

Reachability analysis of linear hybrid systems via block decomposition

Sergiy Bogomolov, Marcelo Forets, Goran Frehse, Kostiantyn Potomkin, and Christian Schilling

Abstract—Reachability analysis aims at identifying states reachable by a system within a given time horizon. This task is known to be computationally expensive for linear hybrid systems. Reachability analysis works by iteratively applying continuous and discrete post operators to compute states reachable according to continuous and discrete dynamics, respectively. In this paper, we enhance both of these operators and make sure that most of the involved computations are performed in low-dimensional state space. In particular, we improve the continuous-post operator by performing computations in high-dimensional state space only for time intervals relevant for the subsequent application of the discrete-post operator. Furthermore, the new discrete-post operator performs low-dimensional computations by leveraging the structure of the guard and assignment of a considered transition. We illustrate the potential of our approach on a number of challenging benchmarks.

Index Terms—Reachability, Hybrid systems, Decomposition.

I. INTRODUCTION

A hybrid system [1] is a formalism for modeling cyber-physical systems. Reachability analysis is a rigorous way to reason about the behavior of hybrid systems.

In this paper, we describe a new reachability algorithm for linear hybrid systems, i.e., hybrid systems with dynamics given by linear differential equations and invariants and guards given by linear inequalities. The key feature of our algorithm is that it performs computations in low-dimensional subspaces, which greatly improves scalability. To this end, we integrate (and modify) a recent reachability algorithm for (purely continuous) LTI systems [2], which we call $Post_C^\square$ in the following, in a new algorithm for linear hybrid systems.

The $Post_C^\square$ algorithm decomposes the calculation of the reachable states into calculations in subspaces (called *blocks*).

S. Bogomolov is with Newcastle University, United Kingdom, (e-mail: sergiy.bogomolov@newcastle.ac.uk).

M. Forets is with Universidad de la República, CURE, Maldonado, Uruguay, (e-mail: mforets@gmail.com).

G. Frehse is with ENSTA ParisTech - U2IS, Palaiseau Cedex, France, (e-mail: goran.frehse@ensta-paris.fr).

K. Potomkin is with Newcastle University, United Kingdom, (e-mail: K.Potomkin2@newcastle.ac.uk).

C. Schilling is with IST Austria, Klosterneuburg, Austria, (e-mail: christian.schilling@ist.ac.at).

Manuscript received April 17, 2020; revised June 17, 2020; accepted July 6, 2020. This article was presented in the International Conference on Embedded Software 2020 and appears as part of the ESWEK-TCAD special issue.

This research was supported in part by the Austrian Science Fund (FWF) under grants S11402-N23 (RISE/SHiNE) and Z211-N23 (Wittgenstein Award), the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 754411, and the Air Force Office of Scientific Research under award number FA2386-17-1-4065. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the United States Air Force.

This decomposition has two benefits. The first benefit is that computations in lower dimensions are generally more efficient and thus the algorithm is highly scalable. The second benefit is that the analysis for different subspaces is decoupled; hence one can effectively skip the computations for dimensions that are of no interest (e.g., for a safety property).

Conceptually, reachability analysis for hybrid systems is performed in a “hybrid loop” that interleaves a continuous-post algorithm and a discrete-post algorithm. If we consider $Post_C^\square$ as a black box, we can plug it into this hybrid loop, which we refer to as $Post_H$ (cf. Section III). However, there are two shortcomings of such a naive integration. First, all operations besides $Post_C^\square$ are still performed in high dimensions, making $Post_H$ still suffer from scalability issues. Second, $Post_H$ does not utilize the decoupling of $Post_C^\square$ at all.

We demonstrate that, unlike in $Post_H$, it is possible to perform all computations in low dimensions (cf. Section IV). For this purpose, we modify $Post_C^\square$ as well as $Post_H$. In common cases, our algorithm does not introduce an additional approximation error. Furthermore, our algorithm makes proper use of the second benefit of $Post_C^\square$ by computing the reachable states only in specific dimensions whenever possible.

We implemented the algorithm in JuliaReach, a toolbox for reachability analysis [3], [4], and we evaluate the algorithm on several benchmark problems, including a 1024-dimensional hybrid system (cf. Section V). Our algorithm outperforms the naive $Post_H$ and other state-of-the-art algorithms by several orders of magnitude.

To summarize, we show how to modify the decomposed reachability algorithm for LTI systems from [2] in order to efficiently integrate it into a new, decomposed reachability algorithm for linear hybrid systems. The key insights are (1) to exploit the decomposed structure of the algorithm to perform all operations in low dimensions and (2) to only compute the reachable states in specific dimensions whenever possible.

Related work

Reachability analysis for linear hybrid systems: There are several tools available for the analysis of the class of systems we consider in this paper. All participating tools in a recent friendly competition [5] follow the architecture of two post operators, one for continuous time and one for discrete time [6], [7], [8], [3], [9], [10]. We implemented our algorithm in JuliaReach [3] and will consider the mature tool SpaceEx [9] in the evaluation in Section V.

Decomposition: Hybrid systems given as a network of components can be explored efficiently in a symbolic way, e.g., using bounded model checking [11]. We consider a decomposition in the continuous state space. Schupp et al. perform such a decomposition by syntactic independence [12], which corresponds to dynamics with block-diagonal matrices (whereas our decomposition is generally applicable).

For purely continuous systems there exist various decomposition approaches. In this work we build on [2] for LTI systems, which decomposes the system into blocks and exploits the linear dynamics to avoid the wrapping effect. Other approaches for LTI systems are based on Krylov subspace approximations [13], time-scale decomposition [14], [15], similarity transformations [16], [17], projectahedra [18], [19], and sub-polyhedra abstract domains [20]. Approaches for nonlinear systems are based on projections with differential inclusions [21], Hamilton-Jacobi methods [22], [23], and hybridization with iterative refinement [24].

Lazy flowpipe computation: The support-function representation of convex sets can be used to represent a flowpipe (a sequence of sets that covers the behaviors of a system) symbolically [25]. Only sets that are of interest, e.g., those that intersect with a linear constraint, need then be approximated [26]. Using our decomposition approach, we can even avoid the symbolic computation in dimensions that are irrelevant to the intersection. Our approach is independent of the set representation, so it can also be applied in analyses based on, e.g., zonotopes [27], [28], [29], [30]. Given a linear switching system with a hyperplanar state-space partition, one approach computes ellipsoidal approximations of the reach set on the partition borders, without computing the full reach set [31].

Intersection of convex sets: Performing intersections in low dimensions allows for efficient computations that are not possible in high dimensions. For example, checking for emptiness of a polyhedron in constraint representation is a feasibility linear program, which can be solved in weakly polynomial time, but solutions in strongly polynomial time are only known in two dimensions [32].

In the context of hybrid-system reachability, computing intersections is a major challenge because it usually requires a conversion from efficient set representations (like zonotopes, support function, or Taylor models) to polytopes and back, often entailing additional approximation. Below we summarize how other approaches tackle the intersection problem.

A coarse approximation of an intersection of a set \mathcal{X} and another set \mathcal{Y} can be obtained by only detecting a nonempty intersection (which is generally easier to do) and then taking the original set \mathcal{X} as overapproximation [33]. In general, the intersection between a polytope and polyhedral constraints (invariants and guards) can be computed exactly, but such an approach is not scalable [34]. Girard and Le Guernic consider hyperplanar constraints where reachable states are either zonotopes, in which case they work in a two-dimensional projection [28], or general polytopes [35], [36]. The tool SpaceEx approximates the intersection of polytopes and general polyhedra using template directions [9]. Frehse and Ray propose an optimization scheme for the intersection of a compact convex set \mathcal{X} , represented by its support function, and a polyhedron

\mathcal{Y} , and this scheme is exact if \mathcal{X} is a polytope [26]. The problem of performing intersections can also be cast in terms of finding separating hyperplanes [37], [38]. Althoff et al. approximate zonotopes by parallelotopes before considering the intersection [29]. For must semantics, Althoff and Krogh use constant-dynamics approximation and obtain a nonlinear mapping [39]. Under certain conditions, Bak et al. apply a model transformation by replacing guard constraints by time-triggered constraints, for which intersection is easy [40].

II. PRELIMINARIES

We introduce some notation. The real numbers are denoted by \mathbb{R} . The origin in \mathbb{R}^n is written $\mathbf{0}_n$. Given two vectors $x, y \in \mathbb{R}^n$, their dot product is $\langle x, y \rangle := \sum_{i=1}^n x_i \cdot y_i$. For $p \geq 1$, the p -norm of a matrix $A \in \mathbb{R}^{n \times n}$ is denoted $\|A\|_p$. The diameter of a set $\mathcal{X} \subseteq \mathbb{R}^n$ is $\Delta_p(\mathcal{X}) := \sup_{x, y \in \mathcal{X}} \|x - y\|_p$. The n -dimensional unit ball of the p -norm is $\mathcal{B}_p^n := \{x \in \mathbb{R}^n \mid 1 \geq \|x\|_p\}$. An n -dimensional half-space is the set $\{ \langle a, x \rangle \leq b \mid x \in \mathbb{R}^n \}$ parameterized by $a \in \mathbb{R}^n, b \in \mathbb{R}$. A polyhedron is an intersection of finitely many half-spaces, and a polytope is a bounded polyhedron.

Given sets $\mathcal{X} \subseteq \mathbb{R}^n$ and $\mathcal{Y} \subseteq \mathbb{R}^m$, scalar $\lambda \in \mathbb{R}$, matrix $A \in \mathbb{R}^{n \times n}$, and vector $b \in \mathbb{R}^n$, we use the following operations on sets: scaling $\lambda\mathcal{X} := \{\lambda x \mid x \in \mathcal{X}\}$, linear map $A\mathcal{X} := \{Ax \mid x \in \mathcal{X}\}$, Minkowski sum $\mathcal{X} \oplus \mathcal{Y} := \{x + y \mid x \in \mathcal{X} \text{ and } y \in \mathcal{Y}\}$ (if $n = m$), affine map $(A, b) \odot \mathcal{X} := A\mathcal{X} \oplus \{b\}$, Cartesian product $\mathcal{X} \times \mathcal{Y} := \{(x, y) \mid x \in \mathcal{X}, y \in \mathcal{Y}\}$, intersection $\mathcal{X} \cap \mathcal{Y} := \{z \mid z \in \mathcal{X}, z \in \mathcal{Y}\}$ (if $n = m$), and convex hull $\text{CH}(\mathcal{X}) := \{\lambda \cdot x + (1 - \lambda) \cdot y \mid x, y \in \mathcal{X}, 0 \leq \lambda \leq 1\}$.

