



HAL
open science

A Generic Ellipsoid Abstract Domain for Linear Time Invariant Systems

Pierre Roux, Romain Jobredeaux, Pierre-Loïc Garoche, Eric Féron

► **To cite this version:**

Pierre Roux, Romain Jobredeaux, Pierre-Loïc Garoche, Eric Féron. A Generic Ellipsoid Abstract Domain for Linear Time Invariant Systems. 2011. hal-00608673v1

HAL Id: hal-00608673

<https://hal.science/hal-00608673v1>

Preprint submitted on 13 Jul 2011 (v1), last revised 26 Jul 2011 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A Generic Ellipsoid Abstract Domain for Linear Time Invariant Systems^{*}

Pierre Roux¹, Romain Jobredeaux², Pierre-Loïc Garoche¹, and Éric Féron²

¹ ONERA – The French Aerospace Lab, Toulouse, FRANCE

² Georgia Institute of Technology, Atlanta, Georgia, USA

Abstract. Embedded system control often relies on linear systems, which admit quadratic invariants. The parts of the code that host linear system implementations need dedicated analysis tools since intervals or linear abstract domains will give imprecise results, if any at all, on these systems. Reference [9] proposes a specific abstraction for digital filters that addresses this issue on a specific class of controllers.

This paper aims at generalizing the idea, relying on existing methods from Control Theory to automatically generate quadratic invariants for linear time invariant systems, whose stability is provable. This class encompasses n -th order digital filters and, in general, controllers embedded in critical systems.

While control theorists only focus on the existence of such invariants, this paper proposes a method to effectively compute tight ones. The method has been implemented and applied to some benchmark systems, giving good results. It also considers floating points issues and validates the soundness of the computed invariants.

Keywords: stable linear systems, ellipsoids, quadratic invariants, Lyapunov functions, semi-definite programming, floating point errors, abstract interpretation.

1 Control-command based critical systems

A wide range of today's real-time embedded systems, especially their most critical parts, rely on a control-command computation core. The control-command of an aircraft, a satellite or a car engine, is processed into a global loop repeated indefinitely during the activity of the controlled device. This loop models the acquisition of new input values via sensors, the update of internal state variables and the generation of new outputs. The acquisition is made either from environmental measurements (wind speed, acceleration, engine RPM, ...) or from human input via the brakes, the accelerator, the stick or wheel control.

Control command theorists are used to model both the environment and the system behavior and, using their own set of tools, add the necessary elements to obtain the target controlled system. After discretizing the system, they mathematically prove its stability and its performance by exhibiting a quadratic form,

^{*} This work has been partially supported by the FNRAE Project CAVALE.

i.e. an ellipsoid, that over-approximates the system behavior with respect to a given input. All these steps are well known by control theory specialists. Reference [15] is a good introduction to these approaches. In this paper, we focus on a class of such systems where control is computed using stable linear systems, i.e. the controller is open-loop stable.

The system control law is then compiled from its description to executable code, like embedded C. This description is usually specified in Matlab Simulink, Scilab Scicos or in a dedicated synchronous language such as Lustre or Scade.

Fig. 1 sketches the loop body of a coupled mass controller, as generated by Matlab. It corresponds to a one step evaluation of the following linear system $x_{k+1} = Ax_k + Bu_k$, where $\|u_k\|_\infty \leq 1$. Vector x_k represents the state of the system at a given time. Matrix A models the system update according to its previous state, while matrix B expresses the effect of the input values u_k .

