Optimal Control for Multi-Mode Systems with Discrete Costs

Dominik Wojtczak

University of Liverpool

AVeRTS workshop
14th July 2017

(based on joint work with Mahmoud A. Mousa and Sven Schewe)

Hybrid Automata

They consist of

- A finite state transition system
- Differential equations in each control state

Example

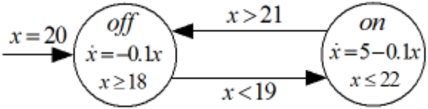

Linear Hybrid Automata

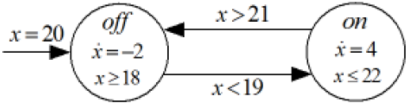

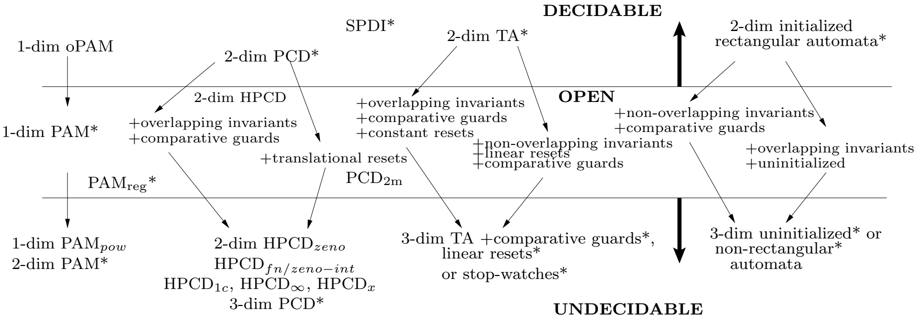
DECIDABLE
OPEN
UNDECIDABLE
1-dim oPAM
1-dim PAM*
PAM_{reg}*
1-dim PAM_{pow}
2-dim PAM*
SPDI*
2-dim PCD*
2-dim HPCD
+overlapping invariants
+comparative guards
+translational resets
PCD_{2m}
2-dim $HPCD_{zeno}$
$HPCD_{fn/zeno-int}$
$HPCD_{1c}$, $HPCD_{\infty}$, $HPCD_x$
3-dim PCD*
2-dim TA*
+overlapping invariants
+comparative guards
+constant resets
+non-overlapping invariants
+linear resets
+comparative guards
3-dim TA +comparative guards*,
linear resets*
or stop-watches*
2-dim initialized
rectangular automata*
+non-overlapping invariants
+comparative guards
+overlapping invariants
+uninitialized
3-dim uninitialized* or
non-rectangular*
automata

Multi-mode Systems: Safe Schedulability

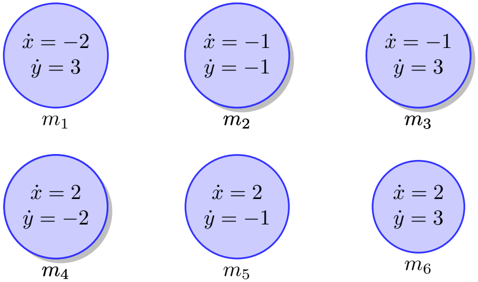

Safe set: $x \in [65, 70], y \in [65, 70]$

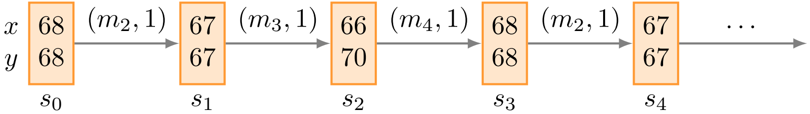

Keywords: modes, schedule, run, and safe schedule

Multi-mode System: Zeno schedule

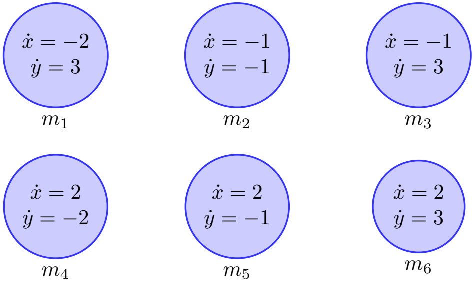

Safe set: $x \in [65, 70], y \in [65, 70]$

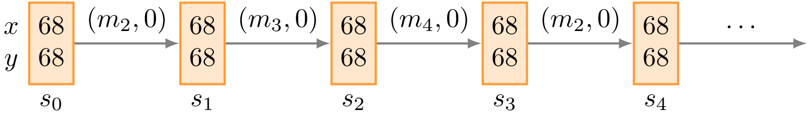

Keywords: Zeno Schedule

Multi-mode Systems: Zeno schedule

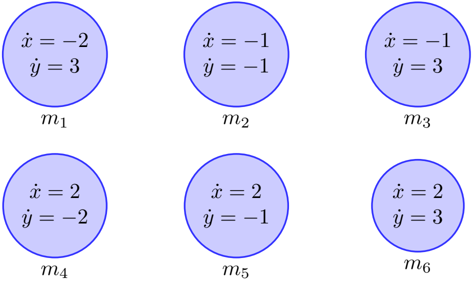

Safe set: $x \in [65, 70], y \in [65, 70]$

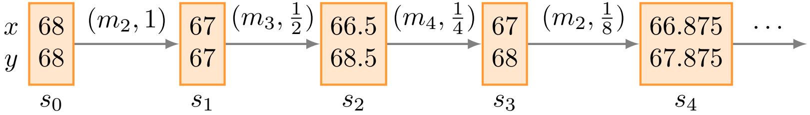

Keyword: Zeno Schedule

Formal Definition

Definition (Constant-Rate Multi-Mode Systems: MMS)

A 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,
- $R : M \to \mathbb{R}^N$ gives for each mode the rate vector,
- $S \subseteq \mathbb{R}^N$ is a bounded convex set of safe states.

- The run of a schedule $\langle (m_1, t_1), (m_2, t_2), \ldots, (m_k, t_k) \rangle$ from s_0 is

  $$s_0, (m_1, t_1), s_1, \ldots, (m_k, t_k), s_k$$

  such that $s_i = s_{i-1} + t_i \cdot R(m_i)$ for all for all $1 \leq i \leq k$.
- A schedule σ is safe at s_0 if all states of the run of σ from s_0 are safe (i.e., $\in S$).

Problems Studied and Results

Safe Schedulability Problem

Given an MMS $\mathcal{H}$ and a starting state s_0 decide whether there exists a non-Zeno safe schedule.

