

PolyAR: A Highly Parallelizable Solver For Polynomial Inequality Constraints Using Convex Abstraction Refinement

Wael Fatnassi* Yasser Shoukry*

* University of California, Irvine, California, USA
(e-mail: {wfatnass, yshoukry}@uci.edu).

Abstract: Numerical tools for constraints solving are a cornerstone to control verification problems. This is evident by the plethora of research that uses tools like linear and convex programming for the design of control systems. Nevertheless, the capability of linear and convex programming is limited and is not adequate to reason about general nonlinear polynomial constraints that arise naturally in the design of nonlinear systems. This limitation calls for new solvers that are capable of utilizing the power of linear and convex programming to reason about general multivariate polynomials. In this paper, we propose PolyAR, a highly parallelizable solver for polynomial inequality constraints. PolyAR provides several key contributions. First, it uses convex relaxations of the problem to accelerate the process of finding a solution to the set of the non-convex multivariate polynomials. Second, it utilizes an iterative convex abstraction refinement process which aims to prune the search space and identify regions for which the convex relaxation fails to solve the problem. Third, it allows for a highly parallelizable usage of off-the-shelf solvers to analyze the regions in which the convex relaxation failed to provide solutions. We compared the scalability of PolyAR against Z3 8.9 and Yices 2.6 on control designing problems. Finally, we demonstrate the performance of PolyAR on designing switching signals for continuous-time linear switching systems.

Copyright © 2021 The Authors. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0>)

Keywords: Polynomial inequalities, Abstraction refinement, convex programming.

1. INTRODUCTION

Advances in constraints programming have opened several venues for control system synthesis and verification. For instance, linear programming and convex optimization are heavily used in a multitude of control system design and analysis tools. Recent surveys Ferreau et al. (2016); Rao (2009) showed that such numerical tools had changed the control system design philosophy.

Nevertheless, linear and convex programming are limited in their ability to problems with specific structures. In several control design and verification problems, constraints are neither linear nor convex. This calls for efficient solvers that can reason about *general multivariate polynomial constraints*. In that regard, Cylindrical Algebraic Decomposition (CAD) has long been one of the most influential algorithms capable of solving general multivariate polynomial constraints. The first CAD algorithm was introduced by Collins (1975) Collins (1975). However, the use of CAD is often limited by the number of variables in the input polynomials, a reflection of its worst-case complexity that grows in a doubly exponential fashion in the number of variables England and Davenport (2016). Due to Collins CAD's high time complexity, there have been improvements to the projection operator of this algorithm Hong (1990); McCallum (1998). However, CAD's modified versions are still limited by the number of variables Hong (1990) and are prone to errors McCallum (1998). The ABSolver and Z3 tools Bauer et al. (2007); De Moura and Bjørner (2008) are two solvers that solve non-linear constraints. However, ABSolver solver may produce incomplete results, and possibly incorrect, due to the local nature of the solver, and Z3 is still affected a lot by the increase in the number of variables in the polynomials. Because of the high complexity and incompleteness of existing ap-

proaches, we introduce PolyAR, a highly parallelizable, efficient, and complete solver that uses the advantage and the simplicity of convex programming to solve general multivariate polynomial inequality constraints. To the best of our knowledge, this approach is new and has not been highlighted before. The main novel contributions of this work can be summarized as follows:

- PolyAR is a highly parallelizable solver that uses a combination of convex programming and abstraction refinement to solve multivariate polynomial inequality constraints.
- PolyAR uses a novel convex abstraction refinement process where the original problem is iteratively relaxed into a series of convex programming problems with the aim to find the solution and prune the search space. Second, it refines such abstraction where it becomes tighter with each iteration of the algorithm. Finally, it examines in parallel all the identified small volume regions left from the abstraction using off-the-shelf solvers (e.g., Z3 and Yices) to search for a solution in these regions.
- We validate our approach by comparing the scalability of the proposed PolyAR solver with respect to the latest versions of state-of-art non-linear real arithmetic solvers, such as Z3 8.9 and Yices 2.6, on synthesizing stabilizing static output feedback controller (SOF) for linear time-invariant (LTI) continuous systems and designing non-linear controller for a Duffing oscillator.

2. PROBLEM FORMULATION

2.1 Notation

We denote by $x = (x_1, x_2, \dots, x_n) \in \mathbb{R}^n$ the set of real-valued variables, where $x_i \in \mathbb{R}$. We denote by $I_n = [\underline{d}_1, \bar{d}_1] \times \dots \times [\underline{d}_n, \bar{d}_n] \subset \mathbb{R}^n$ the n -dimensional region. We denote the space of polynomials with n vari-

* This work was partially sponsored by the NSF awards #CNS-2002405 and #CNS-2013824.

ables and coefficients in \mathbb{R} by $\mathbb{R}[(x_1, x_2, \dots, x_n)]$. We denote by \wedge the Boolean conjunction. A set of the form $L_0^-(f) = \{(x_1, \dots, x_n) | f(x_1, \dots, x_n) \leq 0\}$ ($L_0^+(f) = \{(x_1, \dots, x_n) | f(x_1, \dots, x_n) \geq 0\}$) is called zero sublevel (superlevel) set of f , respectively.