Given two sets $\mathcal{X}, \mathcal{Y} \subseteq \mathbb{R}^n$, the Hausdorff distance is

$$d_H^p(\mathcal{X}, \mathcal{Y}) := \inf_{\varepsilon \in \mathbb{R}} \{ \mathcal{Y} \subseteq \mathcal{X} \oplus \varepsilon \mathcal{B}_p^n \text{ and } \mathcal{X} \subseteq \mathcal{Y} \oplus \varepsilon \mathcal{B}_p^n \}.$$

Let $\mathcal{C}_n \subseteq 2^{\mathbb{R}^n}$ be the set of n -dimensional compact convex sets. For a nonempty set $\mathcal{X} \in \mathcal{C}_n$, the support function $\rho_{\mathcal{X}} : \mathbb{R}^n \rightarrow \mathbb{R}$ is defined as

$$\rho_{\mathcal{X}}(d) := \max_{x \in \mathcal{X}} \langle d, x \rangle.$$

The Hausdorff distance of two sets $\mathcal{X}, \mathcal{Y} \in \mathcal{C}_n$ with $\mathcal{X} \subseteq \mathcal{Y}$ can also be expressed in terms of the support function as

$$d_H^p(\mathcal{X}, \mathcal{Y}) = \max_{\|d\|_p \leq 1} \rho_{\mathcal{Y}}(d) - \rho_{\mathcal{X}}(d).$$

We use β to denote the number of blocks in a partition. Let $\{\pi_j\}_{j=1}^{\beta}$ be a set of (contiguous) projection matrices that partition a vector $x \in \mathbb{R}^n$ into $x = [\pi_1 x, \dots, \pi_{\beta} x]$. Given a set \mathcal{X} and projection matrices $\{\pi_j\}_{j=1}^{\beta}$, we call $\pi_j \mathcal{X}$ a *block* of \mathcal{X} and $\times_j \pi_j \mathcal{X} = \pi_1 \mathcal{X} \times \dots \times \pi_{\beta} \mathcal{X}$ the *Cartesian decomposition* with the block structure induced by projections π_j . We write $\tilde{\mathcal{X}}$ to indicate a decomposed set (i.e., a Cartesian product of lower-dimensional sets). For instance, given a nonempty set $\mathcal{X} \in \mathcal{C}_n$, its box approximation is the Cartesian decomposition into intervals (i.e., one-dimensional blocks). We can bound the approximation error by the radius of \mathcal{X} .

Proposition 1: Let $\mathcal{X} \in \mathcal{C}_n$ be nonempty, $p = \infty$, $r_{\mathcal{X}}^p$ be the radius of the box approximation of \mathcal{X} , and let π_j be appropriate projection matrices. Then $d_H^p(\mathcal{X}, \times_j \pi_j \mathcal{X}) \leq \|r_{\mathcal{X}}^p\|_p$.

A. LTI systems

An n -dimensional LTI system (A, B, \mathcal{U}) , with matrices $A \in \mathbb{R}^{n \times n}$, $B \in \mathbb{R}^{n \times m}$, and input domain $\mathcal{U} \in \mathcal{C}_m$, is a system of ODEs of the form

$$\dot{x}(t) = Ax(t) + Bu(t), \quad u(t) \in \mathcal{U}. \quad (1)$$

We denote the set of all n -dimensional LTI systems by \mathcal{L}_n . From now on, a vector $x \in \mathbb{R}^n$ is also called a (continuous) state. Given an initial state $x_0 \in \mathbb{R}^n$ and an input signal u such that $u(t) \in \mathcal{U}$ for all t , a *trajectory* of (1) is the unique solution $\xi_{x_0, u} : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}^n$ with

$$\xi_{x_0, u}(t) = e^{At}x_0 + \int_0^t e^{A(t-s)}Bu(s)ds.$$

Given an LTI system (A, B, \mathcal{U}) and a set $\mathcal{X}_0 \in \mathcal{C}_n$ of initial states, the *continuous-post operator*, $Post_C$, computes the set of reachable states for all input signals u over \mathcal{U} :

$$Post_C((A, B, \mathcal{U}), \mathcal{X}_0) := \{\xi_{x_0, u}(t) \mid t \geq 0, x_0 \in \mathcal{X}_0, u(s) \in \mathcal{U} \text{ for all } s\}.$$

B. Linear hybrid systems

We briefly introduce the syntax of linear hybrid systems used in this work and refer to the literature for the semantics [41], [1]. An n -dimensional linear hybrid system is a tuple $\mathcal{H} = (Var, Loc, Flow, Inv, Grd, Asgn)$ with variables $Var = \{x_1, \dots, x_n\}$, a finite set of locations Loc , two functions $Flow : Loc \rightarrow \mathcal{L}_n$ and $Inv : Loc \rightarrow \mathcal{C}_n$ that respectively assign continuous dynamics and an invariant to each location, and two functions $Grd : Loc \times Loc \rightarrow \mathcal{C}_n$ and $Asgn : Loc \times Loc \rightarrow \mathbb{R}^{n \times n} \times \mathbb{R}^n$ that respectively assign a guard and an assignment in the form of a deterministic affine map to each pair of locations. If $Grd(\ell, \ell') \neq \emptyset$, we call (ℓ, ℓ') a (discrete) transition.

Let $\mathcal{H} = (Var, Loc, Flow, Inv, Grd, Asgn)$ be a linear hybrid system. A (symbolic) state of \mathcal{H} is a pair $(\ell, \mathcal{X}) \in Loc \times \mathcal{P}^{\mathbb{R}^n}$. The *discrete-post operator*, $Post_D$, maps a symbolic state to a set of symbolic states by means of discrete transitions:

$$Post_D(\ell, \mathcal{X}) := \bigcup_{\ell' \in Loc} \{(\ell', Asgn(\ell, \ell')) \odot (\mathcal{X} \cap Inv(\ell) \cap Grd(\ell, \ell')) \cap Inv(\ell')\} \quad (2)$$

The *reach set* of \mathcal{H} from a set of initial symbolic states \mathcal{R}_0 of \mathcal{H} is the smallest set \mathcal{R} of symbolic states such that

$$\mathcal{R}_0 \cup \bigcup_{(\ell, \mathcal{X}) \in \mathcal{R}} Post_D(\ell, Post_C(Flow(\ell), \mathcal{X})) \subseteq \mathcal{R}. \quad (3)$$

We note that the usual definition of the reach set of a hybrid system is stricter wrt. invariants. However, practical analysis tools use the above definition for efficiency reasons.

Safety properties can be given as a set of symbolic error states that should be avoided and be encoded as the guard of a transition to a new error location. In our implementation we can also check inclusion in the safe states (the complement) and do not require an encoding with additional transitions.

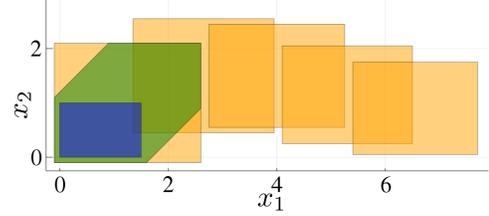


Fig. 1. Starting from the set of initial states \mathcal{X}_0 (blue set), we first compute the set $\mathcal{X}(0)$ by time discretization (green set), then decompose the set into intervals and obtain $\hat{\mathcal{X}}(0)$ (orange box around $\mathcal{X}(0)$), and finally compute the (decomposed) flowpipe $\hat{\mathcal{X}}(1), \dots, \hat{\mathcal{X}}(4)$ by propagating each of the intervals (other orange sets).

III. REACHABILITY ANALYSIS OF LINEAR HYBRID SYSTEMS

Our reachability algorithm for linear hybrid systems is centered around the algorithm from [2], called $Post_C^\square$ for convenience, which implements $Post_C$ for LTI systems in a compositional way. In this section, we first recall the $Post_C^\square$ algorithm. Two important properties of $Post_C^\square$ are that (1) the output is a sequence of *decomposed* sets and that (2) this sequence is computed efficiently in low dimensions.

After explaining the algorithm $Post_C^\square$, we incorporate it in a standard reachability algorithm for linear hybrid systems. However, this standard reachability algorithm will not make use of the above-mentioned properties. This motivates our new algorithm (presented in Section IV), which is a modification of both this standard reachability algorithm and $Post_C^\square$ to make optimal use of these properties.

A. Decomposed reachability analysis of LTI systems

The decomposition-based approach [2] follows a flowpipe-construction scheme using time discretization, which we shortly recall here. Given an LTI system (A, B, \mathcal{U}) and a set of initial (continuous) states \mathcal{X}_0 , by fixing a time step δ we first compute a set $\mathcal{X}(0)$ that overapproximates the reach set up to time δ , a matrix $\Phi = e^{A\delta}$ that captures the dynamics of duration δ , and a set \mathcal{V} that overapproximates the effect of the inputs up to time δ . We obtain an overapproximation of the reach set in time interval $[k\delta, (k+1)\delta]$, for step $k > 0$, with

$$\mathcal{X}(k) := \Phi \mathcal{X}(k-1) \oplus \mathcal{V} = \Phi^k \mathcal{X}(0) \oplus \bigoplus_{j=0}^{k-1} \Phi^j \mathcal{V}.$$

Algorithm $Post_C^\square$ decomposes this scheme. Fixing some block structure, let $\hat{\mathcal{X}}(0) := \mathcal{X}_1(0) \times \dots \times \mathcal{X}_\beta(0)$ be the corresponding Cartesian decomposition of $\mathcal{X}(0)$. We compute a sequence $\hat{\mathcal{X}}(k) := \mathcal{X}_1(k) \times \dots \times \mathcal{X}_\beta(k)$ such that $\mathcal{X}(k) \subseteq \hat{\mathcal{X}}(k)$ for every k . Each low-dimensional set $\mathcal{X}_i(k)$ is computed as

$$\mathcal{X}_i(k) := \bigoplus_{j=1}^{\beta} (\Phi^k)_{ij} \mathcal{X}_j(0) \oplus \bigoplus_{j=0}^{k-1} [(\Phi^j)_{i1} \dots (\Phi^j)_{i\beta}] \mathcal{V}.$$

The above sequences $\mathcal{X}(k)$ resp. $\hat{\mathcal{X}}(k)$ are called flowpipes.

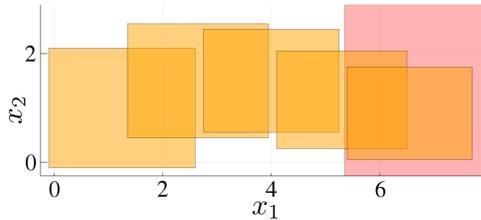


Fig. 2. The flowpipe from Figure 1 together with a guard (red).

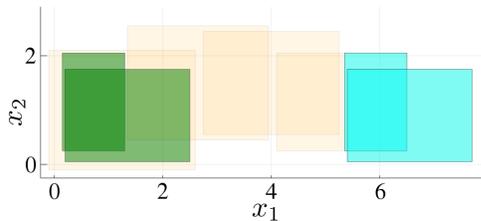


Fig. 3. The assignment shifts the intersection of the flowpipe and the guard from Figure 2 (cyan) to the green sets, which are both contained in the first set of the original flowpipe.

Example: We illustrate the algorithms with a running example throughout the paper. For illustration purposes, the example is two-dimensional (and hence we decompose into one-dimensional blocks, but we emphasize that the approach also generalizes to higher-dimensional decomposition) and we consider a hybrid system with only a single location and one transition (a self-loop). Figure 1 depicts the flowpipe construction for the example.

