

Control Synthesis for Large Collections of Systems with Mode-Counting Constraints

Petter Nilsson

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

Necmiye Ozay

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

ABSTRACT

Given a large homogeneous collection of switched systems, we consider a novel class of safety constraints, called mode-counting constraints, that impose restrictions on the number of systems that are in a particular mode. We propose an approach for synthesizing correct-by-construction switching protocols to enforce such constraints over time. Our approach starts by constructing an approximately bisimilar abstraction of the individual system model. Then, we show that the aggregate behavior of the collection can be represented by a linear system, whose system matrices are induced by the transition graph of the abstraction. Finally, the control synthesis problem with mode-counting constraints is reduced to a cycle assignment problem on the transition graph. One salient feature of the proposed approach is its scalability; the computational complexity is independent of the number of systems involved. We illustrate this approach on the problem of coordinating a large collection of thermostatically controlled loads while ensuring a bound on the number of loads that are extracting power from the electricity grid at any given time.

Keywords

Control synthesis; Control of switched systems; Abstraction; Energy applications

1. INTRODUCTION

Formal methods-based correct-by-construction control synthesis has attracted considerable attention in recent years as a principled means to ensure that the closed-loop behavior of a dynamical system satisfies certain high-level specifications [23]. This synthesis paradigm is particularly convenient to incorporate state and input constraints, and to handle continuous as well as hybrid systems in a unified manner. Unfortunately, many correct-by-construction control synthesis methods suffer from the curse of dimensionality. Therefore, it is crucial to take into account the special structure of the

problem at hand to improve scalability, as demonstrated in recent work in the context of synthesis [20, 3, 5, 4] and in the context of verification [10].

In this paper, we consider the problem of synthesizing controllers to coordinate a collection of N identical switched systems subject to a novel class of safety constraints, called mode-counting constraints. Mode-counting constraints impose restrictions on the number of systems that are in a particular mode at any given time. Although the individual systems are dynamically decoupled, coordination among them is required to ensure that the mode-counting constraints are not violated. In order to address this problem, we first construct an ϵ -approximately bisimilar finite abstraction for an individual system. Then, we construct a linear system that models the dynamics of the population of systems on the nodes of a graph that the abstraction induces. Finally, by restricting the controllers to have a prefix-suffix form, we reduce the control synthesis problem to an (integer) linear program, constraints of which are determined by the abstraction graph, the population dynamics, and the mode-counting constraints. Tightness and complexity of the approach and its various relaxations are discussed. One salient feature of the approach is that its complexity is almost independent of the number of systems N . Examples are provided where N is in the order of tens of thousands.

The motivation for this problem comes from coordination of thermostatically controlled loads (TCLs) to help improve power grid operations [12, 26, 15, 21]. Thermostatically controlled loads include air conditioners, water heaters, refrigerators, etc., that operate within a desired temperature range. The idea in TCL coordination is to use the flexibility within these temperature ranges to track a power signal by appropriately turning the loads on and off. Aggregate TCL population models based on state-space partitioning are proposed in [12, 26, 21], which are similar to the abstractions that we develop in this paper. Yet, the relation between the actual and aggregate models has not been quantified formally in the earlier works as is done in this paper, except in a stochastic setting in [6]. Moreover, in all of these papers, aggregate models are used for control design with the objective to track a power signal. The more traditional control design techniques that are used do not take into account any hard constraints on overall power consumption. Nevertheless, it is well-known that if all air conditioners are turned on at the same time, for instance during some transient phase, it can result in distribution line overload [18]. Similarly, it might be desirable to have at least a certain

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

HSCC'16, April 12-14, 2016, Vienna, Austria

© 2016 ACM. ISBN 978-1-4503-3955-1/16/04...\$15.00

DOI: <http://dx.doi.org/10.1145/2883817.2883831>

amount of loads on at any given time to utilize the TCL collection as a “battery” that stores energy from renewable resources [12]. Motivated by these types of requirements or hard constraints, which can be precisely captured by mode-counting constraints, we propose a novel control synthesis approach for enforcing them.

This paper is structured in the following way. The subsequent Section 2 introduces notation, along with the bisimilarity notion for a general class of transition systems. We proceed with our problem statement in Section 3, followed by a solution approach in Section 4. To motivate the applicability of this approach, two synthesis examples are provided in Section 5. There are several possible extensions of varying difficulty, which we outline in Section 6, before the paper is concluded in Section 7.

2. PRELIMINARIES

In the following paragraphs we introduce some notation that is used throughout the paper. To express a finite set of positive integers, we write $[N] = \{1, \dots, N\}$. The indicator function of a set A is denoted $\mathbb{1}_A(x)$ and is equal to 1 if $x \in A$ and 0 otherwise. The identity function is written as Id . For two sets A and B , we write $A \oplus B = \{a + b : a \in A, b \in B\}$, and $A \ominus B$ is the largest (w.r.t. inclusion) solution to $X \oplus B = A$.

To denote the floor of a number, $\lfloor \cdot \rfloor$ is used. We use the same notation for vectors, where the floor operation is performed component-wise. We write the infinity norm as $\| \cdot \|$, and the ϵ -ball centered at the point x is denoted as $\mathcal{B}(x, \epsilon) = \{y : \|y - x\| \leq \epsilon\}$. The vector of all 1’s is written as $\mathbf{1}$.

Given an ODE $\dot{x} = f_m(x)$, the corresponding flow operator is denoted by $\phi_t^m(x)$ and has the properties that $\phi_0^m(x) = x$, $\frac{d}{dt} \phi_t^m(x) = f_m(\phi_t^m(x))$. So called \mathcal{KL} -functions are related to nonlinear stability theory and are functions $\beta : \mathbb{R}^+ \times \mathbb{R}^+ \rightarrow \mathbb{R}^+$ that are strictly increasing from 0 in the first argument and decreasingly converging to 0 in the second argument.

We also employ the following definition for a *transition system*, which captures systems with both continuous and discrete state spaces.

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

As an intuitive notation for transitions, we write $q_1 \xrightarrow{u} q_2$ to indicate that $(q_1, u, q_2) \in \longrightarrow$. When we speak about trajectories of a transition system, we mean a sequence q_1, q_2, \dots of states in Q , with the property that $q_i \xrightarrow{u_i} q_{i+1}$ for some $u_i \in U$.

We adopt the notion of approximate bisimilarity of transition systems from [23].

Definition 2. Two transition systems $(Q_1, U_1, \xrightarrow{1}, Y_1)$ and $(Q_2, U_2, \xrightarrow{2}, Y_2)$ are ϵ -**approximately bisimilar** if there exists a relation $R \subset Q_1 \times Q_2$ such that for all $(q_1, q_2) \in R$,

- $\|Y_1(q_1) - Y_2(q_2)\| \leq \epsilon$,
- if $q_1 \xrightarrow{u_1} p_1$, there exists $q_2 \xrightarrow{u_2} p_2$ s.t. $(p_1, p_2) \in R$,
- if $q_2 \xrightarrow{u_2} p_2$, there exists $q_1 \xrightarrow{u_1} p_1$ s.t. $(p_1, p_2) \in R$.

3. PROBLEM STATEMENT

We consider a family of identical switched individual systems. The state x_i of system number i obeys the switched differential equation

$$\Sigma : \dot{x}_i(t) = f_{\sigma_i(t)}(x_i(t)), \quad \sigma_i : \mathbb{R} \mapsto [M], \quad (1)$$

where $\sigma_i(t)$ is the mode for system i at time t . We restrict attention to a compact domain of interest $\mathcal{D}^d \subset \mathbb{R}^d$. The *time-sampled* analogue of (1) on \mathcal{D}^d , Σ_τ , is defined as the transition system

$$\Sigma_\tau = (\mathcal{D}^d, [M], \xrightarrow{\tau}, \text{Id}_{\mathbb{R}^d}), \quad (2)$$

where $x_1 \xrightarrow{\tau} x_2$ if and only if $\phi_\tau^m(x_1) = x_2$.

