

Optimal Scheduling for Constant-Rate Multi-Mode Systems

Rajeev Alur
University of Pennsylvania,
Philadelphia, USA
alur@cis.upenn.edu

Ashutosh Trivedi
University of Pennsylvania,
Philadelphia, USA
ashut@cis.upenn.edu

Dominik Wojtczak
University of Liverpool,
Liverpool, UK
d.wojtczak@liv.ac.uk

ABSTRACT

Constant-rate multi-mode systems are hybrid systems that can switch freely among a finite set of modes, and whose dynamics is specified by a finite number of real-valued variables with mode-dependent constant rates. The schedulability problem for such systems is to design a mode-switching policy that maintains the state within a specified safety set. The main result of the paper is that schedulability can be decided in polynomial time. We also generalize our result to optimal schedulability problems with average cost and reachability cost objectives. Polynomial-time scheduling algorithms make this class an appealing formal model for design of energy-optimal policies. The key to tractability is that the only constraints on when a scheduler can switch the mode are specified by global objectives. Adding local constraints by associating either invariants with modes, or guards with mode switches, lead to undecidability, and requiring the scheduler to make decisions only at multiples of a given sampling rate, leads to a PSPACE-complete schedulability problem.

Categories and Subject Descriptors

D.4.7 [Organization and Design]: Real-time systems and embedded systems; B.5.2 [Design Aids]: Optimization, Verification

General Terms

Theory, Verification

Keywords

Multi-Mode Systems, Cyber-Physical Systems, Peak Minimization, Green Scheduling, Hybrid Automata

1. INTRODUCTION

Our study of optimal scheduling on constant-rate multi-mode systems is motivated largely by a series of work by

Nghiem et al. [17, 18] on energy peak demand reduction within a large organization by synchronizing switching decisions of various “heating, ventilation, and air conditioning” (HVAC) systems. The correlation between extreme weather and energy demand peaks is well documented [6], and hence reducing the energy peak demand due to HVAC systems can potentially significantly reduce the total energy peak demand. In [17] Nghiem et al. considered a model of an organization where at any given time the HVAC system of a zone can be in either ON or OFF mode, and in each mode the temperature of the corresponding zone changes with a mode-dependent constant rate. In order to minimize peak energy-usage they studied the following schedulability problem: find a switching schedule of HVAC systems across different zones so as to maintain the temperature in each zone within a given temperature interval, with the restriction that simultaneously at most a fix number of HVAC systems are switched ON. They showed that the schedulability problem can be reduced to testing an inequality involving the rates of temperature change. Our motivation is to explore that to what extent this result can be generalized, and to identify where this result fits into existing literature on schedulability such as real-time scheduling theory [9] and hybrid automata based schedulability analysis [1, 2].

Real-time scheduling is a mature research area [9] with an excellent collection of well-studied algorithms for periodic scheduling, for instance the rate monotonic and the earliest deadline first algorithms. However, as noted by Nghiem et al. [18], generally these algorithms are restricted to tasks whose worst case execution times are fixed and known in advance, and hence they are not directly applicable to energy peak reduction problem as posed in [18, 17].

Another prominent approach [1, 2] to real-time schedulability analysis is via reduction to optimization problems on timed and hybrid automata. Timed automata [3] can model multi-mode systems with a finite set of continuous variables, called clocks, that grow with uniform rate. Clocks can be used to constrain mode-switches and to specify mode-dependent invariants. The decidability of a number of optimization problems [4, 7] on timed automata, and availability of efficient tool support, e.g. Kronos and UPPAAL [15, 20], make them an attractive choice for real-time scheduling. They are, however, not applicable in energy peak reduction problem as the temperature variables in our system grow with non-uniform rates. Hybrid automata generalize timed automata by allowing mode-dependent variable rates, however having two variables with different rates leads to undecidability [11] even for reachability problems. As we see later

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. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Copyright 20XX ACM X-XXXXX-XX-X/XX/XX ...\$10.00.

the key property of our systems that contributes to decidability (and even tractability) of schedulability problems is the absence of structure in the system, i.e. intuitively speaking as long as global safety set is not violated schedulers are allowed to switch among modes without any restriction.

Results. We define our model, a constant-rate multi-mode system (MMS), as a hybrid system with a finite set of modes where dynamics of each mode is specified by a finite set of continuous variables with mode-dependent constant-rates. Given a bounded convex set of safe states (variable valuations) and a starting state, the *safe schedulability problem* for an MMS is to find a non-Zeno schedule that visits only safe states. Another closely related problem is *safe reachability problem* that asks if there exists a schedule to steer the system from a given starting state to a given set of target states while only visiting states from a given safety set.

Our first result concerning safe schedulability problem is that for all the starting states in the interior of the safety set, a safe schedule exists iff there is an assignment of dwell times to modes that allows the system to return to the starting state. Due to constant-rate dynamics of the system this condition can easily be posed as a linear programming (LP) feasibility problem. The algorithm for finding safe schedule is more involved if the starting state is on the boundary of the safety set. A key contribution of the paper is a polynomial-time algorithm for computing a safe schedule for an arbitrary starting state when the safety set is given as a bounded convex polytope.

For safe reachability problem we show that if a state is reachable from a given starting state for some assignment of dwell times to modes, then it is safely reachable for any arbitrary bounded and convex safety set as long as both states lie in the interior of the safety set. We show via an example that this observation is not valid when one of the states lies on the boundary of the safety set. We present a polynomial-time algorithm to find safe schedules for reachability problems when both the starting and the target states lie in the interior of the safety set.

In Section 3 we extend these results to optimization problem with average cost and reachability cost objectives, and present polynomial-time algorithms to solve these problems. Furthermore, we prove that reachability-cost optimal and average-cost optimal strategies always exist, and have a particular simple periodic structure as long as both the target and the starting states are in the interior of the safety set. In Section 4 we show that requiring the scheduler to make decisions only at multiples of a given clock-rate makes the safe schedulability and safe reachability problems complete for PSPACE. We also show that the largest sampling rate for which safe schedulability problem yields a positive answer can be approximated in polynomial space.

Related Work. The work most closely related to our is of Nghiem et al. [17]. The energy peak reduction problem was first posed in [17] and safe schedulability checking for this problem was reduced to checking a simple formula on temperature rates. Authors also presented a lazy scheduling algorithm where scheduler is required to take decisions only at multiples of a given sampling rate. Although the practical motivation of our approach is the same as [17] we have different goals. The central focus of our research is to characterize the complexity of various schedulability problems in

this context. Safe schedulability problem for MMSs generalizes the energy peak reduction problem studied by Nghiem et al. in [17] since MMSs can model HVAC systems with more modes than simply ON and OFF. Moreover, MMSs allow safety set to be an arbitrary bounded convex set as opposed to hyperrectangular sets in [17]. Unlike [17] our algorithm can analyze the safe schedulability problem for starting states on the boundary of the safety set, assuming it is a polytope. Moreover, we also extend our results to safe reachability problem, and optimization problems for average cost and reachability cost objectives. We also establish PSPACE-completeness of finding the optimal sampling rate for safe schedulability of our more general systems.

In [18] Nghiem et al. generalized their work to multi-mode systems with linear dynamics. Using similar restrictions as [18] the ideas presented in this paper can be generalized to handle linear dynamics with some effort. Heymann et al. [12] study checking whether a given hybrid system, under several restrictions, is strongly Zeno and characterize LP feasibility test for this problem. Like Nghiem et al. [17], Heymann et al. only consider hyperrectangular safety states and starting states in the interior of the safety region.

The *practical stabilization* problem studied by Xu and Antsaklis [21] roughly corresponds in our model to an unconstrained reachability followed by a specifically constrained safe schedulability problem where the system cannot leave a ball of radius ϵ when starting from anywhere inside a ball with the same origin and radius δ . The existence and synthesis of a scheduler satisfying such constraints can be solved in polynomial time using our algorithms, while the running time of the method suggested in [21] is exponential in the number of modes. Moreover, the scheduler proposed in [21] is more complicated than ours and requires solving multiple linear programs as opposed to essentially one in our algorithm. Finally, the problem of computing a lower bound on the optimal sampling rate was studied there for systems with two variables and the problem was left open for systems with more variables. We show that even approximation of the optimal sampling rate within a constant error is PSPACE-hard, which makes unlikely the existence of a tractable general procedure for this problem.

