

Optimal Control for Multi-Mode Systems with Discrete Costs

Mahmoud A. A. Mousa, Sven Schewe, and Dominik Wojtczak

University of Liverpool, Liverpool, U.K.

Abstract. This paper studies optimal time-bounded control in multi-mode systems with discrete costs. Multi-mode systems are an important subclass of linear hybrid systems, in which there are no guards on transitions and all invariants are global. Each state has a continuous cost attached to it, which is linear in the sojourn time, while a discrete cost is attached to each transition taken. We show that an optimal control for this model can be computed in NEXPTIME and approximated in PSPACE. We also show that the one-dimensional case is simpler: although the problem is NP-complete (and in LOGSPACE for an infinite time horizon), we develop an FPTAS for finding an approximate solution.

1 Introduction

Multi-mode systems [8] are an important subclass of linear hybrid systems [4], which consist of multiple continuous variables and global invariants for the values that each variable is allowed to take during a run of the system. However, unlike for the full linear hybrid systems model, multi-mode systems have no guards on transitions and no local invariants. In this paper, we study multi-mode systems with discrete costs, which extend linear hybrid systems by adding both continuous and discrete costs to states. Every time a transition is taken (i.e. when the current state changes), the discrete cost assigned to the target state is incurred. The continuous cost is the sum of the products of the sojourn time in each state and the cost assigned to this state. Our aim is to minimise the total cost over a finite-time horizon or a long-time average cost over an infinite time horizon. We exemplify this by applying this model to the optimal control of heating, ventilation, and air-conditioning (HVAC) systems. HVAC systems account for about 50% of the total energy cost in buildings [27], so a lot of energy can be saved by optimising their control. Many simulation programs have been developed to analyse the influence of control on the performance of HVAC system components such as TRNSYS [3], EnergyPlus [1], and the Matlab's IBPT [2]. Our approach has the advantage over the existing control theory techniques that it provides approximation guarantees. Although the actual dynamics of a HVAC system is governed by linear differential equations, one can argue [24, 25, 22] that constant rate dynamic, as in our model, can approximate well such a behaviour.

The simplest subclass of our model is multi-mode systems with a single dimension. It naturally occurs when controlling the temperature in a single room or building to stay in a pleasant range. For this, the system can be in different modes, e.g. the air-conditioning can be switched on or off, or one can choose to switch on an electrical radiator or a gas burner. Each such a configuration can be modelled as mode of our multi-mode system. Modes have start-up cost (gas burners, e.g. may suffer some wear and tear when switched

on) as well as continuous costs.

When keeping an office building in a pleasant temperature range during opening hours, we face a control problem for multi-mode systems with a finite time horizon. We show that finding an optimal schedule in such a case is NP-complete and significantly more challenging than for the infinite time horizon (LogSPACE). However, we devise an FPTAS for the finite time horizon problem.

Heating multiple rooms simultaneously can be naturally modelled by multi-mode systems (with multiple dimensions). In such a scenario, we might have different pleasant temperature ranges in different rooms and the temperatures of the individual rooms may influence each other. Naturally, controlling a multi-dimensional multi-mode systems is more complex than controlling a one-dimensional multi-mode system. We develop a non-deterministic exponential time algorithm for the construction of optimal control, whose complexity is only driven by potentially required high precision in exponentially many mode switches. Allowing for an ε -deviation from the ranges of pleasant temperatures reduces the complexity to PSPACE.

Related work. Our model can be viewed as a weighted extension of the linear hybrid automata model ([5, 17]), but with global constraints only. Even basic questions for the general linear hybrid automata model are undecidable already for three variables and not known to be decidable for two variables [9]. Most of the research for this model has focused on qualitative objectives such as reachability. Various subclasses of hybrid systems with a decidable reachability problem were considered, see e.g. [9] for an overview. In particular, reachability in linear hybrid systems, where the derivative of each variable in each state is constant, can be shown to be decidable for one continuous variable by using the techniques from [19]. In [6], it has been shown that reachability is decidable for timed automata, which are a particular subclass of hybrid automata where the slope of all variables is equal to 1.

In [22] we only studied the one-dimensional case of our model with the simplifying assumption that there is exactly one mode that can bring the temperature down and it is cost-free. In this paper, we drop this assumption and generalise the model to multiple dimensions. In the one-dimensional setting, we manage to prove similar nice algorithmic properties as in [22], i.e. the existence of finitely many patterns for optimal schedules, polynomial constant-factor approximation algorithm and an FPTAS. However, as opposed to the existence of a unique pattern for an optimal schedule in [22], we show that that there can be 44 different patterns when the simplifying assumption is dropped. To show this, we need to devise five safety-preserving and cost-non-increasing operations on schedules, while in [22] it sufficed for each mode to just lump together all timed actions that use this mode. Also, our constant-factor approximation algorithm requires a careful analysis of the interplay between different sections of the normal form for schedules, which results in an $\mathcal{O}(n^7)$ algorithm, while in [22] it sufficed to use one mode all the time and the algorithm ran in linear time.

Multi-mode systems were studied in [8], but with no discrete costs and with infinite time horizons only. They were later extended in [7] to a setting where the rate of change of each variable in a mode belonging to an interval instead of being constant. [28] studied a hybrid automaton model where the dynamics are governed by linear differential equations, but again without switching costs and only with an infinite time

horizon. Both of these papers show that, for any number of variables, a schedule with the optimal long-time average cost can be computed in polynomial time. In [24, 25], the same models without switching costs have been studied over the infinite time horizon, with the objective of minimising the peak cost, rather than the long-time average cost. In [11], long-time average and total cost games have been shown to be decidable for hybrid automata with strong resets, in which all variables are reset to 0 after each discrete transition. The long-time average and total cost optimisation for the weighted timed automata model have been shown to be PSPACE-complete (see e.g. [10] for an overview).

There are many practical approaches to the reduction of energy consumption and peak demand in buildings. One particularly popular one is model predictive control (MPC) [12]. In [26], stochastic MPC was used to minimise the energy consumption in a building. In [21], On-Off optimal control was considered for air conditioning and refrigeration. The drawback of using MPC is its high computational complexity and the fact that it cannot provide any worst-case guarantees. UPPAAL Stratego [15] supports the analysis of the expected cost in linear hybrid systems, but uses a stochastic semantics of these models [16, 14]. I.e. a control strategy induces a stochastic model where the time delay in each state is uniformly or exponentially distributed. This is different to the standard nondeterministic interpretation of the model, which we use in this paper. In [20], an on-line controller synthesis combined with machine learning and compositional synthesis techniques was applied for optimal control of a floor heating system.