Theorem (Alur, Trivedi, W. (HSCC 2012))

Safe Schedulability can be solved in polynomial time for polytopes.

Safe Reachability Problem

Given an MMS $\mathcal{H}$, a starting state $s_0 \in S$, and a target state $s_t \in S$, decide whether there exists a safe schedule that reaches s_t from s_0.

Theorem (Alur, Trivedi, W. (HSCC 2012))

Safe Reachability can be solved in polynomial time if the starting and the target states lie in the interior of the polytope S.

Safe Schedulability Problem: Geometry

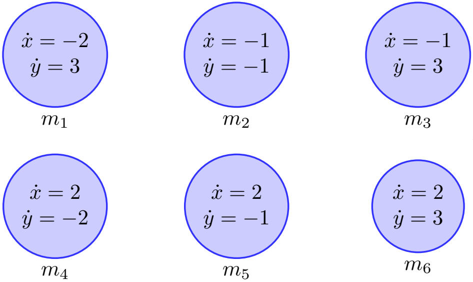

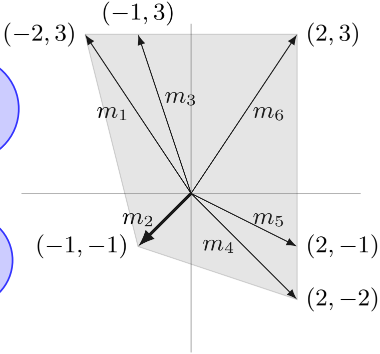

Safe set: $x \in [65, 70], y \in [65, 70]$

Safe Schedulability Problem: Geometry

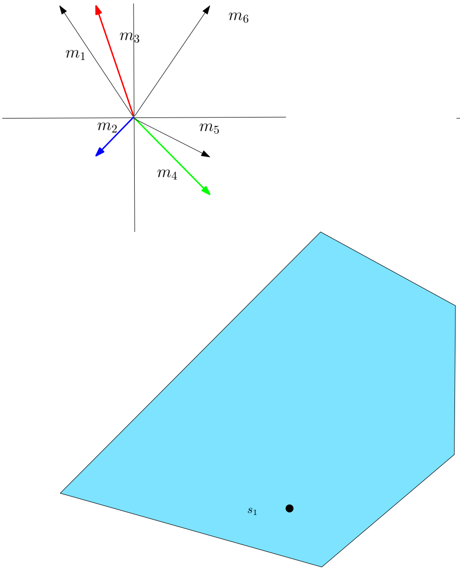

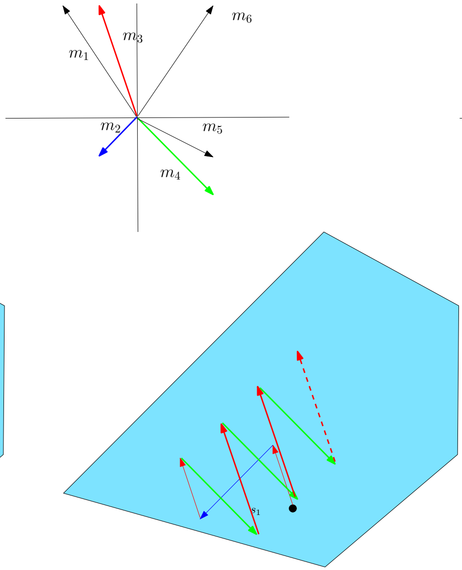

Reachability Problem: Geometry

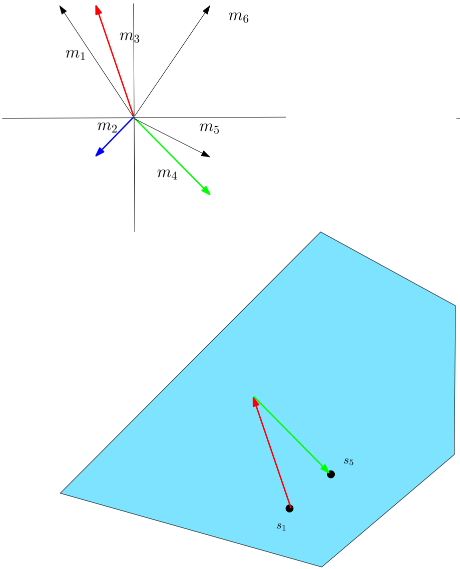

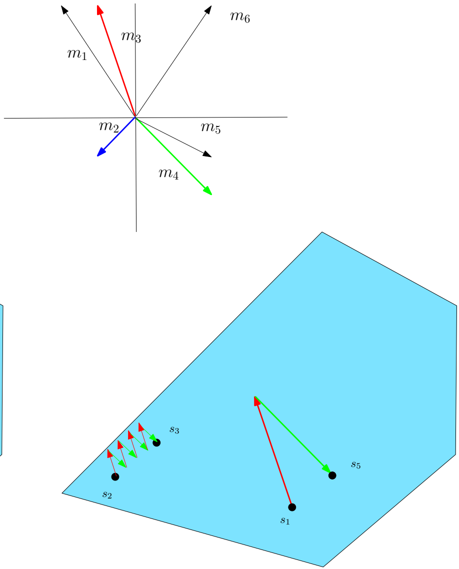

Thumb Rules: Reachability

The following is feasible:

$$s_0 + \sum_{i=1}^{|M|} R(i) \cdot t_i = s_t$$

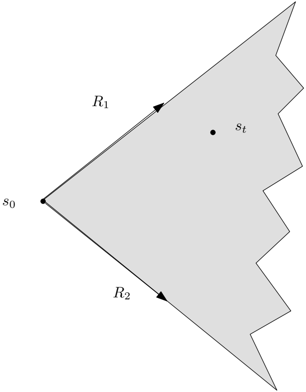

Reachability: Boundary Case

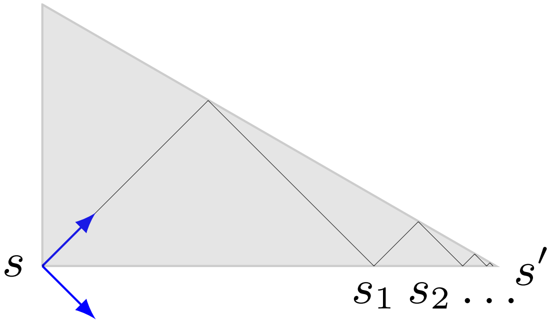

1. Rate vectors are $(1,1)$ and $(1,-1)$
2. Angle at s' is 30^o.
3. $\|s_k, s\| = \|s, s'\| \cdot (\frac{\sqrt{3}-1}{\sqrt{3}+1})^k$.