For a time-sampled system, we define the *mode-counting problem with mode safety constraints* as follows.

Problem 1. Given a family of N plants with states $\{x_i\}_{i \in [N]}$ obeying the dynamics of (2), mode-specific unsafe sets $\{\mathcal{U}_m\}_{m \in [M]}$, and mode-counting bounds $\{\underline{K}_m, \overline{K}_m\}_{m \in [M]}$, synthesize an aggregate switching policy $\{\sigma_i\}_{i \in [N]}$, such that for all $m \in [M]$ and all trajectories $x_i(0), x_i(1), \dots$

$$(x_i(s), \sigma_i(s)) \notin \mathcal{U}_m \times \{m\} \quad \forall i \in [N], s \in \mathbb{N}, \quad (3)$$

$$\underline{K}_m \leq \sum_{i=1}^N \mathbb{1}_{\{m\}}(\sigma_i(s)) \leq \overline{K}_m \quad \forall s \in \mathbb{N}. \quad (4)$$

An instance of this problem is referred to with the tuple $(N, \Sigma_\tau, \{\mathcal{U}_m, \underline{K}_m, \overline{K}_m\}_{m \in [M]})$.

An analogous problem could easily be defined for the continuous-time system (1). The difference is that in the time-sampled system, mode switches can only happen at the sampling instants. Furthermore, a solution where mode safety constraints are violated in between samplings (but satisfied at sample instants) is still a valid solution to Problem 1. If such inter-sample safety violations are unacceptable, the unsafe sets can be expanded by some margin determined by the dynamics to ensure safety satisfaction for all $t \in \mathbb{R}$.

In order to leverage the bisimulation theory presented below, we make the following assumption.

Assumption 1. The system $\dot{x} = f_m(x)$ is forward complete and incrementally stable for all $m \in [M]$. That is, there exist \mathcal{KL} -functions β_m for $m \in [M]$, such that

$$\|\phi_t^m(x) - \phi_t^m(y)\| \leq \beta_m(\|x - y\|, t). \quad (5)$$

4. SOLUTION APPROACH

We propose an abstraction-based solution to Problem 1. In Section 4.1, we describe how the dynamics can be abstracted so that the time-sampled system (2) is ϵ -approximately bisimilar to the abstraction. We can then define the discrete analogue to Problem 1, and show that existence of a solution on the abstraction is equivalent to existence of a solution for the time-sampled system, up to an error margin ϵ .

We then proceed by presenting a way to solve the discrete mode-counting problem. Our method relies on reasoning about the discrete transition graph, we derive some properties of this graph in Section 4.2. Subsequently, in Section 4.3 we give a linear program from which, if feasible, a solution can be extracted. We analyze some aspects of our approach in Section 4.4.

4.1 Abstractions and aggregation

Let Σ be a system of the form (1). For a parameter $\eta > 0$ we define an abstraction function $\alpha_\eta : \mathcal{D}^d \rightarrow \mathcal{D}^d \oplus \mathcal{B}(0, \eta/2)$ as

$$\alpha_\eta(x) = \eta \cdot \left\lfloor \frac{x}{\eta} \right\rfloor + \frac{\eta}{2} \mathbf{1}. \quad (6)$$

This function takes a finite number of values (since \mathcal{D}^d is compact) and is constant on hyper boxes of side η . Using this function, we define a transition system $\Sigma_{\tau, \eta} = (Q_\eta, U, \xrightarrow{\tau, \eta}, Y)$, in the following referred to as the *time-state abstraction* of Σ , as follows:

1. $Q_\eta = \alpha_\eta(\mathcal{D}^d)$,
2. $U = [M]$,
3. $q_1 \xrightarrow{\tau, \eta} q_2$ if and only if $\alpha_\eta(\phi_\tau^m(q_1)) = q_2$,
4. $Y(q) = q$.

In essence, the domain is partitioned into uniform boxes, and every mode is simulated during time τ starting at the center of each box in order to determine transition relations.

Remark 1. The transition system $\Sigma_{\tau, \eta}$ is deterministic. That is, for each state q_1 and action m there exists (at most) one successor state q_2 . In the event that the ODE solution starting in q_1 for mode m exits the domain, i.e. $\phi_\tau^m(q_1) \notin \mathcal{D}^d$, the action m is disabled at the discrete state q_1 .

Using a result from [17], it follows that an abstraction constructed in this way is bisimilar to the time-sampled system (2) if a certain inequality holds.

THEOREM 1 ([17]). *Suppose Assumption 1 holds and let β_m be \mathcal{KL} -functions satisfying (5) for $m \in [M]$. If the inequality $\beta_m(\epsilon, \tau) \leq \epsilon - \eta/2$ holds for all $m \in [M]$, then Σ_τ and $\Sigma_{\tau, \eta}$ are ϵ -approximately bisimilar.*

It is known that the trajectories of ϵ -approximately bisimilar systems remain within distance ϵ of each other, when ‘‘similar’’ control actions are chosen at each time instant (c.f. Definition 2) [8]. This fact is the key to establishing relations between existence of solutions for Σ_τ and existence of solutions for $\Sigma_{\tau, \epsilon}$ later in the text.

Aggregation dynamics:

The abstraction constructed above can be viewed as a graph $G = (V, E)$, where the edges are labeled according to the mode. In the following, we take an arbitrary such *mode-transition graph* (not necessarily obtained from an abstraction) and define a linear system that represents the aggregate dynamics on such a graph.

To this end, we assume that there are N systems which simultaneously move around on G . Assume that the number of nodes in the graph is K , i.e., $|V| = K$, and that they are labeled ν_k for $k \in [K]$. Furthermore, assume that there are M different mode-labels numbered from 1 to M , and let $l_E : E \rightarrow [M]$ be the function which assigns a mode to each edge.

We then introduce $K \times M$ aggregate states labeled w_k^m for $k \in [K], m \in [M]$, that describe the number of individual systems that are *at node ν_k and in mode m* . By also introducing control actions $r_k^{m_1, m_2}$ that represent the number of

systems *at node ν_k that switch from mode m_1 to mode m_2* , the aggregate dynamics $\Sigma_{\tau, \eta}^N$ can be written as

$$(w_k^{m_1})^+ = \sum_{j \in \mathcal{N}_k^{m_1}} \left(w_j^{m_1} + \sum_{m_2 \in [M] \setminus \{m_1\}} (r_j^{m_2, m_1} - r_j^{m_1, m_2}) \right), \quad (7)$$

for $k \in [K], m_1 \in [M]$, where \mathcal{N}_k^m is the set of predecessors of the k 'th node under the action m . That is,

$$\mathcal{N}_k^m = \{i \in [K] : (\nu_i, \nu_k) \in E, l_E(\nu_i, \nu_k) = m\}.$$

We constrain the control actions $r_k^{m_1, m_2}$ so that for all $m_1 \in [M]$,

$$0 \leq \sum_{m_2} r_k^{m_1, m_2} \leq w_k^{m_1}, \quad (8)$$

which ensures the continued positivity of the states. That is, $(w_k^m)^+ \geq 0$ for $k \in [K], m \in [M]$. Furthermore, we have the invariant $\sum_k \sum_m w_k^m = N$ over time, where N is the number of individual systems. In the following, we use the compact notation

$$\mathbf{w}^+ = A\mathbf{w} + B\mathbf{r}, \quad (9)$$

where $\mathbf{w}_i = w_k^m$ if $i = K(m-1) + k$, to denote this system. Here A and B are composed of the incidence matrices of the mode-transition graph for each mode. If A_1, \dots, A_M are the incidence matrices, then A is the block-diagonal matrix with block diagonal given by $[A_1, \dots, A_M]$, and B is a block matrix composed by the same A_m 's. The state space \mathcal{W} and admissible control space \mathcal{R} of this system are

$$\begin{aligned} \mathcal{W} &= \left\{ \mathbf{w} \in \mathbb{N}^{KM} : \sum \mathbf{w} = N \right\}, \\ \mathcal{R} &= \left\{ \mathbf{r} \in \mathbb{N}^{M(M-1)K} : (8) \text{ holds for } (\mathbf{w}, \mathbf{r}) \right\}. \end{aligned}$$

For a given state $\mathbf{w} \in \mathbb{N}^{KM}$ of (7), we can define a *condensed state* $\lambda \in \mathbb{N}^K$ as $\lambda_k = \sum_m w_k^m$, which counts the number of systems at node ν_k irrespective of mode. We introduce a mapping Λ that takes \mathbf{w} to its condensed counterpart, $\mathbf{w} \xrightarrow{\Lambda} [\sum_m w_1^m, \dots, \sum_m w_K^m] \in \mathbb{N}^K$.