We also mention the work of Jha et al. [13, 14] where they synthesize guards for multi-mode systems so as to satisfy certain optimization criteria. Their model is more general than ours as it allows guards on mode-switches, and moreover variables are allowed to have more general dynamics, however authors did not present any complexity or decidability results. Henzinger and Kopke [10] studied the safe reachability problem for hybrid automata [10] where the scheduler is allowed to make decisions at multiples of a given sampling rate, and showed the problem to be PSPACE-complete. Upper bound for the similar problem for MMSs directly follows from their work. Bouyer et al. [8] study the safe schedulability problem on *weighted timed automata* (timed automata extended with a cost variable having mode-dependent constant-rate) with global cost constraints. The safe schedulability problem is to find a schedule that keeps the value of the cost variable within a given interval. MMS and weighted timed automata are incomparable models: the former disallows guards on mode-switches, while the latter disallows more than one variable with mode-dependent rate.

An extended version of this paper with complete proofs is available as a technical report [5].

2. SAFE SCHEDULABILITY

Before we formally introduce constant-rate multi-mode systems, we need to introduce the notation used throughout the rest of the paper. We write \mathbb{N} for the set of natural numbers, \mathbb{R} for the set of real numbers, and \mathbb{Z} for the set of integers. Also, we write $\mathbb{R}_{\geq 0}$ for the set of non-negative reals and $\mathbb{N}_{>0}$ for the set of positive integers.

States of our system will be points in \mathbb{R}^n that is equipped with the standard *Euclidean norm* $\|\cdot\|$. By \bar{x}, \bar{y} we denote points in this state space, by \vec{f}, \vec{v} vectors, while $\bar{x}(i)$ and $\vec{f}(i)$ will denote the i -th coordinate of point \bar{x} and vector \vec{f} , respectively. We denote the distance between \bar{x} and \bar{y} by $\|\bar{x}, \bar{y}\| \stackrel{\text{def}}{=} \|\bar{x} - \bar{y}\|$. For two vectors $\vec{v}_1, \vec{v}_2 \in \mathbb{R}^n$, we write $\vec{v}_1 \circ \vec{v}_2$ to denote their dot product defined as $\sum_{i=1}^n \vec{v}_1(i) \vec{v}_2(i)$. We write $\vec{0}$ for any vector with all its coordinates equal to 0; its exact dimension will depend on the context.

We say that a point \bar{x} is a convex combination of a set of points $X = \{\bar{x}_1, \bar{x}_2, \dots, \bar{x}_k\}$ if there exist $\lambda_1, \lambda_2, \dots, \lambda_k \in [0, 1]$ such that $\sum_{i=1}^k \lambda_i = 1$ and $\bar{x} = \sum_{i=1}^k \lambda_i \bar{x}_i$. We say that the set $S \subseteq \mathbb{R}^n$ is *convex* iff for all $\bar{x}, \bar{y} \in S$ and all $\lambda \in [0, 1]$ we have $\lambda \bar{x} + (1 - \lambda) \bar{y} \in S$ and moreover, S is a *convex polytope* if there exists $k \in \mathbb{N}$, a matrix A of size $k \times n$ and a vector $\vec{b} \in \mathbb{R}^k$ such that $\bar{x} \in S$ iff $A\bar{x} \leq \vec{b}$. Let $B_d(\bar{x}) = \{\bar{y} \in \mathbb{R}^n : \|\bar{x}, \bar{y}\| \leq d\}$ denote a *closed ball* of radius $d \in \mathbb{R}_{>0}$ centered at \bar{x} . We say that a set $S \subseteq \mathbb{R}^n$ is *bounded* if there exists $d \in \mathbb{R}_{>0}$ such that for all $\bar{x}, \bar{y} \in S$ we have $\|\bar{x}, \bar{y}\| \leq d$. The *interior* of a set S , $\text{int}(S)$, is the set of all points $\bar{x} \in S$ for which there exists $d > 0$ s.t. $B_d(\bar{x}) \subseteq S$.

2.1 Constant-Rate Multi-Mode Systems

A constant-rate multi-mode system consists of a finite number of modes and a finite number of real-valued variables whose dynamics is specified by mode-dependent constant rates. Formally,

DEFINITION 1. A *constant-rate multi-mode system (MMS)* is a tuple $\mathcal{H} = (M, n, R)$ where M is a finite nonempty set of modes, n is the number of continuous variables in the system, and $R : M \rightarrow \mathbb{R}^n$ gives for each mode the rate vector whose i -th entry specifies the change in value of the i -th variable per time unit.

For computation purposes, we assume that all real numbers are rational and represented in the standard way by writing down the numerator and denominator in binary.

A *schedule* of an MMS specifies a timed sequence of mode switches. Formally, a *schedule* is defined as a finite or infinite sequences of *timed actions*, where a timed action $(m, t) \in M \times \mathbb{R}_{\geq 0}$ is a tuple consisting of a mode and a time delay. We say that an infinite schedule $\langle (m_1, t_1), (m_2, t_2), \dots \rangle$ is *periodic* if there exists $k \geq 1$ such that for all $i \geq 1$ we have $(m_i, t_i) = (m_{(i \bmod k)+1}, t_{(i \bmod k)+1})$, and it is *ultimately periodic* if it has a suffix that is periodic. We say that an infinite schedule $\langle (m_1, t_1), (m_2, t_2), \dots \rangle$ is *Zeno* if $\sum_{i=1}^{\infty} t_i < \infty$. Zeno schedules require infinitely many mode-switches within a finite time, and hence, are physically unrealizable.

For a (finite or infinite) schedule $\sigma = \langle (m_1, t_1), (m_2, t_2), \dots \rangle$, we write $T_k(\sigma) \stackrel{\text{def}}{=} \sum_{i=1}^k t_i$ for the total time elapsed up to step k of the schedule σ , and we write $T_k^m(\sigma) \stackrel{\text{def}}{=} \sum_{i \leq k: m_i = m} t_i$ for the total time spent in mode m up to step k . For any non-Zeno schedule σ we have that $\lim_{k \rightarrow \infty} T_k(\sigma) = \infty$.

A finite *run* of an MMS \mathcal{H} is a finite sequence of states and timed actions $r = \langle \bar{x}_0, (m_1, t_1), \bar{x}_1, \dots, (m_k, t_k), \bar{x}_k \rangle$ such that

for all $1 \leq i \leq k$ we have that $\bar{x}_i = \bar{x}_{i-1} + t_i \cdot R(m_i)$. For such a run r we say that \bar{x}_0 is the *starting state*, while \bar{x}_k is its *terminal state*. Given a state \bar{x} and a finite schedule $\sigma = \langle (m_1, t_1), (m_2, t_2), \dots, (m_k, t_k) \rangle$, we write $\text{Run}(\bar{x}, \sigma)$ for the (unique) finite run $\langle \bar{x}_0, (m_1, t_1), \bar{x}_1, (m_2, t_2), \dots, (m_k, t_k), \bar{x}_k \rangle$ such that $\bar{x}_0 = \bar{x}$. In this case, we also say that schedule σ leads the system \mathcal{H} from state \bar{x}_0 to state \bar{x}_k . The concept of an *infinite run* and an infinite run $\text{Run}(\bar{x}, \sigma)$ corresponding to an infinite schedule σ is defined in an analogous manner.

Given a set $S \subseteq \mathbb{R}^n$ of safe states, we say that an infinite run $\langle \bar{x}_0, (m_1, t_1), \bar{x}_1, (m_2, t_2), \dots \rangle$ is S -safe if for all $i \geq 0$ we have that $\bar{x}_i \in S$ and $\bar{x}_i + \tau_{i+1} \cdot R(m_{i+1}) \in S$ for all $\tau_{i+1} \in [0, t_{i+1}]$. Notice that if S is a convex set then $\bar{x}_i \in S$ for all $i \geq 0$, implies that $\bar{x}_i + \tau_{i+1} \cdot R(m_{i+1}) \in S$ for all $i \geq 0$ and all $\tau_{i+1} \in [0, t_{i+1}]$. Given a set $S \subseteq \mathbb{R}^n$ of safe states and a starting state $\bar{x} \in \mathbb{R}^n$, we say that an infinite schedule $\sigma = \langle (m_1, t_1), (m_2, t_2), \dots \rangle$ is S -safe at \bar{x} if the corresponding run $\text{Run}(\bar{x}, \sigma)$ is S -safe. The concept of S -safety for finite runs and schedules is defined in a similar manner. Sometimes we simply call a schedule or a run safe when the safety set and the starting state is clear from the context.

We say that a state \bar{x}' is " S -safe reachable" from a state \bar{x} if there exists a finite schedule σ that is S -safe at \bar{x} and leads the system from state \bar{x} to \bar{x}' . The following observations will be useful in some of the proofs later.