Schedulability: Boundary Case

Lemma

For any finite safe schedule σ there exists a finite safe schedule σ' s.t.:

1. *All modes that were ever safe during the run of σ will be safe in the final state of σ', and*
2. *The set of safe modes along the run of σ' will always be increasing.*

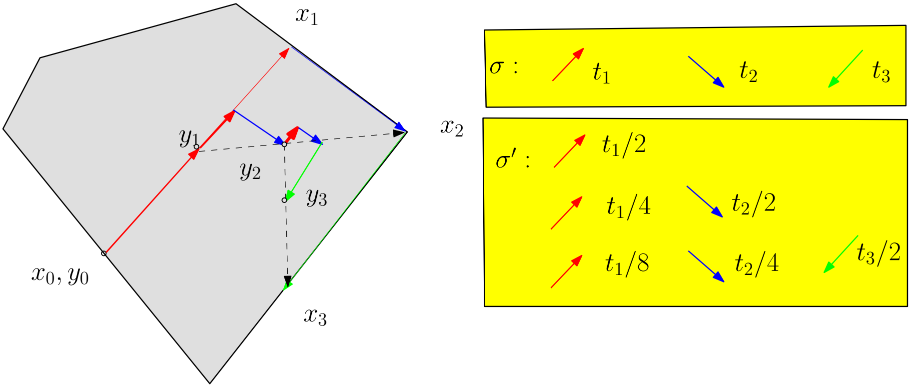

Algorithm: Boundary Case

1. Compute the sequence of set of modes $M_1, M_2, \ldots, M_k$ such that
   - M_1 is the set of safe modes at x_0, and
   - M_i is the set of safe modes at states reachable from x_0 using only modes from M_{i-1}.
2. $M_1 \subset M_2 \subset \cdots \subset M_k$.
3. Modes outside M_k can never be used when starting at x_0.
4. The set M_k can be computed in polynomial time.
5. MMS is schedulable from x_0 if and only if:

$$\sum_{m \in M_k} R(m) \cdot f_m = 0 \text{ and } \sum_{m \in M_k} f_m = 1$$

6. That can, again, be checked in polynomial time.
7. If the system is safe, there exists an ultimately periodic schedule.

Generalisations and Cost Optimal schedules

Generalisations:

- One can add some structure to the real-time system by adding
  - guards on mode-switches
  - mode-dependent invariants
- Both generalisations lead to undecidability of the reachability problem.
- Cost per time unit in each mode – the same complexity via similar analysis.
- .... and what about cost per mode switch?

Multi-mode Systems with Discrete Costs

Definition (Multi-mode system with discrete costs)

MMS with discrete costs is a tuple $\mathcal{H} = (M, N, R, \pi_c, \pi_d, V_{\mathsf{min}}, V_{\mathsf{max}}, V_0)$ where:

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

Costs of Schedules

The cost of a finite schedule $\sigma = \langle (m_1, t_1), (m_2, t_2), \ldots, (m_k, t_k) \rangle$ is defined as $\pi(\sigma) = \sum_{i=1}^{k} \pi_d(m_i) + \pi_c(m_i) t_i$ and its time horizon is $\sum_{i=1}^{k} t_i$.

Finite-time horizon problem

Minimise $\pi(\sigma)$ among schedules σ with a given time horizon t_{max} (in binary).

The limit-average cost of an infinite schedule $\sigma = \langle (m_1, t_1), (m_2, t_2), \ldots \rangle$ is defined as

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

Infinite-time horizon problem

Minimise $\pi_{avg}(\sigma)$ among all schedules with infinite time horizon.

Finite-time horizon problem

Observation

A finite-time horizon problem may not have an optimal schedule.

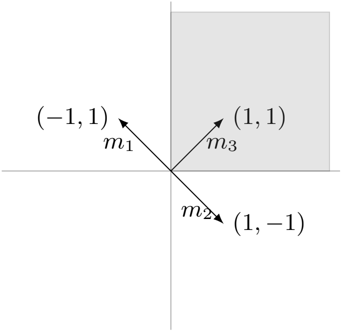

All costs are 0 apart from $\pi_c(m_3) = 1$.

Schedule $\sigma_\epsilon = (m_3, \epsilon), \big((m_1, t), (m_2, t)\big)^l$, where $t' = t_{max} - \epsilon$, $l = \lceil t'/\epsilon \rceil$, and $t = t'/2l$, has time horizon t_{max} and total cost $\epsilon > 0$.

Finite-time horizon problem

Lemma

There exists an ϵ-optimal safe schedule of an exponential length.

Proof (sketch)

1. Remove all modes that can never be used by a safe schedule from V_0 using a similar procedure as in the boundary case for the schedulability problem.
2. Find an easy target state with time horizon t_{max}:
   - minimum number of coefficients at the border using $\mathcal{O}(N)$ LP queries;
   - sufficiently far away from the border using another LP.
3. Split the problem into reaching the mid-point of this schedule by considering $-\mathcal{H}$.
4. All safe modes at V_0 are safe at this mid-point, which makes it easier to reach.

Finite-time horizon problem

Theorem

Checking for the existence of an ϵ-optimal safe schedule with cost at most C is in $\textsc{NExpTime}$.

Proof (sketch)

1. Guess the order of modes used in an ϵ-optimal safe schedule.
2. Write down an exponentially sized LP where the duration of each timed action is a separate variable.
3. Check for the existence of a safe schedule and its minimal total cost.
4. Compare this cost with C.

Theorem

Checking for the existence of an ϵ-optimal ϵ-safe schedule with cost at most C is in $\textsc{PSpace}$.

Proof (sketch)

We can guess on the fly the modes and intermediate points along the schedule if we allow for ϵ deviations from the safe set, and use Savitch's theorem.

Infinite-time horizon problem

However, we know more about the 1-dimensional case.

Theorem

An optimal safe infinite schedule in a 1-dimensional MMS with discrete costs can be computed in deterministic LogSpace.

Finite-time horizon problem in 1-dimension

Let $M^+ = \{m \in M \mid R(m) > 0\}, M^- = \{m \in M \mid R(m) < 0\}$,
$M^0 = \{m \in M \mid R(m) = 0\}$.