2.2 Main Problem

In this paper, we focus on *polynomial inequality constraints* with input ranges as closed boxes defined as follows:

Problem 1.

$$\exists x = (x_1, \dots, x_n) \quad \text{subject to:}$$

$$\bigwedge_{i=1}^n x_i \in [\underline{d}_i, \bar{d}_i], \quad (1)$$

$$\bigwedge_{i=1}^m p_i(x_1, \dots, x_n) \leq 0. \quad (2)$$

where $p_i(x) = p_i(x_1, \dots, x_n) \in \mathbb{R}[(x_1, x_2, \dots, x_n)]$ is a polynomial over variables x_1, \dots, x_n . Without loss of generality, $\bigwedge_{i=1}^m p_i(x) \geq 0$ and $\bigwedge_{i=1}^m p_i(x) = 0$ can be encoded in constraint number (2).

3. ABSTRACTION REFINEMENT OF HIGHER ORDER POLYNOMIALS USING QUADRATIC POLYNOMIALS

Traditional techniques for solving Problem 1 focus on finding all the n roots of the m polynomials and check all the regions between two successive roots to assign a positive/negative sign for each of these regions. Therefore, solving Problem 1 is known to be a doubly combinatorial problem in n with a total running time that is bounded by $(md)^{2^n}$ England and Davenport (2016), where d is the maximum degree among polynomials in P_m .

In problems that are doubly exponential in the input space n , it is beneficial to isolate subsets of the search space in which the solution is guaranteed not to exist. Recall that Problem 1 asks for an x in \mathbb{R}^n for which all the polynomials are negative. Therefore, a solution does not exist in subsets of \mathbb{R}^n at which one of the polynomials is always positive. Similarly, isolating regions of the input space for which some of the polynomials are negative is also beneficial to finding the solution faster.

Our tool's main novelty is to use "convex abstractions" of the polynomials to find subsets of $L_0^+(p)$ and $L_0^-(p)$ efficiently. Indeed such "abstractions" may not be able to identify all regions for which the polynomial is positive or negative, which calls for an "abstraction refinement" process in which these "convex abstractions" become tighter with each iteration of the algorithm.

Figure 1(top) visualizes the proposed abstraction refinement process. Starting from a polynomial $p(x) \in \mathbb{R}[x]$ and an interval $I_n \subset \mathbb{R}^n$, we compute two quadratic polynomials:

$$\begin{aligned} O_1^p(x) &\geq p(x) & \forall x \in I_n, \\ U_1^p(x) &\leq p(x) & \forall x \in I_n. \end{aligned}$$

where O and U stands for Over-approximate and Under-approximate quadratic polynomials, respectively, and the subscript in $O_1^p(x)$ and $U_1^p(x)$ encodes the iteration index of the abstraction refinement process. Computing such upper and lower abstractions can be carried out efficiently using Taylor approximation. Please refer to the example depicted in Figure 1 (top) for a visualization of $O_1^p(x)$ and




Fig. 1. Abstraction Refinement of higher order polynomial using quadratic approximations: (top) first iteration and (bottom) second iteration.

$U_1^p(x)$ for one dimensional higher order polynomial (order ≥ 3) defined in the closed interval $[\underline{d}, \bar{d}] \subset \mathbb{R}$.

The next step is to use the quadratic abstractions to isolate subsets of $L_0^+(p)$ and $L_0^-(p)$. It is particularly direct to show that the zero superlevel set of $U_1^p(x)$ is a subset of $L_0^+(p)$, i.e., $L_0^+(U_1^p) \subseteq L_0^+(p)$. Similarly, the zero sublevel set of $O_1^p(x)$ is a subset of $L_0^-(p)$, i.e., $L_0^-(O_1^p) \subseteq L_0^-(p)$. Thanks to the fact that $O_1^p(x)$ and $U_1^p(x)$ are quadratic polynomials, finding their zero superlevel and zero sublevel sets, respectively, can be computed efficiently. Referring to the example in Figure 1(top), these zero superlevel and sublevel sets are $L_0^+(U_1^p) = [x_1, \bar{d}]$ and $L_0^-(O_1^p) = [\underline{d}, x_0]$, respectively.