Structure of the paper. The paper is organised as follows. We introduce all necessary notation and formally define the model in Section 2. In Section 3, we study the computational complexity of limit-safe and ϵ -safe control in multiple dimensions. In Section 4, we show that in one dimension every schedule can be transformed without increasing its cost into a schedule following one of 44 different patterns. In Section 5, we show that the cost optimisation decision problem in one-dimension with infinite and finite horizon is LOGSPACE and NP-complete, respectively. In Section 6, still for the one-dimension case, we first show a constant factor approximation algorithm and, building on it, develop an FPTAS by a reduction to the 0-1 knapsack problem. Due to the space constraints, some of the proofs and algorithms are only available in the extended version of this paper [23].

2 Preliminaries

Let $\mathbb{0}_N$ and $\mathbb{1}_N$ be N -dimensional vectors with all entries equal to 0 and 1, respectively. By $\mathbb{R}_{\geq 0}$ and $\mathbb{Q}_{\geq 0}$ we denote the sets of all non-negative real and rational numbers, respectively. We assume that $0 \cdot \infty = \infty \cdot 0 = 0$. For a vector v , let $\|v\|$ be its ∞ -norm (i.e. the maximum coordinate in v). We write $v_1 \leq v_2$ if every coordinate vector of vector v_1 is smaller than or equal to the corresponding coordinate in vector v_2 , and $v_1 < v_2$ if, additionally, $v_1 \neq v_2$ holds.

2.1 Formal Definition of Multi-Mode Systems

Motivated by our application of keeping temperature in multiple rooms within comfortable range, we restrict ourselves to safe sets being hyperrectangles, which can be specified by giving its two extreme corner points. A *multi-mode system with discrete costs*, \mathcal{A} , henceforth referred to simply as *multi-mode system*, is formally defined as a

tuple $\mathcal{A} = (M, N, A, \pi_c, \pi_d, V_{\min}, V_{\max}, V_0)$ where:

- M is a finite set of modes;
- $N \geq 1$ is the number of continuous variables in the system;
- $A : M \rightarrow \mathbb{Q}^N$ is the slope of all the variables in a given mode;
- $\pi_c : M \rightarrow \mathbb{Q}_{\geq 0}$ is the cost per time unit spent in a given mode;
- $\pi_d : M \rightarrow \mathbb{Q}_{\geq 0}$ is the cost of switching to a given mode;
- $V_{\min}, V_{\max} \in \mathbb{Q}^N : V_{\min} < V_{\max}$, define the safe set, S , as follows $\{x \in \mathbb{R}^N : V_{\min} \leq x \leq V_{\max}\}$;
- $V_0 \in \mathbb{Q}^N$, such that $V_0 \in S$, defines the initial value of all the variables.

2.2 Schedules, their cost and safety

A *timed action* is a pair $(m, t) \in M \times \mathbb{R}_{\geq 0}$ of a mode m and time delay $t > 0$. A *schedule* σ (of length k) with time horizon t_{\max} is a finite sequence of timed actions $\sigma = \langle (m_1, t_1), (m_2, t_2), \dots, (m_k, t_k) \rangle$, such that $\sum_{i=1}^k t_i = t_{\max}$. A *schedule* σ with infinite time horizon is either an infinite sequence of timed actions $\sigma = \langle (m_1, t_1), (m_2, t_2), \dots, (m_k, t_k), \dots \rangle$, such that $\sum_{i=1}^{\infty} t_i = \infty$ or a finite sequence of timed actions $\sigma = \langle (m_1, t_1), (m_2, t_2), \dots, (m_k, t_k) \rangle$, such that $t_k = \infty$. The *run* of a finite schedule $\sigma = \langle (m_1, t_1), (m_2, t_2), \dots, (m_k, t_k) \rangle$ is a sequence of *states* $\text{run}(\sigma) = \langle V_0, V_1, \dots, V_k \rangle$ such that, for all $0 \leq i \leq k-1$, we have that $V_{i+1} = V_i + t_i A(m_i)$.

A schedule and its run are called *safe* if $V_{\min} \leq V_i \leq V_{\max}$ holds for all $1 \leq i \leq k$. A schedule and its run are called ϵ -*safe* if $V_{\min} - \epsilon \cdot \mathbb{1}_N < V_i < V_{\max} + \epsilon \cdot \mathbb{1}_N$ holds for all $1 \leq i \leq k$. The run of an infinite schedule and its safety and ϵ -safety are defined accordingly.

The *total cost* of a schedule $\sigma = \langle (m_1, t_1), (m_2, t_2), \dots, (m_k, t_k) \rangle$ with a finite time horizon is defined as $\pi(\sigma) = \sum_{i=1}^k \pi_d(m_i) + \pi_c(m_i)t_i$. The *limit-average cost* for a finite schedule $\sigma = \langle (m_1, t_1), (m_2, t_2), \dots, (m_k, t_k) \rangle$ with an infinite time horizon is defined as $\pi_{avg}(\sigma) = \pi_c(m_k)$ and for an infinite schedule $\sigma = \langle (m_1, t_1), (m_2, t_2), \dots \rangle$ it is defined as

$$\pi_{avg}(\sigma) = \limsup_{k \rightarrow \infty} \left(\sum_{i=1}^k \pi_d(m_i) + \pi_c(m_i)t_i \right) / \sum_{i=1}^k t_i$$

A safe finite schedule σ is ϵ -*optimal* if, for all safe finite schedules σ' , we have that $\pi(\sigma') \geq \pi(\sigma) - \epsilon$. A safe finite schedule is *optimal* if it is 0-optimal. A safe infinite schedule σ is *optimal* if, for all safe infinite schedules σ' , we have that $\pi_{avg}(\sigma') \geq \pi_{avg}(\sigma)$.

The following example shows that there may not be an optimal schedule for a multi-mode system with a finite time horizon.

Example 1. Consider a multi-mode system with three modes: M_1, M_2, M_3 . The slope vectors in these modes are $A(M_1) = (1, 1)$, $A(M_2) = (1, -1)$ and $A(M_3) = (-1, 1)$, respectively. The continuous cost of using M_1 is $\pi_c(M_1) = 1$ and all the other costs are 0. Let $V_0 = V_{\min} = \mathbb{0}_2$ and $V_{\max} = \mathbb{1}_2$. Notice that we can only use M_2 or M_3 once we get out of the initial corner V_0 . This can only be done using M_1 . Now let the time

horizon be t_{\max} . Note that the following schedule $\sigma_\epsilon = (M_1, \epsilon), ((M_2, t), (M_3, t))^l$, where $t' = t_{\max} - \epsilon$, $l = \lceil t'/\epsilon \rceil$, and $t = t'/2l$, has time horizon t_{\max} and total cost $\epsilon > 0$. As ϵ can be made arbitrarily small but has to be > 0 , σ_ϵ is an ϵ -optimal schedule for all $\epsilon > 0$, but no optimal schedule exists.