A leap is a subsequence $(m_k, t_k), (m_{k+1}, t_{k+1})$ in a schedule such that $m_k \in M^+$, $m_{k+1} \in M^-$, and $R(m_k)t_k = -R(m_{k+1})t_{k+1} = V_{max} - V_{min}$.

A leap is of type $(i, j) \in M^+ \times M^-$ iff $m_k = i, m_{k+1} = j$.

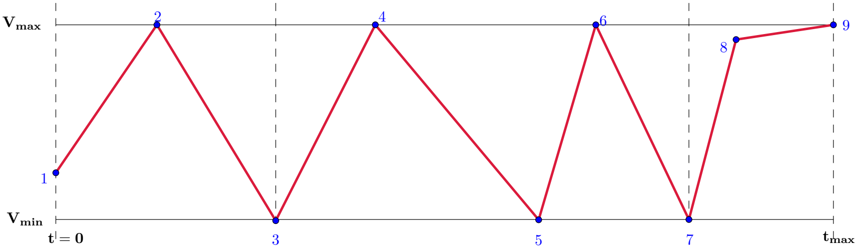

Structure of an optimal schedule

Any schedule longer than 2 can be partitioned into its head, leaps, and tail sections.

- The head section ends once V_{min} is reached.
- The leaps section is the maximal part of the schedule after the head section consisting only of leaps of possibly different types.
- The tail section starts where the leaps section ends. (It may consists of further leaps.)

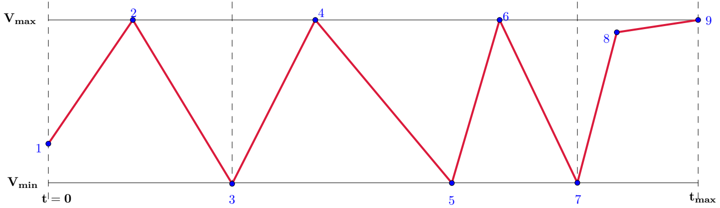

Theorem (Mousa, Schewe, W., 2017)

Any schedule can be transformed without increasing its cost nor compromising its safety into one where the head and tail sections follow one of the following patterns...

10 Possible Head Patterns

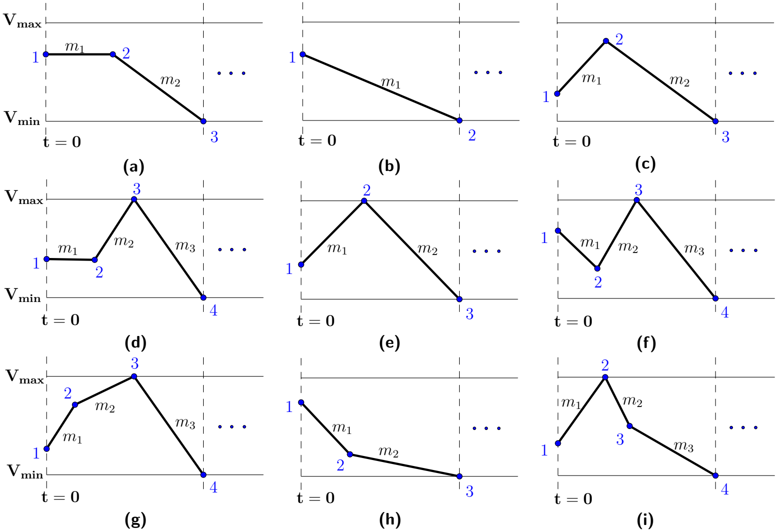

+ the empty one!

10 Possible Tail Patterns

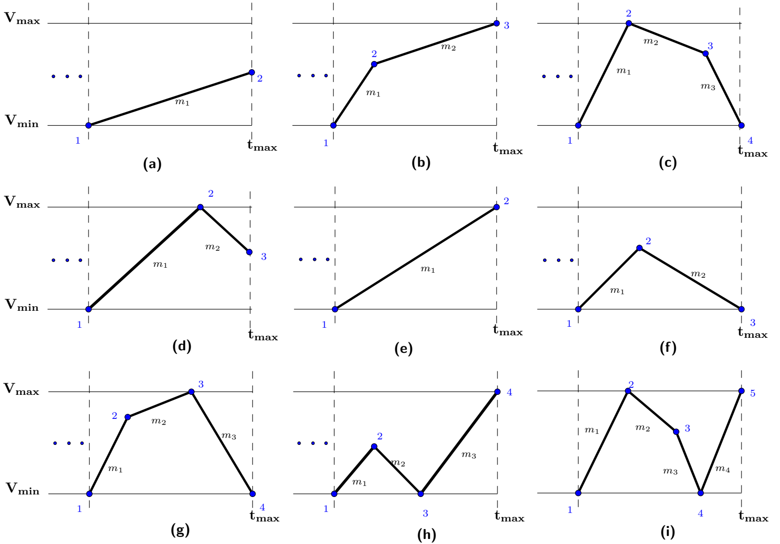

+ the empty one!