PROPOSITION 1. For every MMS \mathcal{H} and a convex safety set S we have that any convex combination $\bar{x}^* = \sum_{i=1}^k \lambda_i \bar{x}_i$ of S -safe reachable states $\bar{x}_1, \bar{x}_2, \dots, \bar{x}_k$ from a given state \bar{x} is also S -safe reachable from \bar{x} . Moreover, if mode m is safe for t amount of time at $\bar{x}_i \in S$, then it is safe for $\lambda_i t$ amount of time at \bar{x}^* .

Two fundamental problems for MMS are the following *safe schedulability* and *safe reachability* problems.

DEFINITION 2 (SAFE SCHEDULABILITY). Given an MMS \mathcal{H} , a bounded convex set $S \subseteq \mathbb{R}^n$, and a state $\bar{x} \in S$, decide if a non-Zeno infinite schedule exists that is S -safe at \bar{x} .

DEFINITION 3 (SAFE REACHABILITY). Given a MMS \mathcal{H} , a bounded convex set $S \subseteq \mathbb{R}^n$, and a pair of states $\bar{x}, \bar{x}' \in S$ decide if \bar{x}' is S -safe reachable from \bar{x} .

We present algorithms to solve safe schedulability and safe reachability problems in Section 2.3 and 2.4, respectively. We next present two examples of posing scheduling problems using constant-rate multi-mode systems.

2.2 Examples

The first example generalizes energy peak demand minimization problem as studied by Nghiem et al. [17].

EXAMPLE 1. Consider an organization with two zones A and B. HVAC units in each zone can be in one of the three modes 0 (OFF), 1 (LOW), and 2 (HIGH). We write the mode of the combined system as $m_{i,j}$ to represent the fact that unit A is in mode i and unit B is in mode j . The rate of temperature change and the energy usage for each zone in each mode is summarized in the following table:

Zones	HIGH	LOW	OFF
A (temp. change rate/ usage)	-2/3	-1/2	2/0.2
B (temp. change/ usage)	-2/3	-1/2	3/0.2

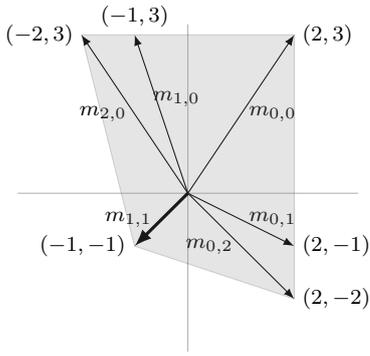


Figure 1: Rate vectors from Example 1.

For instance, if HVAC unit in zone A is in mode LOW then the temperature of the zone A drops 1 unit per second, while the HVAC unit consumes 2 energy units per second. To simplify energy pricing, we assume that the energy cost is equal to energy usage if peak energy usage (sum of the energy usage in all units) at every given point in time is less than or equal to 4 units, otherwise energy cost is 10 times of that standard rate. This assumption is in agreement with the bucket-based pricing [19] used by some energy providers.

It follows that to minimize energy cost, the peak usage, if possible, must not be higher than 4 units at any given time. Hence we model the system as an MMS with modes $m_{0,0}$, $m_{0,1}$, $m_{1,0}$, $m_{0,2}$, $m_{2,0}$, and $m_{1,1}$, because these are the only ones that have peak usage at most 4. The temperature of the zones are the variables of the MMS, while the safety set as the constraint that temperature of both zones should be between 65°F to 75°F . The existence of a safe non-Zeno schedule implies the existence of a switching schedule with energy peak demand less than or equal to 4 units. The rates of the variables in different modes is shown in Figure 1.

Two tank system model, a popular [16, 12] example of hybrid systems, can be modeled as MMS as shown below.

EXAMPLE 2. Consider a system consisting of two leaking tanks and a hose, such that each tank leaks water with some constant rate, and the hose can pump water in either of the tanks with a constant rate. The goal is to find a non-Zeno schedule to keep the water level of both tanks within a given range. It is straightforward to see that this problem can be modeled as a safe schedulability problem on MMS with two variables (water levels of tanks) and three modes that correspond to the positions of the water hose.

2.3 Safe Schedulability Problem

There exists a simple characterization of safe schedulability if the starting state \bar{x} is in the interior of set S . However, if the starting state \bar{x} is on the boundary of safety set S , then the safe schedulability problem is more involved. We treat these two cases separately.

2.3.1 Starting state is in the interior of the safety set

For all starting states in the interior of the safety set a safe non-Zeno schedule exists iff the following constraints are feasible for some vector $(f^{(1)}, f^{(2)}, \dots, f^{(|M|)}) \in \mathbb{R}_{\geq 0}^{|M|}$:

$$\sum_{i=1}^{|M|} R(i)(j) \cdot f^{(i)} = 0 \text{ for } 1 \leq j \leq n \text{ and } \sum_{i=1}^{|M|} f^{(i)} = 1. \quad (1)$$

The first constraint simply states that starting from an arbitrary state if the system spends $f^{(i)}$ fraction of total time in mode i then system comes back to the original state, while the second constraint is required to ensure non-Zenoness. The “if” part is straightforward as for any starting state \bar{x} a satisfying assignment to $f^{(i)}$ can be used to characterize a non-Zeno periodic S -safe schedule that forces all intermediate states to stay within a closed ball of arbitrary non-zero radius centered at the starting state. The “only if” part follows from Farkas’ lemma that states that constraints in (1) are feasible if and only if the following constraints are infeasible for all vectors $(v^{(1)}, v^{(2)}, \dots, v^{(n)}) \in \mathbb{R}^n$:

$$(v^{(1)}, v^{(2)}, \dots, v^{(n)}) \circ R(i) > 0 \text{ for all } 1 \leq i \leq |M| \quad (2)$$

Hence if constraints in (1) are infeasible, then constraints in (2) are feasible, i.e. there exists a vector $(v^{(1)}, v^{(2)}, \dots, v^{(n)})$ such that no matter which mode system stays in, it makes a positive progress along that vector. Since we assume the safety set is bounded, it implies that no safe non-Zeno schedule can exist in this case. This observation also implies that if there exist a safe non-Zeno schedule then there exists one that is periodic.

We say that mode m is S -safe at $\bar{x} \in S$ for $t > 0$ amount of time iff $\bar{x} + tR(m) \in S$ and m is S -safe at \bar{x} if there exists such a $t > 0$. We show that the constraints in (1) give a necessary and sufficient condition for safe schedulability for all starting states where all modes are S -safe. Algorithm 1 returns an S -safe schedule from a given starting state $\bar{x}_0 \in S$ where all modes are S -safe, while Theorem 2 states the correctness and the complexity of the algorithm.

Algorithm 1: Returns an S -safe schedule, if it exists, from a given \bar{x}_0 where all modes are safe.

Input: MMS \mathcal{H} , starting state \bar{x}_0 and $t > 0$ such that all modes of \mathcal{H} are safe at \bar{x}_0 for t amount of time.

Output: A periodic S -safe schedule from \bar{x} or NO if no such a schedule exists.

1 Check whether the following linear program is feasible:

$$\begin{aligned} \sum_{m \in M} R(m) \cdot f^{(m)} &= \vec{0} \\ \sum_{m \in M} f^{(m)} &= 1 \\ f^{(m)} &\geq 0 \text{ for all } m \in M. \end{aligned} \quad (3)$$

if no satisfying assignment exists then

2 \lfloor return NO

3 else

4 \lfloor Find a polynomial size assignment $\{f^{(m)}\}_{m \in M}$.

5 \lfloor return the following periodic schedule with period

$$\lfloor |M|: m_k = (k \bmod |M|) + 1 \text{ and } t_k = f^{(m_k)} \cdot t.$$

THEOREM 2. Given an MMS \mathcal{H} , a bounded convex safety set S , a state $\bar{x}_0 \in S$ and $t > 0$ s.t. all modes are safe at \bar{x}_0 for at least t amount of time, Algorithm 1 returns an S -safe non-Zeno periodic schedule (if exists) in polynomial time.

Notice that Algorithm 1 solves safe reachability problem for all starting states in the interior of the safety set, thanks to the following proposition.

PROPOSITION 3. If $\bar{x} \in \text{int}(S)$ then there exists $t > 0$ such that all modes are safe at \bar{x} for at least t amount of time.