We are now in a position to define the discrete analogue of Problem 1.

Problem 2. Given a discrete-time linear system Γ of the aggregate form (9) built from a mode-transition graph, an initial condensed state λ_0 , unsafe subsets $\{U_m\}_{m \in [M]}$ of $[K]$, and mode-counting bounds $\{\underline{K}_m, \overline{K}_m\}_{m \in [M]}$, synthesize an initial mode assignment $\mathbf{w}(0)$ such that $\Lambda(\mathbf{w}(0)) = \lambda_0$, and a controller $\mathbf{r} : \mathcal{W}^* \rightarrow \mathcal{R}$, such that for all $m \in [M]$ and all $s \in \mathbb{N}$,

$$w_k^m(s) = 0 \quad \forall k \in U_m, \quad (10)$$

$$\underline{K}_m \leq \sum_{i \in [N]} w_i^m(s) \leq \overline{K}_m. \quad (11)$$

We will refer to an instance of this problem as $(N, \Gamma, \{(U_m, \underline{K}_m, \overline{K}_m)\}_{m \in [M]})$.

Since no disturbance is included in this formulation, it is enough to consider open-loop controllers. In the following, we will denote such controllers with $\mathbf{r}(s)$ for $s = 0, 1, \dots$. Now, using the approximate bisimilarity between Σ_τ and $\Sigma_{\tau, \eta}$ we can state the following results.

THEOREM 2. Let $\Sigma_\tau, \Sigma_{\tau,\eta}$ be the time-sampled and time-state abstracted dynamics of a system Σ of the form (1), such that Σ_τ and $\Sigma_{\tau,\eta}$ are ϵ -approximately bisimilar. Let α_η be the abstraction function for $\Sigma_{\tau,\eta}$. Furthermore, let $\Sigma_{\tau,\eta}^N$ be the aggregate sampled dynamics of the form (9) obtained from the mode-transition graph of $\Sigma_{\tau,\eta}$.

Then, if there exists a solution to the instance $(N, \Sigma_{\tau,\eta}^N, \{(\alpha_\eta(\mathcal{U}_m \oplus \mathcal{B}(0, \epsilon)), \underline{K}_m, \overline{K}_m)\}_{m \in [M]})$ of Problem 2, there exists a solution to the instance $(N, \Sigma_\tau, \{(\mathcal{U}_m, \underline{K}_m, \overline{K}_m)\}_{m \in [M]})$ of Problem 1.

PROOF. Assume that a controller $\mathbf{r}(s)$ constitutes a solution to Problem 2, i.e. produces a closed-loop trajectory $\mathbf{w}(s)$ of $\Sigma_{\tau,\eta}^N$ that satisfies (10) - (11). Then we can extract N individual trajectories $\xi_i(0), \xi_i(1), \dots$, of $\Sigma_{\tau,\eta}$. By implementing switching protocols $\{\sigma_i\}_{i \in [N]}$ on the continuous level that agree with the modes of the individual trajectories, it immediately follows that (4) is satisfied also for the resulting trajectories $x_i(0), x_i(1), \dots$, of Σ_τ with the same bounds $\{\underline{K}_m, \overline{K}_m\}_{m \in [M]}$. Furthermore, since $\Sigma_{\tau,\eta}$ and Σ_τ are ϵ -approximately bisimilar, it holds that for all $s \in \mathbb{N}$, $\|\xi_i(s) - x_i(s)\| \leq \epsilon$, therefore $\xi_i(s) \notin \alpha_\eta(\mathcal{U}_m \oplus \mathcal{B}(0, \epsilon))$ implies that $x_i(s) \notin \mathcal{U}_m$ so (3) also holds. \square

Theorem 2 ensures that if a discrete-state solution exists for unsafe sets with an added ϵ -margin, then it can be implemented on the continuous-state level while preserving correctness. Conversely, we can show that within an $(\epsilon + \eta/2)$ -margin, existence of a discrete-state solution is also a necessary condition.

THEOREM 3. Under the same assumptions as in Theorem 2, if there is no solution to the instance $(N, \Sigma_{\tau,\eta}^N, \{(\alpha_\eta(\mathcal{U}_m \oplus \mathcal{B}(0, \epsilon + \eta/2)), \underline{K}_m, \overline{K}_m)\}_{m \in [M]})$ of Problem 2, there is no solution to the instance $(N, \Sigma_\tau, \{(\mathcal{U}_m, \underline{K}_m, \overline{K}_m)\}_{m \in [M]})$ of Problem 1.

PROOF. Assume that a solution exists to the instance $(N, \Sigma_\tau, \{(\mathcal{U}_m, \underline{K}_m, \overline{K}_m)\}_{m \in [M]})$ of Problem 1, but that no solution exists to the instance $(N, \Sigma_{\tau,\eta}^N, \{(\alpha_\eta(\mathcal{U}_m \oplus \mathcal{B}(0, \epsilon + \eta/2)), \underline{K}_m, \overline{K}_m)\}_{m \in [M]})$ of Problem 2.

Like in the previous result, the continuous-state switching protocol can be implemented on the discrete-state system, and the resulting trajectories will deviate at most ϵ due to approximate bisimilarity. Thus a continuous-state solution with unsafe sets \mathcal{U}_m can be implemented on the discrete-state system and will be correct for unsafe sets $\mathcal{U}_m \oplus \mathcal{B}(0, \epsilon) \supset \alpha_\eta(\mathcal{U}_m \oplus \mathcal{B}(0, \epsilon + \eta/2))$, which is a contradiction. \square

4.2 Mode-transition graph properties

We now state some results that connect the properties of the underlying mode-transition graph to the aggregate dynamics. Let a system Γ of the form (9) obtained from a mode-transition graph $G = (V, E)$ with M different modes be given, such that every node $\nu \in V$ has at most¹ M outgoing edges. We proceed by stating some results about this graph. Assume that the nodes are labeled with integers 1 through K , i.e., $V = \{\nu_k : k \in [K]\}$.

We recall some standard definitions from graph theory. A *path* in G is a list of edges $(\nu_1, \nu_2), (\nu_2, \nu_3), \dots, (\nu_{j-1}, \nu_j)$. If the first and last nodes are equal, i.e. $\nu_j = \nu_1$, the path is a *cycle*. A cycle is *simple* if it visits every node at most one

¹Some modes can be disabled at certain nodes.

time. For a subset of nodes $D \subset V$ (or the corresponding subgraph, which we use interchangeably), it is said to be *strongly connected* if for each pair $\nu_1, \nu_2 \in D$, there exists a path from ν_1 to ν_2 . Any directed graph can be decomposed into strongly connected components. The *period* of a subgraph D is the greatest common divisor of all cycles in D . A subgraph D is called *aperiodic* if it has period one.