Rearrange Operation

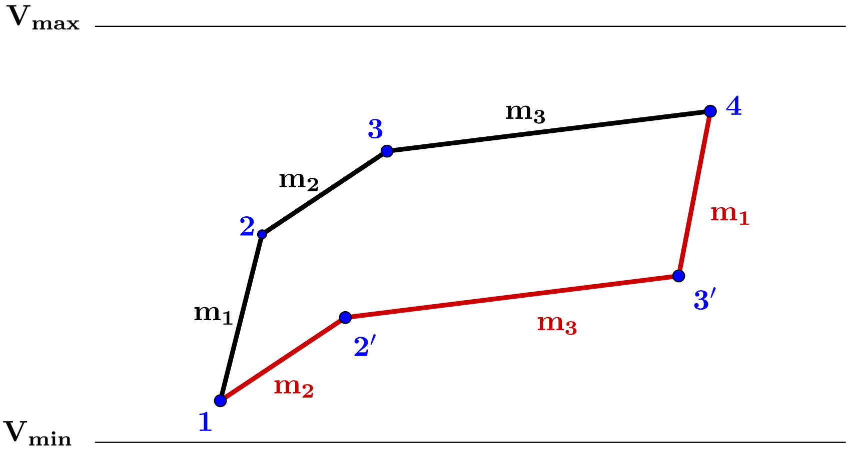

Shift Operation

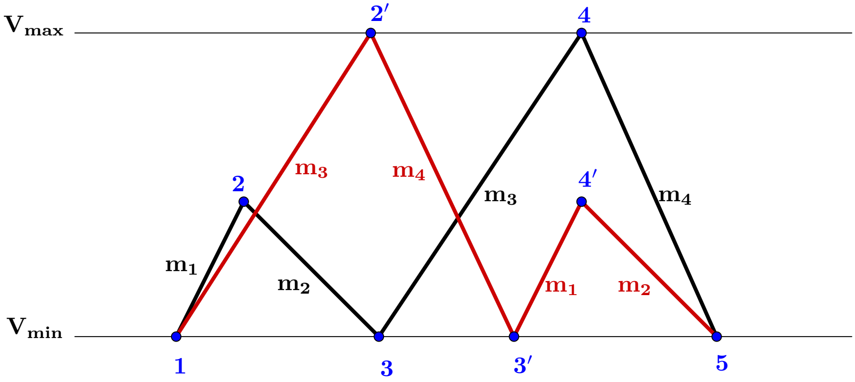

Shift-down operation

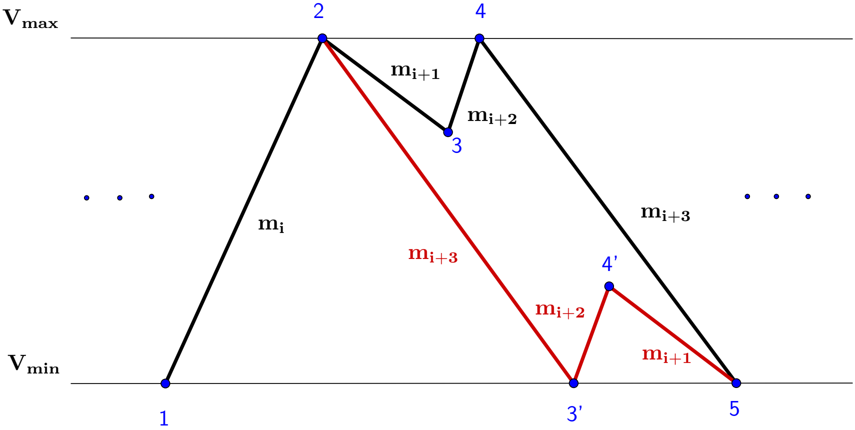

Wedge Operation

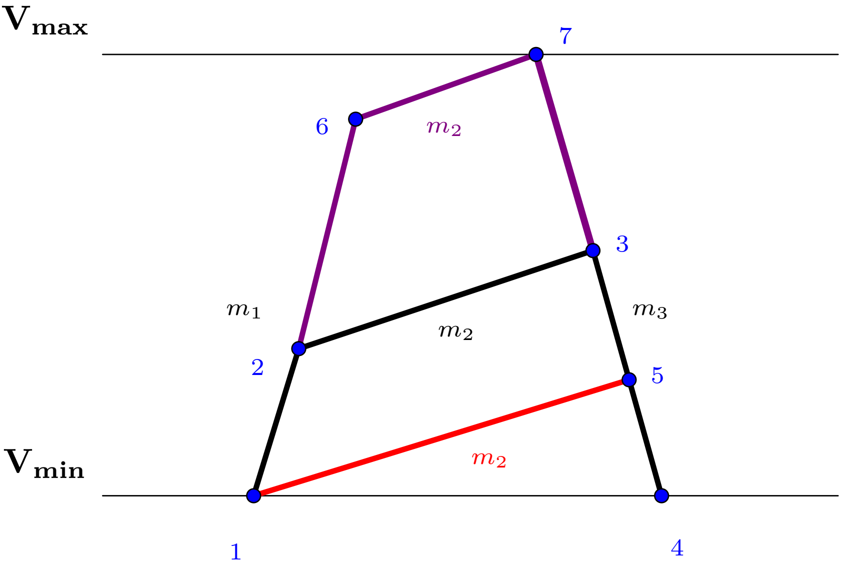

Shrink and Stretch operations

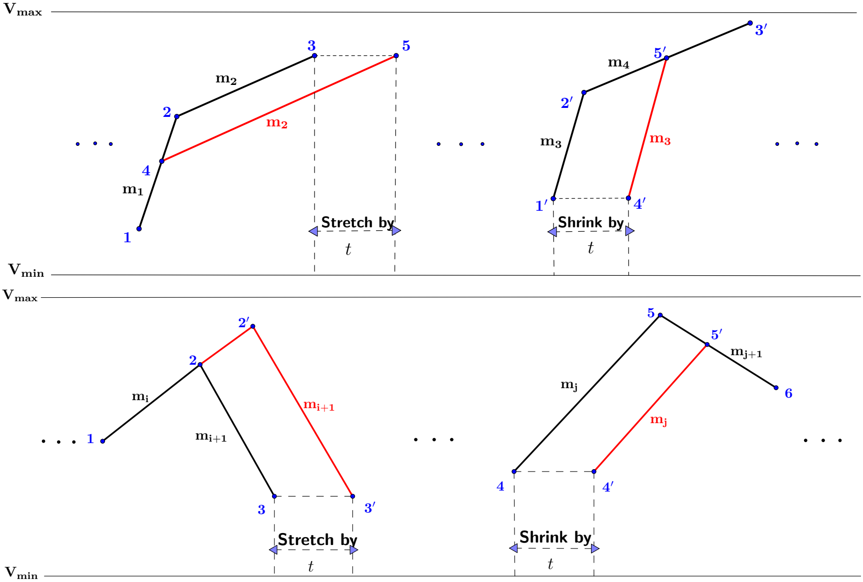

These can also be applied to the last timed action and the ones with $m \in M^0$.