It is clear from Figure 1(top) that the abstractions $O_1^p(x)$ and $U_1^p(x)$ fails to identify all subsets of $L_0^-(p)$ and $L_0^+(p)$. Therefore, the next step is to compute tighter over and under approximations of $p(x)$. Such a refinement process can be carried out by removing the zero superlevel and the zero sublevel sets, i.e., $L_0^+(U_1^p)$ and $L_0^-(O_1^p)$, identified using the previous abstraction and computing new over and lower approximation, as shown in Figure 1(bottom). The process of abstraction refinement can continue until the remaining subsets of the search space, in which case we call them ambiguous regions, and with some abuse of notation, denoted them by $L_0^{+/-}(p)$, are *small* enough to be analyzed using off-the-shelf solvers. More details about the proposed abstraction refinement process are given in the next section.

4. ALGORITHM ARCHITECTURE

In this section, we describe the different steps used by our solver PolyAR to solve Problem 1. More details about the implementation of the algorithm is given in this technical report Fatnassi and Shoukry (2021).

Our design methodology for the PolyAR tool aims to reduce the number of the required abstraction refinement and tries to find a solution early on in the process. To that end, the tool starts by computing a set of convex

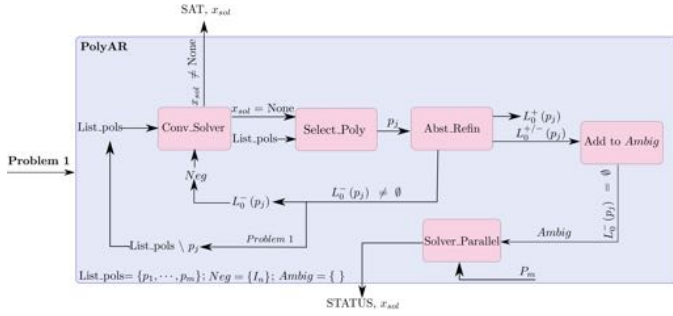


Fig. 2. Framework of PolyAR.

(quadratic or linear) polynomials $O_0^{p_i}(x), i = 1, \dots, m$, that over approximate the original polynomials. The next step is to solve a convex feasibility problem aiming to find a solution that satisfies the constraints:

$$\exists x = (x_1, \dots, x_n) \quad \text{s.t.} \quad O_0^{p_i}(x) \leq 0, \quad i = 1, \dots, m.$$

Indeed, if such a convex problem is feasible, the tool terminates and returns the solution found by the convex feasibility problem above (**Conv_Solver**, Figure 2). If not, then the tool selects one polynomial p_j (**Select_Poly**, Figure 2) to perform the abstraction refinement process. Indeed, several heuristics can be applied to select which polynomial will be selected. In the PolyAR tool, we opt-out to select the polynomial with the highest Lipschitz constant. Our intuition is that the higher the Lipschitz constant, the harder to obtain a tight over-approximation that can be used to find the solution.

Once a polynomial p_j is selected, the next step is to apply the abstraction refinement process on p_j (**Abst_Refin**, Figure 2). The objective of the abstraction refinement process is to identify subsets of the positive regions $L_0^+(p_j)$ and negative regions $L_0^-(p_j)$. Indeed, such abstraction refinement may not be able to identify all positive and negative regions, and hence a remaining portion of the search space may not be identified to belong to either $L_0^+(p_j)$ or $L_0^-(p_j)$ in which case it belongs to the ambiguous region $L_0^{+/-}(p_j)$. The abstraction refinement process of the polynomial p_j ensure that the volume of such ambiguous regions is below a certain user defined threshold.

The process of using the convex solver to find the solution and abstracting one polynomial continues. Since a solution of Problem 1 needs to lie in a negative region for all the polynomials, we confine the tool attention to the negative regions identified by the abstraction refinement in the previous iterations to accelerate the process of searching for the solution.

While excluding the positive regions identified in previous iterations does not affect the tool (since a solution is guaranteed not to exist in such regions), excluding the ambiguous regions from the next iterations may affect the correctness of the tool. Therefore, the last step in the PolyAR tool is to examine all the identified ambiguous regions using off-the-shelf solvers (e.g., Z3 and Yices) to search for a solution in these regions (**Solver_Parallel**, Figure 2). Because the volume of these ambiguous regions is smaller than a user-defined threshold, the execution time of running off-the-shelf tools on such small volume regions is shorter than solving the original problem. This reflects that the number of roots for each polynomial is limited in small regions. Moreover, searching for a solution in these ambiguous regions can be highly parallelized, leading to an extra level of efficiency.

4.1 Early Termination Using **Conv_Solver**:

The objective of the **Conv_Solver** is to search for a solution to Problem 1 using the information of (i) a set of closed convex regions Neg identified by the previous iterations of the abstraction refinement process and (ii) a list of polynomials ($List_pols$) that have not yet been processed by the abstraction refinement process.