PROOF. Since $\bar{x} \in \text{int}(S)$, there exists $d > 0$ such that $B_d(\bar{x}) \subseteq S$. We can just set t to be $\min_{m \in M} d / \|R(m)\| > 0$, because then for any $m \in M$ we have $\|tR(m)\| \leq t\|R(m)\| \leq d$ and so $\bar{x} + tR(m) \in B_d(\bar{x}) \subseteq S$. \square

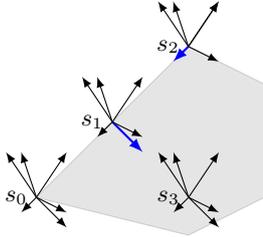
The constraints (1) and (2) give us two ways to test safe schedulability of an MMS for starting states in the interior of the safety set. From (1) it follows that there exists a safe non-Zeno schedule if and only if the origin lies in the convex hull of the points corresponding to the rate vector of each mode. While from (2) it follows that there is no safe non-Zeno schedule if there exists a vector such that the angles between that vector and every other rate vector are all less than 90° .

EXAMPLE 3. Using Figure 1 one can easily check that for the MMS from Example 1 there exists a safe non-Zeno schedule for all starting states in the interior of the safety set, as the convex hull (the shaded area) of the points corresponding to the rate vectors includes the origin, and also there is no vector which makes an angle less than 90° with all the vectors. It can also be easily verified that if we remove mode $m_{1,1}$ then there is no safe non-Zeno schedule.

2.3.2 Starting state is on the boundary of safety set

Feasibility of constraint 1 is not a sufficient condition for safe schedulability if the starting state lies on the boundary of the safe region as shown in the following example.

EXAMPLE 4. In the following figure we revisit Example 1 and draw the safety set (shaded region) S and four states s_0, s_1, s_2 and s_3 inside the safety set. The state s_3 is inside the interior of the set S , while other states are on the boundary of the set S .



From Theorem 2 it follows that from the starting state s_3 there exists a safe non-Zeno schedule. Now let us consider the state s_0 . It is clear that no mode is safe at s_0 as making any infinitesimally small progress along one of the rate vectors leads outside of the safety set S . Hence, for the starting state s_0 there is no safe non-Zeno schedule. On the other hand, from state s_1 choosing the vector $(2, -2)$ corresponding to mode $m_{0,2}$ leads into the interior of the safety region from where there exists a safe non-Zeno schedule. Similarly, from state s_2 first choosing the vector $(-1, -1)$ corresponding to mode $m_{1,1}$, and then choosing vector $(2, -2)$ corresponding to mode $m_{0,2}$ leads into the interior of the safe set.

The algorithm for the safe schedulability problem for boundary starting states follows from Theorem 7 presented in Section 3 and crucially depends on our results for the safe reachability problem presented next.

2.4 Safe Reachability Problem

Given an MMS $\mathcal{H} = (M, n, R)$, a bounded convex safety set S , and a starting state $\bar{x} \in S$, and a target state $\bar{x}' \in S$, the safe reachability problem is to decide whether \bar{x}' is S -safe reachable from \bar{x} . When all modes of \mathcal{H} are safe at \bar{x} and \bar{x}' , a safe schedule from \bar{x} to \bar{x}' exists iff the following constraints are satisfied for some vector $\vec{t} = (t^{(1)}, t^{(2)}, \dots, t^{(|M|)}) \in \mathbb{R}_{\geq 0}^{|M|}$:

$$\bar{x}(j) + \sum_{i=1}^{|M|} R(i)(j) \cdot t^{(i)} = \bar{x}'(j) \text{ for all } 1 \leq j \leq n \quad (4)$$

This constraint expresses that using a combination of rate vectors it is possible to reach \bar{x}' from \bar{x} . If the set of constraints in (4) is feasible and \bar{x}, \bar{x}' are in the interior of the safe set—or more generally all modes are safe at \bar{x} and \bar{x}' —then a satisfying assignment to $t^{(i)}$ can be used to make progress towards the vector $\bar{x}' - \bar{x}$ by scaling the $t^{(i)}$ -s appropriately without leaving the safety set. Repeating this process, \bar{x}' can be reached from \bar{x} in finitely many steps.

Algorithm 2 returns an S -safe schedule, if exists, that leads the system from \bar{x} to \bar{x}' given all modes are safe at both states. Theorem 4 establishes the correctness and the computational complexity of the algorithm.

Algorithm 2: Returns an S -safe schedule, if exists, from states \bar{x} and \bar{x}' when all modes are safe at \bar{x} and \bar{x}' .

Input: MMS \mathcal{H} , two points \bar{x}, \bar{x}' and $t > 0$ such that all modes are safe at \bar{x} and \bar{x}' for time t .

Output: NO, if no S -safe schedule from \bar{x} to \bar{x}' exists, and a periodic such schedule, otherwise.

1 Check whether the following linear program is feasible:

$$\begin{aligned} \bar{x} + \sum_{m \in M} R(m)t^{(m)} &= \bar{x}' \text{ and} \\ t^{(m)} &\geq 0 \text{ for all } m \in M. \end{aligned} \quad (5)$$

2 if no satisfying assignment exists then

3 return NO

4 else

5 Find a polynomial size assignment $\{t^{(m)}\}_{m \in M}$.

6 Let l be the smallest natural number greater than $\sum_{m \in M} t^{(m)} / t$.

7 return the following schedule of length $l|M|$ with period $|M|$: $m_k = (k \bmod |M|) + 1$ and $t_k = t^{(m_k)} / l$ for $k = 1, 2, \dots, l|M|$.

THEOREM 4. Given an MMS \mathcal{H} , a bounded convex set $S \subseteq \mathbb{R}^n$, a starting state $\bar{x} \in S$ and a target state $\bar{x}' \in S$, and a $t > 0$ such that all modes are safe at \bar{x} and \bar{x}' for at least t amount of time, Algorithm 2 returns an S -safe schedule (if exists) from \bar{x} to \bar{x}' in polynomial time.

PROOF. Assume that the linear constraints are feasible, and consider the schedule constructed in Algorithm 2. Let us denote that schedule by σ and consider the run $r = \text{Run}(\bar{x}, \sigma)$ corresponding to this schedule from state \bar{x} . Let us denote the state visited after the k -th step of r by \bar{x}_k , and also look at the states when the schedule σ is extended by $|M|$ extra steps. Notice that because σ is periodic with period $|M|$, for any $i < l$ and $j \leq |M|$ the following holds:

$$\bar{x}_{i|M|+j} = (1 - i/l)\bar{x}_j + (i/l)\bar{x}_{l|M|+j}.$$

So it suffices to prove that the points \bar{x}_j and $\bar{x}_{|M|+j}$ are S -safe for $j \leq |M|$. Let us denote for $j = 0, \dots, |M|$ by a_j the expression $1 - \sum_{m=1}^j t^{(m)} / (t \cdot l)$. Notice that

$$a_j \geq 1 - \sum_{m=1}^j t^{(m)} / \sum_{m \in M} t^{(m)} \geq 0,$$

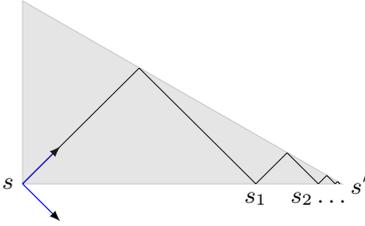
because $t \cdot l \geq \sum_{m \in M} t^{(m)}$. On the other hand, $\bar{x}_j = \bar{x} + \sum_{m=1}^j R(m)t^{(m)} / l = a_j \bar{x} + \sum_{m=1}^j (t^{(m)} / (t \cdot l)) (\bar{x} + tR(m))$. However, all modes were assumed to be safe at \bar{x} for t amount of time and so $\bar{x}_j \in S$ for all $j \leq M$, because it is a convex combination of states from the convex set S . In exactly the same manner we show $\bar{x}_{|M|+j} \in S$ for all $j \leq |M|$. This concludes the proof that the schedule returned by Algorithm 1 is S -safe.

The safe reachability problem can be solved in polynomial time, since the main computation of the algorithm involves solving a linear program. Moreover, although the schedule can be of exponential length, it can be represented compactly in polynomial space, because it has a period $|M|$. \square

2.4.1 General Case

Feasibility of the set of constraints in (5) does not guarantee safe reachability of state \bar{x}' from state \bar{x} if both states are on the boundary of the safety set. Moreover, in such a case the state \bar{x}' may not be reachable in finitely many steps as shown in the following example.

EXAMPLE 5. *In the following figure we present a MMS with two variables and two modes (1, 1) and (1, -1), and the starting state is s , while the target state is s' .*



The safety set is shown as the shaded area such that the angle at state s' is 30° . It is easy to see that the constraint in Equation 5 is feasible for this example, however, as we can see from the figure, the distance between states s_k and s' is equal to $\|s, s'\| \cdot \left(\frac{\sqrt{3}-1}{\sqrt{3}+1}\right)^k$. Hence, although \bar{x}' can be approached arbitrarily close using the two available modes, it is impossible to reach that state using a non-Zeno run.