Proof by Example

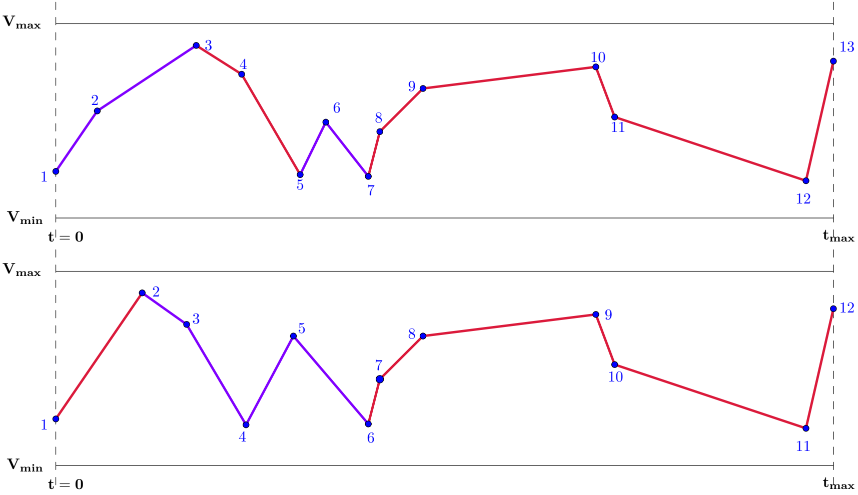

Proof by Example

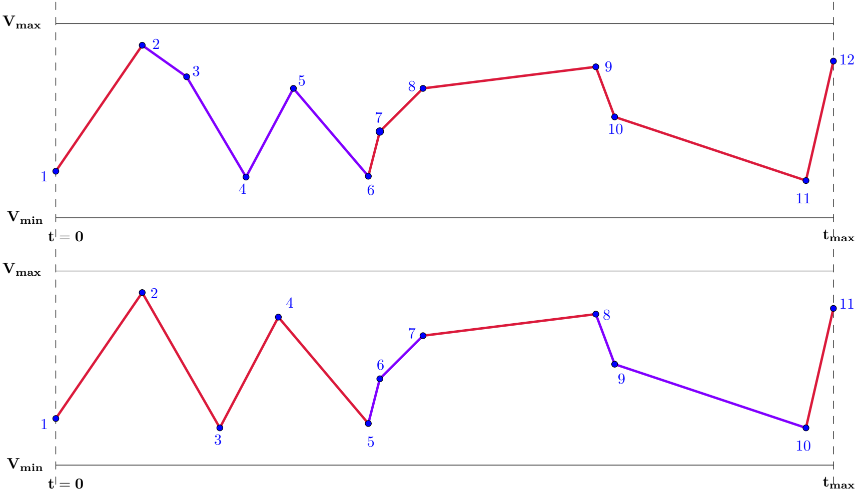

Proof by Example

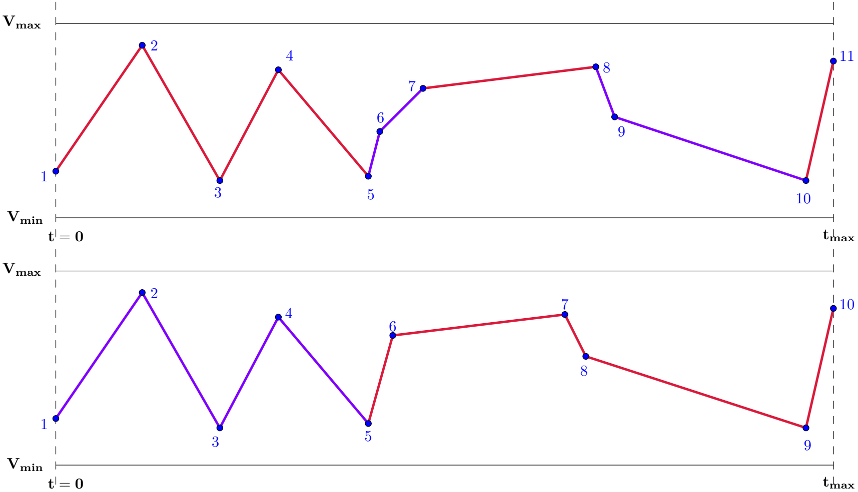

Proof by Example

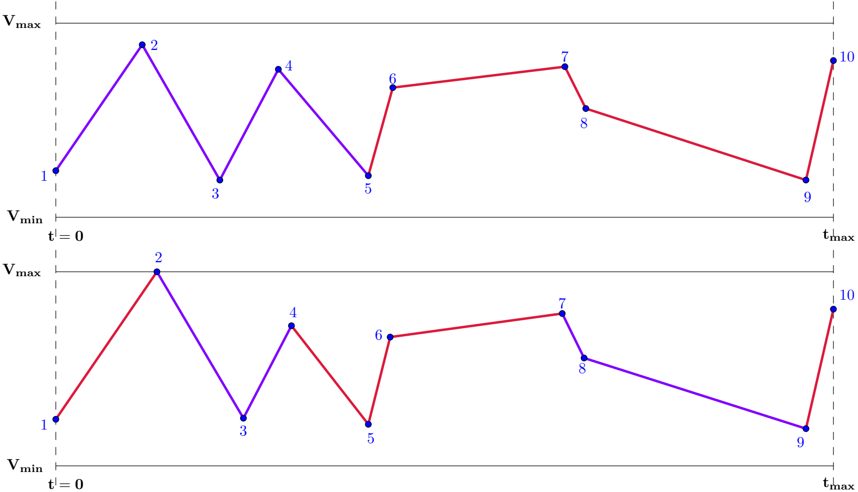

Proof by Example

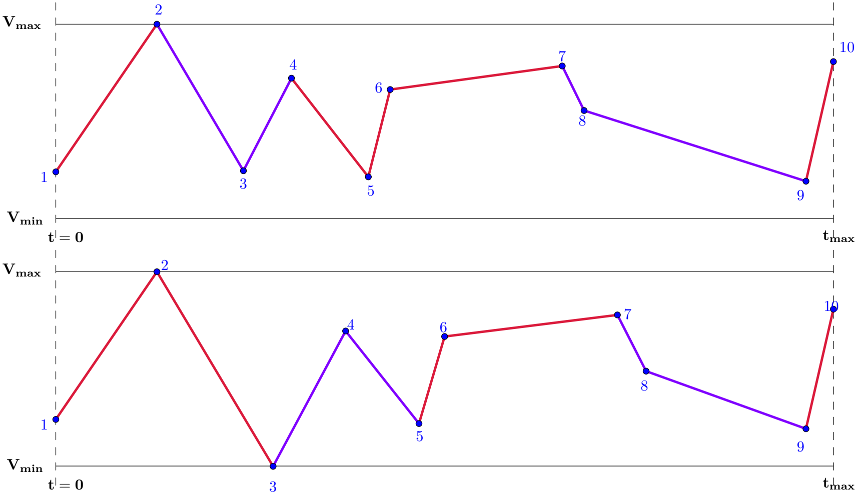

Proof by Example

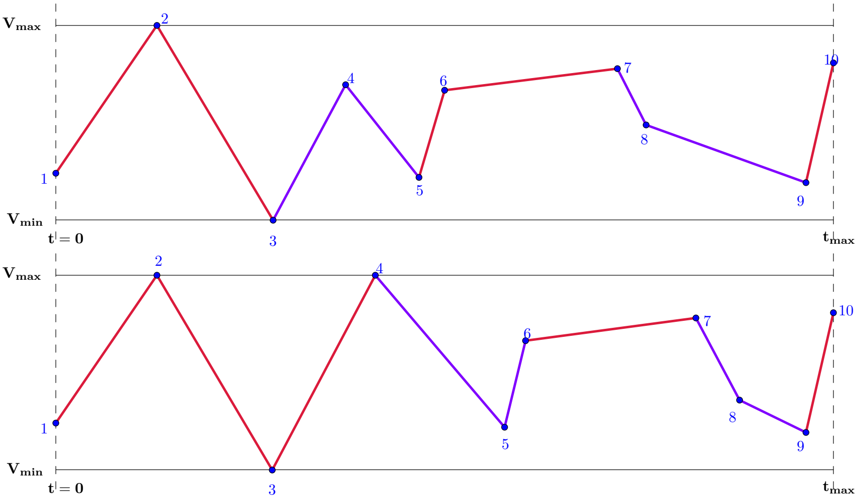

Proof by Example

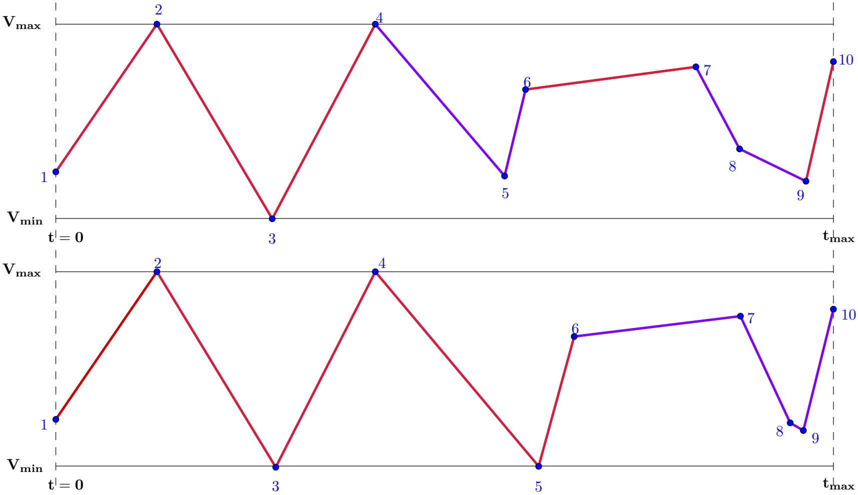

Proof by Example

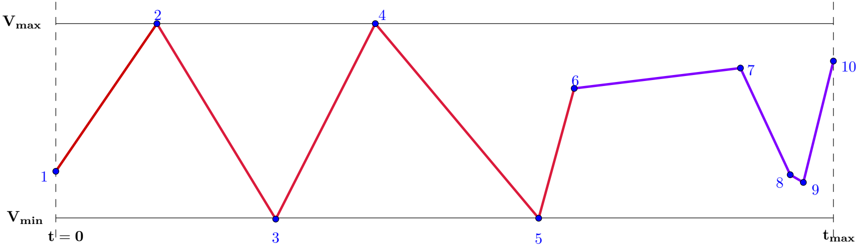

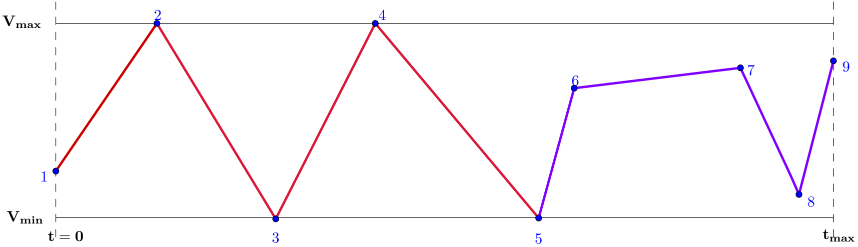

Proof by Example

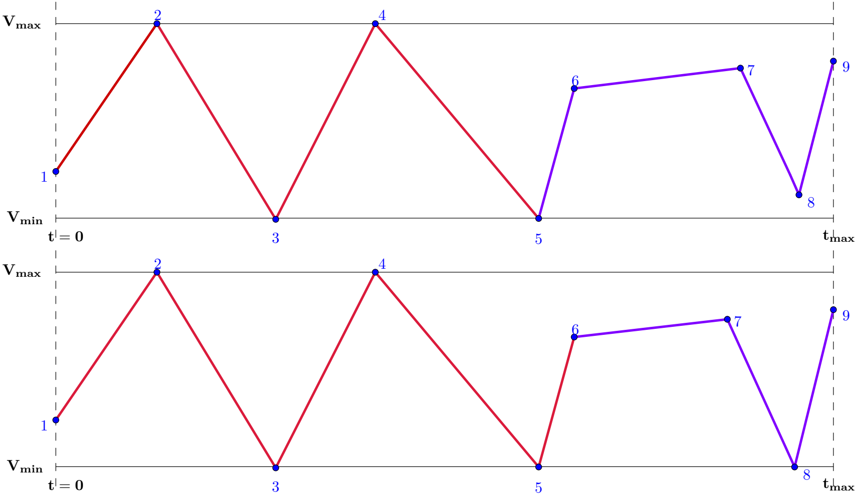

Proof by Example

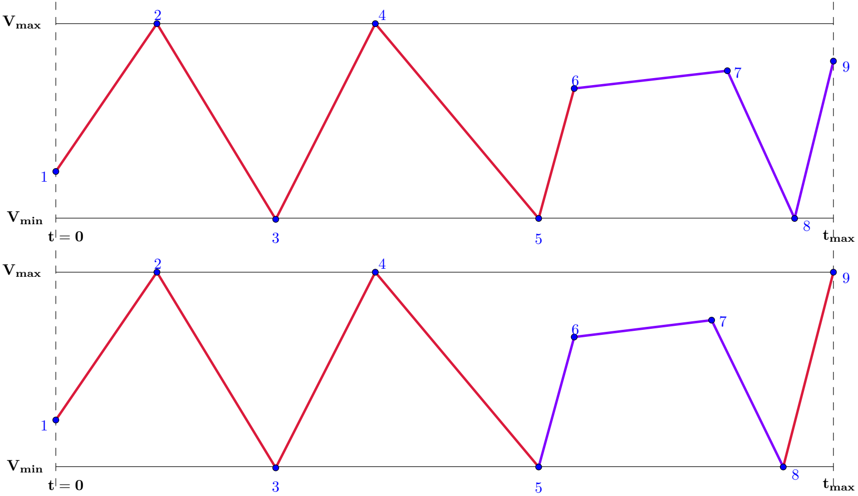

Proof by Example

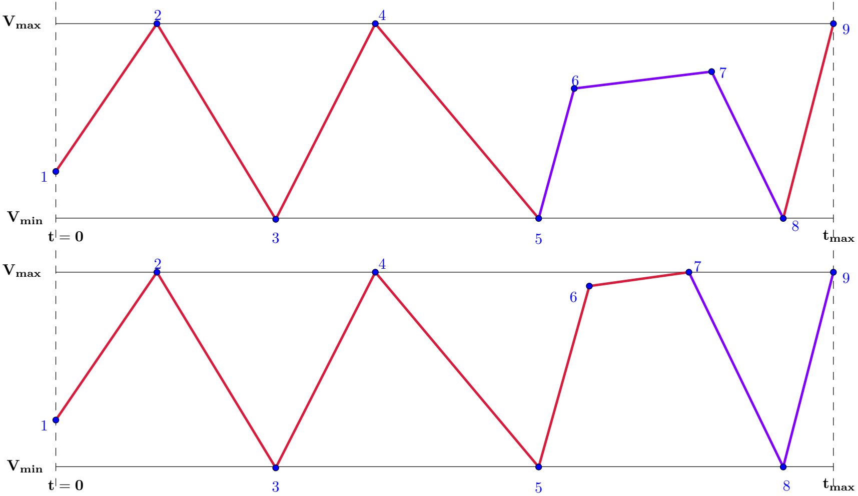

Proof by Example

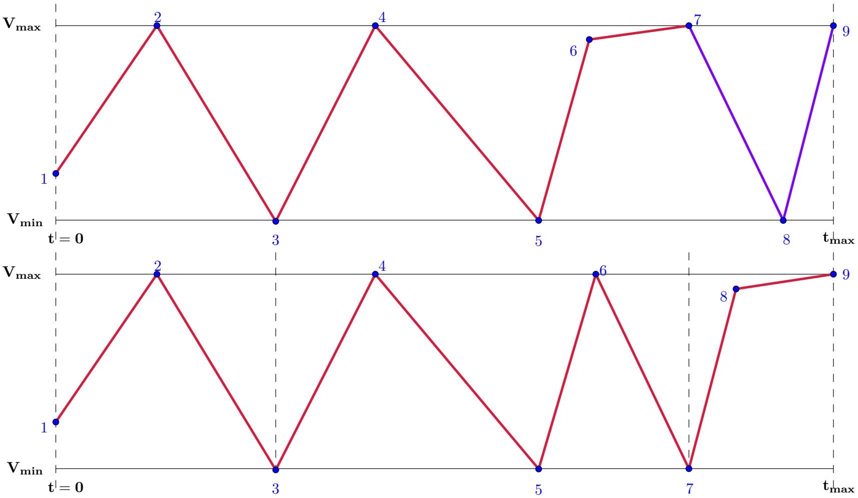

Structure of an optimal schedule

There are 100 different combined patterns for the head and tail sections.

However, many of these combinations can be further reduced using one of the just defined operations.

Intuition: there can be only one point of flexibility in a given schedule.

If there are two then one of them can be removed using the shrink - stretch operation combination.

In the end, we obtain 44 different combined patterns that cannot be further reduced and their combined length is at most 5.

Corollary

Optimal schedules for 1-dimensional MMS with a finite-time horizon always exist.

