\mathcal{I}}.$$

Thus, $\boldsymbol{\xi}$ belongs to $\gamma_\eta(\Theta_{\text{safe}} \ominus \mathcal{B}(0, \delta_1))$.

By the property of controlled invariant set, $u \in U(x)$ and x' exist which satisfy $x' = Ax + Bu \in X'$. Let the policy $\boldsymbol{\mu} = (\mu_1, \dots, \mu_N)$ be consistent with the modes of the individual subsystem's transition induced by u, and $\boldsymbol{\xi}$ moves to $\boldsymbol{\xi'}$ under $\boldsymbol{\mu}$ (*i.e.* $\boldsymbol{\xi} \xrightarrow[\eta,\times N]{\boldsymbol{\mu}} \boldsymbol{\xi'}$). Then, $\boldsymbol{\mu}$ satisfies (10) and $\boldsymbol{\xi'}$ is equal to $\Theta_{\text{map}}(X')$. Thus, $\boldsymbol{\xi'}$ is an element of $\Theta_{\text{map}}(X')$.

Therefore, $\Theta_{\text{map}}(X')$ is a controlled invariant set of S_η. □

Proofs for Section IV :

Let $L_q(\beta')$ be q-step shifted cycle assignment from $\beta' = (\beta'_1, \dots, \beta'_l)$ which is defined as follows.

$$L_{j+lq'}(\beta') = (\beta'_{j+1}, \dots, \beta'_l, \beta'_1, \dots, \beta'_j) \quad \forall j \in [l],\ q' \in \mathbb{N}_0,$$

Lemma 4. *Suppose that $\boldsymbol{\beta} = (\beta_1, \dots, \beta_n) \in \mathbb{N}_0^{\sum_j l_j}$ is an element of Ω. For any integer q, $\boldsymbol{L}_q(\boldsymbol{\beta}) = (L_q(\beta_1), \dots, L_q(\beta_n))$ also belongs to Ω.*

Proof of Lemma 4. For any integer s, $\sum_{l=1}^{l_j} L_s(\beta_j)_l = \sum_{l=1}^{l_j} \beta_{jl}$ holds. Therefore, the following is easily obtained.

$$\sum_{j=1}^{n}\sum_{l=1}^{l_j} L_s(\beta_j)_l = \sum_{j=1}^{n}\sum_{l=1}^{l_i} \beta_j(l) = N \tag{34}$$

In addition, the following inequalities hold.

$$\begin{aligned} &\sum_{j=1}^{n} \sum_{\substack{l \text{ s.t.} \\ \tilde{\mu}_{C_j,l}=m}} L_s(\beta_j)_{(-q+l \mod l_j)+1} \\ &= \sum_{j=1}^{n} \sum_{\substack{l \text{ s.t.} \\ \tilde{\mu}_{C_j,l}=m}} \beta_j((-q+l-s \mod l_j)+1) \end{aligned} \tag{35}$$

Using (14) to (35), the following is obtained

$$\begin{aligned} \underline{N}_m \le \sum_{j=1}^{n} \sum_{\substack{l \text{ s.t.} \\ \tilde{\mu}_{C_j,l}=m}} L_s(\beta_j)_{(-q+l \mod l_j)+1} \le \overline{N}_m \\ \forall q \in [\text{lcm}(l_1, \dots, l_n)],\ \forall m \in [M], \end{aligned} \tag{36}$$

From (34) and (36), $\boldsymbol{L}_k(\boldsymbol{\beta})$ satisfies all the conditions in the definition (15), and belongs to Ω. □

Proof of Theorem 3. We first show that X_{cyc} is a subset of X_{safe}. For any element x of X_{cyc}, there exists $\boldsymbol{\beta} = (\beta_1, \dots, \beta_n) \in \Omega$ which satisfies $x = \sum_{j=1}^{n} \Phi_{C_j,\tilde{\mu}_j}(\beta_j)$. From the definition of Ω in (15) and the mapping function (16), the following is obtained.

$$\begin{aligned} \sum_{k=1}^{K} x_{m,k} &= \sum_{k=1}^{K}\sum_{j=1}^{n} \Phi_{C_j,\tilde{\mu}_j}(\beta_j)_{m,k} \\ &= \sum_{k=1}^{K}\sum_{j=1}^{n} \sum_{(l:\tilde{\nu}_{C_j,l}=\nu_i,\tilde{\mu}_{jl}=m)} \beta_j(l) \\ &= \sum_{j=1}^{n} \left(\sum_{i=1}^{K} \sum_{(l:\tilde{\nu}_{C_j,l}=\nu_i,\tilde{\mu}_{jl}=m)} \beta_j(l) \right) \\ &= \sum_{j=1}^{n} \sum_{l:\tilde{\mu}_{jl}=m} \beta_j(l) \end{aligned}$$

Since β_j satisfy the inequalities in (14), we obtain the following.

$$\underline{N}_m \le \sum_{k=1}^{K} x_{k,m} \le \overline{N}_m \quad \forall m \in [M] \tag{37}$$

Moreover, the following holds by using $\tilde{\nu}_{C_j,l} \neq \nu_k$ for all $k \in \tilde{I}$

$$x_{m,k} = \sum_{j=1}^{n} \sum_{(l:\tilde{\nu}_{C_j,l}=\nu_k,\tilde{\mu}_{jl}=m)} \beta_j(l) = 0 \quad \forall k \in \tilde{\mathcal{I}}. \tag{38}$$

By (37) and (38), x is an element of X_{safe}. Therefore, X_{cyc} is a subset of X_{safe}.