To ensure infinite horizon correctness we will focus on solution trajectories that consist of a finite prefix phase, and a periodic suffix phase defined on cycles. The purpose of the prefix phase is to steer the system to a “nice” state from which there is a periodic solution. Before explaining such solutions, we first define a concept of controllability on a subset of nodes D for the aggregate dynamics Γ . Similarly as for controllability of linear systems on a subspace, controllability of Γ on D means that the system can be steered between any two aggregate states with support on D .

Definition 3. A subset of nodes $D \subset V$ is **completely controllable** for Γ if for any two condensed states λ^1, λ^2 with support² on D such that $\sum \lambda^1 = \sum \lambda^2$, there exists a finite horizon T , controls $\{\mathbf{r}(s)\}_{s=0}^{T-1}$, and states $\{\mathbf{w}(s)\}_{s=0}^T$ such that $\Lambda(\mathbf{w}(0)) = \lambda^1, \Lambda(\mathbf{w}(T)) = \lambda^2$, and $\mathbf{w}(s+1) = \mathbf{A}\mathbf{w}(s) + \mathbf{B}\mathbf{r}(s)$ for $s = 0, \dots, T-1$.

As the following result shows, the aggregate system Γ of the form (9) has a lot of “control freedom” on strongly connected, aperiodic components.

THEOREM 4. If a strongly connected component D is aperiodic, it is completely controllable for Γ .

PROOF. It is known that the incidence matrix A_D of an aperiodic, strongly connected graph D is primitive [22], i.e., there exists an integer T such that all entries of A_D^T are positive. This means that for each vertex pair (ν_j, ν_l) , there exists a path of length T that connects them. Thus, by sending p_{jl} systems along paths $\nu_j \rightarrow \nu_l$ such that $\sum_i p_{ji} = \lambda_j^1$ and $\sum_j p_{jl} = \lambda_l^2$, the condensed state at time T will be equal to λ^2 . We can define a control $\mathbf{r}(s)$ that realizes these paths by switching the correct number of systems at each node over time. \square

In the case of periodicity, it is not possible to reach every condensed state. In fact, the periodicity will preserve the parity structure of the initial state. However, within this restriction, the system is still controllable in the following sense. If a strongly connected component D has period P , its nodes can be labeled with a function $L_P : D \rightarrow \{0, 1, \dots, P-1\}$ such that a node ν_1 with $L_P(\nu_1) = p$ only has edges to nodes ν_2 with $L_P(\nu_2) = i + 1 \pmod{P}$. Let D_0, \dots, D_{P-1} be the subsets of nodes induced by the equivalence relation $\nu_1 \sim \nu_2$ iff $L_P(\nu_1) = L_P(\nu_2)$.

COROLLARY 1. The subsets of nodes D_i as constructed above are completely controllable for Γ .

PROOF. We can connect the nodes in D_i with edges that correspond to paths of length P in D . By construction, the resulting graphs are aperiodic, so the previous result applies. \square

²That a condensed state λ has support on D means that $\lambda_k = 0$ for $\nu_k \notin D$.




Figure 1: Illustration of the assignment $\alpha = [3, 2, 3, 4, 4]$ on a cycle C of length 5, with two modes 1 (blue) and 2 (red). On the left, the mode-1-count is $3+4+4 = 11$, while the mode-2-count is $3+2=5$. After circulating the assignment one step, as displayed to the right, the mode-1-count is $2+3+4 = 9$ and the mode-2-count is $4+3 = 7$. Over all possible circulations, the minimal mode-1-count is 8, and the maximal mode-1-count is 11, so $\bar{\Psi}^1(C, \alpha) = 11$ and $\underline{\Psi}^1(C, \alpha) = 8$. Similarly, $\bar{\Psi}^2(C, \alpha) = 8$ and $\underline{\Psi}^2(C, \alpha) = 5$.

In order to treat the periodic suffix part of a solution, we introduce the following notation pertaining to a cycle $C = (\nu_1, \nu_2), \dots, (\nu_{J-1}, \nu_1)$ of length given by $|C| = J - 1$.

Definition 4. The **outgoing mode** $\Xi_C(\nu_j)$ is the mode of the outgoing edge at ν_j , that is,

$$\Xi_C(\nu_j) = \begin{cases} l_E(\nu_j, \nu_1) & \text{if } j = |C|, \\ l_E(\nu_j, \nu_{(j+1)}) & \text{otherwise.} \end{cases}$$

Definition 5. An **assignment** to a cycle C is a function $\alpha : [|C|] \rightarrow \mathbb{R}^+$.

Definition 6. An **integer assignment** to a cycle C is an assignment α to C such that $\alpha(j)$ is an integer for $j \in [|C|]$.

Definition 7. The **maximal mode- m -count** for a cycle C with assignment α is

$$\bar{\Psi}^m(C, \alpha) = \max_{k \in [|C|]} \sum_{i: \Xi_C(\nu_i) = m} \alpha((k+i) \bmod |C|).$$

For a given assignment, the maximal mode- m -count denotes the maximal number of systems that are simultaneously in mode m when the assignment α circulates around C .

Definition 8. The **minimal mode- m -count** for a cycle C with assignment α is

$$\underline{\Psi}^m(C, \alpha) = \min_{k \in [|C|]} \sum_{i: \Xi_C(\nu_i) = m} \alpha((k+i) \bmod |C|).$$

These functions are illustrated in Figure 1 for an example cycle-assignment pair. Finally, we define a function $\Phi_C : \mathbb{R}^{|C|} \rightarrow \mathbb{R}^{|V|}$ that for a cycle C maps the values of a cycle assignment α to the corresponding nodes in the graph.

$$\Phi_C(\alpha)_k = \begin{cases} \alpha(j) & \text{if } \nu_j \text{ in } C \text{ corresponds to } \nu_k \text{ in } V, \\ 0 & \text{otherwise.} \end{cases}$$

4.3 Prefix-suffix strategies as a linear program

We are now ready to define the type of strategies that we consider in this paper.

Definition 9. A control strategy for a condensed initial state λ_0 is of **prefix-suffix** type if it consists of an initial mode assignment $\mathbf{w}(0)$ s.t. $\Lambda(\mathbf{w}(0)) = \lambda_0$, a finite number of inputs $\mathbf{r}(0), \dots, \mathbf{r}(T-1)$, and a set of cycles $\{C_j\}_{j \in J}$ with assignments $\{\alpha_j\}_{j \in J}$ such that the cycles are populated with their respective cycle assignments at time T .

For given initial positions $\lambda_0 \in \mathbb{N}^K$, mode-counting bounds $\{\underline{K}_m, \bar{K}_m\}_{m \in [M]}$, a given set of cycles $\{C_j\}_{j \in J}$, and a horizon T , the following linear feasibility program searches for a prefix-suffix control strategy that solves Problem 2.

find $\alpha_1, \dots, \alpha_J$ cycle assignments,

$$\mathbf{r}(0), \dots, \mathbf{r}(T-1),$$

$$\mathbf{w}(0), \dots, \mathbf{w}(T),$$

$$\text{s.t. } \underline{K}_m \leq \sum_{k \in [K]} w_k^m(s) \leq \bar{K}_m, \quad s = 0, \dots, T, \quad (12a)$$

$$\underline{K}_m \leq \sum_j \bar{\Psi}^m(C_j, \alpha_j), \quad (12b)$$

$$\sum_j \bar{\Psi}^m(C_j, \alpha_j) \leq \bar{K}^m, \quad (12c)$$

$$\Lambda(\mathbf{w}(T)) = \sum_j \Phi_{C_j}(\alpha_j), \quad (12d)$$

$$\mathbf{w}(s+1) = \mathbf{A}\mathbf{w}(s) + \mathbf{B}\mathbf{r}(s), \quad s = 0, \dots, T-1, \quad (12e)$$

$$\Lambda(\mathbf{w}(0)) = \lambda_0, \quad (12f)$$

$$\sum_{m_2} r_j^{m_1, m_2} = w_j^{m_1} \text{ for all } j \in \bigcup_{i \in U_{m_1}} \mathcal{N}_i^{m_1}, \quad (12g)$$

$$r_j^{m_2, m_1} = 0 \text{ for all } m_2 \in [M], j \in U_{m_1}, \quad (12h)$$

$$\text{control constraints (8)}. \quad (12i)$$

We briefly describe the purpose of each constraint. Firstly, (12a) assures that mode-counting constraints are satisfied in the prefix phase, i.e., up to time T . Similarly, (12b)-(12c) restrict mode-counting in the suffix phase by ensuring that the sums of maximal and minimal mode-counts over all cycles are within the bounds. Eq. (12d) connects the prefix phase to the suffix phase by ensuring that the condensed state at time T agrees with the sum of all cycle assignments, while (12e) propagates the dynamics up to time T , and (12f) implies that the initial state $\mathbf{w}(0)$ must condense to the given initial condition λ_0 . The mode-safety constraints are taken care of through (12g)-(12h).