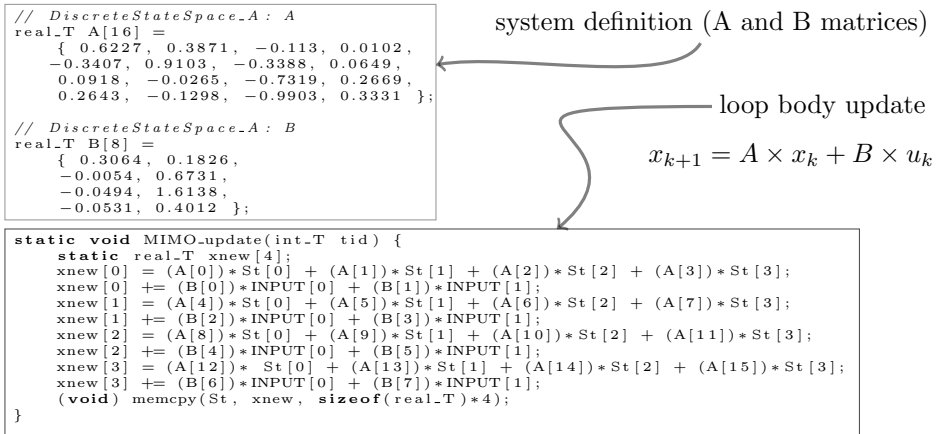


Fig. 1: System update for a coupled mass system controller generated by Matlab.

Once such controller source code is generated, it is embedded in the controlled device, eg. an aircraft. Critical embedded systems are then a major target for static analysis in order to ensure their good behavior. The success story of Astrée [7] illustrates such needs: it targets the analysis of the control command of the Airbus A380, and was used to formally prove the absence of any runtime error on the 300kloc of the controller source code. It relies on the theory of abstract interpretation [5,6] to compute a sound overapproximation of all possible values of the program variables in any reachable states. Then it is able to ensure that this over-approximation does not reach any possible bad state like overflows, division by zero, or invalid pointer dereferencing. In [7], the authors enumerate the different abstractions used to compute this over-approximation: intervals, octagons, ... Among them, we focus here on digital filters abstractions represented by ellipsoid abstract domains.

In [9], Feret proposes an analysis dedicated to stable linear filters in control command programs. These short pieces of code correspond to the kind of systems illustrated in Fig. 1, restricted to only one input argument, i.e. the matrix B is a column vector, and specific types of matrices A , i.e. companion matrices.

Most of the abstract domains available in actual tools only represent linear properties and this kind of non linear invariant is a real need. For example, *no interval invariant exists* for the example of Fig. 1. Thus analyzing it with intervals will give the $(-\infty; +\infty)$ over-approximation, whereas quadratic invariants will bound all parameters in $[-5; 5]$.

In this paper, we propose to generalize the approach of [9] by considering any inherently stable linear system. In particular, we

- characterize quadratic forms, invariants of the linear system analyzed, with techniques inspired by the control theory community;
- propose an open implementation of the analysis that handles floating point rounding errors;
- validate the result using a sound external solver.

This paper focuses on the quadratic invariants (ellipsoids) computation. Any reader interested in the reduction of the quadratic form obtained with the other domains should refer to [9].

The paper is structured as follow: Section 2 introduces the reader to the notion of stability based on Lyapunov invariants. Section 3 presents our global approach and the steps of our algorithm. Sections 4, 5 and 6 detail the main steps while Section 7 covers floating point issues and soundness. Finally concrete results and related work are presented in Sections 8 and 9.

2 Introduction to Lyapunov stability theory

One common way to establish stability of a discrete, time-invariant closed (i.e. with no inputs) system described in state space form, (i.e $x_{k+1} = f(x_k)$) is to use what is called a Lyapunov function. It is a function $V : \mathbb{R}^n \rightarrow \mathbb{R}$ which must satisfy the following properties

$$V(0) = 0 \wedge \forall x \in \mathbb{R}^n \setminus \{0\}, V(x) > 0 \wedge \lim_{\|x\| \rightarrow \infty} V(x) = \infty \quad (1)$$

$$\forall x \in \mathbb{R}^n, V(f(x)) - V(x) \leq 0. \quad (2)$$

It is shown for example in [14] that exhibiting such a function proves the so-called Lyapunov stability of the system, meaning that its state variables will remain bounded through time. Equation (2) expresses the fact that the function $k \mapsto V(x_k)$ decreases, which, combined with (1), shows that the state variables remain in the bounded sublevel-set $\{x \in \mathbb{R}^n | V(x) \leq V(x_0)\}$ at all instants $k \in \mathbb{N}$.