B. Reachability analysis of linear hybrid systems

We now discuss a standard reachability algorithm for hybrid systems. Essentially, this algorithm interleaves the operators $Post_C$ and $Post_D$ following (3) until it finds a fixpoint. Here we use $Post_C^\square$ as the continuous-post operator.

We first compute a flowpipe $\hat{\mathcal{X}} = \hat{\mathcal{X}}(0), \dots, \hat{\mathcal{X}}(N)$ using $Post_C^\square$ as described above. Then we use $Post_D$ to take a discrete transition. According to (2), we want to compute $((A, b) \odot (\hat{\mathcal{X}} \cap \mathcal{I}_1 \cap \mathcal{G})) \cap \mathcal{I}_2$, where (A, b) is a deterministic affine map \mathcal{I}_1 and \mathcal{I}_2 are the source and target invariant, and \mathcal{G} is the guard. Frehse and Ray showed that for such maps the term can be simplified to

$$(A, b) \odot (\hat{\mathcal{X}} \cap \mathcal{G}^*) \quad (4)$$

where the set \mathcal{G}^* can be statically precomputed [26], which is usually easy because the sets \mathcal{I}_1 , \mathcal{G} , and \mathcal{I}_2 are given as polyhedra in constraint representation (also called H-representation, as opposed to the (vertex) V-representation). Hence we ignore invariants in the rest of the presentation for simplicity.

Example: We continue with the flowpipe from Figure 2. The guard \mathcal{G} is a half-space that is constrained in dimension x_1 and unconstrained in dimension x_2 . Only the last two sets in the flowpipe intersect with the guard. The assignment here is a translation in dimension x_1 . The resulting intersection, before and after the translation, is depicted in Figure 3.

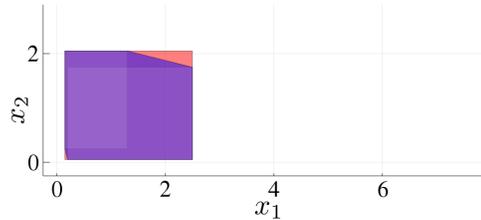


Fig. 4. Approximation of the union of the green sets from Figure 3 using the convex hull (purple) and the decomposed convex hull (red).

Finally, the algorithm checks for a fixpoint, i.e., for inclusion of the symbolic states we computed with $Post_D$ in previously-seen symbolic states.

Example: The green set in Figure 3 shows that the two sets we obtained from $Post_D$ are contained in $\hat{\mathcal{X}}(0)$ computed before. (Recall that in this example we only consider a single location; hence the inclusion of the sets implies inclusion of the symbolic states.)

The steps outlined above describe one iteration of the standard reachability algorithm. Each symbolic state for which the fixpoint check was negative spawns a new flowpipe. Since this can lead to a combinatorial explosion, one typically applies a technique called *clustering* (cf. [9]), where symbolic states are merged after the application of $Post_D$. Here we consider clustering with a convex hull.

Example: Assume that the fixpoint check above was negative for both sets that we tested. In Figure 4 we show the convex hull of the sets in purple.

Up to now, we have seen a standard incorporation of an algorithm for the continuous-post operator $Post_C$ (for which we used $Post_C^\square$) into a reachability algorithm for hybrid systems. Observe that $Post_C$ was used as a black box. Consequently, we could not make use of the properties of the specific algorithm $Post_C^\square$. In particular, apart from $Post_C^\square$, we performed all computations in high dimensions. In the next section, we describe a new algorithm that instead performs all computations in low dimensions.

IV. DECOMPOSED REACHABILITY ANALYSIS

We now present a new, decomposed reachability algorithm for linear hybrid systems. The algorithm uses a modified version of $Post_C^\square$ for computing flowpipes and has two major performance improvements over the algorithm seen before.

Recall that $Post_C^\square$ computes flowpipes consisting of decomposed sets. The first improvement is to exploit the decomposed structure to perform all other operations (intersection, affine map, inclusion check, and convex hull) in low dimensions.

The second improvement is to compute flowpipes in a sparse way. Roughly speaking, we are only interested in those dimensions of a flowpipe that are relevant to determine an intersection with a guard. We only need to compute the other dimensions of the flowpipe if we detect such an intersection.

We defer proofs in this section to Appendix A.

The algorithm starts as before: given an initial (symbolic) state, we compute $\mathcal{X}(0)$ (discretization) and decompose the

Algorithm 1: Function `reach`.

Input: $\mathcal{D} = (\Phi, \mathcal{V}(\cdot), \mathcal{X}(0))$: discrete system
 N : total number of steps
 $blocks$: list of constrained block indices
 $other_blocks$: list of unconstrained block indices
 $constraints$: linear constraints of the outgoing transitions

```

1  $all\_blocks \leftarrow blocks \cup other\_blocks$ ;
2  $\hat{\mathcal{X}}(0) \leftarrow decompose(\mathcal{X}(0), all\_blocks)$ ;
3  $P \leftarrow identity(dim(\Phi))$ ;
4  $Q \leftarrow \Phi$ ;
5  $\hat{\mathcal{V}}_{tmp} \leftarrow []$ ;
6 for  $b_i \in blocks$  do
7    $\hat{\mathcal{V}}_{tmp}[b_i] \leftarrow \{0_{dim(b_i)}\}$ ;
8 end
9 for  $k = 1$  to  $N - 1$  do
10   $\hat{\mathcal{X}}_{tmp} \leftarrow []$ ;
11  for  $b_j \in all\_blocks$  do
12     $\hat{\mathcal{V}}_{tmp}[b_j] \leftarrow \hat{\mathcal{V}}_{tmp}[b_j] \oplus P[b_i, :] \odot \mathcal{V}$ ;
13  end
14   $\hat{\mathcal{X}}_{tmp} = reach\_blocks(\hat{\mathcal{X}}(0), \hat{\mathcal{X}}_{tmp}, Q, \hat{\mathcal{V}}_{tmp}, blocks, all\_blocks)$ ;
15  if  $!isdisjoint(\hat{\mathcal{X}}_{tmp}, constraints)$  then
16    // compute high-dimensional sets only if
17    // intersection with guards is nonempty
18     $\hat{\mathcal{X}}_{tmp} = reach\_blocks(\hat{\mathcal{X}}(0), \hat{\mathcal{X}}_{tmp}, Q, \hat{\mathcal{V}}_{tmp}, other\_blocks,$ 
19     $all\_blocks)$ ;
20  end
21   $\hat{\mathcal{X}}(k) \leftarrow \hat{\mathcal{X}}_{tmp}$ ;
22   $P \leftarrow Q$ ;
23   $Q \leftarrow Q \cdot \Phi$ ;
24 end
25 return  $[\hat{\mathcal{X}}(0), \hat{\mathcal{X}}(1), \dots, \hat{\mathcal{X}}(N - 1)]$ ;

```

function $reach_blocks(\hat{\mathcal{X}}(0), \hat{\mathcal{X}}_{tmp}, Q, \hat{\mathcal{V}}_{tmp}, blocks, all_blocks)$

```

26 for  $b_i \in blocks$  do
27    $\hat{\mathcal{X}}_{tmp}[b_i] \leftarrow \{0_{dim(b_i)}\}$ ;
28   for  $b_j \in all\_blocks$  do
29      $\hat{\mathcal{X}}_{tmp}[b_j] \leftarrow \hat{\mathcal{X}}_{tmp}[b_j] \oplus Q[b_i, b_j] \odot \hat{\mathcal{X}}(0)[b_j]$ ;
30   end
31    $\hat{\mathcal{X}}_{tmp}[b_i] \leftarrow \hat{\mathcal{X}}_{tmp}[b_i] \oplus \hat{\mathcal{V}}_{tmp}[b_i]$ ;
32 end
33 return  $\hat{\mathcal{X}}_{tmp}$ ;

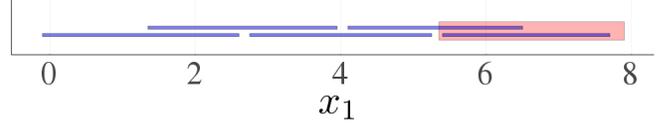
```

set to obtain $\hat{\mathcal{X}}(0)$. Next we want to compute a flowpipe, and this is where we deviate from the previous algorithm.

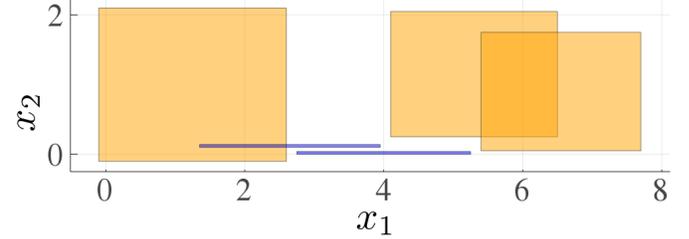
A. Computing a sparse flowpipe

We modify $Post_C^\square$ in order to control the dimensions of the flowpipe. Recall that the black-box version of $Post_C^\square$ computed the flowpipe $\hat{\mathcal{X}}(k) = \mathcal{X}_1(k) \times \dots \times \mathcal{X}_\beta(k)$ for $k = 1, \dots, N$, i.e., in all dimensions. This is usually not necessary for detecting an intersection with a guard. We will discuss this formally below, but want to establish some intuition first. Recall the running example from before. The guard was only constrained in dimension x_1 . This means that the bounds of the sets $\hat{\mathcal{X}}(k) = \mathcal{X}_1(k) \times \mathcal{X}_2(k)$ in dimension x_2 are irrelevant. Consequently, we do not need to compute the sets $\mathcal{X}_2(k)$ at all (at least for the moment). We only compute those dimensions of a flowpipe that are necessary to determine intersection with the guards. Identifying these dimensions and projecting the guards accordingly has to be performed only once per transition and is often just a syntactic procedure.

In particular, Algorithm 1 describes the new implementation of $Post_C^\square$ to compute a sparse flowpipe. The algorithm starts



(a) The flowpipe from Figure 1 in dimension x_1 only consists of intervals (blue). The linear constraint \mathcal{G}_1 (red) is the guard \mathcal{G} projected to x_1 . For better visibility, we draw the sets thicker and add a slight offset to some of the intervals.



(b) Illustration of the effective flowpipe computation from Figure 1. Only the first set $\hat{\mathcal{X}}(0)$ and the sets that intersect with the guard ($\hat{\mathcal{X}}(3)$ and $\hat{\mathcal{X}}(4)$) are computed in high dimensions.

Fig. 5. Illustration of the mixed sparse and high-dimensional flowpipe construction.

the same way as the original algorithm in [2] with a decomposed set in line 2 and iteratively computes a sparse flowpipe only for the dimensions of interest (line 14). However, if we detect an intersection with a guard constraint (line 15), we compute the full-dimensional flowpipe for the corresponding time frame. The computation of the inputs $\hat{\mathcal{V}}_{tmp}$ remains the same as in the original algorithm.