Our approach is to compute a convex over-approximation of the polynomials in $List_pols$ using Taylor approximation. Let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be two times differentiable in open interval in \mathbb{R}^n , then $f(x)$ can be bounded by:

$$\begin{aligned} T_1(x) - M_1 &\leq f(x) \leq T_1(x) + M_1, \\ T_2(x) - M_2 &\leq f(x) \leq T_2(x) + M_2, \end{aligned}$$

where $T_1(x)$ and $T_2(x)$ are the first and the second order Taylor polynomials, and M_1 and M_2 are the upper bound on the higher order terms.

Next, we check the convexity of the obtained second Taylor approximation and use it to compute the over-approximation function O^{p_i} whenever it is convex. Otherwise, we use the first Taylor approximation instead. Finally, for each *region* in the set of negative regions (Neg), we solve the following convex feasibility problem:

$$x_{Sol} := \arg \min_{x \in region} 1 \quad \text{s.t.} \quad O^{p_i}(x) \leq 0, \quad i \in List_pols. \quad (3)$$

4.2 Abstraction Refinement Using **Abst_Refin**:

Given the set of negative regions Neg identified by the previous abstraction refinement process along with the polynomial p_j selected by the **Select_Poly** algorithm, the objective of the **Abst_Refin** algorithm is to find subsets of the zero sublevel sets of p_j that lie inside Neg . The output of this algorithm are subsets of $L_0^-(p_j)$ and $L_0^+(p_j)$. The remainder of Neg is then considered to be part of the ambiguous regions $L_0^{+/-}(p_j)$. To do so, for every *region* in Neg , the tool initiates a list of ambiguous regions $List_Ambig_reg$, which will contain all the ambiguous regions from the abstraction refinement. Next, it selects one element from these ambiguous regions and performs the abstraction refinement on this region iteratively until the volume of the remaining ambiguous region is smaller than a user-defined threshold. During the iterative abstraction refinement, all the identified zero sublevel and superlevel subsets are stored in the sets $L_0^-(p_j)$ and $L_0^+(p_j)$, respectively.

While the zero sublevel (superlevel) sets of the quadratic over-approximation (under-approximation) are ellipsoid or hyperboloid in general, we opt to represent all the subsets of $L_0^-(p_j)$ and $L_0^+(p_j)$ as n -dimensional hypercubes. This choice reflects the fact that off-the-shelf solvers (e.g., Z3 and Yices) can exploit the geometry of hypercubes to accelerate their computations. The process of finding these hypercubes is visualized in Figure 3 and can be summarized as follows:

Step 1: Computing the largest polytope subset of $L_0^-(p_j)$ and $L_0^+(p_j)$ Given the over-approximation O^{p_j} computed using Taylor polynomials (detailed in Section 4.1) and a convex ambiguous region ($Ambig_reg$), we start by computing a set of $n + 1$ vertices v_1^N, \dots, v_{n+1}^N that are inscribed in the ambiguous region $Ambig_reg$. Each vertex can be computed by solving the following convex optimization problem:

$$v_i^N = \arg \min_{v_i \in Ambig_reg} (l_i^T v_i) \quad \text{s.t.} \quad O^{p_j}(v_i) \leq 0, \quad (4)$$




Fig. 3. Polytopic under-approximation of a 2-dimensional ellipse sublevel set $L_0^-(p_j)$. \mathcal{P}^N presents the under-approximate polytope inscribed in $L_0^-(p_j)$, and \mathcal{B}^N represents the axis-aligned box of maximum volume inscribed in \mathcal{P}^N .

where l_i is a user defined normal vector (or template) (see Figure 3 for graphical representation of such normal vectors). Using these vertices, we can obtain the polytope \mathcal{P}^N as:

$$\mathcal{P}^N = \text{Convex_Hull}(v_1^N, \dots, v_{n+1}^N).$$

Thanks to the constraints in the optimization problem (4) along with the convexity of $L_0^-(p_j)$, it is direct to conclude that the polytope \mathcal{P}^N satisfy $\mathcal{P}^N \subset L_0^-(p_j)$. We compute the polytope $\mathcal{P}^P \subset L_0^+(p_j)$ in a similar fashion using the under-approximation U^{p_j} .

Step 3: Under approximate the polytopes with axis aligned boxes: To compute the largest axis-aligned hypercube \mathcal{B}^N inscribed inside the polytope \mathcal{P}^N , we solve the following convex optimization problem Behroozi (2019):

$$\begin{aligned} & \arg \max_{(l_1^N, u_1^N, \dots, l_n^N, u_n^N) \in \mathbb{R}^{2n}} \sum_{k=1}^n \log(u_k^N - l_k^N) \\ \text{s.t. } & \sum_{k=1}^n (p_{ik}^{N,+} u_k^N - p_{ik}^{N,-} l_k^N) \leq c_i^N, \quad i = 1, \dots, n_p, \end{aligned} \quad (5)$$

where $(l_1^N, u_1^N, \dots, l_n^N, u_n^N) \in \mathbb{R}^{2n}$ is the representation of the box \mathcal{B}^N with (l_k, u_k) is the lower/upper limit of the box in the k th dimension, $p_{ik}^{N,+} = \max\{p_{ik}^N, 0\}$, $p_{ik}^{N,-} = \max\{-p_{ik}^N, 0\}$, and p_{ik}^N, c_i^N are the rows of the half-space matrix/vector representation of the polytope \mathcal{P}^N .