In the case of Linear Time Invariant systems (of the form $x_{k+1} = Ax_k$, with $A \in \mathbb{R}^{n \times n}$), one can always look for V as a quadratic form in the state variables of the system: $V(x) = x^T P x$ with $P \in \mathbb{R}^{n \times n}$ a symmetric matrix such that

$$P \succ 0 \quad (3)$$

$$A^T P A - P \preceq 0 \quad (4)$$

where “ $P \succ 0$ ” means that the matrix P is positive definite, i.e. for all non-zero vector $x, x^T P x > 0$.

Now, to account for the presence of an external input to the system (which is usually the case with controllers: they use data collected from sensors to generate their output), the model is usually extended into the form

$$x_{k+1} = Ax_k + Bu_k, \|u_k\|_\infty \leq 1. \quad (5)$$

To study this difference equation as precisely as possible, another model, expressing the behavior of the controlled system (the plant), is usually introduced. The two systems taken together form a closed system with no inputs which can be analyzed by looking for a P matrix matching the criteria mentioned before. Such an analysis is referred to as 'closed loop stability analysis'. Here we seek not to model the plant, instead we only require for $\|u\|_\infty$ to remain bounded³. Then, through a slight reinforcement of Equation (4) into

$$A^T P A - P \prec 0 \quad (6)$$

we can still guarantee that the state variables of (5) will remain in a sub-level set $\{x \in \mathbb{R}^n \mid x^T P x \leq \lambda\}$ (for some $\lambda > 0$), which is an ellipsoid in this case. This approach only enables us to study control laws that are inherently stable, i.e stable when taken separately from the plant they control. Nevertheless a wide range of controllers remain that can be analyzed, and this encompasses in particular all those handled by Astrée. In addition, inherent stability is required in a context of critical applications.

These stability proofs have the very nice side effect that they provide a quadratic invariant on the state variables, which can be used at the code level to find bounds on the program variables. Furthermore, there are many P matrices that fulfill the equations described above. This gives some flexibility as to the choice of such a matrix: by adding relevant constraints on P , one can obtain increasingly better bounds.

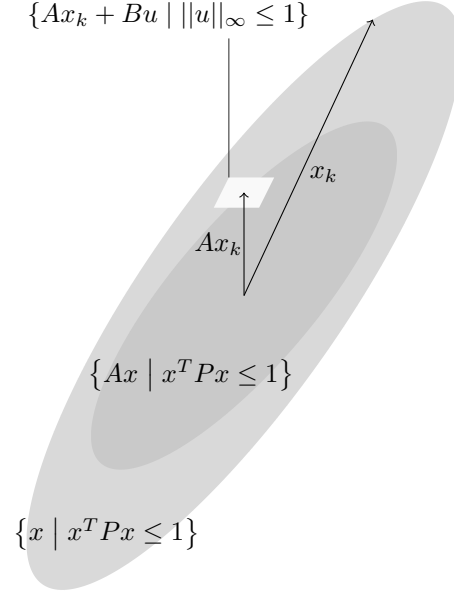


Fig. 2: Illustration of the stability concepts: if x_k is in the light gray ellipse, then, after a time step, Ax_k is in the dark gray ellipse, which is exactly what is expressed by Equation (4). The white box represents the potential values of x_{k+1} after adding the effect of the bounded input u_k . We see here the necessity that the dark gray ellipse be strictly included in the light gray one, which is the stronger condition expressed by Equation (6).

³ While we could consider different bounds for each component of the input u , we will only deal with $\|u\|_\infty \leq 1$ for simplicity of the exposition.