Example: As discussed, we only compute the flowpipe $\mathcal{X}_1(1), \dots, \mathcal{X}_1(4)$ for the first block (in dimension x_1), i.e., a sequence of intervals, which is depicted in Figure 5(a). Projecting the guard to x_1 , we obtain a ray \mathcal{G}_1 . As expected, we observe an intersection with the guard for the same time steps as before (namely steps $k = 3$ and $k = 4$).

B. Decomposing an intersection

Computing the intersection of two n -dimensional sets \mathcal{X} and \mathcal{G} in low dimensions is generally beneficial for performance; yet it is particularly interesting if one of the sets is a polytope that is not represented by its linear constraints. Common cases are the V-representation or zonotopes represented by their generators, which are used in several approaches [27], [28], [29], [30]. To compute the (exact) intersection of such a polytope \mathcal{X} with a polyhedron \mathcal{G} in H-representation, \mathcal{X} needs to be converted to H-representation first. A polytope with m vertices can have $O\left(\binom{m-n/2}{m-n}\right)$ linear constraints [42]. (For two polytopes in V-representation in general position there is a polynomial-time intersection algorithm [43], but this assumption is not practical.) A zonotope with m generators can have $O\left(m\binom{m}{n-1}\right)$ linear constraints [29]. If \mathcal{G} is a polytope in H-representation, checking disjointness of \mathcal{X} and \mathcal{G} can also be solved more efficiently in low dimensions, e.g., for m linear constraints in $O(m)$ for $n \leq 3$ [44].

Next we discuss how to apply low-dimensional reasoning to the intersection $\hat{\mathcal{X}} \cap \mathcal{G}$ of a decomposed set $\hat{\mathcal{X}}$ and a polyhedron

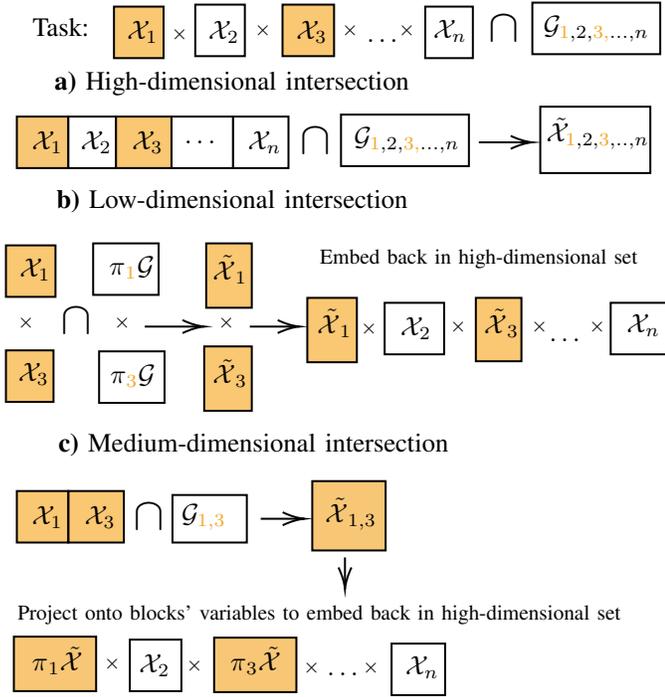


Fig. 6. Illustration of the different intersection algorithms. The task (top row) is to compute the intersection of a decomposed set $\hat{\mathcal{X}}$ with a guard \mathcal{G} that is constrained in the blocks 1 and 3 (marked in orange). We write $\tilde{\mathcal{X}}$ for the new set after an intersection operation with set \mathcal{X} . **a)** The traditional high-dimensional intersection does not make use of the decomposed structure of $\hat{\mathcal{X}}$. **b)** The low-dimensional intersection projects the guard to each block, computes the block-wise intersection, and embeds the result in high dimensions. **c)** The medium-dimensional intersection projects the guard to the constrained dimensions, computes the intersection with the Cartesian product of the corresponding sets (here: \mathcal{X}_1 and \mathcal{X}_3), projects the result to the original block structure, and finally embeds this result in high dimensions.

\mathcal{G} , or respectively detect emptiness of the intersection $\hat{\mathcal{X}} \cap \mathcal{G}$ (which can often be achieved more efficiently). The key idea is to exploit that $\hat{\mathcal{X}}$ is decomposed. For ease of discussion, we consider the case of two blocks (i.e., $\hat{\mathcal{X}} = \mathcal{X}_1 \times \mathcal{X}_2$). Below we discuss the two cases that \mathcal{G} is decomposed or not.

Intersection of two decomposed sets: We first consider the case that \mathcal{G} is also decomposed and agrees with $\hat{\mathcal{X}}$ on the block structure, i.e., $\mathcal{G} = \hat{\mathcal{G}} = \mathcal{G}_1 \times \mathcal{G}_2$ and $\mathcal{X}_1, \mathcal{G}_1 \subseteq \mathbb{R}^{n_1}$ for some n_1 . Clearly

$$\hat{\mathcal{X}} \cap \hat{\mathcal{G}} = (\mathcal{X}_1 \times \mathcal{X}_2) \cap (\mathcal{G}_1 \times \mathcal{G}_2) = (\mathcal{X}_1 \cap \mathcal{G}_1) \times (\mathcal{X}_2 \cap \mathcal{G}_2)$$

because the Cartesian product and intersection distribute; thus

$$\hat{\mathcal{X}} \cap \hat{\mathcal{G}} = \emptyset \iff (\mathcal{X}_1 \cap \mathcal{G}_1 = \emptyset) \vee (\mathcal{X}_2 \cap \mathcal{G}_2 = \emptyset). \quad (5)$$

Now consider the second disjunct in (5) and assume that \mathcal{G}_2 is universal. We get $\mathcal{X}_2 \cap \mathcal{G}_2 = \emptyset \iff \mathcal{X}_2 = \emptyset$. In our context, $\hat{\mathcal{X}}$ (and hence \mathcal{X}_2) is nonempty by construction. Hence (5) simplifies to

$$\hat{\mathcal{X}} \cap \hat{\mathcal{G}} = \emptyset \iff \mathcal{X}_1 \cap \mathcal{G}_1 = \emptyset,$$

so we *never* need to compute \mathcal{X}_2 to determine whether the intersection is empty. In practice, the set \mathcal{G}^* from (4) takes the

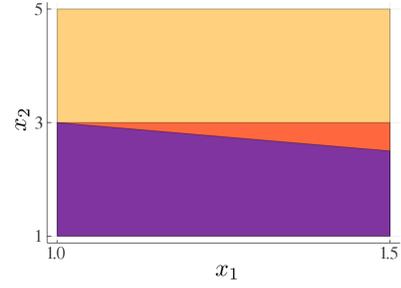


Fig. 7. Example of performing the different intersection algorithms. We intersect a three-dimensional hyperrectangle with the ranges $1 \leq x_1 \leq 5$, $1 \leq x_2 \leq 5$, $1 \leq x_3 \leq 5$ and a polyhedral guard with the linear constraints $x_1 + x_2 \leq 4$ and $x_1 \leq 1.5$ (observe that dimension x_3 is unconstrained). The exact intersection is the purple set, obtained using the high-dimensional intersection. The red set corresponds to the medium-dimensional intersection. The orange set corresponds to low-dimensional intersection.

role of $\hat{\mathcal{G}}$ and is often only constrained in *some* dimensions (and hence decomposed and universal in all other dimensions). We illustrate this algorithm in Figure 6b).

Intersection of a decomposed and a non-decomposed set: If \mathcal{G} is not decomposed in the same block structure as \mathcal{X} , we can still decompose it, at the cost of an approximation error. Let π_1 and π_2 be suitable projection matrices. Then

$$\hat{\mathcal{X}} \cap \mathcal{G} \subseteq (\mathcal{X}_1 \cap \pi_1 \mathcal{G}) \times (\mathcal{X}_2 \cap \pi_2 \mathcal{G})$$

and hence

$$\begin{aligned} \hat{\mathcal{X}} \cap \mathcal{G} = \emptyset &\iff (\mathcal{X}_1 \cap \pi_1 \mathcal{G} = \emptyset) \vee (\mathcal{X}_2 \cap \pi_2 \mathcal{G} = \emptyset) \\ &\iff \mathcal{X}_1 \cap \pi_1 \mathcal{G} = \emptyset. \end{aligned} \quad (6)$$

From (6) we obtain 1) a sufficient test for emptiness of $\hat{\mathcal{X}} \cap \mathcal{G}$ in terms of only \mathcal{X}_1 and 2) a more precise sufficient test in terms of \mathcal{X}_1 and \mathcal{X}_2 in low dimensions. If both tests fail, we can either fall back to the (exact) test in high dimensions or conservatively assume that the intersection is nonempty.

The precision of the above scheme highly depends on the structure of $\hat{\mathcal{X}}$ and \mathcal{G} . If several (but not all) blocks of \mathcal{G} are constrained, instead of decomposing \mathcal{G} into the low-dimensional block structure, one can alternatively compute the intersection for medium-dimensional sets to avoid an approximation error; we apply this strategy in the evaluation (Section V). If \mathcal{G} is compact, the following proposition shows that the approximation error is bounded by the maximal entry in the diameters of $\hat{\mathcal{X}}$ and \mathcal{G} , and this bound is tight. This strategy is illustrated in Figure 6c).

Proposition 2: Let $\hat{\mathcal{X}} = \times_j \mathcal{X}_j \in \mathcal{C}_n$, $\mathcal{G} \in \mathcal{C}_n$, $\hat{\mathcal{X}} \cap \mathcal{G} \neq \emptyset$, $\hat{\mathcal{G}} := \times_j \pi_j \mathcal{G}$ for appropriate projection matrices π_j corresponding to \mathcal{X}_j , and $p = \infty$. Then

$$d_H^p(\hat{\mathcal{X}} \cap \mathcal{G}, \hat{\mathcal{X}} \cap \hat{\mathcal{G}}) \leq \max_j \min(\|\Delta_p(\mathcal{X}_j)\|_p, \|\Delta_p(\pi_j \mathcal{G})\|_p).$$

Example: Consider Figure 5(b). Previously we have already identified the intersection with the flowpipe for time steps $k = 3$ and $k = 4$. The resulting sets are $\hat{\mathcal{X}}(k) \cap \mathcal{G} = \mathcal{X}_1(k) \cap \mathcal{G}_1 \times \mathcal{X}_2(k)$, where \mathcal{G}_1 was the projection of \mathcal{G} to x_1 . We emphasize that we compute the intersections in low

dimensions, that we need not compute $\mathcal{X}_2(1)$ and $\mathcal{X}_2(2)$ at all, and that in this example all computations are exact (i.e., we obtain the same sets as in Figure 3).