4.3 Highly Parallelizable Analysis of Ambiguous Regions using `Solver_Parallel`

Once all the ambiguous regions are identified, the next step is to analyze all of them using off-the-shelf solvers. In particular, PolyAR supports the use of the latest versions Z3 8.9 and Yices 2.6 solvers. Thanks to the fact that all the ambiguous regions are hypercubes, both these solvers can exploit the geometry of the region to accelerate their computations. Also, thanks to the fact that the volume of all ambiguous regions is lower than a user-defined threshold, the CAD algorithm can run efficiently. To that end, PolyAR tool runs multiple instances of Z3 or Yices to analyze all these ambiguous regions in parallel.

5. EXTENSION TO SMT SOLVING

We extend the PolyAR solver described in the previous sections to account for combinations of Boolean and Polynomial inequality constraints of the form:

$$\exists(b_1, \dots, b_o, x_1, \dots, x_n) \in \mathbb{B}^o \times \mathbb{R}^n,$$

subject to:

$$p_i(x_1, \dots, x_n) \leq 0, \quad i = 1, \dots, m \quad (6)$$

$$x_k \in [\underline{d}_k, \bar{d}_k], \quad k = 1, \dots, n \quad (7)$$

$$\varphi_j(b_1, \dots, b_o) \leftrightarrow \text{TRUE}, \quad j = 1, \dots, r \quad (8)$$

$$b_l \leftrightarrow (p_{l+m}(x_1, \dots, x_n) \leq 0), \quad l = 1, \dots, h \quad (9)$$

where $\varphi_j(b_1, \dots, b_o)$ is any combinations of Boolean and pseudo-Boolean predicates.

We can create a Satisfiability Modulo Theory (SMT) solver by combining a SAT solver for Boolean and pseudo-Boolean constraints and a theory solver (PolyAR) for interval and polynomial constraints on real numbers by following the lazy SMT paradigm Barrett and Tinelli (2018). The SAT solver solves the combination of Boolean and pseudo-Boolean constraints using the David-Putnam-Logemann-Loveland (DPLL) algorithm and suggests satisfying assignments for the Boolean variables b and thus suggesting which polynomial constraints should jointly satisfied (or unsatisfied). The theory solver (PolyAR) checks the validity of the given assignments and provides an explanation of the conflict, i.e., an *UNSAT certificate*, whenever a conflict is found. Each certificate is a new Boolean constraint that will be used by the SAT solver to prune the search space.

While in the lazy SMT paradigm, the PolyAR solver needs to be executed multiple times with a different set of polynomial constraints, we modify the PolyAR solver to perform all the abstraction refinement for all the polynomials as a pre-processing step. This eliminates the need to re-compute the same abstraction refinement every time the PolyAR solver is executed.

Whenever the SAT solver assigns one of the Boolean variables b_l in (9) to zero, then the PolyAR solver needs to guarantee that the corresponding polynomial p_{l+m} satisfy $p_{l+m}(x) > 0$ or equivalently $-p_{l+m}(x) \leq 0$. To eliminate the need to apply the convex abstraction refinement process for both $p_{l+m}(x)$ and $-p_{l+m}(x)$, the PolyAR solver computes the negative and positive boxes (\mathcal{B}^N and \mathcal{B}^P) only for $p_{l+m}(x)$ and flips their usage for $-p_{l+m}(x)$.

6. NUMERICAL RESULTS

We compare the performance of PolyAR to the state-of-the-art solvers Z3 8.9 and Yices 2.6. The objective of this comparison is to study the performance on:

- Problems that appear naturally in parametric controller synthesis. In particular, we focus on the problem of designing stabilizing SOF controllers for LTI systems Bahavarnia et al. (2020).
- Problems that appear in non-parametric controller synthesis for non-linear systems. In particular, we focus on the problem of designing a controller for the nonlinear Duffing oscillator Fotiou et al. (2006).
- Additionally, we demonstrate the performance of PolyAR on designing a hybrid switching system; a problem which state-of-the-art tools are incapable of handling.

All the experiments were executed on an Intel Core i7 2.6-GHz processor with 16 GB of memory.

6.1 Static Output Feedback Controller Synthesis for Linear Time Invariant Systems

In this subsection, we assess the scalability of the PolyAR solver compared to state-of-the-art solvers on control synthesis problems. In particular, we consider the problem of synthesizing a parametric controller for the following continuous LTI system:

$$\dot{x} = Ax + Bu, \quad y = Cx,$$

Table 1: Experiment results for SOF design. The timeout is set by 3600 s.