Note that in Example 1, for any $\epsilon > 0$, there exists an optimal ϵ -safe schedule σ with total cost 0: $\sigma_0 = \langle ((M_2, t), (M_3, t))^l \rangle$ where l is defined as in Example 1. Our aim is to find an “abstract schedule” that, for any given $\epsilon > 0$, can be used to construct in polynomial time an ϵ -safe ϵ -optimal schedule.

Let $M^* = \{m \in M \mid \pi_d(m) = 0\}$ be the subset of modes without discrete costs. Note that, as shown in [8], the cost and safety of a schedule with M^* modes only, depends only on the total amount of time spent in each of the M^* modes. We therefore lump together any sequence of timed actions that only use M^* modes and define an *abstract timed action (over M^*)* as a function $\mathbf{t} : M^* \rightarrow \mathbb{R}_{\geq 0}$. A finite *abstract schedule* with time horizon t_{\max} (of length k) is a finite sequence $\tau = \langle \mathbf{t}_1, (m_1, t_1), \mathbf{t}_2, (m_2, t_2), \dots, (m_{k-1}, t_{k-1}), \mathbf{t}_k \rangle$ such that $\forall_i m_i \in M \setminus M^*$ and $\sum_{i \leq k, m \in M^*} \mathbf{t}_i(m) + \sum_{i < k} t_i = t_{\max}$. The run of the abstract schedule τ is a sequence $\langle V_0, V_1, \dots, V_{2k+1} \rangle$ such that, for all $i \leq k$, we have $V_{2i} = V_{2i-1} + A(m_i)t_i$ and $V_{2i+1} = V_{2i} + \sum_{m \in M^*} A(m)\mathbf{t}_i(m)$. We say that an abstract schedule is *limit-safe* if its run is safe. The total cost of an abstract schedule τ is defined as

$$\sum_{i \leq k, m \in M^*} \pi_c(m, \mathbf{t}_i(m)) + \sum_{i < k} (\pi_d(m_i) + \pi_c(m_i)t_i).$$

Note that any safe schedule can be turned into a limit-safe abstract schedule with the same cost by simply replacing any maximal subsequence of consecutive timed actions that only use M^* modes by a single abstract timed action. A limit-safe abstract schedule σ is optimal if the total cost of all other limit-safe abstract schedules is higher than $\pi(\sigma)$. The following statement justifies the name “limit-safe”.

Proposition 1. *Given a limit-safe abstract schedule τ and $\epsilon > 0$, we can construct in polynomial time an ϵ -safe schedule σ such that $\pi(\tau) = \pi(\sigma)$.*

Proof. Let $M^* = \{m_1, m_2, \dots, m_j\}$. To obtain σ from τ , we replace each abstract timed action $\{(m, t_m) \mid m \in M^*\}$ by a sequence $((m_1, t_{m_1}/l), \dots, (m_j, t_{m_j}/l))^l$ for a sufficiently large $l \in \mathbb{N}$.

Sufficiently large means that, for $t^* = \sum_{m \in M^*} t_m$, $l > t^* \cdot \max_{m \in M^*} \|A(m)\|/\epsilon$. This choice guarantees that $\sum_{m \in M^*} \|A(m)\| \cdot t_m/l < \epsilon$. Thus, when the abstract action $\{(m, t_m) \mid m \in M^*\}$ joins two states V_{2i}, V_{2i+1} along the run $\langle V_0, V_1, \dots, V_{2k+1} \rangle$ of τ , we know that this concrete schedule will cover the l -th part of V_{2i}, V_{2i+1} after every sequence $(m_1, t_{m_1}/l), (m_2, t_{m_2}/l), \dots, (m_j, t_{m_j}/l)$. As the safe set is convex, the start and end points of this sequence are safe points. Also, $\sum_{m \in M^*} \|A(m)\| \cdot t_m/l < \epsilon$ implies that the points in the middle are ϵ -safe. \square

Example 1 continues. *An example limit-safe abstract schedule of length 1 is $\tau = \{(m_1, t_{\max}/2), (m_2, t_{\max}/2)\}$. Based on τ we can construct an ϵ -safe schedule $\langle ((m_1, t_{\max}/2l), (m_2, t_{\max}/2l))^l \rangle$ where l is any integer greater than t_{\max}/ϵ .*

2.3 Structure of optimal schedules

We show here that it later suffices to consider only schedules with a particular structure.

Definition 1. We call a finite schedule σ angular if there are no two consecutive timed actions $(m_i, t_i), (m_{i+1}, t_{i+1})$ in σ such that $A(m_i) = A(m_{i+1})$.

We show that while looking for an (ϵ) -safe (ϵ) -optimal finite schedule, we can restrict our attention to angular schedules only.

Proposition 2. For every finite (ϵ) -safe schedule with time horizon t_{max} there exists an angular safe schedule with the same or lower cost.

Henceforth, we assume that all finite schedules are angular. Let $M^0 = \{m \mid A(m) = 0\}$, which we will also refer to as *zero-modes*.

Proposition 3. For every finite safe schedule with time horizon t_{max} there exists a safe schedule with the same or lower cost, in which at most one zero-mode is used at the very beginning.

Henceforth, we assume that all finite schedules use at most one zero-mode timed action and only at the very beginning.

2.4 Approximation algorithms

We study approximation algorithms for the total cost minimisation problem in multi-mode systems. We say that an algorithm is a *constant factor approximation algorithm* with a *relative performance* ρ iff, for all inputs x , the cost of the solution that it computes, $f(x)$, satisfies $OPT(x) \leq f(x) \leq (1 + \rho) \cdot OPT(x)$, where $OPT(x)$ is the optimal cost for the input x . We are particularly interested in polynomial-time approximation algorithms. A polynomial-time approximation scheme (PTAS) is an algorithm that, for every $\rho > 0$, runs in polynomial-time and has relative performance ρ . Note that the running time of a PTAS may depend in an arbitrary way on ρ . Therefore, we typically strive to find a fully polynomial-time approximation scheme (FPTAS), which is an algorithm that runs in polynomial-time in the size of the input and $1/\rho$.

The 0-1 Knapsack problem is a well-known NP-complete optimisations problem, which possess multiple FPTASes (see e.g. [18]). In this problem we are given a knapsack with a fixed volume and a list of items, each with an integer volume and value. The aim is to pick a subset of these items that together do not exceed the volume of the knapsack and have the maximum total value.

3 Complexity of limit-safe and ϵ -safe finite control

As our one-dimensional model strictly generalises the simple linear hybrid automata considered in [22], we immediately obtain the following result.

Theorem 1 (follows from [22], Theorem 3). Given (one-dimensional) multi-mode system \mathcal{A} , constants t_{max} and C (both in binary), checking whether there exists a safe schedule in \mathcal{A} with time horizon t_{max} and total cost at most C is NP-hard.

In the rest of this section we fix a (multi-dimensional) multi-mode system \mathcal{A} and time horizon t_{\max} .