3 Overall method

3.1 Separate shape and ratio

We keep the same overall representation as Feret [9,10], representing an ellipsoid by a pair (P, λ) where $P \in \mathbb{R}^{n \times n}$ is a symmetric positive definite matrix giving the *shape* of the ellipsoid and $\lambda \in \mathbb{R}$ a scalar giving its *ratio*. The represented ellipsoid is then the set of all $x \in \mathbb{R}^n$ such that $x^T P x \leq \lambda$, i.e. the concretization function γ is given by $\gamma : (P, \lambda) \mapsto \{x \in \mathbb{R}^n \mid x^T P x \leq \lambda\}$. To avoid having multiple representations for the same ellipsoid⁴ we can normalize P for instance by requiring its largest coefficient to be 1. The underlying lattice also remains the same. In particular the join of two abstract values (P, λ) and (P', λ') is $(P, \max(\lambda, \lambda'))$ if $P = P'$ and \top otherwise.

This seemingly strange choice at first sight allows us to decompose the computation in two successive steps

1. first determine the shape of the ellipsoid by choosing a well suited matrix P ;
2. then find the smallest possible ratio λ such that $x \in \gamma(P, \lambda)$ is an invariant.

Various methods for both steps are detailed and compared in Sections 4 and 5.

3.2 Instrumentation: use of semidefinite programming

To perform the aforementioned computations we rely heavily on semidefinite programming [4,12]. These tools allow us to compute in polynomial time a solution to a linear matrix inequality (LMI) while minimizing a linear objective function. A LMI is an inequality of the form

$$A_0 + \sum_{i=1}^k y_i A_i \succeq 0$$

where the A_i are known matrices, the y_i are the unknowns and “ $P \succeq 0$ ” means that the matrix P is positive semidefinite, i.e. $x^T P x \geq 0$ for all vector x . Indeed we can easily have unknown matrices since a matrix $A \in \mathbb{R}^{n \times n}$ can be expressed as $\sum_{i=1, j=1}^{n, n} A_{i,j} E^{i,j}$, where $E^{i,j}$ is the matrix with zeros everywhere except a one at line i and column j . Likewise multiple LMIs can be grouped into one since $A \succeq 0 \wedge B \succeq 0$ is equivalent to $\begin{pmatrix} A & 0 \\ 0 & B \end{pmatrix} \succeq 0$.

We will also have to deal with some implications which will be achieved by transforming them into a LMI thanks to the following theorem.

Theorem 1 (S-Procedure). *For any $P \in \mathbb{R}^{n \times n}$, $a, a' \in \mathbb{R}^n$ and $b, b' \in \mathbb{R}$, the two following conditions are equivalent*

1. $\forall x \in \mathbb{R}^n, x^T P x + 2a^T x + b \geq 0 \Rightarrow x^T P' x + 2a'^T x + b' \geq 0$

⁴ For instance $(P, 2\lambda)$ and $(\frac{P}{2}, \lambda)$ represent exactly the same ellipsoid.

$$2. \exists \tau \in \mathbb{R}, \tau > 0 \wedge \begin{pmatrix} P' & a' \\ a'^T & b' \end{pmatrix} - \tau \begin{pmatrix} P & a \\ a^T & b \end{pmatrix} \succeq 0$$

Proof. Soundness (2 \Rightarrow 1) is obvious. A proof of completeness (1 \Rightarrow 2) can be found in [15].

4 Choosing the shape of the ellipsoid

As was presented in Sect. 2, any positive definite matrix P satisfying the Lyapunov equation

$$A^T P A - P \prec 0 \tag{7}$$

will yield a proof of stability and provide *some* bound on the variables. However, additional constraints on P can be introduced that make it possible to obtain better results than others.

While in Control Theory the existence of such ellipsoids is sufficient to prove stability of the system, we are here interested in characterizing it concretely. Investigating heuristically multiple possible shapes allows us to find one which is more adequate, i.e. more precise, with respect to the analyzed system.

The following subsections describe three different types of additional constraints on P and their respective advantages.