3. OPTIMAL SCHEDULABILITY

In the previous section we presented algorithms to compute an arbitrary schedule out of all possible safe schedules. However, for most practical control problems when more than one safe schedule exists, it is often desirable to use a schedule that is optimal in according to some quantitative objective. In this section we study a priced extension of MMS where every mode is associated with a cost incurred per time unit when the system is that mode, and the natural generalization of safe schedulability and safe reachability problems to optimal average cost and optimal reachability cost problems for such extension.

DEFINITION 4. *A priced constant-rate multi-mode system is a tuple $\mathcal{H} = (M, n, R, \pi)$ where (M, n, R) is a MMS and*

$\pi : M \rightarrow \mathbb{R}$ is a price function such that $\pi(m)$ characterizes the price per-time unit of staying in mode m .

The price function can be extended to define costs for finite and infinite schedules. The cost of a finite schedule $\sigma = \langle (m_1, t_1), (m_2, t_2), \dots, (m_k, t_k) \rangle$ is defined as the sum of the costs of its timed actions, i.e., $Cost(\sigma) \stackrel{\text{def}}{=} \sum_{i=1}^k \pi(m_i) \cdot t_i$. The cost of an infinite schedule $\sigma = \langle (m_1, t_1), (m_2, t_2), \dots \rangle$ is defined as the average cost per time-unit, i.e.

$$AvgCost(\sigma) = \limsup_{k \rightarrow \infty} \frac{\sum_{i=1}^k \pi(m_i) \cdot t_i}{\sum_{i=1}^k t_i}.$$

Notice that we could define $AvgCost(\sigma)$ with \liminf instead of \limsup . However, because the aim is to find a schedule with the minimum average cost, it makes more sense to pick \limsup in order to minimize the maximal recurring average cost along a run. Also, as we will see later, this assumption is crucial in the proof of a key theorem of this paper.

Let $\Sigma(\bar{x}, \bar{x}', S)$ be the set of finite schedules that are S -safe at \bar{x} and lead the system from \bar{x} to \bar{x}' . The *optimal reachability cost* $Cost_*^S(\bar{x}, \bar{x}')$ for a starting state \bar{x} , a target state \bar{x}' , and a safety set S is defined as:

$$Cost_*^S(\bar{x}, \bar{x}') = \inf \{ Cost(\sigma) : \sigma \in \Sigma(\bar{x}, \bar{x}', S) \}.$$

We say that a schedule σ is reachability-cost optimal for a starting state \bar{x} , a target state \bar{x}' and a safety set S , if $\sigma \in \Sigma(\bar{x}, \bar{x}', S)$ and $Cost(\sigma) = Cost_*^S(\bar{x}, \bar{x}')$.

Let $\Sigma(\bar{x}, S)$ be the set of non-Zeno infinite schedules that are S -safe at \bar{x} . The optimal average cost of a state \bar{x} and a safety set S is defined as

$$AvgCost_*^S(\bar{x}) = \inf \{ AvgCost(\sigma) : \sigma \in \Sigma(\bar{x}, S) \}.$$

We say that a schedule σ is average-cost optimal for a starting state \bar{x} and a safety set S if $\sigma \in \Sigma(\bar{x}, S)$ and $AvgCost(\sigma) = AvgCost_*^S(\bar{x})$.

Two fundamental problems for priced MMS are the following optimal reachability and optimal average schedulability problems.

DEFINITION 5 (OPTIMAL REACHABILITY). *Given a priced MMS \mathcal{H} , a bounded convex set $S \in \mathbb{R}^n$, a starting state $\bar{x} \in S$, and a target state $\bar{x}' \in S$, compute the optimal reachability cost $Cost_*^S(\bar{x}, \bar{x}')$ and find, if exists, a reachability-cost optimal schedule.*

DEFINITION 6 (OPTIMAL AVERAGE SCHEDULABILITY). *Given a priced MMS \mathcal{H} , a bounded convex set $S \in \mathbb{R}^n$, and a state $\bar{x} \in S$ compute the optimal average cost $AvgCost_*^S(\bar{x})$ and find, if exists, an average-cost optimal schedule.*

In Sections 3.1 and 3.2 we present algorithms to solve optimal reachability and optimal average schedulability problems, respectively.

3.1 Optimal Reachability Problem

Let us fix a priced MMS $\mathcal{H} = (M, n, R, \pi)$, a starting state \bar{x} , a target state \bar{x}' , and a bounded convex safety set S for this section. As we noticed in Example 5, if the points \bar{x} and \bar{x}' are on the boundary of the safety set, optimal schedules may not always exist. We now show that using Algorithm 3, we can solve in polynomial time the optimal reachability problem between any two points in the interior of S .

Algorithm 3: Given all modes are safe for time $t > 0$ at \bar{x}, \bar{x}' and at least one S -safe schedule from \bar{x} to \bar{x}' exists, the algorithm returns a cost-optimal such schedule.

Input: MMS \mathcal{H} , two points \bar{x}, \bar{x}' and $t > 0$ such that all modes of \mathcal{H} are safe at \bar{x} and \bar{x}' for time t .

Output: NO, if no S -safe schedule from \bar{x} to \bar{x}' exists, and an optimal periodic such schedule, o/w.

- 1 Check whether the following linear programming problem with variables $\{t^{(m)}\}_{m \in M}$ has a solution.

$$\begin{aligned} & \text{Minimize } \sum_{m \in M} \pi(m)t^{(m)} \text{ subject to:} \\ & \bar{x} + \sum_{m \in M} R(m)t^{(m)} = \bar{x}' \text{ and} \\ & t^{(m)} \geq 0 \text{ for all } m \in M. \end{aligned}$$

- 2 if no satisfying assignment exists then

- 3 \lfloor return NO

- 4 else

- 5 \lfloor Find a polynomial size assignment $\{t^{(m)}\}_{m \in M}$.
 - 6 \lfloor Let l be the smallest natural number greater or equal to $\sum_{m \in M} t^{(m)}/t$.
 - 7 \lfloor return the following schedule of length $l|M|$ with period $|M|$: $m_k = (k \bmod |M|) + 1$ and $t_k = t^{(m_k)}/l$ for $k = 1, 2, \dots, l|M|$.
-

THEOREM 5. Given an MMS \mathcal{H} , a bounded convex set $S \subseteq \mathbb{R}^n$, states \bar{x}, \bar{x}' , and $t > 0$ such that all modes of \mathcal{H} are safe at \bar{x} and \bar{x}' for at least t amount of time, Algorithm 3 returns a reachability-cost optimal schedule (if exists) in polynomial time.

PROOF. It follows from the proof of Theorem 4 that the schedule, σ_* , returned by Algorithm 3 is S -safe. We now show σ_* is also reachability-cost optimal.

Assume that there is at least one S -safe schedule from \bar{x} to \bar{x}' . Let σ be an arbitrary such schedule and let k be its length. Notice that of course we have

$$\bar{x}' = \bar{x} + \sum_{m \in M} R(m)T_k^m(\sigma),$$

and for all $m \in M$ we have that $T_k^m(\sigma) \geq 0$, since recall that $T_k^m(\sigma)$ is the total time spent in mode m up to step k .

Therefore, by setting $t^{(m)} \mapsto T_k^m(\sigma)$ for all $m \in M$, the linear constraints in Algorithm 1 become satisfied. Furthermore, we have $\text{Cost}(\sigma) = \sum_{m \in M} \pi(m) \cdot T_k^m(\sigma)$, which is exactly the same as the value of the objective function for such an assignment. Hence, $\text{Cost}(\sigma) \geq \text{Cost}(\sigma_*)$ because σ_* has the minimal value of the objective function among all assignments that satisfy the linear constraints in Algorithm 1. However, we picked σ arbitrarily and so we have that $\text{Cost}(\sigma_*) = \text{Cost}_S^*(\bar{x}, \bar{x}')$. \square

3.2 Optimal Average Schedulability Problem

Let us fix a priced MMS $\mathcal{H} = (M, n, R, \pi)$ and a bounded convex set S . We first present an algorithm for computing the average-optimal schedule for the case when all the modes are safe at the starting state. From Proposition 3 such starting states include all states in the interior of the safe set S .

3.2.1 Starting state is in the interior of the safety set

We show that the optimal average-cost does not depend on the starting state as long as all modes are S -safe at that state. Moreover, there always exists a period schedule that has the optimal average-cost among all S -safe schedules.