Example	Times (seconds)					
	Z3 8.9	Yices 2.6	PolyAR+Z3 (1 thread)	PolyAR+Z3 (max threads)	PolyAR+Yices (1 thread)	PolyAR+Yices (max threads)
1	timeout	timeout	timeout	7.552	2.405	2.442
2	timeout	timeout	83.776	114.453	timeout	3.766
3	timeout	timeout	23.551	23.970	timeout	8.725
4	timeout	timeout	0.718	0.729	0.416	0.432
5	timeout	timeout	3.636	3.768	0.621	0.498
# Problems Solved	0	0	4	5	3	5
Total Time (seconds)	timeout	timeout	111.681	150.472	3.442	15.863

where $x \in \mathbb{R}^{n_A}$ is the system state, $u \in \mathbb{R}^{n_B}$ is the system control input, $y \in \mathbb{R}^{n_C}$ is the system output, and the matrices $A \in \mathbb{R}^{n_A \times n_A}$, $B \in \mathbb{R}^{n_A \times n_B}$ and $C \in \mathbb{R}^{n_C \times n_A}$ are the system matrices. We are interested in designing a static output feedback controller of the form:

$$u = Ky,$$

such that the resulting closed loop system:

$$\dot{x} = (A + BKC)x,$$

is stable, i.e., the matrix $A + BKC$ is Hurwitz.

We follow the steps detailed in Bahavarnia et al. (2020) to pose the problem of designing the static output feedback controller as a set of polynomial constraints using the Routh-Hurwitz stability criteria. The Routh-Hurwitz stability criteria result in a set of n_A polynomials in the elements of the controller matrix K . We consider five instances of the controller synthesis problem. Details of the parameters of each of the five instances are detailed in Fatnassi and Shoukry (2021).

For each of the five instances, we generate random system matrices from a zero-mean normal distribution and feed them to four versions of our solver PolyAR:

- PolyAR + Z3 (1 thread): This version uses one instance of Z3 to analyze all the ambiguous regions.
- PolyAR + Z3 (max threads): This version uses a separate instance of Z3 to analyze each of the ambiguous regions. All Z3 instances are running in parallel.
- PolyAR + Yices (1 thread): This version uses one instance of Yices to analyze all the ambiguous regions.
- PolyAR + Yices (max thread): This version uses a separate instance of Yices to analyze each of the ambiguous regions. All Yices instances are running in parallel.

We compare the execution times of these four solvers with Z3 8.9 and Yices 2.6. Table 1 shows the execution time for all the solvers. As evident by the results in Table 1, off-the-shelf solvers are incapable of solving all the five examples and they time out after one hour. On the other hand, and thanks to the abstraction refinement process, the PolyAR solver is able to solve all the instances in a few seconds, leading to 240X speed up in the total execution time in the PolyAR+Yices (max threads) case, evidence of the scalability of the proposed approach.

6.2 Non-Linear Controller Design for a Duffing Oscillator

In this subsection, we assess the scalability of PolyAR solver compared to state-of-the-art solvers on synthesizing a non-parametric controller for a Duffing oscillator reported by Fotiou et al. (2006). The dynamics of the oscillator is given by the higher-order differential equation:

$$y^{(n)}(t) + \dots + y^{(2)}(t) + 2\zeta y^{(1)}(t) + y(t) + y(t)^3 = u(t), \quad (10)$$

where $y \in \mathbb{R}$ is the continuous state variable and $u \in \mathbb{R}$ is the control input. The parameter ζ is the damping coefficient. The objective of the control is to regulate the state to the origin. To derive the discrete-time model, forward difference approximation is used (with sampling period of $h = 0.05$ time units). The resulting state space

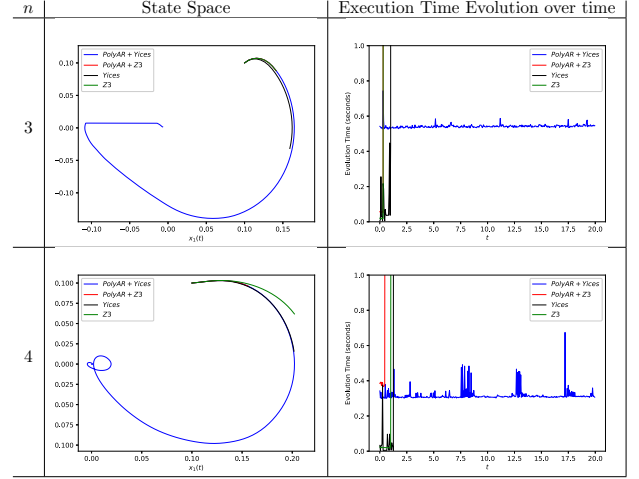


Fig. 4. Results of controlling the Duffing oscillator with different n (left) evolution of the states $x_1(k)$ and $x_2(k)$ for the solvers in the state-space, (right) evolution of the execution time of solvers during the 20 seconds. The timeout is equal to 1s. Trajectories are truncated once the solver exceeds the timeout limit.