The maximal and minimal mode-counts for a given assignment α to a cycle C can be represented by the maximal and minimal entries of the product $Y_C^m \alpha$, where Y_C^m is the circulant $(0, 1)$ -matrix s.t.

$$[Y_C^m]_{ij} = \begin{cases} 1, & \text{if } \Xi_C(\nu_{j-(i-1) \bmod |C|}) = m, \\ 0, & \text{otherwise.} \end{cases}$$

To illustrate, the cycle C in Figure 1 has matrices

$$Y_C^1 = \begin{bmatrix} 0 & 0 & 1 & 1 & 1 \\ 1 & 0 & 0 & 1 & 1 \\ 1 & 1 & 0 & 0 & 1 \\ 1 & 1 & 1 & 0 & 0 \\ 0 & 1 & 1 & 1 & 0 \end{bmatrix}, \quad Y_C^2 = \begin{bmatrix} 1 & 1 & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 1 & 1 \\ 1 & 0 & 0 & 0 & 1 \end{bmatrix}.$$

Thus, the constraints $\underline{\Psi}^m(C, \alpha) \geq \underline{K}_m$, $\bar{\Psi}^m(C, \alpha) \leq \bar{K}_m$, can be enforced by the linear vector inequalities

$$\underline{K}_m \mathbf{1} \leq Y_C^m \alpha \leq \bar{K}_m \mathbf{1}.$$

The feasibility program (12) can be solved either as a normal linear program (LP) feasibility problem or as an integer linear program (ILP) feasibility problem. Since the size of it can be large in practice (for instance due to a fine-grained abstraction, see paragraph on complexity below), the ILP version may be impractical. Furthermore, the number of individual systems N may affect the difficulty of the ILP,

since a larger N increases the number of possible integer points. In the next section we discuss how feasible solutions to the ILP are related to feasible solutions of the LP. By construction, the following result holds.

PROPOSITION 1. *If a feasible integer solution of (12) exists, a solution to Problem 2 can be extracted.*

The formulation of (12) can be altered to capture different requirements. For instance, if the counts of certain modes are not important for the application, they can be omitted. An objective function can also be added to search for the tightest possible bounds, but in practice ILP solvers are typically more efficient in finding feasible points than in proving optimality.

How to select the cycle set.

The number of cycles in a graph is infinite, so in order to provide a finite number of cycles as input to (12), a selection must be made. As shown in Section 4.4 below, it is enough to consider the set of simple cycles in the LP case, whereas non-simple cycles may enable tighter bounds in the ILP case. However, even the number of simple cycles can be exponential in the size of the graph. For this reason, some randomized cycle selection is a reasonable choice in practice.

How to select the horizon T .

If mode-counting constraints during the prefix phase are disregarded, Theorem 4 implies the existence of a time T in which any assignment can be achieved. This time is upper bounded by a quadratic polynomial for general graphs, although more specialized bounds exist [13]. However, mode-counting constraints may exacerbate the upper bounds, so there is a trade-off between a short prefix horizon T and tight mode-counting constraints during the prefix phase.

Complexity of linear program.

The number of variables and constraints in the linear program above are

$$\mathcal{O}\left(M^2KT + \sum_{j \in J} |C_j|\right), \quad \mathcal{O}\left(MKT + \sum_{j \in J} |C_j|^2\right),$$

respectively. Taking the abstraction step into account, $K = \mathcal{O}((1/\eta)^d)$, where η is the state discretization parameter and d the state-space dimension for an individual system. Notably, the complexity does not depend on the number of systems N , which makes our approach attractive for large collections of relatively simple homogeneous systems.

4.4 Analysis

We now give sufficient conditions for the existence of a feasible solution to (12).

THEOREM 5. *If there exists an (integer) solution to Problem 2, then there exists a finite horizon T such that the LP version of (12) is feasible when solved for the cycle set consisting of all simple cycles.*

The proof of this result is divided into two parts. First, we show in Lemma 1 that it is sufficient to consider solutions consisting of a transient and a periodic (i.e., cyclic) part. Secondly, we show in Lemma 2 that an assignment for a non-simple cycle can be decomposed into (non-integer) assignments for simple cycles, while preserving the same bounds. Together, these results constitute a proof to Theorem 5.

We also give a result that bounds the maximal deviation in suffix phase mode-counting when a non-integer solution is rounded to an integer solution.

LEMMA 1. *Suppose that a correct strategy \mathbf{r}^* induces a behavior such that there are times T_1 and $T_2 > T_1$ for which $\Lambda(\mathbf{w}(T_1)) = \Lambda(\mathbf{w}(T_2))$. Then there is a prefix-suffix strategy that achieves the same performance.*

In particular, if \mathbf{r}^* is an integer strategy the lemma applies, since it will necessarily result in a behavior that satisfies the assumptions in this lemma due to the finiteness of the number of possible condensed states.

PROOF. We show that the graph flows induced by \mathbf{r}^* on the time interval $[T_1, T_2]$ can be achieved with cycle assignments. To this end, we define a static flow on a graph in a higher dimension, decompose it into cyclic flows, and project the cyclic flow onto the original graph.

Let $G = (V, E)$ be the system graph, and define a new graph $H = (V_H, E_H)$. The node set $V_H = \underbrace{V \times V \times \dots \times V}_{T_2 - T_1 \text{ times}}$ contains $T_2 - T_1$ copies of each vertex in V , and copies of $\nu \in V$ are labeled ν_t for $t \in [T_1, T_2]$. The set of edges is defined as

$$E_H = \{(\nu_t, \tilde{\nu}_{t+1}) : t \in \{T_1, \dots, T_2 - 1\}, (\nu, \tilde{\nu}) \in E\} \cup \{(\nu_{T_2}, \nu_{T_1}) : \nu \in V\}.$$

A static flow is induced on H by \mathbf{r}^* , obtained by letting the flow along $(\nu_t, \tilde{\nu}_{t+1})$ be the number that traverses the edge $(\nu, \tilde{\nu}) \in E$ at time t , and letting the flow along (ν_{T_2}, ν_{T_1}) be equal to the number of systems at ν at time T_1 . By construction, this flow is balanced at each node (i.e. inflows equal outflows).

By the flow decomposition theorem [1, Theorem 3.5], we can then find cycles in H that achieve the static edge flow. By projecting these cycles onto a single copy of V , we obtain cycles and assignments that when circulated mimic the performance of \mathbf{r}^* on the interval $[T_1, T_2]$.

We can therefore define a prefix-suffix strategy by taking as the prefix part inputs $\mathbf{r}^*(s)$ up to time $s = T_1 - 1$, together with a suffix part consisting of the cycles and assignments found above. \square

The preceding result implies the following convergence statement, which provides a converse result to Proposition 1.