Next, we show the recurrence property of X_{cyc}. Let $\boldsymbol{\beta'}$ be the one-step shifted cycle assignment from $\boldsymbol{\beta}$ (i.e. $\boldsymbol{\beta'} = (L_1(\beta_1), \dots, L_1(\beta_n))$). By the Lemma 4, it is obvious that $\boldsymbol{\beta'}$ is also an element of Ω, and thus $x' = \sum_j \Phi_{C_j,\tilde{\mu}_j}(\beta'_j)$ belongs to X_{cyc}. Since transition from x to x' can be conducted by one-step cycle shifting, it is obvious that there is an input $u \in U(x)$ driving x to x' (i.e. $x' = Ax + Bu$) which shows the recurrence property of X_{cyc}.

Therefore, X_{cyc} is a controlled invariant set in X_{safe}. □

Proofs for section V:

Proof of Theorem 4. It can be shown by principle of mathematical induction.

1) When $t = 0$,
   Since X_{cyc} is a controlled invariant set, there exist a sequence of states $x^\tau \in X_{\text{cyc}}$ ($\tau \in \{0, \dots, h\}$), inputs $u^\tau \in U(x^\tau)$ ($\tau \in \{0, \dots, h-1\}$), and $(\beta'_1, \dots, \beta'_n) \in \Omega$ which satisfy

$$\begin{aligned} &x^0 = x(0) \\ &x^{\tau+1} = Ax^\tau + Bu^\tau \qquad \forall \tau \in \{0, \dots, h-1\} \\ &x^\tau \in X_{\text{cyc}} \qquad \forall \tau \in \{0, \dots, h\} \\ &x^h = \sum_{j=1}^{n} \Phi_{C_j,\tilde{\mu}_j}(\beta'_j). \end{aligned}$$

   If we substitute x^τ for $x^{\tau|0}$, u^τ for $u^{\tau|0}$, and β'_j for β_j, all the constraints of the program hold. Therefore, this program has a feasible solution, and an optimal solution $x^{\tau|0} = x^{\tau|0}_*$, $u^{\tau|0} = u^{\tau|0}_*$, and $\beta_j = \beta_{j*}$ exist.

Now assume that there exists $\tau' \in \{0, \dots, h\}$ such that $x_*^{\tau'|0}$ does not belong to $X^{\max}$. Since every state outside of $X^{\max}$ cannot be driven into $X^{\max}$ by any input, $x_*^{\tau|0}$ for any $\tau \in \{\tau', \dots, h\}$ does not belong to $X^{\max}$.
However, from $x_*^{h|0} = \sum_{j=1}^{n} \Phi_{C_j,\tilde{\mu}_j}(\beta'_{j*})$ where $(\beta_{1*}, \dots, \beta_{n*})$ belongs to Ω, $x_*^{h|0}$ is an element of X_{cyc} . Since X_{cyc} is a subset of $X^{\max}$, $x_*^{h|0}$ belongs to $X^{\max}$ which is a contradiction.
By the contradiction, $x_*^{\tau|0}$ belongs to $X^{\max}$ for any $\tau \in \{0, \dots, h\}$. Thus, the state at the next time step $x(1) = x_*^{1|0}$ belongs to $X^{\max}$. Therefore, the statement has been proved for the case of $t = 0$.

2) Assume that the statement holds when $t = t' \in \mathbb{N}_0$, which is equivalent to the statement that the program at $t = t'$ is feasible and $x(t'+1)$ belongs to $X^{\max}$.
Let the optimal solution obtained at $t = t'$ be $x_*^{t'+\tau|t'}$, $u_*^{t'+\tau|t'}$. Since $x_*^{t'+h|t'}$ is an element of X_{cyc}, there exists $u^{t'+h} \in U(x_*^{t'+h|t'})$ such that $x^{t'+h+1} = Ax_*^{t'+h|t'} + Bu^{t'+h}$ belongs to X_{cyc}. If we substitute the variables of the program at $t = t'+1$ as follows,

$$
\begin{aligned}
x^{t'+\tau+1|t'+1} &= x_*^{t'+\tau+1|t'} \quad \tau \in \{0, \dots, h-1\} \\
x^{t'+h+1|t'+1} &= x^{t'+h+1} \\
u^{t'+\tau+1|t'+1} &= x_*^{t'+\tau+1|t'} \quad \tau \in \{0, \dots, h-2\} \\
u^{t'+h|t'+1} &= u^{t'+h},
\end{aligned}
$$

the constraints of the program are satisfied. Therefore, the program at $t = t'+1$ is feasible.
Moreover, we can prove that $x(t'+2) = x_*^{t'+2|t'+1}$ from this program at $t = t'+1$ belongs to $X^{\max}$ by taking the similar procedure as 1).
Therefore, the statement holds at $t = t'+1$.

By 1) and 2), the statement is proven. □