model with discrete state vector $x = [x_1, x_2, \dots, x_n]^T \in \mathbb{R}^n$ and input $u \in \mathbb{R}$ is:

$$\begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix}^+ = \begin{bmatrix} 1 & h & 0 & \dots & 0 \\ 0 & 1 & h & 0 & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ -h-2\zeta h-h & \dots & -h & 1 & -h \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_2 \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ \vdots \\ h \end{bmatrix} u + \begin{bmatrix} 0 \\ 0 \\ \vdots \\ -hx_1^3 \end{bmatrix}. \quad (11)$$

The previous equation is written in the form of $x(k+1) = Ax(k) + Bu(k) + E(x)$, which includes a nonlinear term $E(x) = [0, \dots, -hx_1^3(k)]^T$. Our objective is to design a non-parametric controller. To that end, we encode the controller as the solution of a feasibility problem of several constraints that capture the system dynamics, state/input constraints, and stability constraints as discussed below.

First, to enforce the stability of the resulting non-parametric controller, we consider the candidate quadratic Lyapunov function $V(x) = x^T P x$ with the symmetric positive definite matrix P is a solution of the discrete-time Lyapunov equation $APA^T + P + Q = 0$ and is a positive definite matrix. Thanks to the fact that $E(x)$ satisfies $\lim_{\|x\| \rightarrow 0} \frac{\|E(x)\|}{\|x\|} = 0$ along with the Lyapunov's indirect method in Khalil (2002), one can directly conclude that $V(x)$ is indeed a Lyapunov function. For simplicity, we pick $Q = I_n$, where I_n is the identity matrix of size n .

Moreover, to ensure the smoothness of the resulting controller signals, we add additional filters in the form of high order polynomial $L(x, u) \leq 0$. In addition, we consider the state-constraints of the form $\|x(k)\|_\infty \leq 0.6$. The final non-parametric controller is then encoded as the solution of a feasibility problem (details of the exact constraints can be found in Fatnassi and Shoukry (2021)).

We consider three instances of the controller synthesis problem for the Duffing oscillator with the following parameters:

- $n = 3$, $\zeta = 1.0$, $x(0) = [0.1, 0.1, 0.1]^T$, $L(x(k), u(k)) = x_1^5(k) + x_2^5(k) + x_3^5(k) + u^5(k)$, which results in 8 polynomial constraints with 4 variables and max polynomial order of 5.
- $n = 4$, $\zeta = 1.75$, $x(0) = [0.1, 0.1, 0.01, 0.1]^T$, $L(x(k), u(k)) = x_1^4(k) + x_2^4(k) + x_3^4(k) + x_4^4(k) -$




Fig. 5. The trajectory that starts from an initial state $x(0) = [40, 30]^T$ and reaching a final state $x(3) \in \text{Goal}$ while avoiding the obstacles. The goal and the obstacles are represented with a red and yellow rectangle, respectively.

$u^4(k)$, which results in 10 polynomial constraints with 5 variables and max polynomial order of 4.

We feed the resultant polynomial inequality constraint to PolyAR+Yices, PolyAR+Z3, Yices, and Z3. We solve the feasibility problem for $n = 3$, and $n = 4$. We set the timeout to be 1s. Figure 4 (left) shows the state-space evolution of the controlled Duffing oscillator for different solvers for number of variables n of 2, 3, and 4. Figure 4 (right) shows the evolution of the execution time of the solvers during the 20 seconds. As it can be seen from Fig. 4, our solver PolyAR+ Yices succeeded to find a control input u that regulates the state to the origin for all n . However, off-the-shelf solvers are incapable of solving all the three instances and they early time out after one second out of the simulated 20 seconds.

6.3 Designing Switching Signals for Continuous-Time Linear Switching Systems

In this subsection, we show how to use the PolyAR solver to successfully design a controller for a continuous-time linear switching system. In particular, we consider the following switching dynamics:

$$\dot{x} = A_{\sigma(t)}x, \quad \sigma(t) = \{1, 2, 3\},$$

with $x(t) \in \mathcal{X} \subset \mathbb{R}^2$ is the system state at time t and the matrices A_1, A_2 , and $A_3 \in \mathbb{R}^{2 \times 2}$ represents three modes for the switching system. Consider the state space in Figure 5. The objective is to design a switching signal $\sigma(t)$ that can steer the state of the system to the goal set $\text{Goal} \subset \mathcal{X}$ while avoiding entering the obstacle set $\text{Obstacle} \subset \mathcal{X}$. For simplicity, we confine our attention to step-wise switching signals $\sigma(t)$. That is, we assume the switching signal $\sigma(t)$ will be constant for some amount of time t_1, t_2, \dots, t_L . Our objective is then to design the switching times and the associated system mode that leads to the satisfaction of the reach-avoid specifications. To that end, we define a set of Boolean variables b_{ij} such that b_{ij} is equal to 1 whenever the j th mode is active during t_i . Given the initial condition of the system $x(0)$, we can use these Boolean variables to encode the problem of designing the switching signal as an SMT problem (details of the exact SMT constraints can be found in Fatnassi and Shoukry (2021)).