4.1 Minimizing condition number

Graphically, the condition number of a positive definite matrix expresses a notion similar to that addressed by eccentricity for ellipses in dimension 2. It measures how 'close' to a circle (or its higher dimension equivalent) the resulting ellipsoid will be. Multiples of the identity matrix, which all represent a circle, have a condition number of 1. Thus one idea of constraint we can impose on P is to have its condition number as close to 1 as possible. One reason is that 'flat' ellipsoids can yield a very bad bound on one of the variables. This is done [3] by minimizing a new variable, r , in the following matrix inequality

$$I \preceq P \preceq rI.$$

This constraint, along with the others (Lyapunov equation, positive definiteness, ...), can be expressed as an LMI, which is solved using the semi definite programming techniques mentioned in Section 2.

4.2 Preserving the shape

Another approach [22] is to minimize $r \in (0, 1)$ in the following inequality

$$A^T P A - rP \preceq 0.$$

Intuitively, this corresponds to finding the shape of ellipsoid that gets 'preserved' the best when the update $x_{k+1} = Ax_k$ is applied. This is the choice implicitly

made in [9] for a particular case of 2×2 matrices A . With this technique however, the presence of a quadratic term rP in the equation prevents the use of usual LMI solving tools 'as is'. To overcome this we chose an approach where we try a value for r and refine it by dichotomy. Only a few steps are required to obtain a good approximation of the optimal value.

4.3 All in one

The two previous methods were based only on A , completely abstracting B away, which could lead to rather coarse abstractions. We try here to take both A and B into account by finding the smallest possible P such that

$$\forall x, \forall u, \|u\|_\infty \leq 1 \Rightarrow x^T P x \leq 1 \Rightarrow (Ax + Bu)^T P (Ax + Bu) \leq 1$$

which, using the S procedure, amounts to the existence of $\tau_i > 0$ such that

$$\begin{pmatrix} -A^T P A & -A^T P B e_i \\ -e_i^T B^T P A & 1 - e_i^T B^T P B e_i \end{pmatrix} - \tau_i \begin{pmatrix} -P & 0 \\ 0 & 1 \end{pmatrix} \succeq 0$$

for all the vertices e_i of the hypercube of dimension p , the number of inputs. The rationale behind this formula is explained in section 5.2. This is not an LMI since τ and P are both variables but a reasonably good solution can be found by trying various values of τ between some $\tau_{min} \in (0, 1)$, which can be found by dichotomy, and 1.

4.4 Comparison and combination

There is no proof that one method always performs better than the others, and, for each method, there exists examples where it performs better than the other two, see Sect. 8. It appears, however, that the third method, albeit a little more costly, yields the best bounds in general. In fact the cost is also debatable since, despite being costlier, it does not require the search for the ratio, a necessary step for the first two methods described in Sections 4.1 and 4.2.

In any case, the methods are not exclusive of each other and can be combined: the resulting (sound) value will be the intersection of the projection of each obtained ellipsoids. Having multiple, not-always-comparable values will only yield more precise results.

5 Finding a stable ratio

Now that we have chosen a matrix P , we need to find a ratio λ such that $x^T P x \leq \lambda$ is an invariant for the whole system $x_{k+1} = Ax_k + Bu_k$ with a bounded input u that satisfies $\|u\|_\infty \leq 1$. The existence of such a λ is guaranteed by the choice of P as a solution of the Lyapunov inequality (7). Those λ are exactly the fixpoints of the function mapping λ_k to the maximum of λ_k and the least λ_{k+1} such that

$$\forall x_k \in \mathbb{R}^n, u_k \in \mathbb{R}^p, \|u_k\|_\infty \leq 1 \Rightarrow x_k^T P x_k \leq \lambda_k \Rightarrow x_{k+1}^T P x_{k+1} \leq \lambda_{k+1} \quad (8)$$

where $x_{k+1} = Ax_k + Bu_k$. We are of course interested in the least fixpoint.

5.1 Initial ratio λ_0

Since the system starts in state x_0 , we initialize λ_0 as $x_0^T P x_0$. If instead of a simple point the initial conditions are only known to lie in a polyhedron, we just have to take the maximum of $x^T P x$ among all vertices x of the polyhedron.