THEOREM 6. Given a MMS \mathcal{H} , a bounded convex set $S \subseteq \mathbb{R}^n$, $t > 0$, an initial state $\bar{x}_0 \in S$ such that all modes of \mathcal{H} are S -safe at \bar{x}_0 for t amount of time, Algorithm 4 returns a periodic average-cost optimal schedule in polynomial time.

PROOF. We first show that for any S -safe schedule, $\sigma = \langle (m_1, t_1), (m_2, t_2), \dots \rangle$, from \bar{x}_0 we can construct a periodic schedule with period $|M|$ whose average-cost is not greater than σ 's.

Let $f_k^{(m)} = T_k^m(\sigma)/T_k(\sigma)$ represents the fraction of the time spent by σ in mode m up to step k ; note that $f_k^{(m)} \in [0, 1]$, and $\sum_{m \in M} f_k^{(m)} = 1$ for all k . Also, from the definitions,

$$\bar{x}_k = \bar{x}_0 + \sum_{m \in M} R(m)T_k^m(\sigma) = \bar{x}_0 + T_k(\sigma) \sum_{m \in M} R(m)f_k^{(m)},$$

for any k , and $\text{AvgCost}(\sigma) = \limsup_{k \rightarrow \infty} \sum_{m \in M} \pi(m)f_k^{(m)}$. The definition of lim sup stipulates that we can pick a subsequence of the sequence $\langle \sum_{m \in M} \pi(m)f_k^{(m)} \rangle_{k=1}^{\infty}$ that converges to $\text{AvgCost}(\sigma)$. In other words, there exists an increasing integer sequence i_1, i_2, \dots such that $\text{AvgCost}(\sigma) = \lim_{k \rightarrow \infty} \sum_{m \in M} \pi(m)f_{i_k}^{(m)}$. Let us now look at the sequence of vectors $\langle \vec{f}_k \in [0, 1]^M \rangle_{k=1}^{\infty}$ where we set $\vec{f}_k(m) = f_{i_k}^{(m)}$. Since this sequence is bounded, by the Bolzano-Weierstrass theorem, there exists an increasing integer sequence j_1, j_2, \dots such that $\lim_{k \rightarrow \infty} \vec{f}_{j_k}$ exists and let us denote this limit by \vec{f} . We next prove by contradiction that

$$\sum_{m \in M} R(m)\vec{f}(m) = \vec{0}. \quad (6)$$

Assume that $\sum_{m \in M} R(m)\vec{f}(m) \neq \vec{0}$. Then for some variable $1 \leq v \leq n$ we have $c := \sum_{m \in M} R(m)(v)\vec{f}(m) \neq 0$ and wlog assume this value to be positive. From the definition of \vec{f} , for any $\varepsilon > 0$ we can pick N such that for all $k > N$ and $m \in M$ we have $|\vec{f}_{j_k}(m) - \vec{f}(m)| < \varepsilon$. Now, notice that

$$\begin{aligned} \bar{x}_{i_{j_k}}(v) - \bar{x}_0(v) &= T_{i_{j_k}}(v) \sum_{m \in M} R(m)(v)f_{i_{j_k}}^{(m)} \\ &= T_{i_{j_k}} \sum_{m \in M} R(m)(v)\vec{f}_{j_k}(m) \\ &= T_{i_{j_k}} \sum_{m \in M} R(m)(v) \cdot (\vec{f}(m) + (\vec{f}_{j_k}(m) - \vec{f}(m))) \\ &\geq T_{i_{j_k}} \sum_{m \in M} (R(m)(v) \cdot \vec{f}(m) - R_{\max}\varepsilon) \\ &= T_{i_{j_k}}(c - |M|R_{\max}\varepsilon) \end{aligned}$$

where $R_{\max} := \max_{m \in M, w \in V} |R(m)(w)|$. If we now set ε to be $c/(2|M|R_{\max})$ for $R_{\max} \neq 0$ and $\varepsilon = 0$ otherwise, then $\bar{x}_{i_{j_k}}(v) - \bar{x}_0(v) \geq \frac{1}{2}T_{i_{j_k}}c$, where the right-hand side tends to ∞ , because $\lim_{k \rightarrow \infty} i_{j_k} = \infty$ and $c > 0$. Therefore, we have $\lim_{k \rightarrow \infty} \|\bar{x}_{i_{j_k}} - \bar{x}_0\| = \infty$, which would imply that either S is not bounded or r is not S -safe; a contradiction.

Let the periodic schedule $\sigma' = \langle (m'_1, t'_1), (m'_2, t'_2), \dots \rangle$ be s.t. $m'_k = (k \bmod |M|) + 1$ and $t'_k = \vec{f}(m_k) \cdot t$ for each $k \geq 1$. Also, let $r' = \langle \bar{x}_0, (m'_1, t'_1), \bar{x}'_1, \dots \rangle$ be the corresponding run

from \bar{x}_0 . It is straightforward to see that σ' is non-Zeno. We show that σ' is S -safe at \bar{x}_0 . Notice that from (6) it follows

$$\bar{x}'_{|M|} = \bar{x}_0 + \sum_{m \in M} R(m)t\vec{f}(m) = \bar{x}_0.$$

So it suffices to show that the finite prefix of length $|M|$ of r' is S -safe. However, this prefix is exactly the same schedule that we would construct in the proof of Proposition 1 for the convex combination of points $\{\bar{x}_0 + tR(m) \mid m \in M\}$, where $\vec{f}(m)$ is picked as the coefficient of the point $\bar{x}_0 + tR(m)$ in this combination. All these points are trivially reachable by a finite path from \bar{x}_0 , because all modes were assumed to be safe for t amount of time at \bar{x}_0 .

We now show that σ' has the average-cost not greater than the original schedule σ .

$$\begin{aligned} \text{AvgCost}(\sigma) &= \limsup_{k \rightarrow \infty} \sum_{m \in M} \pi(m) f_k^{(m)} \\ &\geq \limsup_{k \rightarrow \infty} \sum_{m \in M} \pi(m) f_{j_k}^{(m)} = \sum_{m \in M} \pi(m) \lim_{k \rightarrow \infty} f_{j_k}^{(m)} \\ &= \sum_{m \in M} \pi(m) \vec{f}(m) = \text{AvgCost}(\sigma'). \end{aligned}$$

The first inequality follows from the fact that removing elements from a sequence can only lower its limsup value, while the second equality follows as the limsup of a sum of bounded converging sequences is equal to the sum of its limits. This proves that the optimal average-cost among all S -safe schedules is equal to the infimum of the average-cost over S -safe periodic runs with period $|M|$.

Notice that Algorithm 4 differs from Algorithm 1 only at line 1 where an objective function is added to the linear program (7). The linear constraints guarantee the periodic schedule constructed at the end of the algorithm to be S -safe, while the objective function guarantees its average-cost to be the lowest among all S -safe schedules. \square

3.2.2 General Case

In this section we show how to handle arbitrary starting states as long as the safety set is a convex polytope.

THEOREM 7. *Given any MMS \mathcal{H} , bounded convex polytope $S \subseteq \mathbb{R}^n$, and an initial state $\bar{x}_0 \in S$, Algorithm 5 returns in polynomial time an ultimately periodic S -safe schedule with the minimum average-cost.*

PROOF. The algorithm first computes an increasing sequence of sets of modes $M_1 \subset M_2 \subset \dots$, where M_1 is the set of modes safe at \bar{x}_0 , and for $i \geq 1$ the set M_{i+1} consists of all modes safe at states S -safe reachable from \bar{x}_0 using only modes in M_i . Let k be the smallest number such that $M_{k+1} = M_k$. Of course $k \leq |M|$, because M_{i+1} has to have at least one more mode than M_i and the total number of modes is $|M|$. We will show that no mode outside of M_k can ever become safe during a S -safe schedule starting at \bar{x}_0 . First, we need the following lemma that can transform any S -safe schedule into a S -safe schedule with the set of safe modes never decreasing as the new schedule progresses.

LEMMA 8. *Any S -safe finite schedule σ from $\bar{x}_0 \in S$ can be modified to a S -safe finite schedule σ' of length polynomial in σ in a way that all modes that were safe at any state along σ will be safe at the terminal state of σ' and no safe mode along σ' can become unsafe as σ' progresses.*

Algorithm 4: Given all modes are safe for time $t > 0$ at \bar{x} and an S -safe non-Zeno schedule exists from \bar{x} , the algorithm returns average-cost optimal such schedule.

Input: MMS \mathcal{H} , initial point $\bar{x} \in \text{int}(S)$, and $t > 0$ such that all modes of \mathcal{H} are safe for time t .

Output: NO, if no S -safe non-Zeno schedule exists from \bar{x} , and a periodic such schedule with the minimum average-cost, otherwise.