Now consider the case that \mathcal{G} is not decomposed in the same structure as $\widehat{\mathcal{X}}$, e.g., the hyperplane $x_1 = x_2$. One option is to decompose \mathcal{G} to the blocks of $\widehat{\mathcal{X}}$, i.e., $\mathcal{G}_1 := \pi_1\mathcal{G}$, $\mathcal{G}_2 := \pi_2\mathcal{G}$ and compute the block-wise intersection $(\mathcal{X}_1(k) \cap \mathcal{G}_1) \times (\mathcal{X}_2(k) \cap \mathcal{G}_2)$. Here \mathcal{G}_1 and \mathcal{G}_2 are universal, so we obtain a coarse approximation of the intersection (namely $\widehat{\mathcal{X}}(k)$ itself). Alternatively, computing the intersection $\widehat{\mathcal{X}}(k) \cap \mathcal{G}$ is exact but computationally more demanding. If we assume that the system has higher dimension, e.g., 10, then computing the intersection in two dimensions (i.e., with a 2D projection $(\mathcal{X}_1(k) \times \mathcal{X}_2(k)) \cap \pi\mathcal{G}$) is still exact and yet cheaper than computing the intersection in full dimensions (i.e., $\widehat{\mathcal{X}}(k) \cap \mathcal{G}$).

We exemplify the possible difference between the intersection algorithms in Figure 7. There we compare the three algorithms visualized in Figure 6. The high-dimensional intersection is the most precise as expected because it does not suffer from a projection error, however, the result is a high-dimensional set, instead of a decomposed high-dimensional set. In addition, this algorithm intersects the sets in three dimensions (including x_3) while the other algorithms ignore the unconstrained dimensions. The medium-dimensional intersection corresponds to the box approximation of the true intersection. We observe that if we cannot decompose the polyhedron to the same block structure as the decomposed set without a projection error (e.g. projecting $x_1 + x_2 \leq 4$ onto x_1 and x_2 individually results in one-dimensional universes), then the low-dimensional intersection can produce a large approximation error. Thanks to the second linear constraint $x_1 \leq 1.5$ and the fact that we treat each half-space separately, we are still able to compute a nontrivial intersection.

C. Decomposing an affine map

The next step after computing the intersection with the guard is the application of the assignment. We consider an affine map $A\widehat{\mathcal{X}} \oplus \{b\}$ with $A \in \mathbb{R}^{n \times n}$ and $b \in \mathbb{R}^n$. Affine-map decomposition has already been presented as part of the operator $Post_C^\square$ [2]:

$$A\widehat{\mathcal{X}} \oplus \{b\} \subseteq \times_i \bigoplus_j A_{ij} \mathcal{X}_j \oplus \{b_i\}$$

where A_{ij} is the block in the i -th block row and the j -th block column. We recall an error estimation.

Proposition 3: [2, Prop. 3] Let $\mathcal{X} = \times_{j=1}^\beta \mathcal{X}_j \in \mathcal{C}_n$ be nonempty, $A \in \mathbb{R}^{n \times n}$, $q_j := \arg \max_i \|A_{ij}\|_p$ (the index of the block with the largest matrix norm in the j -th block column) so that $\alpha_j := \max_{i \neq q_j} \|A_{ij}\|_p$ is the second largest matrix norm in the j -th block column. Let $\alpha_{\max} := \max_j \alpha_j$ and $\Delta_{\text{sum}} := \sum_j \Delta_\infty(\mathcal{X}_j)$. Then

$$\begin{aligned} & d_H^p(A\mathcal{X}, \times_i \bigoplus_j A_{ij} \mathcal{X}_j) \\ &= \max_{\|d\|_p \leq 1} \sum_{i,j} \rho_{\mathcal{X}_j}(A_{ij}^T d_i) - \rho_{\mathcal{X}_j}(\sum_k A_{kj}^T d_k) \\ &\leq (\beta - 1) \sum_j \alpha_j \Delta_\infty(\mathcal{X}_j) \leq \frac{n}{2} \alpha_{\max} \Delta_{\text{sum}}. \end{aligned}$$

In particular, if only one block per block column of matrix A is nonzero, the approximation is exact [2]. For example, consider a two-block scenario and a block-diagonal matrix A , i.e., $A_{12} = A_{21} = 0$. Then

$$\begin{aligned} & \begin{pmatrix} A_{11} & 0 \\ 0 & A_{22} \end{pmatrix} \mathcal{X}_1 \times \mathcal{X}_2 \oplus \{b_1\} \times \{b_2\} \\ &= (A_{11}\mathcal{X}_1 \oplus \{b_1\}) \times (A_{22}\mathcal{X}_2 \oplus \{b_2\}). \end{aligned}$$

In practice, affine maps with such a structure are unrealistic for the $Post_C^\square$ operator but typical for assignments. Prominent cases include resets, translations, and scalings, for which A is even diagonal and hence block diagonal for any block structure.

Example: Recall that, after computing the intersections, we ended up with the same sets as in Figure 3. In our example, the assignment is a translation in dimension x_1 . Hence, as mentioned above, the application of the decomposed assignment is also exact. In particular, the translation only affects $\mathcal{X}_1(k)$ and we obtain the same result as in Figure 3.

D. Inclusion check for decomposed sets

Our algorithm is now fully able to take transitions. Observe that *all* sets ever occurring in scheme (3) using the algorithm are decomposed. The following proposition gives an exact low-dimensional fixpoint check under this condition.

Proposition 4: Let $\widehat{\mathcal{X}} = \times_j \mathcal{X}_j \in \mathcal{C}_n$, $\widehat{\mathcal{G}} = \times_j \mathcal{G}_j \in \mathcal{C}_n$ be nonempty sets with identical block structure. Then

$$\widehat{\mathcal{X}} \subseteq \widehat{\mathcal{G}} \iff \bigwedge_j \mathcal{X}_j \subseteq \mathcal{G}_j.$$

E. Decomposing a convex hull

As the last part of the algorithm, we decompose the computation of the convex hull. We exploit that all sets in the same flowpipe share the same block structure.

Proposition 5: Let $\widehat{\mathcal{X}} = \times_j \mathcal{X}_j \in \mathcal{C}_n$, $\widehat{\mathcal{G}} = \times_j \mathcal{G}_j \in \mathcal{C}_n$ be nonempty sets with identical block structure. Then

$$\text{CH}(\widehat{\mathcal{X}} \cup \widehat{\mathcal{G}}) \subseteq \times_j \text{CH}(\mathcal{X}_j \cup \mathcal{G}_j).$$

For the decomposition operations proposed before (intersection, affine map, and inclusion), there are common cases where the approximations were exact. In these cases it is always beneficial to perform the decomposed operations instead of the high-dimensional counterparts. The decomposition of the convex hull, however, always incurs an approximation error, which we can bound by the radius of the box approximation and by the block-wise difference in bounds.

Proposition 6: Let $\widehat{\mathcal{X}} = \times_j \mathcal{X}_j \in \mathcal{C}_n$, $\widehat{\mathcal{G}} = \times_j \mathcal{G}_j \in \mathcal{C}_n$ be nonempty sets with identical block structure and let r^∞ be the radius of the box approximation of $\text{CH}(\widehat{\mathcal{X}} \cup \widehat{\mathcal{G}})$. Then

$$\begin{aligned} & d_H^p(\text{CH}(\widehat{\mathcal{X}} \cup \widehat{\mathcal{G}}), \times_j \text{CH}(\mathcal{X}_j \cup \mathcal{G}_j)) \\ &\leq \min \left(\|r^\infty\|_\infty, \max_{\|d\|_p \leq 1} \sum_j |\rho_{\mathcal{X}_j}(d_j) - \rho_{\mathcal{G}_j}(d_j)| \right). \end{aligned}$$

Example: Figure 4 shows the decomposed convex hull of the sets $\widehat{\mathcal{X}}(3) \cap \mathcal{G}$ and $\widehat{\mathcal{X}}(4) \cap \mathcal{G}$ after applying the translation. Since each block is one-dimensional in our example, we obtain the box approximation.

F. Discussion

The combination of the decomposed set operations outlined above results in a reachability algorithm that is sound. This immediately follows from the fact that all operations compute an overapproximation.

Proposition 7 (Soundness): The analysis using Algorithm 1 and the aforementioned procedures to apply the decomposed set operations (intersection, affine map, inclusion check, and convex hull) results in an overapproximation of the reach set.

When using the above algorithm, the question about the choice of the decomposition arises. Unfortunately, there is generally no automatic way to find the “best” decomposition.

Regarding LTI systems, different decompositions generally result in different flowpipes [2]. Unless the LTI system consists of fully decoupled sub-components (which is usually not the case), the only “best” decomposition is the degenerate case of a single big block.

Regarding hybrid systems given a fixed decomposition of the LTI systems, there is indeed a unique “best” decomposition: a single block for all constrained dimensions. Finding this “best” decomposition is trivial. However, this block may be too big in terms of computational cost.

We can generally say that a finer partition in the decomposition (i.e., using more and smaller blocks) only ever degrades precision. However, while lower-dimensional set operations have lower computational cost, a decrease in precision may actually affect the rest of the algorithm in a negative way. For instance, we may not be able to prove that a transition cannot be taken. Thus statements about computational cost are generally not possible in the case of hybrid systems.

In the next section we shall see that in practical use cases the decomposition is always beneficial and that using the block size of the constrained dimensions is a good heuristic.

Our reachability algorithm benefits greatly from a low number of constrained dimensions, especially in high-dimensional models. High-dimensional models are scarcely available, but those models that we found supported the hypothesis that this number is typically low. In particular, the ARCH-COMP competition [9] represents the state of the art in reachability analysis and considers a scalable model of a powertrain from [39] where only one dimension is constrained. In the next section we consider another scalable model with just three constrained dimensions.

V. EVALUATION

We implemented the ideas presented in Section IV in JuliaReach [4], [3]. The code is publicly available [4]. We performed the experiments presented in this section on a Mac notebook with an Intel i5 CPU@3.1 GHz and 16 GB RAM.

A. Benchmark descriptions

We evaluate our implementation on a number of benchmarks taken from the HyPro model library [45], from the ARCH-COMP 2019 competition [5], and a scalable model from [9]. To demonstrate the qualitative performance of our approach, we verify a safety property for each benchmark,

which requires precise approximations in each step of the algorithm. We briefly describe the benchmarks below.

Linear switching system. This five-dimensional model taken from [45] is a linear hybrid system with five locations of different controlled, randomly-generated continuous dynamics stabilized by an LQR controller. The discrete structure has a ring topology with transitions determined heuristically from simulations. The safety property for this system is $x_1 > -1.2$.

Spacecraft rendezvous. This model with five dimensions represents a spacecraft steering toward a passive target in space [46]. We use a linearized version of this model with three locations. We consider two scenarios, one where the spacecraft successfully approaches the target and another one where a mission abort occurs at $t = 120$ min. For the safety properties we refer to [5].

Platooning. This ten-dimensional model with two locations represents a platoon of three vehicles with communication loss at deterministic times [47]. The safety property enforces a minimum distance d between the vehicles: $\bigwedge_{x \in \{x_1, x_4, x_7\}} x \geq -d$. We consider both a time-bounded setting with $d = 42$ and a time-unbounded setting with $d = 50$; note that in the latter setting a fixpoint must be found.

Filtered oscillator. This scalable model consists of a two-dimensional switched oscillator (dimensions x and y) and a parametric number of filters (here: 64–1024) which smooth x [9]. We fixed the maximum number of transitions to five by adding a new variable. The safety property is $y < 0.5$.