In our experiments, we pick the horizon $L = 3$ and the modes $A_1 = \begin{bmatrix} -1, 2 \\ -2, -2 \end{bmatrix}$, $A_2 = \begin{bmatrix} -1, 3 \\ -3, -1 \end{bmatrix}$, and $A_3 = \begin{bmatrix} 0, 2 \\ -2, 0 \end{bmatrix}$, and we start with an initial state $x(0) = [40, 30]^T$. We restrict the states to be inside $\mathcal{X} \in [-100, 100]$. We feed the resultant polynomial inequality constraints to

PolyAR+Yices, and as it can be seen from Fig.5, the solver succeeded to find the right modes ($b_{11} = b_{22} = b_{33} = 1$) and the necessary times $t_1 = 0.391s$, $t_2 = 0.5s$, and $t_3 = 0.25s$ that ensures that $x(3)$ reaches a Goal while the intermediate states $x(2), x(1)$ avoid Obstacles. In addition, we remark that the trajectory between $x(0)$ and $x(2)$ is making its way to the equilibrium point $[0, 0]^T$. This is explained by the fact that the matrices A_1 and A_2 are stables. Our solver computes the necessary time t_3 that ensures that the final state $x(3) \in \text{Goal}$ and does not converge to the equilibrium point.

REFERENCES

- Bahavarnia, M., Shoukry, Y., and Martins, N.C. (2020). Controller Synthesis subject to Logical and Structural Constraints: A Satisfiability Modulo Theories (SMT) Approach. In *2020 American Control Conference (ACC)*, 5281–5286. doi:10.23919/ACC45564.2020.9147460.
- Barrett, C. and Tinelli, C. (2018). Satisfiability modulo theories. In *Handbook of Model Checking*, 305–343. Springer.
- Bauer, A., Pister, M., and Tautschnig, M. (2007). Tool-support for the analysis of hybrid systems and models. In *2007 Design, Automation Test in Europe Conference Exhibition*, 1–6.
- Behroozi, M. (2019). Largest Inscribed Rectangles in Geometric Convex Sets. *CoRR*, abs/1905.13246.
- Collins, G.E. (1975). Quantifier elimination for real closed fields by cylindrical algebraic decomposition. In *Automata Theory and Formal Languages 2nd GI Conference Kaiserslautern, May 20–23, 1975*, 134–183. Springer.
- De Moura, L. and Bjørner, N. (2008). Z3: An efficient SMT solver. In *International conference on Tools and Algorithms for the Construction and Analysis of Systems*, 337–340.
- England, M. and Davenport, J.H. (2016). The complexity of cylindrical algebraic decomposition with respect to polynomial degree. In *International Workshop on Computer Algebra in Scientific Computing*, 172–192. Springer.
- Fatnassi, W. and Shoukry, Y. (2021). Polyar: A highly parallelizable solver for polynomial inequality constraints using convex abstraction refinement. *arXiv preprint arXiv:2101.04655*.
- Ferreau, H.J., Almér, S., Peyrl, H., Jerez, J.L., and Domahidi, A. (2016). Survey of industrial applications of embedded model predictive control. In *2016 European Control Conference (ECC)*, 601–601. IEEE.
- Fotiou, I.A., Rostalski, P., Parrilo, P.A., and Morari, M. (2006). Parametric optimization and optimal control using algebraic geometry methods. *International Journal of Control*, 79(11), 1340–1358.
- Hong, H. (1990). An Improvement of the Projection Operator in Cylindrical Algebraic Decomposition. In *Proceedings of the International Symposium on Symbolic and Algebraic Computation, ISSAC '90*, 261–264. Association for Computing Machinery, New York, NY, USA. doi:10.1145/96877.96943. URL <https://doi.org/10.1145/96877.96943>.
- Khalil, H.K. (2002). *Nonlinear systems; 3rd ed.* Prentice-Hall, Upper Saddle River, NJ.
- McCallum, S. (1998). An improved projection operation for cylindrical algebraic decomposition. In *Quantifier Elimination and Cylindrical Algebraic Decomposition*, 242–268. Springer.
- Rao, A.V. (2009). A survey of numerical methods for optimal control. *Advances in the Astronautical Sciences*, 135(1), 497–528.