5.2 One iteration

Given some λ_k , we want to compute the least λ_{k+1} satisfying equation (8). By a convexity argument⁵, it is enough to have the following for every vertex $e_i, i \in \llbracket 1, 2^p \rrbracket$ of the hypercube⁶ $\{u_k \mid \|u_k\|_\infty \leq 1\}$ of dimension p

$$\forall x_k \in \mathbb{R}^n, x_k^T P x_k \leq \lambda_k \Rightarrow (A x_k + B e_i)^T P (A x_k + B e_i) \leq \lambda_{k+1}.$$

Using the S-procedure⁷ we get the equivalent formulation

$$\forall i \in \llbracket 1, 2^p \rrbracket, \exists \tau_i, \tau_i \geq 0 \wedge \begin{pmatrix} -A^T P A & -A^T P B e_i \\ -e_i^T B^T P A & \lambda_{k+1} - e_i^T B^T P B e_i \end{pmatrix} - \tau_i \begin{pmatrix} -P & 0 \\ 0 & \lambda_k \end{pmatrix} \succeq 0$$

which is an LMI in λ_{k+1} and the τ_i which is solved by minimizing λ_{k+1} .

We can notice that, by a symmetry argument, we can forget about half of the e_i as depicted on figure 3, page 9.

5.3 Iterating to fixpoint and classical widening

Now we can compute Kleene iterates but it will be slow to converge to a fixpoint. To accelerate, we can use a widening with thresholds, which allows us to find a value for λ up to a factor q of the least one by using a sequence of powers of q as thresholds.

5.4 An alternative to classical widening

When looking for a good postfixpoint we are indeed looking for a small λ satisfying the following equation

$$\forall i \in \llbracket 1, 2^p \rrbracket, \tau_i > 0 \wedge \begin{pmatrix} -A^T P A & -A^T P B e_i \\ -e_i^T B^T P A & \lambda - e_i^T B^T P B e_i \end{pmatrix} - \tau_i \begin{pmatrix} -P & 0 \\ 0 & \lambda \end{pmatrix} \succeq 0. \quad (9)$$

This is not an LMI because of the τ_i but if we used the method described in Sect. 4.2 for choosing the shape of P , we have obtained⁸ a parameter $r \in (0, 1)$ such that $\tau_i \in [r, 1]$. Computing the smallest λ satisfying the following LMI then directly gives a postfixpoint

$$\forall i \in \llbracket 1, 2^p \rrbracket, \begin{pmatrix} -A^T P A & -A^T P B e_i \\ -e_i^T B^T P A & \lambda - e_i^T B^T P B e_i \end{pmatrix} - \frac{r+1}{2} \begin{pmatrix} -P & 0 \\ 0 & \lambda \end{pmatrix} \succeq 0.$$

⁵ See figure 2 for a graphical illustration of this.

⁶ A major drawback of the approach is that the number of vertices is exponential in the number of inputs p . We could design a cheaper abstraction but it would be coarser, in addition the number of inputs p often remains reasonable.

⁷ See theorem 1.

⁸ Otherwise we can still recompute such a parameter r .

5.5 Refining a postfixpoint by dichotomy

Once we have found a postfixpoint λ_{max} using widening with thresholds, we can refine it through decreasing iterations with narrowing but this usually does not lead quickly to anything close to the least fixpoint. However, an interesting property of this least fixpoint λ_{min} is that $\lambda_{k+1} \leq \lambda_k$ exactly when $\lambda_k \geq \lambda_{min}$, then enabling to efficiently and tightly overapproximate it by a dichotomy testing satisfiability of the LMI 9 for values of λ between zero⁹ and λ_{max} .

6 Back to intervals

While quadratic forms precisely overapproximate the set of reachable states of linear systems subject to bounded inputs, they are hardly usable as such in conjunction with other abstractions. We can solve LMIs to project the obtained ellipsoid and get bounds on the variables, $x_i \in [-a, a]$ with a the least value such that

$$\begin{pmatrix} 0 & -\frac{e_i}{2} \\ -\frac{e_i^T}{2} & a \end{pmatrix} - \tau \begin{pmatrix} -P & 0 \\ 0 & \lambda \end{pmatrix} \succeq 0$$

where e_i is the i th vector of the canonical base.