B. Tool descriptions

We compare our implementation in JuliaReach to two other algorithms available in the same tool. All three algorithms use the same decomposition-based continuous-post operator $Post_C^\square$ [3], so the main difference between these algorithms is the intersection operation with discrete jumps, which allows for a direct comparison of the approach presented in this paper. The existing implementations can be seen as instances of the algorithm in Section III. Furthermore, we compare the implementation to two algorithms available in SpaceEx [9], which is an efficient and mature verification tool for linear hybrid systems. We summarize the different approaches below.

Deco. This algorithm implements the decomposed approach presented in this paper. To compute the (low-dimensional) intersections, we use a polyhedra library [48].

LazySupp. This algorithm uses a (lazy) support-function-based approximation of the intersection operation using the simple heuristic $\rho_{X \cap Y}(\ell) \leq \min(\rho_X(\ell), \rho_Y(\ell))$. This heuristic is fast but not precise enough to verify the safety property of the *filtered oscillator* model.

LazyOptim. In contrast to the coarse intersection in LazySupp, this algorithm uses a more precise implementation of the intersection operation based on line search [26].

SpaceEx LGG. This is an efficient implementation of the algorithm by Le Guernic and Girard [35]. The algorithm uses a support-function representation and can hence check low-dimensional conditions efficiently, but operations such as intersections work in high dimensions.

SpaceEx STC. The STC algorithm is an extension of the LGG algorithm with automatic time-step adaptation [49].

TABLE I. For each benchmark we report the number of dimension and the number of constrained dimensions (*Dim. (constr.)*), the step size of the time discretization (*Step*) and the run time of the different algorithms as described in Section V-B (e.g., “Deco” refers to the decomposed algorithm presented in this paper) in seconds (where the fastest solution is marked in bold face). For benchmark instances with parenthesized computation time, the safety property could not be proven because the overapproximation was too coarse, in which case the computation terminated as soon as the violation was detected. “TO” and “OOM” indicate a timeout of 5×10^4 seconds and an out-of-memory error, respectively.

Benchmark	Dim. (constr.)	Step	Deco	LazySupp	LazyOptim	SpaceEx LGG	SpaceEx STC
linear_switching	5 (1)	0.0001	2.50×10^0	1.27×10^1	2.81×10^1	2.60×10^1	2.30×10^1
spacecraft_noabort	5 (4)	0.04	5.30×10^0	3.42×10^0	2.19×10^2	1.18×10^0	3.50×10^{-1}
spacecraft_120	5 (5)	0.04	5.30×10^0	2.10×10^0	4.30×10^1	1.91×10^0	8.10×10^{-1}
platoon_bounded	10 (4)	0.01	1.30×10^{-1}	1.60×10^{-1}	5.69×10^0	5.55×10^0	1.60×10^0
platoon_unbounded	10 (4)	0.03	1.08×10^0	1.16×10^0	4.96×10^1	3.46×10^1	6.50×10^1
filtered_osc64	67 (3)	0.01	2.81×10^0	$(7.43 \times 10^0)^\dagger$	5.63×10^2	2.04×10^1	3.25×10^1
filtered_osc128	131 (3)	0.01	7.95×10^0	$(4.29 \times 10^1)^\dagger$	1.79×10^3	1.69×10^2	4.67×10^3
filtered_osc256	259 (3)	0.01	2.80×10^1	$(9.19 \times 10^1)^\dagger$	9.99×10^4	8.70×10^3	OOM
filtered_osc512	515 (3)	0.01	1.13×10^2	$(4.73 \times 10^2)^\dagger$	TO	TO	TO
filtered_osc1024	1027 (3)	0.01	5.09×10^2	$(5.11 \times 10^3)^\dagger$	TO	TO	TO

[†]The safety property could not be proven by this tool because the overapproximation was too coarse.

All tools use a support-function representation of sets. We use template polyhedra with box constraints to overapproximate these sets, which roughly corresponds to a decomposition into one-dimensional blocks. This approximation is fast.

For the algorithms implemented in JuliaReach, we use one-dimensional block structures for all models. For SpaceEx we use the options given in the benchmark sources. For algorithms that use a fixed time step, we fix this parameter to the same value for each benchmark. The SpaceEx STC algorithm does not have such a parameter, so we instead fix the parameter “flowpipe tolerance,” which controls the approximation error, to the following values: *Linear switching system*: 0.01, *Spacecraft rendezvous*: 0.2, *Platooning*: 1, *Filtered oscillator*: 0.05.

C. Experimental results

The results are presented in Table I. We generally observe a performance boost of the Deco algorithm for models with more than five dimensions, and only a minor overhead for “small” models. This demonstrates the general scalability improvement by performing operations, especially the intersection, in low dimensions. Since all models have a small number of constrained dimensions in their guards and invariants, choosing an appropriate block structure results in very low-dimensional sets for computing the intersections, for which concrete polyhedral computations are very efficient and most precise. We note that such an intersection computation does not scale with the dimension, and so other algorithms must resort to approximation techniques.

Moreover, we found that our approach scales more favorably compared to the high-dimensional approaches when decreasing the time step δ (cf. Section III-A). We demonstrate this observation for the *filtered oscillator* model in Table II and explain it as follows. Recall that we only compute those sets in the flowpipe in high dimensions for which we have detected an intersection in low dimensions. With a smaller time step, the total number of sets increases and hence the savings due to our approach become more dominant. To give an example, for the *filtered oscillator* with time step 0.0005, out of the 9,661 sets in total we only computed 1,400 sets in high dimensions.

TABLE II. Evaluation for different time steps on the *filtered oscillator* model with 64 filters (“filtered_osc64” in Table I) and all algorithms (see Section V-B) that allow varying the time step. The last row shows the relative change of the fourth row over the first row as base line.

Step	Deco	LazySupp	LazyOptim	SpaceEx LGG
0.01	2.8×10^0	7.4×10^0	5.6×10^2	2.0×10^1
0.005	4.1×10^0	1.3×10^1	9.5×10^2	3.6×10^1
0.001	1.7×10^1	6.5×10^1	4.7×10^3	1.6×10^2
0.0005	3.2×10^1	1.3×10^2	8.4×10^3	3.2×10^2
$\times 20.0$	$\times 11.4$	$\times 17.5$	$\times 15.0$	$\times 16.0$

Consequently, making the time step 20 times smaller only makes the run time 11.4 times slower for our algorithm as opposed to at least 15 times slower for other algorithms.

Furthermore, while in general our algorithm may induce an additional approximation error with the block structure, in the evaluation it was always precise enough to prove the safety properties. In Table I we report the total amount of constrained dimensions in guards, invariants, and safety properties for each benchmark instance. These are the dimensions that determine our low-dimensional flowpipes. For some of the models, the constrained dimensions differ between the invariants/guards and the safety property, but a one-dimensional block structure was still sufficient. Note that we need to compute intersections only with invariants and guards; for safety properties we just need to check inclusion.

In the *linear switching* model, only one dimension is constrained, so our algorithm, together with one-dimensional block structure, is appropriate, especially because a small time step is required to prove the property. In the *platooning* model, the safety property constrains three dimensions but the invariants and guards constrain just one dimension.

The invariants and guards of the models *spacecraft* and *filtered oscillator* constrain two dimensions, so the natural choice for the decomposition is to keep these dimensions in the same block. However, we chose to decompose into one-dimensional blocks for better run-time comparison in

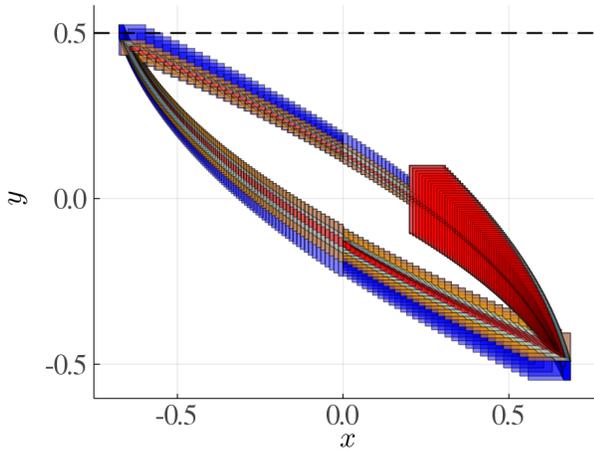


Fig. 8. Four flowpipes for the model *filtered oscillator* using a time step of 0.01. The dark blue flowpipe (low-dimensional intersection algorithm) and the orange flowpipe (medium-dimensional intersection algorithm) are obtained for a one-dimensional block structure. The light blue flowpipe is obtained for a two-dimensional block structure and octagon approximation. The red flowpipe is obtained by the tool SpaceEx (which uses high-dimensional analysis). The dashed line corresponds to the safety property $y \geq 0.5$.

Table I. In the implementation, we follow the strategy to perform the intersection in two dimensions and then project back to one dimension (cf. Section IV-B). If instead we decide to employ a two-dimensional block structure, we can further gain precision by using different template polyhedra, e.g., with octagon constraints. In Figure 8 we visualize this alternative for the *filtered oscillator*. As expected, the one-dimensional analysis is less precise as it corresponds to a box approximation. However, we note that a one-dimensional block structure also inherently reduces the precision of $Post_C^\square$, so the additional approximation error does not only stem from the handling of discrete transitions by to our approach. In addition, one can observe that the approximation using a one-dimensional block structure with low-dimensional intersection is coarser compared with the medium-dimensional strategy and we cannot prove the safety property anymore.

D. Scaling the number of constrained dimensions

We now investigate the effect of increasing the number of constrained dimensions in two experiments.

In the first experiment we focus on the intersection operation. We fix a hypercube \mathcal{X} of dimension 1,024, of radius 4 and centered in the origin, and a half-space \mathcal{G} $x_1 + \dots + x_k \leq 2$ for parameter k that controls the number of constrained dimensions. The half-space properly cuts the hypercube for any k . In Figure 9 we show the run times for different intersection algorithms. We note that these algorithms compute different approximations of the true intersection: the “DecoLow” and “DecoMedium” algorithms implement the respective ideas from Figure 6; the “LazySupp” and “LazyOptim” algorithms use the ideas as in the reachability algorithms of the same name (described in Section V-B); the “Concrete” algorithm uses a polyhedra library (note that the H-representation of a

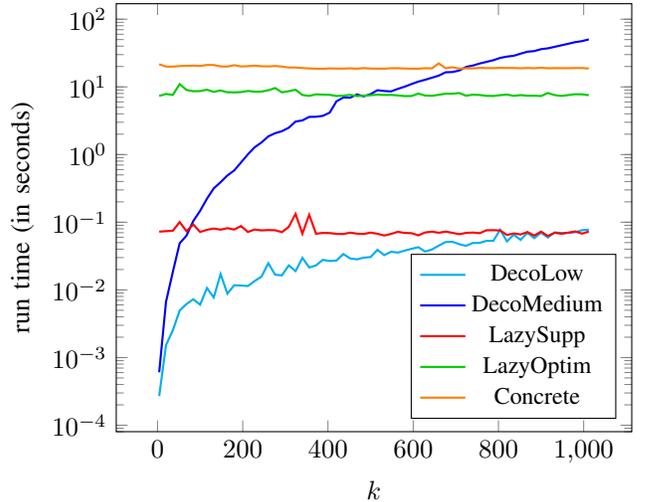


Fig. 9. Scaling the number of constrained dimensions k for an intersection $\mathcal{X} \cap \mathcal{G}$ where \mathcal{X} is a hypercube and \mathcal{G} is a half-space.