Theorem 2. *If a limit-safe abstract schedule exists in \mathcal{A} , then there exists one of exponential length and it can be constructed in polynomial time.*

Proof (sketch). Before we formally prove this theorem, we need to introduce first a bit of terminology. We call a mode m *safe for time* $t > 0$ at $V \in S := \{x \in \mathbb{R}^N : V_{\min} \leq x \leq V_{\max}\}$ if $V + A(m)t \in S$. Also, m is *safe* at V if there exists $t > 0$ such that m is safe for time t at V . We say that a *coordinate of a state*, $V \in S$, is *at the border* if that coordinate in V is equal to the corresponding coordinate in V_{\min} or V_{\max} .

Our algorithm first removes from M all modes that will never be safe to use in a limit-safe schedule (and it can be found in the extended version of this paper [23]). This is an adaptation of [8, Theorem 7] where an algorithm was given for finding safe modes that can ever be used in a schedule with no time horizon. The main difference here is that the modes in M^* can always be used in a limit-safe abstract schedule even if they are not safe to use. We find here a sequence of sets of modes $M^* = M_0 \subset M_1 \subset M_2 \subset \dots$ such M_{i+1} is the set of modes that are safe at a state reachable from V_0 via a limit-safe abstract schedule that only uses modes from M_i . Note that at some step $k \leq |M|$ this sequence will stabilise, i.e. $M_k = M_{k+1}$. Similarly as in the proof of [8, Theorem 7], we can show that no mode from $M \setminus M_k$ can ever be used by a limit-safe abstract schedule. As a result, we can remove all these modes from M .

Next, we remove all modes that cannot be part of a limit-safe abstract schedule with time horizon t_{\max} . For this, for each m , we formulate a very similar linear programme (LP) as above (again, more details in [23]), where we ask for the time delay of m to be positive and the total time delay of all the modes to be t_{\max} . By a simple adaptation of the proof of [8, Theorem 4], if this LP is not satisfiable then m can be removed from \mathcal{A} .

Next, we look for the easiest possible target state V_{end} that can potentially be reached using a limit-safe abstract schedule from V_0 with time horizon t_{\max} . For this, V_{end} has to have the least number of coordinates at the border of the safe set. Note that this is well-defined, because if V and V' are two points reachable from V_0 via a limit-safe abstract schedules τ and τ' with time horizon t_{\max} , respectively, then $\tau/2$ (i.e. divide all abstract and timed actions delays in τ by 2) followed by $\tau'/2$, is also a limit-safe abstract schedule with time horizon t_{\max} , which reaches $(V + V')/2$. However, $(V + V')/2$ has a coordinate at the border iff both V and V' have it as well. This shows that there is a state with a minimum number of coordinates at the border.

To find the coordinates that need to be at the border we will use the following LP. We have a variable x_i for each dimension $i \leq N$ and a constraint that requires x_i to be less or equal to the i -th coordinate of $V_{\max} - V_{\text{end}}$ **and** $V_{\text{end}} - V_{\min}$. We also add that $\sum_{m \in M} t_m = t_{\max}$ and $V_{\text{end}} = V_0 + \sum_{m \in M} t_m \cdot A(m)$, with the objective *Maximise* $\sum_i x_i$. If the value of the objective is > 0 , we will get to know a new coordinate that does not have to be at the border. We then remove it from the LP and run it again. Once the objective is 0, then all the remaining coordinates, I , have to be at the border and the solution to this LP tells us, at which border the solution has to be located (it cannot possibly be at the border of both V_{\min} and V_{\max} as then we could reach the middle).

Next, in order to bound the length of a limit-safe abstract schedule by an exponential

in the size of the input, we not only need a state with the minimum number of coordinates at the border, but also sufficiently far way from the border. Otherwise, we may need super-exponentially many timed actions to reach it. In order to find such a point, we replace all x_i -s in the previously defined LP by a single variable x which is smaller or equal to all the coordinates of $V_{\max} - V_{\text{end}}$ and $V_{\text{end}} - V_{\min}$ from I . We then set the objective to *Maximise* x , which will give us a suitable easy target state V_{end} .

Now, consider \mathcal{A}' , which is the same as \mathcal{A} but with all slopes negated (i.e. $A'(m) = -A(m)$ for all $m \in M$). We claim that V_{end} is reachable from V_0 using a limit-safe abstract schedule τ iff $(V_0 + V_{\text{end}})/2$ is reachable from V_0 in \mathcal{A} with time horizon $t_{\max}/2$ and $(V_0 + V_{\text{end}})/2$ is reachable from V_{end} in \mathcal{A}' with time horizon $t_{\max}/2$; this again follows by considering $\tau/2$. Note that a coordinate of $(V_0 + V_{\text{end}})/2$ is at the border iff it is at the border in both V_0 and V_{end} .

This way we reduced our problem to just checking whether a limit-safe abstract schedule exists from one point to another more permissive point (i.e. where the set of safe modes is at least as big) within a given time horizon. The algorithm that solves this problem is provided in the extended version. It again reuses the same constructions as above, e.g. by constructing a sequence of sets of modes $M^* = M_0 \subset M_1 \subset \dots \subset M_k$, and its correctness follows similarly as before. We now need to invoke this algorithm twice: to check that $(V_0 + V_{\text{end}})/2$ is reachable from V_0 with time horizon $t_{\max}/2$ and that $(V_0 + V_{\text{end}})/2$ is reachable from V_{end} with time horizon $t_{\max}/2$ in \mathcal{A}' . If at least one of these calls return NO, then no limit-safe abstract schedule from V_0 to V_{end} can exist. Otherwise, let σ and σ' be the schedules returned by these two calls, respectively. Then the concatenation of σ with the reverse of σ' is a limit-safe abstract schedule that reaches V_{end} from V_0 with time horizon t_{\max} . \square

Theorem 3. *Finding an optimal limit-safe abstract schedule in \mathcal{A} can be done in non-deterministic exponential time.*

Proof. The limit-safe abstract schedule constructed in Theorem 2 has an exponential length. To establish a nondeterministic exponential upper bound, we can guess the modes (and the order in which they occur). With them, we can produce an exponentially sized linear program, which encodes that the run of the abstract schedule is safe and minimises the total cost incurred. \square

Theorem 3 and Proposition 1 immediately give us the following.

Corollary 1. *If a limit-safe abstract schedule exists in \mathcal{A} , then for any $\epsilon > 0$ an ϵ -safe schedule with the same cost can be found in nondeterministic exponential time.*

Moreover, from Theorem 2 and the fact that in the case of multi-mode systems with no discrete costs all abstract schedules have length 1, we get the following.

Corollary 2. *Finding an optimal limit-safe abstract schedule for multi-mode systems with no discrete costs can be done in polynomial time.*

We can reduce the computational complexity in the general model if we are willing to sacrifice optimality for ϵ -optimality.