This is not limited to intervals, the same thing can be done for octagons [16] or more generally linear [20] or even quadratic [1,13] templates.

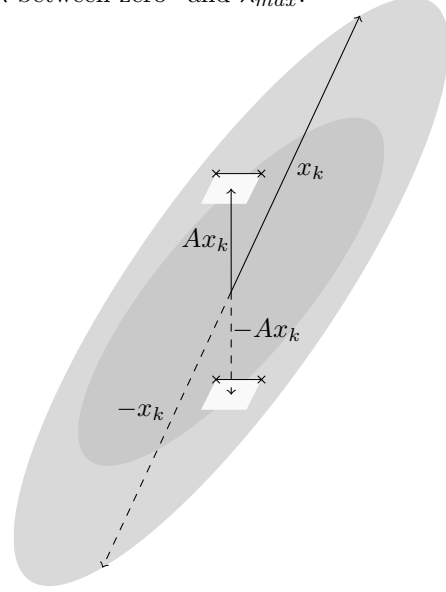


Fig. 3: We can forget half of the vertices of the white box as they will be taken into account on the opposite side.

7 Floating point issues and soundness

Two fundamentally different issues with floating point numbers must be considered

the analyzed system contains floating point computations with rounding errors making it behave differently from the way it would if the same computations were done with real numbers, this is discussed in section 7.1;

the implementation of the abstract domain is also carried out with floating point computations for the sake of efficiency, this usually works well in practice but can give erroneous results, hence the need for some a posteriori validation, see section 7.2 for further details.

⁹ Or, better, the last prefixpoint encountered during the widening iterations.

7.1 Taking rounding errors into account

The sum of two floating point values is, in all generality, not representable as a floating point value and must consequently be rounded. The accumulation of rounding errors can potentially lead to far different results from the ones expected with real numbers, thus floating point computations must be taken into account in our analysis [17].

The rounding errors can be of two different types :

- for normalized numbers represented with a fixed number of bits, we get a relative error: $\text{round}(a + b) \in [(1 - \epsilon)(a + b), (1 + \epsilon)(a + b)]$;
- for denormalized numbers (i.e ones very close to 0), we get an absolute error: $\text{round}(a + b) \in [a + b - \omega, a + b + \omega]$.

A common and easy solution to take both possible errors into account is to sum them which in practice leads only to a very slight overapproximation: $\text{round}(a + b) \in [(1 - \epsilon)(a + b) - \omega, (1 + \epsilon)(a + b) + \omega]$. Although only addition is illustrated here, the method works exactly the same way for any other floating point operation. The actual values of ϵ and ω depend on the characteristics of the considered floating point system. For instance we will take $\epsilon = 2^{-23}$ and $\omega = 2^{-149}$ for single precision¹⁰.

Combining these elementary errors we get a simple postprocessing for each iteration of Sect. 5.2 to soundly overapproximate rounding errors¹¹.

7.2 Checking soundness of the result

Because the LMI solver is implemented with floating point computations, we have no guarantee on the results it provides. Hence the need to check them. This amounts to checking that a given matrix is actually positive definite.

We currently do this using Sylvester’s criterion, stating that a matrix $A \in \mathbb{R}^{n \times n}$ is positive definite if and only if all its leading principal minors are positive, and interval arithmetic to check their positiveness.

This is suited for small values of n but doesn’t scale well since computing a determinant requires $\Omega(n!)$ arithmetic operations. While it remains an easy early working solution, there is probably room for improvement on this point and the use of an interval Cholesky decomposition [11] looks promising.

8 Experimental results

All the elements presented in this paper have been implemented as an autonomous linear system analysis engine. The tool is composed of three parts:

¹⁰ Type `float` in C.

¹¹ Further details can be found in the comments of our implementation.