COROLLARY 2. *Suppose Problem 2 has an (integer) solution. Then, there exists an integer $L < \infty$ and a finite horizon T such that the ILP version of (12) formulated with the set of all cycles of length smaller than or equal to L , has a feasible solution.*

Next, we show that it is sufficient to consider simple cycles, but this comes at the cost of possible non-integer assignments. In the integer case, it is in fact *not* sufficient to consider simple cycles only, as the example in Figure 2 shows.

LEMMA 2. *Suppose $C = C_1 \cup C_2$ is a cycle that visits a node ν_1 twice, so that it can be decomposed into two cycles $C_1 = (\nu_1, \nu_2), \dots, (\nu_{J_1}, \nu_1)$ and $C_2 = (\nu_1, \nu_{i+1}), \dots, (\nu_{J_2}, \nu_1)$.*

Let α be an assignment to C . If α satisfies

$$\underline{K}_m \leq \underline{\Psi}^m(C, \alpha), \quad \overline{\Psi}^m(C, \alpha) \leq \overline{K}^m, \quad (13)$$




Figure 2: Illustration showing how simple cycles may fail to realize mode-counting bounds that are possible with non-simple cycles. Considering $(\nu_1, \nu_2, \nu_3, \nu_4, \nu_5, \nu_1, \nu_6, \nu_7)$ as a single (non-simple) cycle, it can be assigned $[1, 1, 1, 1, 0, 0, 0, 0]$ to achieve a constant mode-count of 4 for both mode 1 (blue) and mode 2 (red). However, there is no way to create separate integer assignments for the simple cycles $(\nu_1, \nu_2, \nu_3, \nu_4, \nu_5)$ and (ν_1, ν_6, ν_7) , which achieve the same constant mode-counting over time.

then there exist assignments α_1, α_2 (not necessarily integral) for C_1, C_2 such that

$$\sum_{i \in [|C_1|]} \alpha_1(i) + \sum_{i \in [|C_2|]} \alpha_2(i) = \sum_{i \in [|C|]} \alpha(i), \quad (14)$$

$$\underline{K}_m \leq \Psi^m(C_1, \alpha_1) + \Psi^m(C_2, \alpha_2), \quad (15)$$

$$\overline{\Psi}^m(C_1, \alpha_1) + \overline{\Psi}^m(C_2, \alpha_2) \leq \overline{K}_m. \quad (16)$$

PROOF. Note that (13) implies that for $k \in [|C|]$,

$$\underline{K}_m \leq \sum_{i: \Xi_C(\nu_i)=m} \alpha((k+i) \bmod |C|) \leq \overline{K}_m.$$

Summing these inequalities over $k = 1, \dots, |C|$ yields, since each $\alpha(i)$ will appear exactly $|C|_m := |\{\nu_i \in C : \Xi_C(\nu_i) = m\}|$ times,

$$|C| \underline{K}_m \leq |C|_m \sum_{i \in [|C|]} \alpha(i) \leq |C| \overline{K}_m. \quad (17)$$

Now consider the constant ‘‘averaging’’ assignments

$$\alpha_1(i) = \alpha_2(j) = \frac{\sum_{i \in [|C|]} \alpha(i)}{|C|}.$$

We obviously have

$$\sum_{i \in [|C_1|]} \alpha_1(i) + \sum_{i \in [|C_2|]} \alpha_2(i) = \underbrace{\frac{|C_1| + |C_2|}{|C|}}_{=1} \sum_{i \in [|C|]} \alpha(i),$$

so (14) is satisfied. Furthermore, for any k_1, k_2 ,

$$\begin{aligned} & \sum_{i: \Xi_{C_1}(\nu_i)=m} \alpha_1((i+k_1) \bmod |C_1|) \\ & + \sum_{j: \Xi_{C_2}(\nu_j)=m} \alpha_2((j+k_2) \bmod |C_2|) \\ & = \underbrace{(|C_1|_m + |C_2|_m)}_{=|C|_m} \frac{\sum_{i \in [|C|]} \alpha(i)}{|C|}. \end{aligned}$$

From (17) we can conclude that this expression is in the interval $[\underline{K}_m, \overline{K}_m]$. It follows that

$$\begin{aligned} \underline{K}_m & \leq \Psi^m(C_1, \alpha_1) + \Psi^m(C_2, \alpha_2) \\ & = \overline{\Psi}^m(C_1, \alpha_1) + \overline{\Psi}^m(C_2, \alpha_2) \leq \overline{K}_m, \end{aligned}$$

so (15)-(16) hold. \square

The next result shows how the mode-counting bounds are affected when the suffix part of a non-integer strategy is rounded to integer cycle assignments. A proof can be found in the Appendix.

PROPOSITION 2. For any integer N_C and a cycle C , there exists an integer assignment α_{int} , with $\sum_i \alpha_{int}(i) = N_C$, such that for the average assignment $\alpha_{avg}(i) := N_C/|C|$,

$$\begin{aligned} \underline{\Psi}^m(C, \alpha_{int}) & \geq \underline{\Psi}^m(C, \alpha_{avg}) \\ & - \min \left(\frac{|C|_m}{|C|} z, \left(1 - \frac{|C|_m}{|C|}\right) (|C| - z) \right), \\ \overline{\Psi}^m(C, \alpha_{int}) & \leq \overline{\Psi}^m(C, \alpha_{avg}) \\ & + \min \left(\frac{|C|_m}{|C|} (|C| - z), \left(1 - \frac{|C|_m}{|C|}\right) z \right), \end{aligned}$$

where $z = N_C \bmod |C|$ and $|C|_m := |\{\nu_i \in C : \Xi_C(\nu_i) = m\}|$.

Remark 2. Looser, but less cumbersome, bounds can be obtained by noting that

$$\begin{aligned} & \min \left(\frac{|C|_m}{|C|} z, \left(1 - \frac{|C|_m}{|C|}\right) (|C| - z) \right) \\ & = |C| \min \left(\frac{|C|_m}{|C|} \frac{z}{|C|}, \left(1 - \frac{|C|_m}{|C|}\right) \left(1 - \frac{z}{|C|}\right) \right) \\ & \leq \frac{|C|}{4}, \end{aligned}$$

since $\max_{a,b \in [0,1]} \min(ab, (1-a)(1-b)) = 1/4$. The same bound holds for the second inequality due to symmetry.

We conjecture that the bounds in Proposition 2 can be tightened substantially by a clever integer assignment algorithm, as opposed to the worst-case analysis used to obtain the bounds above. This question, together with the issue of rounding also the prefix part of a non-integer strategy, will be subject to future research.

5. EXAMPLES

We provide two examples, one numerical example where the individual systems have two states, and one that applies the method to a scheduling problem of thermostatically controlled loads. In both examples, we use the Gurobi solver [9] to obtain LP and ILP solutions of (12)³.

5.1 Numerical example

As a first example, consider the aggregate dynamics where the state of each individual system is governed by the switched two-dimensional nonlinear ODE $\dot{x}_i(t) = f_{\sigma_i(t)}(x_i(t))$, where for $j = 1, 2$,

$$f_j : \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} \mapsto \begin{bmatrix} -(x_1 - \bar{x}_1^j) + x_2 \\ -(x_1 - \bar{x}_1^j) - x_2 - x_2^3 \end{bmatrix}, \quad \begin{matrix} \bar{x}_1^1 = 1, \\ \bar{x}_1^2 = -1. \end{matrix} \quad (18)$$

We want to solve Problem 1 with $N = 10000$ systems with a desired mode-1-count of exactly 6750 systems, without having systems enter an unsafe set $\mathcal{U}_1 = \mathcal{U}_2 = [-0.3, 0.3] \times [-0.2, 0.2] \subset \mathbb{R}^2$.

Using abstraction parameters $\tau = 0.5, \eta = 0.05$ on the domain $\{(x, y) : -2 \leq x \leq 2, -1 \leq y \leq 1\}$, we obtain

³Our software is available at github.com/pettni/mode-count.