Approximation algorithms

Definition

An algorithm has a **relative performance** ρ if 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.

Definition

A fully polynomial-time approximation scheme (FPTAS) is an algorithm that runs in polynomial-time in the size of the input and $1/\rho$.

Constant-approximation algorithm

Theorem

A 3-approximate optimal schedule can be found in $\mathcal{O}(|M|^7)$.

Proof (sketch)

1. Consider all schedules of length less than 3 in $\mathcal{O}(|M|^2)$.
2. Iterate over all possible combined patterns for schedules of length more than 2.
3. Note that no pattern uses more than 5 different modes for its head and tail section (so $44 \cdot |M|^5$ possibilities in total).
4. Picking among them the cheapest schedule that only uses leaps of the same type ($|M|^2$ possibilities) gives us a 3-approximate optimal schedule.

FPTAS

Theorem

There is an FPTAS for the optimal cost with finite time horizon problem for MMS with discrete costs in 1-dimension.

Proof (sketch)

- We first call the 3-approximation algorithm and reduce the problem to a $0-1$ Knapsack problem instance.
- We iterate over all possible combined schedule patterns and their modes ($44 \cdot |M|^5$ possibilities).
- In each case, the FPTAS algorithm is a bit different.

For More Details

Rajeev Alur, Ashutosh Trivedi, and Dominik Wojtczak.