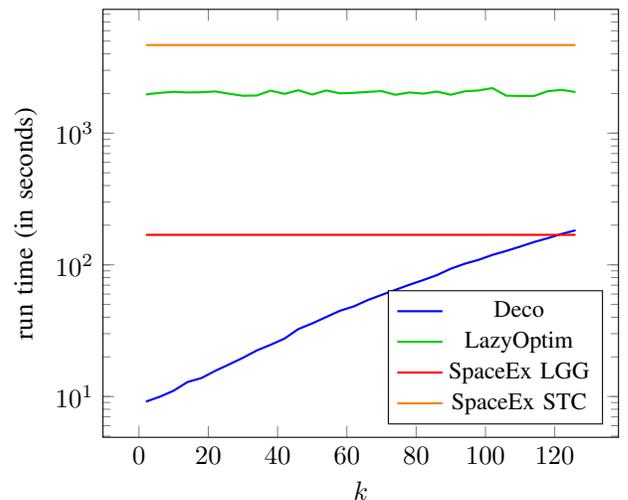


Fig. 10. Scaling the number of constrained dimensions k for the modified *filtered_osc128* benchmark.

hypercube of dimension n has only $2n$ constraints). We can see that the high-dimensional intersection algorithms are not affected by the number of constrained dimensions k whereas the decomposition-based algorithms scale gracefully with k .

In the second experiment we consider the full reachability setting. For that we modify the *filtered oscillator* benchmark with 128 filters (*filtered_osc128* in Table I) by adding small nonzero entries to k previously unconstrained dimensions in the invariants and guards. We consider all the reachability algorithms that are precise enough to verify the safety property (which is still satisfied in our modified benchmark instances), i.e., all the algorithms from Table I except for *LazySupp*. Note that for SpaceEx algorithms we extrapolate the results we report for this benchmark instance in Table I to a range of values of parameter k as the reachability algorithms implemented in SpaceEx do not depend on the number of constrained dimensions. The run times are plotted in Figure 10 as a function of k . Again the high-dimensional algorithms are not

affected by the number k whereas the decomposition-based algorithm scales well with k (note the log scale). For $k \approx 120$ (close to $n = 131$) we see a cross-over with the *SpaceEx LGG* run time, which is expected due to the overhead of the decomposition.

VI. CONCLUSION

We have presented a schema that integrates a reachability algorithm based on decomposition for LTI systems in the analysis loop for linear hybrid systems. The key insight is that intersections with polyhedral constraints can be efficiently detected and computed (approximately or often even exactly) in low dimensions. This enables the systematic focus on appropriate subspaces and the potential for bypassing large amounts of flowpipe computations. Moreover, working with sets in low dimensions allows to use precise polyhedral computations that are infeasible in high dimensions.

Future work: An essential step in our algorithm is the fast computation of a low-dimensional flowpipe for the detection of intersections. In the presented algorithm, we recompute the flowpipe for the relevant time frames in high dimensions using the same decomposed algorithm with the same time step. However, this is not necessary. We could achieve higher precision by using different algorithmic parameters or even a different, possibly non-decomposed, algorithm (e.g., one that features arbitrary precision [25]). This is particularly promising for LTI systems because one can avoid recomputing the homogeneous (state-based) part of the flowpipe [26].

In the benchmarks considered in the experimental evaluation, it was not necessary to change the block structure when switching between locations. In general, different locations may constrain different dimensions, so tracking the “right” dimensions may be necessary to maintain precision. While it is easy to merge different blocks, subsequent computations would become more expensive. Hence one may also want to split blocks again for optimal performance. Since this comes with a loss in precision, heuristics for rearranging the block structure, possibly in a refinement loop, are needed.

REFERENCES

- [1] T. A. Henzinger, “The theory of hybrid automata,” in *LICS*, 1996, pp. 278–292.
- [2] S. Bogomolov, M. Forets, G. Frehse, A. Podelski, C. Schilling, and F. Viry, “Reach set approximation through decomposition with low-dimensional sets and high-dimensional matrices,” in *HSCC*, 2018, pp. 41–50.
- [3] S. Bogomolov, M. Forets, G. Frehse, K. Potomkin, and C. Schilling, “JuliaReach: a toolbox for set-based reachability,” in *HSCC*, 2019, pp. 39–44.
- [4] “JuliaReach,” <https://github.com/JuliaReach>, 2020.
- [5] M. Althoff, S. Bak, M. Forets, G. Frehse, N. Kochdumper, R. Ray, C. Schilling, and S. Schupp, “ARCH-COMP19 category report: Continuous and hybrid systems with linear continuous dynamics,” in *ARCH19*, 2019, pp. 14–40.
- [6] M. Althoff, “An introduction to CORA 2015,” in *ARCH*, 2015, pp. 120–151.
- [7] S. Schupp, E. Ábrahám, I. B. Makhlof, and S. Kowalewski, “Hypro: A C++ library of state set representations for hybrid systems reachability analysis,” in *NFM*, 2017, pp. 288–294.
- [8] S. Bak and P. S. Duggirala, “Hylaa: A tool for computing simulation-equivalent reachability for linear systems,” in *HSCC*, 2017, pp. 173–178.
- [9] G. Frehse, C. L. Guernic, A. Donzé, S. Cotton, R. Ray, O. Lebeltel, R. Ripado, A. Girard, T. Dang, and O. Maler, “SpaceEx: Scalable verification of hybrid systems,” in *CAV*, 2011, pp. 379–395.
- [10] A. Gurung, R. Ray, E. Bartocci, S. Bogomolov, and R. Grosu, “Parallel reachability analysis of hybrid systems in xspeed,” *Int. J. Softw. Technol. Transf.*, vol. 21, no. 4, pp. 401–423, 2019.
- [11] L. Bu and X. Li, “Path-oriented bounded reachability analysis of composed linear hybrid systems,” *STTT*, no. 4, pp. 307–317, 2011.
- [12] S. Schupp, J. Nellen, and E. Ábrahám, “Divide and conquer: Variable set separation in hybrid systems reachability analysis,” in *QAPL@ETAPS*, 2017, pp. 1–14.
- [13] Z. Han and B. H. Krogh, “Reachability analysis of large-scale affine systems using low-dimensional polytopes,” in *HSCC*, 2006, pp. 287–301.
- [14] A. L. Dontchev, “Time-scale decomposition of the reachable set of constrained linear systems,” *MCSS*, no. 3, pp. 327–340, 1992.
- [15] E. V. Goncharova and A. I. Ovseevich, “Asymptotics for singularly perturbed reachable sets,” in *LSSC*, 2009, pp. 280–285.
- [16] S. Kaynama and M. Oishi, “Overapproximating the reachable sets of LTI systems through a similarity transformation,” in *ACC*, 2010, pp. 1874–1879.
- [17] —, “Complexity reduction through a Schur-based decomposition for reachability analysis of linear time-invariant systems,” *Int. J. Control*, no. 1, pp. 165–179, 2011.
- [18] M. R. Greenstreet and I. Mitchell, “Reachability analysis using polygonal projections,” in *HSCC*, 1999, pp. 103–116.
- [19] C. Yan and M. R. Greenstreet, “Faster projection based methods for circuit level verification,” in *ASP-DAC*, 2008, pp. 410–415.
- [20] Y. Seladji and O. Bouissou, “Numerical abstract domain using support functions,” in *NASA Formal Methods*, 2013, pp. 155–169.
- [21] E. Asarin and T. Dang, “Abstraction by projection and application to multi-affine systems,” in *HSCC*, 2004.
- [22] I. M. Mitchell and C. Tomlin, “Overapproximating reachable sets by Hamilton-Jacobi projections,” *J. Sci. Comput.*, no. 1-3, pp. 323–346, 2003.
- [23] M. Chen, S. L. Herbert, and C. J. Tomlin, “Exact and efficient Hamilton-Jacobi guaranteed safety analysis via system decomposition,” in *ICRA*, 2017, pp. 87–92.
- [24] X. Chen and S. Sankaranarayanan, “Decomposed reachability analysis for nonlinear systems,” in *RTSS*, 2016, pp. 13–24.
- [25] A. Girard and C. L. Guernic, “Efficient reachability analysis for linear systems using support functions,” *IFAC Proceedings Volumes*, no. 2, pp. 8966–8971, 2008.
- [26] G. Frehse and R. Ray, “Flowpipe-guard intersection for reachability computations with support functions,” in *ADHS*, no. 9, 2012, pp. 94–101.
- [27] A. Girard, “Reachability of uncertain linear systems using zonotopes,” in *HSCC*, 2005, pp. 291–305.
- [28] A. Girard and C. L. Guernic, “Zonotope/hyperplane intersection for hybrid systems reachability analysis,” in *HSCC*, 2008, pp. 215–228.
- [29] M. Althoff, O. Stursberg, and M. Buss, “Computing reachable sets of hybrid systems using a combination of zonotopes and polytopes,” *Nonlinear Analysis: Hybrid Systems*, no. 2, pp. 233–249, 2010.
- [30] M. Althoff and G. Frehse, “Combining zonotopes and support functions for efficient reachability analysis of linear systems,” in *CDC*, 2016, pp. 7439–7446.
- [31] A. O. Hamadeh and J. M. Goncalves, “Reachability analysis of continuous-time piecewise affine systems,” *Automatica*, no. 12, pp. 3189–3194, 2008.
- [32] D. S. Hochbaum and J. Naor, “Simple and fast algorithms for linear and integer programs with two variables per inequality,” *SIAM J. Comput.*, no. 6, pp. 1179–1192, 1994.
- [33] N. S. Nedialkov and M. V. Mohrenschildt, “Rigorous simulation of hybrid dynamic systems with symbolic and interval methods,” in *ACC*, 2002, pp. 140–147.
- [34] A. Chutinan and B. H. Krogh, “Computational techniques for hybrid system verification,” *IEEE Trans. Automat. Contr.*, no. 1, pp. 64–75, 2003.
- [35] C. L. Guernic and A. Girard, “Reachability analysis of hybrid systems using support functions,” in *CAV*, 2009, pp. 540–554.
- [36] C. Le Guernic, “Reachability analysis of hybrid systems with linear continuous dynamics,” Ph.D. dissertation, Université Grenoble 1 - Joseph Fourier, 2009.
- [37] G. Frehse, S. Bogomolov, M. Greitschus, T. Strump, and A. Podelski, “Eliminating spurious transitions in reachability with support functions,” in *HSCC*, 2015, pp. 149–158.