- 1 Check whether the following linear programming problem with variables $\{f^{(m)}\}_{m \in M}$ has a solution.

$$\begin{aligned} &\text{Minimize } \sum_{m \in M} \pi(m) \cdot f^{(m)} \text{ subject to} \\ &\sum_{m \in M} R(m) \cdot f^{(m)} = \vec{0} \\ &\sum_{m \in M} f^{(m)} = 1 \\ &f^{(m)} \geq 0 \text{ for all } m \in M. \end{aligned} \quad (7)$$

- 2 if no satisfying assignment exists then

- 3 \lfloor return NO

- 4 else

- 5 \lfloor Find a polynomial size assignment $\{f^{(m)}\}_{m \in M}$.

- 6 \lfloor return the following periodic schedule with period $|M|$: $m_k = (k \bmod |M|) + 1$ and $t_k = f^{(m_k)} \cdot t$.

Now, let σ be an arbitrary S -safe schedule starting at \bar{x}_0 . First, using Lemma 8, we construct a new schedule σ' based on σ with the property which implies the existence of a finite list of states, $\bar{x}_1, \bar{x}_2, \dots$, along the run $\text{Run}(\bar{x}_0, \sigma')$, such that for any $i \geq 0$ the set of safe modes at \bar{x}_{i+1} is strictly greater than at \bar{x}_i and does not change at the states between \bar{x}_i and \bar{x}_{i+1} in $\text{Run}(\bar{x}_0, \sigma')$. Hence, for any $i \geq 0$ the point \bar{x}_{i+1} has to be reachable from \bar{x}_0 using modes safe at \bar{x}_i only. From the definition of the sets M_i and an easy induction on i , the set of modes safe at \bar{x}_i has to be a subset of M_{i+1} . Therefore, the set of safe modes at the terminal state of σ' is also a subset of M_i for some i . However, the set of safe modes at the terminal state of σ' was supposed to contain all the modes safe along $\text{Run}(\bar{x}_0, \sigma)$ and so any mode safe along a S -safe schedule from \bar{x}_0 has to belong to M_k .

We now show that a very specific schedule always exists.

LEMMA 9. *There exist states $\bar{x}_1, \bar{x}_2, \dots, \bar{x}_k \in S$ such that for all $1 \leq i \leq k$ we have that state \bar{x}_i is S -safe reachable from \bar{x}_{i-1} using only modes in M_i , and all modes in M_i are safe at \bar{x}_{i-1} .*

PROOF. Pick any mode $q \in M_k$ and let $I(q)$ be the lowest index such that $q \in M_{I(q)}$. Next, pick any finite S -safe schedule such that mode q is safe at some state along the run of this schedule from \bar{x}_0 and that is the first state of that run where a mode belonging to $M_{I(q)} \setminus M_{I(q)-1}$ is safe. Such a schedule has to exist from the definition of the sets M_i . Let σ^q be the result of the transformation of this schedule into a new one using Lemma 8. For $1 \leq i \leq I(q) - 1$, denote by \bar{x}_i^q the first state along the run $\text{Run}(\bar{x}_0, \sigma^q)$ when a mode from $M_{i+1} \setminus M_i$ becomes safe, some of these states may coincide, and for all $i \geq I(q)$ let the terminal state of $\text{Run}(\bar{x}_0, \sigma^q)$ be

assigned to \bar{x}_i^q . Notice that mode q is already safe at state \bar{x}_{i-1}^q , because of the way we picked σ^q .

Let us define $\bar{x}_i := \frac{1}{|M_k|} \sum_{q \in M_k} \bar{x}_i^q$ for all i . From Proposition 1, we know that $\bar{x}_i \in S$ and any mode safe at \bar{x}_i^q is safe at \bar{x}_i . But this means that all modes in M_i are safe at \bar{x}_{i-1} , because any mode $q \in M_i$ is safe at \bar{x}_j^q for all $j \geq i-1$. Furthermore, we know that \bar{x}_i^q is S -safe reachable from \bar{x}_{i-1}^q in $Run(\bar{x}_0, \sigma^q)$ using only modes from M_i , because only these modes are safe before this run reaches \bar{x}_i^q . Therefore, it has to be $\bar{x}_i^q - \bar{x}_{i-1}^q = \sum_{m \in M_i} R(m)t_m^q$ for some $t_m^q \geq 0$. But this means that $\bar{x}_i - \bar{x}_{i-1} = \frac{1}{|M_k|} \sum_{q \in M_k} \bar{x}_i^q - \bar{x}_{i-1}^q = \sum_{m \in M_i} R(m) \left(\frac{1}{|M_k|} \sum_{q \in M_k} t_m^q \right)$. Moreover, all modes from M_i are safe both at \bar{x}_i and \bar{x}_{i-1} . Hence, based on the S -safe reachability characterization given in Theorem 4, we get that \bar{x}_i is reachable from \bar{x}_{i-1} via a S -safe schedule. \square

Algorithm 5 computes the sets of modes M_i iteratively. Having computed M_i for all $i \leq j$, it checks for every mode $q \in M \setminus M_j$ whether there exists a sequence of states $\bar{x}_1, \dots, \bar{x}_j$ with the properties as in Lemma 9 and mode q being safe at \bar{x}_j . Lemma 9 guarantees that all safe modes from M_{j+1} will be found this way. We can use Theorem 4 to characterize the S -safe reachability of \bar{x}_i from \bar{x}_{i-1} via a set of linear constraints. Combining all of these for $i = 1, 2, \dots, j$, and some additional checks regarding the safety of the states and modes used, we obtain the same linear program as in step 6 of Algorithm 5. In the end, the algorithm finds a polynomial size solution to the whole system of constraints while maximizing t , the minimum amount of time each mode is safe for. Algorithm 3 can then be used to find a polynomial length schedule from \bar{x}_i to \bar{x}_{i+1} for each i .

After that, Algorithm 5 calls Algorithm 4 to find the optimal average-cost S -safe schedule σ from \bar{x}_{k-1} that uses only modes from M_k . If no such schedule exists, then there is not any S -safe schedule from \bar{x}_0 either. However, we know that no other mode can be used by a S -safe schedule and so we can assume that the set of all modes is M_k . In such a case, σ has the minimum average-cost among all S -safe schedules. Moreover, a finite prefix of a schedule does not influence its average-cost, so the schedule consisting of the polynomial length S -safe schedules leading the system from \bar{x}_0 to \bar{x}_{k-1} followed by σ is an ultimately periodic S -safe schedule which has the minimum average-cost. Finally, Algorithm 5 runs in polynomial time, because at most $|M|$ iterations of the loop is needed to find the sets of modes M_1, \dots, M_k and it calls other algorithms that were already shown to run in polynomial time.

4. DISCRETE-TIME SCHEDULABILITY

In the previous section, we showed that there exist optimal schedules for the average cost and reachability cost criteria with time delays in each mode expressible as rational numbers of the size polynomial in the size of the MMS. Also, the optimal average cost schedule can be made periodic with period $|M|$ and the time delays for the optimal cost one are also periodic with period $|M|$. It follows that such schedules are implementable by a discrete controller with the sampling rate Δ being the least common multiple of all these $|M|$ rational time delays. Such a number will also be a polynomial sized rational number, however it may be exponentially smaller than any of the time delays used.

Algorithm 5: Finds a S -safe non-Zeno run from a given $\bar{x}_0 \in S$ with the minimal average cost, where S is a bounded convex polytope.

Input: MMS \mathcal{H} , $l \times n$ matrix A and vector $\vec{b} \in \mathbb{R}^l$ that together define a bounded convex polytope S such that $\bar{x} \in S$ iff $A\bar{x} \leq \vec{b}$, and initial point $\bar{x}_0 \in S$.

Output: NO if no S -safe non-Zeno schedule exists from \bar{x}_0 , and an average-cost optimal ultimately periodic such schedule, otherwise.