Optimal Scheduling for Constant-Rate Multi-Mode Systems.
In *Proc. of HSCC*, pages 75–84. ACM, 2012.

Dominik Wojtczak.
Optimal Control for Linear-Rate Multi-mode Systems.
In *Proc. of FORMATS*, volume 8053 of *LNCS*, pages 258–273. Springer, 2013.

Mahmoud A. A. Mousa, Sven Schewe, and Dominik Wojtczak.
Optimal Control for Simple Linear Hybrid Systems.
In *Proc. of TIME*, pages 12–20. IEEE Computer Society, 2016.

Mahmoud A. A. Mousa, Sven Schewe, and Dominik Wojtczak.
Optimal Control for Multi-Mode Systems with Discrete Costs.
In *Proc. of FORMATS*, page (to appear). Springer, 2017.

Summary

- Multi-mode systems are an expressive subclass of hybrid automata with decidable or even tractable analysis.
- **Without discrete costs**:
  - Polynomial-time algorithms for optimal cost reachability and optimal average-cost schedulability;
  - Adding either local invariants or guards lead to undecidability.
- **With discrete switching costs**:
  - an optimal schedule may not exist;
  - in NEXPTIME (and NP-hard) for an ϵ-optimal finite-time horizon problem, or in PSPACE for its ϵ-safe ϵ-optimal version;
  - the decidability of the (ϵ)-optimal infinite-time horizon problem is not known;
  - in the 1-dimensional case:
    - optimal schedules always exist;
    - NP-complete for a finite-time horizon problem, but has FPTAS;
    - infinite-time horizon version is in LOGSPACE.

THANKS!