Figure 3: Illustration of the suffix part of the solution of (12) for (18). The cycles found by the synthesis procedure are highlighted in different colors. The unsafe set is shown in red, with an added margin to compensate for bisimulation approximation errors.

an abstraction with 3483 states that is 0.15-approximately bisimilar to its time-sampled counterpart. Consequently, we need to add a margin of 0.15 to the unsafe set to ensure safety⁴. With a horizon of 10τ (i.e., $T = 10$) and using 100 randomly generated cycles, an ILP solution can be found in about 16 seconds on a 3.4 GHz iMac. The suffix part of the solution consists of three cycle assignments, and the mode-1-count is guaranteed to be exactly 6750 over time. Figure 3 shows a visualization of the cycles found by the solver.

The cycles have the following mode-profiles:

$$\begin{aligned} C_1 &: 1, 1, 1, 1, 1, 1, 2, 2, 2, 1, \\ C_2 &: 2, 1, 1, 1, 1, 2, 2, 2, 2, 2, \\ C_3 &: 2, 2, 2, 1, 1, 1, 1, 2, 2, 2, \end{aligned}$$

and the corresponding cycle assignments found by the solver are

$$\begin{aligned} \alpha_1 &= [917, 917, 917, 917, 917, 917, 917, 917, 917, 917], \\ \alpha_2 &= [1, 1, 1, 1, 1, 1, 1, 1, 1, 1], \\ \alpha_3 &= [0, 164, 0, 164, 0, 164, 0, 164, 0, 164]. \end{aligned}$$

It follows that the mode-1-count is constant for all three cycles, and is equal to 6419, 3, and 328, respectively.

5.2 Application example

Next, we apply the proposed algorithm to a TCL coordination problem, where the objective is to control the aggregate load of a family of domestic air condition units. The dynamics of a single unit are

$$\dot{\theta}_i = -a(\theta_i - \theta_a) - bP_i^m,$$

where $P_i^m = 0$ when unit i is **off** and $P_i^m = 5.6$ when unit i is **on**. The parameters⁵ are taken from [15].

We assume that all units are set to the desired temperature $\theta_0 = 22.5^\circ\text{C}$, and that deviations up to 1°C are tolerated (larger deviations constitute unsafe states). The objective is to control the aggregate load, i.e. the number of TCL's that are **on**, while simultaneously guaranteeing that the state θ_i of every TCL will remain in the interval $[21.5, 23.5]$.

⁴Safety is guaranteed only at sample instants.

⁵ $\theta_a = 32^\circ\text{C}$, $a = 0.25$, $b = 1.25$




Figure 4: Densities (kernel-smoothened for visualization) of TCL populations over time for the desired mode-on-count 3600 (upper) and 3200 (lower). Red color indicates that a larger number of TCL's are in that interval. In the initial phases, control is being applied to place the TCL's on cycles. After this is achieved at $t = 1$, the TCL's circulate in their respective cycles which results in a more regular behavior. As can be seen, when the desired mode-on-count is higher, the temperatures group in the lower part of the allowed temperature range. The reason is that when the temperature difference is large, the room is warmed up at a higher rate, thus the air condition units can be turned on during a larger percentage of the time, on average.

For this example, we created an abstraction with parameters $\tau = 0.05$, $\eta = 0.0195$, which is approximately bisimilar to the original system with accuracy $\epsilon = 0.1$. We then solve (12) for a population of 10000 TCL's with randomized initial states, using a set of 100 randomized cycles and horizon $T = 20$. Since the number of variables in the linear program is quite large, we utilize a two-step approach to obtain an integer solution. First, (12) is solved as a non-integer linear program, which results in non-integer assignments. Secondly, these assignments are rounded to integer assignments and controls needed to reach these integer assignments are found by solving (12) as an ILP with fixed assignments. This reduces the number of variables and constraints which makes the ILP more tractable, however, the rounding of cycle assignments may lead to bound violations, as alluded to in Proposition 2.

We synthesized control strategies for two sets of desired mode-on-count bounds, one around 3600 (high) and another one around 3200 (low), using the same randomized initial

Table 1: Guaranteed mode-on-counting bounds for the TCL application.

Desired mode-count	low	high
Prefix phase bounds	[2500, 2564]	[3696, 4300]
Suffix phase bounds	[3180, 3217]	[3595, 3604]




Figure 5: mode-on-count over time for desired mode-on-count 3600 (red) and 3200 (blue). As can be seen, the mode-on-counts are close to the desired values in the suffix phase.

conditions. To allow more flexibility, we relaxed the mode-counting constraints during the prefix phase. With this setup, we obtained control strategies satisfying the bounds in Table 1. In Figure 4, the densities of TCL’s in the temperature spectrum are shown over time for both desired mode-on-counts. These densities are of continuous-state trajectories, which are guaranteed to stay within 0.1-distance of their discrete counterparts. Figure 5 shows mode-on-counts over time for both solutions.

6. DISCUSSION AND EXTENSIONS

In this section we discuss possible extensions of the proposed framework and potential application areas where mode-counting like constraints can be relevant.

Our approach can be readily applied to a “heterogeneous” collection of systems if there is a limited number of types of systems and the systems are still homogeneous within each type. This is for instance the case with TCLs where one can categorize air conditioners or refrigerators as different types [6]. This will require computing an ϵ -approximate state-time abstraction and an aggregate linear system of the form (9) per type. The resulting feasibility problem will contain these different abstractions and aggregate dynamics as constraints, together with additional mode-counting constraints that couple them. Another straightforward extension is to include state-dependent mode-counting constraints or mode-counting like constraints, where mode is not necessarily the switching signal but a region of the state-space, and mode-counting like constraints restrict the number of systems in these regions of the state-space over time. Such mode-counting like constraints are relevant in application areas like air traffic control or swarm robotics. Another interesting extension is to consider richer constraints on the mode-counts beyond a bound, for instance to synthesize controllers that enforce a linear temporal logic (LTL) property defined on the mode-counts. Such constraints can be incorporated in the proposed framework using mixed integer programming encodings of LTL [11, 24].

There are some limitations of the proposed framework that require less trivial extensions. For instance, there could be small deviations from the model (1) between individual systems that will lead to an uncertain system model. Similarly, all the states might not be available for measurement, or state measurements might be noisy or delayed [15]. Also, the individual system dynamics might not be incrementally stable. Although there are techniques for constructing ab-

stractions that take into account such imperfections [25, 16, 14], the resulting abstract transition system is usually non-deterministic. Therefore, such an extension will require feedback control strategies for which one can potentially use reactive synthesis or a robust ILP formulation [19]. However, the resulting reactive synthesis problem is challenging to solve due to the size of the aggregate (uncertain in this case) linear system (9). A better characterization of the trade-offs between solving the problem (12) as a linear program versus an integer linear program is also important to explore, in order to enable some of these extensions in a computationally tractable way.

Considering the structural properties of the problem, what enables massive aggregation of individual dynamics is the permutation invariance of dynamics due to homogeneity and the permutation invariance of mode-counting constraints. Therefore, it will be interesting to consider other classes of systems with symmetries (see, e.g., [7]) where one can utilize techniques similar to those proposed in this paper to achieve scalability.

Finally, our work is also related to scheduling based methods used in intersection collision avoidance [4] and air traffic control [2]. A crucial difference is that these works consider finite-horizon objectives, whereas we require control strategies that are valid on an infinite horizon. Therefore, we believe the proposed work can be used in these application domains when infinite horizon guarantees are required.

7. CONCLUSIONS

We have considered a new class of constraints, called mode-counting constraints, and proposed an approach to synthesize controllers that can enforce this type of constraints for a homogeneous collection of switched systems. The proposed approach utilizes the particular structure of the problem, the specific form of the control objective and the homogeneity of the systems, in a way that allows handling very large collections of systems. The efficacy of the approach has been demonstrated both with a numerical example and with an application to a TCL coordination problem that involves ten thousand systems.