Theorem 4. *If a limit-safe abstract schedule exists, then finding an ϵ -safe ϵ -optimal strategy can be done in deterministic polynomial space.*

Proof. When reconsidering the linear programme from the end of the proof of Theorem 3, we can guess the intermediate states in polynomial space (and thus guess and output the schedule) as long as all states along the run (including the time passed so far) are representable in polynomial space.

Otherwise we use the opportunity to deviate by up to ϵ from the safe set by increasing or decreasing the duration of each timed action up to some $\delta > 0$, in order to keep the intermediate values representable in space polynomial in $|\mathcal{A}|$ and ϵ . However, we apply these changes in a way that the overall time remains t_{\max} . Clearly this is possible, because within $\delta/2$ of the actual time point of each state along the run, there is a value whose number of digits in the standard decimal notation is at most equal to the sum of the number of digits in $\delta/2$ and t_{\max} . Picking any such point for every interval would induce a schedule with the required property and they can be simply guessed one by one.

The final imprecision introduced by this operation is at most $b \cdot \delta \cdot \max_{m \in M} |A(m)|$, where b is a bound on the number of timed actions in a limit-safe schedule, which is exponential in $|\mathcal{A}|$. If we choose $\delta = \epsilon / (b \cdot \max_{m \in M} |A(m)|)$, then we will get the required precision.

Although our algorithm is nondeterministic, due to Savitch's theorem, it can be implemented in deterministic polynomial space. \square

4 Structure of Finite Control in One-dimension

We show in this section that any finite safe schedule in one-dimension can be transformed without increasing its cost into a safe schedule, which follows one of finitely many regular patterns. The crucial component of this normal form will be a “leap” that we define below. We first introduce some notation. Let $M^+ = \{m \mid A(m) > 0\}$ and $M^- = \{m \mid A(m) < 0\}$. Recall that $M^0 = \{m \mid A(m) = 0\}$. We will call a mode, m , an *up mode*, *down mode*, or *zero-mode* if $m \in M^+$, $m \in M^-$, or $m \in M^0$, respectively. Similarly, the *trend* of a timed action (m, t) is *up*, *down*, *flat* if m is an up, down, zero-mode, respectively. For any subsequence of timed actions $\sigma' = \langle (m_i, t_i), \dots, (m_j, t_j) \rangle$ in a schedule σ , whose run is $\text{run}(\sigma) = \langle V_0, V_1, \dots, V_k \rangle$, we say that σ' *starts at state* v and *ends at state* v' iff $v = V_{i-1}$ and $v' = V_j$. We use the same terminology for a single timed action (in this case this subsequence has length 1).

Definition 2. *A partial leap is a pair of consecutive timed actions $(m_i, t_i), (m_{i+1}, t_{i+1})$ in a safe schedule such that $m_i \in M^+$, $m_{i+1} \in M^-$, and $A(m_i)t_i + A(m_{i+1})t_{i+1} = 0$, i.e. the state of a multi-mode system does not change after any partial leap. A partial leap is complete if $A(m_i)t_i = V_{\max} - V_{\min}$. We will simply refer to complete leaps as leaps.*

There are $|M^+ \times M^-|$ types of leaps. A leap is of type $(m, m') \in M^+ \times M^-$ iff $m_i = m$ and $m_{i+1} = m'$. Let Δt_m and $\Delta \pi_m$ denote the time and cost it takes for an up mode m to get from V_{\min} to V_{\max} or a down mode m to get from V_{\max} to V_{\min} . Note that $\Delta t_m = |(V_{\max} - V_{\min})/A(m)|$ and $\Delta \pi_m = \pi_d(m) + \pi_c(m) \cdot \Delta t_m$. By $\Delta t_{m, m'}$ and $\Delta \pi_{m, m'}$ we denote the time duration and the cost of a leap of type $(m, m') \in M^+ \times M^-$, respectively. Note that $\Delta t_{m, m'} = \Delta t_m + \Delta t_{m'}$ and $\Delta \pi_{m, m'} = \Delta \pi_m + \Delta \pi_{m'}$.

Any safe schedule σ can be decomposed into three sections that we will call its *head*, *leaps*, and *tail*. The *head section* ends after the first timed action that ends at V_{\min} . The *leaps section* contains only leaps of possibly different types following the head section. Finally, the *tail section* starts after the last leap in the leaps section has finished. Note that any of these sections can be empty and the tail section can in principle contain further leaps. We show here that, for any safe schedule of length at least three, there exists another safe one with the same or a smaller cost, whose head and tail sections follow one of the 10 patterns presented in Figure 3 and Figure 4, respectively, where *partial up/down* means that the next state is not at the border. For each of these patterns, there exists an example which shows that an optimal safe schedule may need to use such a pattern and hence it is necessary to consider it. In order to prove this, we first need to define several cost-nonincreasing and safety-preserving operations that can be applied to safe schedules. These will later be applied in Theorem 5 to transform any safe schedule into one of the just mentioned regular patterns. These operations are easy to explain via a picture, but cumbersome to define formally. Therefore, the formal definitions can be found in the extended version of this paper [23] and we present here only the intuition behind them.

Let σ be any safe finite schedule. Following Proposition 2 and 3, we can assume that σ is angular and only contains at most one timed action with a zero-mode, and if it contains one, this action occurs at the very beginning. Unless explicitly stated, the operations below are defined for timed actions with up or down trend only.




Fig. 1: On the left, the rearrange operation applied to three timed actions 1-2-3 with modes m_1, m_2, m_3 results in 1'-2'-3' with modes m_2, m_3, m_1 . On the right, the shift operation is being applied to a partial leap 1-2-3 which will be moved after the (complete) leap 3-4-5.

The first operation that we need is the *rearrange* operation, which simply changes the order of any subsequence of timed actions with the same trend. The next one is the *shift* operation. It cuts any subsequence of timed actions that start and end at the same state, V , and pastes this subsequence after any timed action that ends at V . The effect of these two operations can be seen in Figure 1.

Next is the *shift-down* operation. We can see an example of applying this operation in Figure 2. Intuitively, it can rearrange any subsequence of timed actions that start and end at the same state and move them after any timed action that ends at V_{\min} . The most complicated operation we define is the *wedge* operation. It acts on three consecutive timed actions in a safe schedule and simultaneously shrinks the middle action while extending the other two, or stretches the middle action while shrinking the other two. We can see its behaviour in Figure 2. Intuitively, it moves the timed action m_2 parallelly up




Fig. 2: On the left, an example of applying the shift-down operation to timed actions m_{i+1}, m_{i+2} . These actions are rearranged to move after point 5, which becomes point 3' (i.e. following timed action m_{i+3}). On the right, an example of applying the wedge operation to three timed actions m_1, m_2, m_3 . This operation is a (parallel) translation of the action m_2 , which changes the time duration of each of these actions. After this operation either the m_2 line touches V_{\min} , which would remove m_1 from the schedule, or the m_2 line touches V_{\max} , which would change a state along the run of the schedule to be at the border.

or down, until either the timed action m_1 is removed or m_2 ends at V_{\max} . The direction depends on the cost gradient, but as the cost delta function of this operation is linear, one of these directions is cost-nonincreasing.