- 1 $M_0 := \emptyset; k := -1;$
- 2 **repeat**
- 3 $k := k + 1;$
- 4 $M_{k+1} := M_k;$
- 5 **foreach** mode $q \in M \setminus M_k$ **do**
- 6 **if** the following set of linear constraints is satisfiable for some assignment to the variables $t, \{t_1^{(m)}\}_{m \in M_1}, \{t_2^{(m)}\}_{m \in M_2}, \dots, \{t_k^{(m)}\}_{m \in M_k}$

$$t > 0$$

For all $i = 1, \dots, k$

$$t_i^{(m)} \geq 0 \text{ for all } m \in M_i$$

$$\bar{x}_i = \bar{x}_{i-1} + \sum_{m \in M_i} R(m)t_i^{(m)}$$

$$A\bar{x}_i \leq \vec{b}$$

$$A(\bar{x}_{i-1} + R(m)t) \leq \vec{b} \text{ for all } m \in M_i$$

$$A(\bar{x}_k + R(m)t) \leq \vec{b} \text{ for all } m \in M_k \cup \{q\} \quad (8)$$
- 7 **then**

$$M_{k+1} := M_{k+1} \cup \{q\};$$
- 8 **until** $M_{k+1} = M_k$;
- 9 Compute a polynomial sized solution to the linear program in step 6 but with the constraints (8) removed and the objective function *Maximize* t .
- 10 Call Algorithm 4 with the set of modes M_k , starting point \bar{x}_{k-1} and the safe time bound t .
- 11 **if** the call returned NO **then**
- 12 **return** NO
- 13 **else**
- 14 Let σ be the periodic schedule returned otherwise.
- 15 Repeatedly call Algorithm 3 to find S -safe schedules of polynomial length from \bar{x}_0 to \bar{x}_1 using only modes in M_1 , from \bar{x}_1 to \bar{x}_2 using only modes in M_2, \dots , from \bar{x}_{k-2} to \bar{x}_{k-1} using only modes in M_{k-1} .
- 16 **return** the ultimately periodic schedule created by composing the S -safe finite schedules from \bar{x}_i to \bar{x}_{i+1} for $i = 0, 1, \dots, k-2$ with the average-cost optimal S -safe periodic schedule σ from \bar{x}_{k-1} .

Such a sampling rate may not be acceptable in practice, e.g., the switching frequency is too high, or the dwell time for some modes is too low. In this section we discuss discrete-time schedulers that are allowed to switch modes only at times multiple of a given sampling rate. We present results on deciding whether a given Δ suffices for schedulability or optimality (Theorem 10), and how to find the maximum sampling rate among all feasible Δ values (Theorem 11).

Let $\mathcal{H} = (M, n, R)$ be an MMS and let Δ be a given sampling rate. We say that a schedule $\sigma = \langle (m_1, t_1), (m_2, t_2), \dots \rangle$ is Δ -clocked if for all $i \geq 1$ there is $d_i \in \mathbb{N}_{>0}$ such that $t_i = d_i \cdot \Delta$. Given a bounded convex set S and states $\bar{x}, \bar{x}' \in \mathbb{R}^n$, the optimal average cost $AvgCost_{\Delta}^S(\bar{x})$ and optimal reachability cost $Cost_{\Delta}^S(\bar{x}, \bar{x}')$ over all Δ -clocked S -safe schedules starting from \bar{x} are defined in straightforward manner.

THEOREM 10. *Given a priced MMS $\mathcal{H} = (M, n, R, \pi)$, a bounded polytope S , a sampling rate Δ , and states $\bar{x}_0, \bar{x}_d \in \mathbb{R}^n$ the discrete average cost and the discrete reachability cost problems are PSPACE-complete.*

PSPACE-membership of both problems is shown via discretization of the state space of \mathcal{H} . Since the set S is given as a bounded polytope, the size of the discretization can be shown to be at most exponential in the size of \mathcal{H} and Δ . We prove PSPACE-hardness by a reduction from the acceptance problem for the linear bounded automata (LBAs).

Lower bound on Δ can be computed using Theorems 6 and 7, while the upper bound can be obtained from the diameter of the safety set. Given a target average-cost, using a straightforward binary search algorithm we can approximate the maximum sampling rate for which the optimal average cost does not exceed the target. Since the total number of iterations is polynomial and in each iteration the optimality can be checked in PSPACE, the following theorem follows.

THEOREM 11. *Given a MMS \mathcal{H} , a bounded polytope S , a starting state s , a budget B , and $\varepsilon > 0$, the maximum sampling rate for which the optimal average-cost is not greater than B can be approximated within ε in PSPACE.*

5. CONCLUSION

We have proposed a model for constant-rate multi-mode systems (MMSs) to analyze hybrid systems with variables having mode-dependent constant-rates and no constraints on mode-switching. For this model, we have developed polynomial time algorithms to solve safe schedulability and safe reachability problems, as well as their corresponding optimization problems. From a practical perspective, a number of quantitative analysis problems for hybrid systems, in particular, energy peak demand minimization problem proposed recently by Nghiem et al., can be formalized as optimal schedulability problems for MMSs. Our analysis algorithms reduce the problem to linear programming, and are tractable. As such, MMSs are a new promising, natural and expressive subclass of hybrid systems. There are, however, some natural optimization problems on MMSs where the linear programming formulation breaks down. For instance, if we allow different prices for mode-switches then the choice of the mode switching sequence plays a crucial role in the average cost of a schedule, and the linear programming characterization presented in this paper does not work. For this setting, the exact complexity of optimal schedulability problem remains open.

Acknowledgments

We thank Madhur Behl, Rahul Mangharam, Truong Nghiem, and George Pappas for fruitful discussions on scheduling for energy optimization problems. This research was partially supported by NSF awards CNS 0931239, CNS 1035715, CCF 0915777, and EPSRC grant EP/G050112/2.

6. REFERENCES

- [1] Y. Abdeddaim, E. Asarin, and O. Maler. Scheduling with timed automata. *TCS*, 354:272–300, March 2006.
- [2] Y. Abdeddaim and Oded Maler. Preemptive job-shop scheduling using stopwatch automata. In *TACAS*, volume 2280 of *LNCS*, pages 39–53. Springer, 2002.
- [3] R. Alur and D. Dill. A theory of timed automata. *TCS*, 126(2):183–235, 1994.
- [4] R. Alur, S. La Torre, and G. J Pappas. Optimal paths in weighted timed automata. In *HSCC*, pages 49–62. Springer, 2001.
- [5] R. Alur, A. Trivedi, and D. Wojtczak. Optimal scheduling for constant-rate multi-mode systems. Technical Report MS-CIS-12-01, CIS, UPenn, 2012.
- [6] D. Belzer, M. Scott, and R. Sands. Climate change impacts on U.S. commercial building consumption: An analysis using sample survey data. *Energy Sources*, 18(2):77–201, 1996.
- [7] P. Bouyer, E. Brinksma, and K. G. Larsen. Staying alive as cheaply as possible. In *HSCC*, volume 2993 of *LNCS*, pages 203–218. Springer, 2004.
- [8] P. Bouyer, U. Fahrenberg, K. G. Larsen, N. Markey, and J. Srba. Infinite runs in weighted timed automata with energy constraints. In *FORMATS*, pages 33–47, 2008.
- [9] Giorgio C. Buttazzo. *Hard Real-time Computing Systems*. Springer-Verlag, 2004.
- [10] T. A. Henzinger and P. W. Kopke. Discrete-time control for rectangular hybrid automata. *TCS*, 221(1-2):369–392, 1999.
- [11] T. A. Henzinger, P. W. Kopke, A. Puri, and P. Varaiya. What’s decidable about hybrid automata? *Journal of Comp. and Sys. Sciences*, 57:94–124, 1998.
- [12] M. Heymann, L. Feng, G. Meyer, and S. Resmerita. Analysis of zeno behaviors in a class of hybrid systems. *IEEE Trans. on Auto. Ctrl.*, 50:376–383, 2005.
- [13] Susmit Jha, Sumit Gulwani, Sanjit Seshia, and Ashish Tiwari. Synthesizing switching logic for safety and dwell-time requirements. In *International Conference on Cyber-physical Systems*, pages 22–31, 2010.
- [14] Susmit Jha, Sanjit A. Seshia, and Ashish Tiwari. Synthesizing switching logic to minimize long-run cost. *CoRR*, abs/1103.0800, 2011.
- [15] Kronos. <http://www-verimag.imag.fr/TEMPORISE/kronos/>.
- [16] J. Lygeros. *Lecture Notes on Hybrid Systems*. Cambridge, 2003.
- [17] T. X. Nghiem, M. Behl, R. Mangharam, and G. J. Pappas. Green scheduling of control systems for peak demand reduction. In *IEEE CDC*, December 2011.
- [18] T. X. Nghiem, M. Behl, G. J. Pappas, and R. Mangharam. Green scheduling: Scheduling of control systems for peak power reduction. *2nd International Green Computing Conference*, July 2011.
- [19] TRFund Energy Study. Understanding PECO’s general service tariff.
- [20] UPPAAL. <http://www.uppaal.com/>.
- [21] X. Xu and P. J. Antsaklis. Practical stabilization of integrator switched systems. In *American Control Conference*, pages 2767–2772, 2003.