As discussed in Section 6, there is a wide range of possibilities for future directions; both in terms of potential application areas where mode-counting like constraints are relevant, and in terms of extensions to the type of systems that can be handled or the type of mode-counting constraints that can be enforced. We will explore these directions as part of our future work.

Acknowledgments

The authors would like to thank Johanna Mathieu for enlightening discussions on the TCL coordination problem that motivated this research. The work of PN is supported by NSF grant CNS-1239037 and the work of NO is supported in part by NSF grant CNS-1446298.

8. REFERENCES

- [1] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin. *Network flows*. Prentice Hall, 1993.
- [2] A. Bayen, C. Tomlin, Y. Ye, and J. Zhang. Milp formulation and polynomial time algorithm for an aircraft scheduling problem. In *Proc. of IEEE CDC*, pages 5003–5010, 2003.

- [3] A. Borri, D. V. Dimarogonas, K. H. Johansson, M. D. Di Benedetto, and G. Pola. Decentralized symbolic control of interconnected systems with application to vehicle platooning. In *4th IFAC Workshop on Distributed Estimation and Control in Networked Systems*, pages 285–292, 2013.
- [4] A. Colombo and D. Del Vecchio. Least restrictive supervisors for intersection collision avoidance: A scheduling approach. *IEEE Trans. on Automatic Control*, 60(6):1515–1527, June 2015.
- [5] S. Coogan and M. Arcak. Efficient finite abstraction of mixed monotone systems. In *Proc. of HSCC*, pages 58–67, 2015.
- [6] S. Esmail Zadeh Soudjani and A. Abate. Aggregation and control of populations of thermostatically controlled loads by formal abstractions. *IEEE Trans. on Control Systems Technology*, 23(3):975–990, 2015.
- [7] E. Frazzoli, M. Dahleh, and E. Feron. Maneuver-based motion planning for nonlinear systems with symmetries. *IEEE Trans. on Robotics*, 21(6):1077–1091, 2005.
- [8] A. Girard and G. J. Pappas. Approximation Metrics for Discrete and Continuous Systems. *IEEE Trans. on Automatic Control*, 52(5):782–798, 2007.
- [9] Gurobi Optimization, Inc. Gurobi optimizer reference manual, 2015.
- [10] T. T. Johnson and S. Mitra. A small model theorem for rectangular hybrid automata networks. In *FMOODS/FORTE*, pages 18–34. Springer, 2012.
- [11] S. Karaman, R. G. Sanfelice, and E. Frazzoli. Optimal control of mixed logical dynamical systems with linear temporal logic specifications. In *Proc. of IEEE CDC*, pages 2117–2122, 2008.
- [12] S. Koch, J. L. Mathieu, and D. S. Callaway. Modeling and control of aggregated heterogeneous thermostatically controlled loads for ancillary services. In *Proc. of PSCC*, pages 1–7, 2011.
- [13] M. Lewin. On exponents of primitive matrices. *Numerische Mathematik*, 18(2):154–161, 1971.
- [14] J. Liu and N. Ozay. Abstraction, discretization, and robustness in temporal logic control of dynamical systems. In *Proc. of HSCC*, 2014.
- [15] J. L. Mathieu, S. Koch, and D. S. Callaway. State estimation and control of electric loads to manage real-time energy imbalance. *IEEE Trans. on Power Systems*, 28(1):430–440, 2013.
- [16] O. Mickelin, N. Ozay, and R. M. Murray. Synthesis of correct-by-construction control protocols for hybrid systems using partial state information. In *Proc. of ACC*, 2014.
- [17] G. Pola and P. Tabuada. Symbolic Models for Nonlinear Control Systems: Alternating Approximate Bisimulations. *SIAM J. Control Optim.*, 48(2):719–733, 2009.
- [18] P. Pourbeik, P. S. Kundur, and C. W. Taylor. The anatomy of a power grid blackout. *IEEE Power and Energy Magazine*, 4(5):22–29, 2006.
- [19] V. Raman, A. Donzé, D. Sadigh, R. M. Murray, and S. A. Seshia. Reactive synthesis from signal temporal logic specifications. In *Proc. of HSCC*, pages 239–248, 2015.
- [20] M. Rungger, M. Mazo Jr, and P. Tabuada. Specification-guided controller synthesis for linear systems and safe linear-time temporal logic. In *Proc. of HSCC*, pages 333–342, 2013.
- [21] B. M. Sanandaji, H. Hao, and K. Poolla. Fast regulation service provision via aggregation of thermostatically controlled loads. In *Proc. of HICSS*, pages 2388–2397, 2014.
- [22] E. Seneta. *Non-negative matrices and Markov chains*. Springer, 2006.
- [23] P. Tabuada. *Verification and control of hybrid systems: a symbolic approach*. Springer, 2009.
- [24] E. M. Wolff, U. Topcu, and R. M. Murray. Optimization-based trajectory generation with linear temporal logic specifications. In *Proc. of ICRA*, pages 5319–5325, 2014.
- [25] M. Zamani, G. Pola, M. Mazo Jr, and P. Tabuada. Symbolic models for nonlinear control systems without stability assumptions. *IEEE Trans. on Automatic Control*, 57:1804–1809, 2012.
- [26] W. Zhang, K. Kalsi, J. Fuller, M. Elizondo, and D. Chassin. Aggregate model for heterogeneous thermostatically controlled loads with demand response. In *Proc. of PES General Meeting*, pages 1–8, 2012.

APPENDIX

PROOF OF PROPOSITION 2. Let $N_C = |C| \times r + z$, where $0 \leq z < |C|$. Consider the assignment α_{int} such that

$$\alpha_{int} = \begin{cases} r + 1, & i = 1, \dots, z, \\ r, & i = z + 1, \dots, |C|. \end{cases}$$

Then $\sum_i \alpha_{int}(i) = r|C| + z = N_C$. From a worst-case analysis, it follows that the max- and min-counts satisfy:

$$\begin{aligned} \underline{\Psi}(C, \alpha_{int}) &\geq r \min(|C|_m, |C| - z) \\ &\quad + (r + 1) \max(0, |C|_m - (|C| - z)) \\ &= r|C|_m + \max(0, |C|_m - (|C| - z)), \\ \bar{\Psi}(C, \alpha_{int}) &\leq (r + 1) \min(|C|_m, z) + r \max(0, |C|_m - z) \\ &= r|C|_m + \min(|C|_m, z). \end{aligned}$$

It follows that,

$$\begin{aligned} &\underline{\Psi}(C, \alpha_{avg}) - \underline{\Psi}(C, \alpha_{int}) \\ &\leq |C|_m \frac{N_C}{|C|} - (r|C|_m + \max(0, |C|_m - (|C| - z))) \\ &= |C|_m \left(\frac{|C|r + z}{|C|} \right) - (r|C|_m + \max(0, |C|_m - (|C| - z))) \\ &= |C|_m \frac{z}{|C|} - \max(0, |C|_m - (|C| - z)) \\ &= \min \left(\frac{|C|_m}{|C|} z, \left(1 - \frac{|C|_m}{|C|} \right) (|C| - z) \right). \end{aligned}$$

Similarly,

$$\begin{aligned} &\bar{\Psi}(C, \alpha_{int}) - \bar{\Psi}(C, \alpha_{avg}) \leq |C|_m r + \min(|C|_m, z) - |C|_m \frac{N_C}{|C|} \\ &= |C|_m r + \min(|C|_m, z) - |C|_m r - |C|_m \frac{z}{|C|} \\ &= \min \left(\frac{|C|_m}{|C|} (|C| - z), \left(1 - \frac{|C|_m}{|C|} \right) z \right). \quad \square \end{aligned}$$