Finally, we define the *resize* operation that will be used the most in our procedure. The resize operation requires one parameter $t \in \mathbb{R}$ and can act on any two consecutive timed actions in a safe schedule. Intuitively, if $t < 0$, this operation decreases the total time of this pair of timed actions by $|t|$ while changing only the middle state between these two timed actions along the run of the schedule. If $t > 0$, this operation increases the duration of this pair of timed actions by t while again changing only the state between them along the run. If $t > 0$ then we will also refer to this operation as the *stretch* operation and if $t < 0$ as the *shrink* operation with parameter $-t > 0$. If the stretch and shrink operations are simultaneously applied with the same parameter t to two non-overlapping pairs of timed actions, the result is a safe schedule with the same time horizon as before, but with a possibly different total cost. We will call a *flexi* any subsequence of length 2 in a safe schedule such that both shrink and stretch operations can be applied to it for some $t > 0$ without compromising its safety. A simultaneous application of these two operations to flexis is demonstrated in Figure 5 and 6.

Consider two non-overlapping flexis at positions i and j in a safe schedule σ . Let $\sigma' = \text{resize}(\sigma, i, t)$ be the resulting schedule of applying the resize operation with parameter t to the i -th and $i + 1$ -th timed actions in σ and $\text{resize-domain}(\sigma, i)$ be the maximal closed interval from which t can be picked to ensure that σ' is safe. Similarly, let $\sigma'' = \text{resize}(\sigma, j, -t)$ and $\sigma''' = \text{resize}(\text{resize}(\sigma, i, t), j, -t)$. Note that σ''' has the same time horizon as σ and it is safe as long as $t \in \text{resize-domain}(\sigma, i) \cap \text{resize-domain}(\sigma, j)$ and let us denote this closed interval by I . Furthermore, $\pi(\sigma''') - \pi(\sigma) = \pi(\sigma') - \pi(\sigma) + \pi(\sigma'') - \pi(\sigma)$ because the two flexis did not overlap. As it is shown in the extended version, both $\pi(\sigma') - \pi(\sigma)$ and $\pi(\sigma'') - \pi(\sigma)$ are linear functions

in t in the interior of I . As a result, $\pi(\sigma''') - \pi(\sigma)$ is also a linear function in t and so its minimum value is achieved at one of the endpoints of I . Also, at such an endpoint, one of the time actions in these two flexis will disappear and as a result the total cost would be reduced even further. It follows, that there is an endpoint of I such that selecting it as t will not increase the cost of the schedule, but it will remove a flexi from σ . As the zero-mode timed action and the last timed action in a schedule can have flexible time delay, we can also define the resize operation for them in a similar way. As a result, we can apply the resize operation with parameter t to any of these (including a flexi) and with parameter $-t$ to the other. Reasoning as above, there is a value for t such that the cost of the resulting schedule does not increase, the schedule remains safe, and at least one of the timed actions is removed from σ or one more state along the run of σ becomes V_{\min} or V_{\max} .




Fig. 3: Ten possible head patterns: (a) flat+down (b) down (c) partial-up+down (d) flat+up+down (e) up+down (f) partial-down+up+down (g) partial-up+up+down (h) partial-down+down (i) up+partial-down+down and (j) empty (not depicted).

Theorem 5. For every safe schedule σ in a one-dimensional multi-mode system there exists a safe schedule σ' whose head section matches one of the patterns in Figure 3, tail section matches one of the patterns in Figure 4, and $\pi(\sigma') \leq \pi(\sigma)$ holds. Furthermore, it suffices to consider only 44 combinations of these head and tail patterns, and the length of all of them is at most five.

Proof. We will repeatedly apply combination of shrink and stretch operations to flexis until we remove all non-overlapping ones. Note that after each such an application either a timed action is removed or one more state along the run of σ becomes equal to V_{\max} or V_{\min} . We claim that the following steps will transform σ into a suitable σ' :




Fig. 4: Ten possible tail patterns: (a) partial-up (b) partial-up+up (c) up+partial-down+down (d) up+partial-down (e) up (f) partial-up+down (g) partial-up+up+down (h) partial-up+down+up (i) up+partial-down+down+up and (j) empty (not depicted).

1. as long as there are at least one pair of non-overlapping flexis then shrink one and stretch the other until a timed action is removed or a new state at the border is created;
2. once there is only one flexi left or two overlapping ones, use the shift or shift-down operation to move them to the end of the schedule;
3. if the first timed action is flat, pair it with the remaining flexi to remove one of them using the shrink-stretch operation combination;
4. if the last state of $run(\sigma)$ is not at the border and a flexi or flat timed action remains after the previous step, they should be paired with each other for the shrink-stretch operation combination;
5. if two overlapping flexis exist, use the wedge operation to resolve them;
6. finally, if the tail section still does not follow any of the patterns, apply the shift-down operation to the (unique) segment that starts and ends at V_{\max} .

A graphical representation of this procedure when applied to an example schedule is provided in the extended version. It is easy to see that the first step of this procedure will stop eventually because σ has a finite number of timed actions and states along its run. The rest of the steps of this procedure just try to reduce the number of possibilities for the head and tail sections. Note that, apart from the initial state, there can be only one state, along the run of the resulting σ' , which is not at the border. This is because otherwise a shrink-stretch or wedge operation could still be applied. Drawing all possible patterns with one point not at the border and eliminating the ones that are inter-reducible

using one of these operations, results in Figure 3 for the head section and Figure 4 for the tail section.

If we try to combine all these head and tail pattern together then this would result in $10 \cdot 10 = 100$ possible combinations. However, as just mentioned, there can be only one point not at the border or a zero-mode timed action in a schedule so these combinations of head and tail patterns can be reduced further. In particular, any head pattern can be combined with tail patterns (e) and (j), but only (b), (e), (j) head patterns can be combined with the remaining tail ones. Therefore, there are $10 \cdot 2 + 3 \cdot 8 = 44$ combined patterns and it is easy to check that none of them has length larger than five (this is important for the computational complexity stated in Theorem 8). \square




Fig. 5: Shrink and stretch operations being applied to two up-up flexis. The 1-2-3 one is stretched by t , which results in 1-4-5, and 1'-2'-3' is shrunk by t , which results in 4'-5'-3'. Note that 3 and 5 (also, 1' and 4') are the same states but shifted in time. In fact, all states along the run of the schedule stay the same apart from 2 and 2', and as a result the schedule stays safe.




Fig. 6: Shrink and stretch operations being applied to two up-down flexis.