- [38] S. Bogomolov, G. Frehse, M. Giacobbe, and T. A. Henzinger, “Counterexample-guided refinement of template polyhedra,” in *TACAS*, 2017, pp. 589–606.
- [39] M. Althoff and B. H. Krogh, “Avoiding geometric intersection operations in reachability analysis of hybrid systems,” in *HSCC*, 2012, pp. 45–54.
- [40] S. Bak, S. Bogomolov, and M. Althoff, “Time-triggered conversion of guards for reachability analysis of hybrid automata,” in *FORMATS*, 2017, pp. 133–150.
- [41] R. Alur, C. Courcoubetis, T. A. Henzinger, and P. Ho, “Hybrid automata: An algorithmic approach to the specification and verification of hybrid systems,” in *Hybrid Systems*, 1992, pp. 209–229.
- [42] P. McMullen, “The maximal number of faces of a convex polytope,” *Mathematika*, pp. 179–184, 1970.
- [43] K. Fukuda, T. M. Liebbling, and C. Lütolf, “Extended convex hull,” *Comput. Geom.*, no. 1-2, pp. 13–23, 2001.
- [44] L. Barba and S. Langerman, “Optimal detection of intersections between convex polyhedra,” in *SODA*, 2015.
- [45] “HyPro Benchmark Repository,” <https://ths.rwth-aachen.de/research/projects/hypro/benchmarks-of-continuous-and-hybrid-systems/>, 2020.
- [46] N. Chan and S. Mitra, “Verifying safety of an autonomous spacecraft rendezvous mission,” in *ARCH*, 2017.
- [47] I. B. Makhlof and S. Kowalewski, “Networked cooperative platoon of vehicles for testing methods and verification tools,” in *ARCH*, 2014, pp. 37–42.
- [48] B. Legat, R. Deits, O. Evans, G. Goretkin, T. Koolen, J. Huchette, D. Oyama, M. Forets, guberger, R. Schwarz, H. Ferrolho, E. Saba, and C. Coleman, “JuliaPolyhedra/Polyhedra.jl: v0.6.5,” 2020.
- [49] G. Frehse, R. Kateja, and C. L. Guernic, “Flowpipe approximation and clustering in space-time,” in *HSCC*, 2013, pp. 203–212.

APPENDIX

Proof of Proposition 1: Recall that

$$\begin{aligned} & d_H^p(\mathcal{X}, \times_j \pi_j \mathcal{X}) \\ &= \inf_{\varepsilon \in \mathbb{R}} \left\{ \times_j \pi_j \mathcal{X} \subseteq \mathcal{X} \oplus \varepsilon \mathcal{B}_p^n \text{ and } \mathcal{X} \subseteq \times_j \pi_j \mathcal{X} \oplus \varepsilon \mathcal{B}_p^n \right\}. \end{aligned}$$

We prove that $\|r_{\mathcal{X}}^p\|_p$ is an upper bound on ε by showing the two inclusions. One inclusion is trivial because $\mathcal{X} \subseteq \times_j \pi_j \mathcal{X}$. For the other inclusion, observe that the center of the box approximation of \mathcal{X} , $c_{\mathcal{X}}^p$, lies in \mathcal{X} . Recall that $p = \infty$ and that the box approximation of \mathcal{X} , $\{c_{\mathcal{X}}^p\} \oplus \|r_{\mathcal{X}}^p\|_p \mathcal{B}_p^n$, is the worst-case decomposition of \mathcal{X} , subsuming all other decompositions.

$$\times_j \pi_j \mathcal{X} \subseteq \{c_{\mathcal{X}}^p\} \oplus \|r_{\mathcal{X}}^p\|_p \mathcal{B}_p^n \subseteq \mathcal{X} \oplus \|r_{\mathcal{X}}^p\|_p \mathcal{B}_p^n \quad \blacksquare$$

Proof of Proposition 2:

$$\begin{aligned} & d_H^p(\widehat{\mathcal{X}} \cap \widehat{\mathcal{G}}, \widehat{\mathcal{X}} \cap \widehat{\mathcal{G}}) \\ &= \max_{\|d\|_p \leq 1} \rho_{\widehat{\mathcal{X}} \cap \widehat{\mathcal{G}}}(d) - \rho_{\widehat{\mathcal{X}} \cap \widehat{\mathcal{G}}}(d) \\ &= \max_{\|d\|_p \leq 1} \max_{x \in \widehat{\mathcal{X}} \cap \widehat{\mathcal{G}}} \langle d, x \rangle - \max_{y \in \widehat{\mathcal{X}} \cap \widehat{\mathcal{G}}} \langle d, y \rangle \\ &= \max_{\|d\|_p \leq 1} \max_{x \in \widehat{\mathcal{X}} \cap \widehat{\mathcal{G}}} \min_{y \in \widehat{\mathcal{X}} \cap \widehat{\mathcal{G}}} \langle d, x - y \rangle \end{aligned}$$

Choosing d accordingly we take absolutes in the dot product.

$$\begin{aligned} &= \max_{\|d\|_p \leq 1} \max_{x \in \widehat{\mathcal{X}} \cap \widehat{\mathcal{G}}} \min_{y \in \widehat{\mathcal{X}} \cap \widehat{\mathcal{G}}} \langle d, |x - y| \rangle \\ &\leq \max_{\|d\|_p \leq 1} \max_{x, y \in \widehat{\mathcal{X}} \cap \widehat{\mathcal{G}}} \langle d, |x - y| \rangle \end{aligned}$$

We conservatively bound $|x - y|$ by the diameter of $\widehat{\mathcal{X}} \cap \widehat{\mathcal{G}}$.

$$\begin{aligned} &\leq \max_{\|d\|_p \leq 1} \langle d, \Delta_p(\widehat{\mathcal{X}} \cap \widehat{\mathcal{G}}) \rangle \\ &\leq \min \left(\max_{\|d\|_p \leq 1} \langle d, \Delta_p(\widehat{\mathcal{X}}) \rangle, \max_{\|d\|_p \leq 1} \langle d, \Delta_p(\widehat{\mathcal{G}}) \rangle \right) \\ &\leq \max_j \min(\|\Delta_p(\mathcal{X}_j)\|_p, \|\Delta_p(\pi_j \mathcal{G})\|_p) \quad \blacksquare \end{aligned}$$

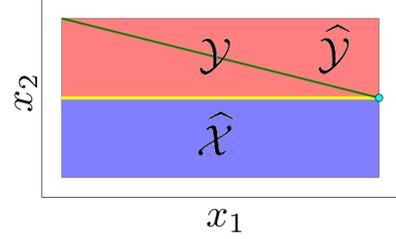


Fig. 11. The rectangle $\widehat{\mathcal{X}}$ (blue) and the line segment \mathcal{G} (green) intersect in a single point (cyan). The Cartesian decomposition $\widehat{\mathcal{G}}$ of \mathcal{G} (red) intersects with $\widehat{\mathcal{X}}$ on a whole facet (yellow). Observe that the length of this intersection facet is determined by the width of $\widehat{\mathcal{X}}$ and \mathcal{G} in dimension x_1 , independent of the height in dimension x_2 .

Figure 11 shows that the bound of Proposition 2 is tight.

Proof of Proposition 4:

$$\begin{aligned} & \widehat{\mathcal{X}} \subseteq \widehat{\mathcal{G}} \iff \forall x \in \widehat{\mathcal{X}} : x \in \widehat{\mathcal{G}} \\ & \iff \forall x_1 \in \mathcal{X}_1, \dots, x_\beta \in \mathcal{X}_\beta : \bigwedge_{j=1}^{\beta} x_j \in \mathcal{G}_j \\ & \iff \bigwedge_{j=1}^{\beta} \forall x_j \in \mathcal{X}_j : x_j \in \mathcal{G}_j \iff \bigwedge_{j=1}^{\beta} \mathcal{X}_j \subseteq \mathcal{G}_j \quad \blacksquare \end{aligned}$$

Proof of Proposition 5: Let $d \in \mathbb{R}^n$. We show $\rho_{\text{CH}(\widehat{\mathcal{X}} \cup \widehat{\mathcal{G}})}(d) \leq \rho_{\times_j \text{CH}(\mathcal{X}_j \cup \mathcal{G}_j)}(d)$.

$$\begin{aligned} & \rho_{\text{CH}(\widehat{\mathcal{X}} \cup \widehat{\mathcal{G}})}(d) = \max(\rho_{\widehat{\mathcal{X}}}(d), \rho_{\widehat{\mathcal{G}}}(d)) \\ &= \max \left(\sum_j \rho_{\mathcal{X}_j}(d_j), \sum_j \rho_{\mathcal{G}_j}(d_j) \right) \\ &\leq \sum_j \max(\rho_{\mathcal{X}_j}(d_j), \rho_{\mathcal{G}_j}(d_j)) \\ &= \sum_j \rho_{\text{CH}(\mathcal{X}_j \cup \mathcal{G}_j)}(d_j) = \rho_{\times_j \text{CH}(\mathcal{X}_j \cup \mathcal{G}_j)}(d) \quad \blacksquare \end{aligned}$$

Proof of Proposition 6: The first bound is due to Proposition 1.

$$\begin{aligned} & d_H^p(\text{CH}(\widehat{\mathcal{X}} \cup \widehat{\mathcal{G}}), \times_j \text{CH}(\mathcal{X}_j \cup \mathcal{G}_j)) \\ &= \max_{\|d\|_p \leq 1} \rho_{\times_j \text{CH}(\mathcal{X}_j \cup \mathcal{G}_j)}(d) - \rho_{\text{CH}(\widehat{\mathcal{X}} \cup \widehat{\mathcal{G}})}(d) \\ &= \max_{\|d\|_p \leq 1} \left(\sum_j \rho_{\text{CH}(\mathcal{X}_j \cup \mathcal{G}_j)}(d_j) \right) - \max(\rho_{\widehat{\mathcal{X}}}(d), \rho_{\widehat{\mathcal{G}}}(d)) \\ &= \max_{\|d\|_p \leq 1} \left(\underbrace{\sum_j \max(\rho_{\mathcal{X}_j}(d_j), \rho_{\mathcal{G}_j}(d_j))}_{=: \varphi(d_j)} \right) - \max \left(\sum_j \rho_{\mathcal{X}_j}(d_j), \sum_j \rho_{\mathcal{G}_j}(d_j) \right) \\ &= \max_{\|d\|_p \leq 1} \min \left(\sum_j \varphi(d_j) - \rho_{\mathcal{X}_j}(d_j), \sum_j \varphi(d_j) - \rho_{\mathcal{G}_j}(d_j) \right) \\ &= \max_{\|d\|_p \leq 1} \min \left(\sum_j \max(0, \rho_{\mathcal{G}_j}(d_j) - \rho_{\mathcal{X}_j}(d_j)), \sum_j \max(0, \rho_{\mathcal{X}_j}(d_j) - \rho_{\mathcal{G}_j}(d_j)) \right) \\ &\leq \max_{\|d\|_p \leq 1} \sum_j |\rho_{\mathcal{X}_j}(d_j) - \rho_{\mathcal{G}_j}(d_j)| \quad \blacksquare \end{aligned}$$