5 Complexity of Optimal Control in One-dimension

We start with considering the easy case of infinite time horizons, before turning to the interesting case of finite time horizons.

5.1 Infinite time horizon

First let us consider the case $M^0 = \emptyset$. If also $M^+ \times M^- = \emptyset$ then there are no safe schedules with infinite horizon at all. Otherwise, let $(i', j') = \operatorname{argmin}_{(i,j) \in M^+ \times M^-} \Delta\pi_{i,j} / \Delta t_{i,j}$.

Let us pick any mode $m^- \in M^-$ and denote $t^- := (V_{\min} - V_0)/A(m^-)$. Consider the infinite schedule σ , which starts with the timed action (m^-, t^-) followed by infinitely many complete leaps of type (i', j') . Obviously, at all times $t = t^- + k \cdot \Delta t_{i', j'}$ where $k \in \mathbb{N}$, σ is more expensive by at most $\pi_d(m^-) + \pi_c(m^-)t^-$ from the cheapest schedule with time horizon t . Consequently, as $k \rightarrow \infty$, this shows that the limit superior of the average cost cannot be smaller than $\Delta\pi_{i', j'}/\Delta t_{i', j'}$. At the same time, σ realises this long-time average.

If $M^0 \neq \emptyset$, then let $m' = \min_{m \in M^0} \pi_c(m)$ be the zero-mode with the lowest continuous cost to run. We claim that if $\pi_c(m') < \Delta\pi_{i', j'}/\Delta t_{i', j'}$ or $M^+ \times M^- = \emptyset$ then an optimal safe schedule is simply (m', ∞) , whose limit-average cost is $\pi_c(m')$, and otherwise σ defined above is an optimal safe schedule. This is because, if $\pi_c(m') < \Delta\pi_{i', j'}/\Delta t_{i', j'}$, then, at any time point of σ where a leap of some type (i, j) is used, removing this leap and increasing the time m' is used for by $\Delta t_{i, j}$ reduces the total cost up to this time point.

Taking into account that $\operatorname{argmin}_{(i, j) \in M^+ \times M^-} \Delta\pi_{i, j}/\Delta t_{i, j}$ can be computed using logarithmic space (because multiplication, division and comparison can be [13]) we get the following theorem.

Theorem 6. *An optimal safe infinite schedule for one-dimensional multi-mode systems can be computed in deterministic LOGSPACE.*

5.2 Finite Time Horizon

Due to Theorem 1, we already know that the decision problem for optimal schedules in one-dimensional multi-mode systems is at least NP-hard. Here, we show that the problem is NP-complete by showing that an optimal schedule exists and that each section of an optimal schedule can be guessed.

Note that the existence of an optimal schedule for the one-dimensional case sets it apart from the general case. In Example 1, we have shown that optimal schedules are not even guaranteed to exist for two-dimensional multi-mode systems.

Theorem 7. *For any one-dimensional multi-mode systems \mathcal{A} and $t_{\max} \geq 0$, there exists an optimal schedule with time horizon t_{\max} , and checking for the existence of an optimal schedule with cost $\leq C$ is NP-complete. (When t_{\max} and C are given in binary.)*

Proof. First, we can simply iterate over all schedules of length one and directly calculate their costs. Next, we can iterate over pairs of modes, m_1 and m_2 , and for each of them solve a linear program (LP) which will give us the cheapest schedule of length two using these two modes. This LP finds the cheapest partition of t_{\max} between the two modes and has the following form: Minimise $\pi_c(m_1)t_1 + \pi_c(m_2)(t_{\max} - t_1) + \pi_d(m_1) + \pi_d(m_2)$

$$\text{Subject to: } 0 \leq t_1 \leq t_{\max}, \quad V_{\min} \leq V_0 + A(m_1)t_1 \leq V_{\max} \quad \text{and}$$

$$V_{\min} \leq V_0 + A(m_1)t_1 + A(m_2)(t_{\max} - t_1) \leq V_{\max}.$$

This can be done in $\mathcal{O}(|\mathcal{A}|^2)$ time.

Now, for schedules of length at least three, we showed in Section 4 that any such a schedule can be transformed without increasing its cost into one that can be split into three sections: the head section, the leaps section, and the tail section (some of which

may be empty). Due to Theorem 5, there are 44 combined patterns for the tail and head sections. Note that, when considering only the cost of the whole schedule, it suffices for us to know the number of leaps of each type in the leaps section and not their precise order. Notice that a schedule with time horizon t_{\max} can contain at most $\lfloor t_{\max}/\Delta\pi_{i,j} \rfloor$ leaps of type (i, j) . The size of this number is polynomial in the size of the input \mathcal{A} . There are $\mathcal{O}(|M|^2)$ types of leaps so the number of leaps of each type and the combined pattern of the schedule can be guessed non-deterministically with polynomially many bits. This guess uniquely determines the cost of the schedule. This is because, after the total time of the leaps section is deducted from t_{\max} , we get the exact time the head and tail section have to last for. Each combined pattern has at most one of the following: a flexi, a zero-mode, or the last state not at the border. The time remaining will determinate exactly (if at all possible) the value of this single flexible point along this schedule. Now, computing the cost of the resulting schedule and checking whether it is lower than C can be done in polynomial time. This shows that the problem is in NP. It also shows that optimal schedules exist, because there are only finitely many options to choose from. \square

6 Approximate Optimal Control in One-Dimension

We first show an approximation algorithm with a 3-relative performance for the cost minimisation problem in one-dimensional multi-mode systems, which runs in $\mathcal{O}(|\mathcal{A}|^7)$ time. Our algorithm tries all possible patterns for an optimal schedule and for the leaps section always picks leaps of the same type. It then adds, if necessary or for cost efficiency, a partial leap to the leaps section and minimises the total cost of the just constructed schedule by optimising the time duration of this partial leap. This constant approximation algorithm is crucial for showing the existence of an FPTAS for the same problem in the next subsection.

Theorem 8. *Computing a safe schedule with total cost at most three times larger than the optimal one for one-dimensional multi-mode system \mathcal{A} can be done in $\mathcal{O}(|\mathcal{A}|^7)$ time.*

We now show that the cost minimisation problem for one dimensional multi-mode systems is in FPTAS by a polynomial time reduction to the 0-1 Knapsack problem, for which many FPTAS algorithms are available (see e.g. [18]). This is similar to the FPTAS construction in [22], but differs in how the modes with fractional duration are handled. First we iterate over all possible schedules of length at most two and find the cheapest one in polynomial time. Next, thanks to Theorem 5, all optimal schedules longer than two can be transformed into one of 44 different patterns. Each of these patterns results in a slightly different FPTAS formulation. An FPTAS for the general model consists of all of these individual FPTASes executed one after another. The details of the proof are provided in the extended version.

Theorem 9. *Solving the optimal control problem for multi-mode systems with relative performance ρ takes $\mathcal{O}(\text{poly}(1/\rho)\text{poly}(\text{size of the instance}))$ time and is therefore in FPTAS.*

Acknowledgement

This work was supported by EPSRC EP/M027287/1 grant “Energy Efficient Control”.

References

1. EnergyPlus: Building energy simulation program. <https://energyplus.net/>.
2. IBPT: International Building Physics Toolbox in Simulink. <http://www.ibpt.org/>.
3. TRaNsient SYstems Simulation Program. <http://sel.me.wisc.edu/trnsys/>.
4. Rajeev Alur, Costas Courcoubetis, Nicolas Halbwachs, Thomas A. Henzinger, P.-H. Ho, Xavier Nicollin, Alfredo Olivero, Joseph Sifakis, and Sergio Yovine. The algorithmic analysis of hybrid systems. *Theoretical computer science*, 138(1):3–34, 1995.
5. Rajeev Alur, Costas Courcoubetis, Thomas Henzinger, and Pei Ho. Hybrid automata: An algorithmic approach to the specification and verification of hybrid systems. In *Hybrid Systems*, pages 209–229. Springer Berlin / Heidelberg, 1993.
6. Rajeev Alur and David L. Dill. A theory of timed automata. *Theoretical Computer Science*, 126(2):183–235, April 1994.
7. Rajeev Alur, Vojtech Forejt, Salar Moarref, and Ashutosh Trivedi. Safe schedulability of bounded-rate multi-mode systems. In *HSCC*, pages 243–252. ACM, 2013.
8. Rajeev Alur, Ashutosh Trivedi, and Dominik Wojtczak. Optimal Scheduling for Constant-Rate Multi-Mode Systems. In *Proc. of Hybrid Systems: Computation and Control 2012*.
9. Eugene Asarin, Venkatesh P. Mysore, Amir Pnueli, and Gerardo Schneider. Low dimensional hybrid systems – decidable, undecidable, don’t know. *Information and Computation*, 211:138–159, February 2012.
10. Patricia Bouyer. Weighted Timed Automata: Model-Checking and Games. *Electronic Notes in Theoretical Computer Science*, 158:3–17, May 2006.
11. Patricia Bouyer, Thomas Brihaye, Marcin Jurdziński, Ranko Lazić, and Michał Rutkowski. Average-Price and Reachability-Price Games on Hybrid Automata with Strong Resets. In Franck Cassez and Claude Jard, editors, *Formal Modeling and Analysis of Timed Systems*, number 5215 in Lecture Notes in Computer Science, pages 63–77. Springer Berlin Heidelberg, September 2008.
12. Eduardo F Camacho and Carlos Bordons Alba. *Model predictive control*. Springer Science & Business Media, 2013.
13. Andrew Chiu, George I. Davida, and Bruce E. Litow. Division in logspace-uniform NC^1 . *ITA*, 35(3):259–275, 2001.
14. Alexandre David, Dehui Du, Kim Guldstrand Larsen, Axel Legay, and Marius Mikučionis. Optimizing Control Strategy Using Statistical Model Checking. In Guillaume Brat, Neha Rungta, and Arnaud Venet, editors, *NASA Formal Methods*, number 7871 in Lecture Notes in Computer Science, pages 352–367. Springer Berlin Heidelberg, May 2013.
15. Alexandre David, Peter Gjøøl Jensen, Kim Guldstrand Larsen, Marius Mikučionis, and Jakob Haahr Taankvist. Uppaal Stratego. In Christel Baier and Cesare Tinelli, editors, *Tools and Algorithms for the Construction and Analysis of Systems*, number 9035 in Lecture Notes in Computer Science, pages 206–211. Springer Berlin Heidelberg, April 2015.
16. Alexandre David, Kim G. Larsen, Axel Legay, Marius Mikučionis, and Zheng Wang. Time for Statistical Model Checking of Real-Time Systems. In Ganesh Gopalakrishnan and Shaz Qadeer, editors, *Computer Aided Verification*, number 6806 in Lecture Notes in Computer Science, pages 349–355. Springer Berlin Heidelberg, July 2011.
17. Thomas A. Henzinger. The theory of hybrid automata. In *Proc. of the 11th IEEE LICS’96 Symposium*, Washington, DC, USA, 1996.
18. Hans Kellerer, Ulrich Pferschy, and David Pisinger. *Knapsack Problems*. Springer Berlin Heidelberg, Berlin, Heidelberg, 2004.
19. François Laroussinie, Nicolas Markey, and Ph Schnoebelen. Model checking timed automata with one or two clocks. In *CONCUR 2004-Concurrency Theory*, pages 387–401. Springer, 2004.

20. Kim G Larsen, Marius Mikučionis, Marco Muniz, Jiří Srba, and Jakob Haahr Taankvist. Online and compositional learning of controllers with application to floor heating. In *International Conference on Tools and Algorithms for the Construction and Analysis of Systems*, pages 244–259. Springer, 2016.
21. Bin Li and Andrew G. Alleyne. Optimal on-off control of an air conditioning and refrigeration system. In *American Control Conference (ACC), 2010*, pages 5892–5897. IEEE, 2010.
22. Mahmoud A. A. Mousa, Sven Schewe, and Dominik Wojtczak. Optimal control for simple linear hybrid systems. In *23rd International Symposium on Temporal Representation and Reasoning (TIME)*, pages 12–20. IEEE, 2016.
23. Mahmoud A. A. Mousa, Sven Schewe, and Dominik Wojtczak. Optimal Control for Multi-Mode Systems with Discrete Costs. *CoRR*, ??????????????????, 2017.
24. Truong X. Nghiem, Madhur Behl, Rahul Mangharam, and George J. Pappas. Green scheduling of control systems for peak demand reduction. In *Decision and Control and European Control Conference (CDC-ECC), 2011 50th IEEE Conference on*, pages 5131–5136. IEEE, 2011.
25. Truong X. Nghiem, George J. Pappas, and Rahul Mangharam. Event-based green scheduling of radiant systems in buildings. In *American Control Conference (ACC), 2013*, pages 455–460. IEEE, 2013.
26. Frauke Oldewurtel, Andreas Ulbig, Alessandra Parisio, Göran Andersson, and Manfred Morari. Reducing peak electricity demand in building climate control using real-time pricing and model predictive control. In *Decision and Control (CDC), 2010 49th IEEE Conference on*, pages 1927–1932. IEEE, 2010.
27. Luis Pérez-Lombard, José Ortiz, and Christine Pout. A review on buildings energy consumption information. *Energy and buildings*, 40(3):394–398, 2008.
28. Dominik Wojtczak. Optimal Control for Linear-Rate Multi-mode Systems. In Víctor Braberman and Laurent Fribourg, editors, *Formal Modeling and Analysis of Timed Systems*, number 8053 in Lecture Notes in Computer Science, pages 258–273. Springer Berlin Heidelberg, August 2013.