	Method	t_1	λ	λ_{∇}	Bounds	t_2	Valid. t_3
Ex. 1 From [10, slides] n=2, 1 input	I	0.07	fp 131072 \perp	105352	[140.4; 189.9]	0.48 \perp	0.04 \perp
	P	0.16	fp 128.0 τ 96.8	96.0	[22.2; 26.5]	0.35 0.28	0.03 0.03
	U	0.35	$1 + \epsilon$		[16.2; 17.6]	0.20	0.03
Ex. 2 From [10, slides] n=4, 1 input	I	0.09	fp 2048 τ 1376	1372	[18.1; 25.2; 24.3; 33.7]	0.49 0.40	0.07 0.07
	P	0.27	fp 8.0 τ 6.4	4.2	[6.3; 7.7; 2.2; 3.4]	0.35 0.28	0.07 0.07
	U	0.85	$1 + \epsilon$		[1.7; 2.0; 2.2; 2.5]	0.21	0.07
Ex. 3 Discretized lead-lag controller n=2, 1 input	I	0.08	fp 262144 τ \perp	204654	[391.8; 21.6]	0.53 \perp	0.03 \perp
	P	0.18	fp 2048 τ 1632	1283	[36.2; 36.1]	0.47 0.35	0.03 0.04
	U	0.35	$1 + \epsilon$		[31.7; 21.4]	0.20	0.03
Ex. 4 Linear quadratic gaussian regulator n=3, 1 input	I	0.09	fp 16.0 τ 10.9	10.3	[1.2; 0.9; 0.5]	0.39 0.33	0.04 0.05
	P	0.19	fp 1.0 τ 1.1	0.7	[0.9; 0.9; 0.9]	0.29 0.25	0.05 0.04
	U	0.50	$1 + \epsilon$		[0.7; 0.5; 0.4]	0.22	0.04
Ex. 5 Observer based controller for a coupled mass system n=4, 2 inputs	I	0.09	fp 512.0 τ 323.0	304.6	[9.8; 8.9; 11.0; 16.8]	0.51 0.43	0.13 0.12
	P	0.25	fp 32.0 τ 28.6	24.3	[5.7; 5.6; 6.4; 10.1]	0.42 0.33	0.12 0.12
	U	0.96	$1 + \epsilon$		[5.0; 5.0; 4.8; 4.7]	0.22	0.12
Ex. 6 Butterworth low-pass filter n=5, 1 input	I	0.10	fp 128.0 τ 113.1	102.4	[7.5; 8.7; 6.1; 7.0; 6.5]	0.46 0.38	0.20 0.20
	P	0.31	fp 8.0 τ 7.7	7.1	[3.6; 5.0; 4.7; 8.1; 8.9]	0.39 0.29	0.20 0.20
	U	1.17	$1 + \epsilon$		[2.2; 1.1; 1.9; 2.0; 2.9]	0.24	0.20
Ex. 7 Dampened oscillator from [1] n=2, no input	I	0.07	fp 353.6 τ 353.6	353.6	[1.7; 2.1]	0.22 0.26	0.03 0.03
	P	0.15	fp 3.0 τ 3.0	3.0	[2.0; 2.0]	0.22 0.20	0.03 (\perp) 0.03 (\perp)
	U	0.34	$1 + \epsilon$		[1.5; 1.5]	0.16	0.03
Ex. 8 Harmonic oscillator from [1] n=2, no input	I	0.08	fp 22.9 τ 22.9	22.9	[1.5; 1.5]	0.24 0.23	0.03 0.03
	P	0.24	fp 2.0 τ 2.0	2.0	[1.5; 1.5]	0.24 0.22	0.03 (\perp) 0.03 (\perp)
	U	0.31	$1 + \epsilon$		[1.5; 1.5]	0.16	0.14

Table 1: Result of the experiments: quadratic invariants computation. Times are expressed in seconds, t_1 is the time spent to compute the shape of the ellipsoid, t_2 is the time spend to find the appropriate ratio λ and project the resulting invariant on intervals and t_3 is the time needed to validate the stability of the resulting ellipsoid with Gappa. I, P and U are respectively the methods of Sections 4.1, 4.2 and 4.3. λ_{∇} denotes the refined value of λ by dichotomy.

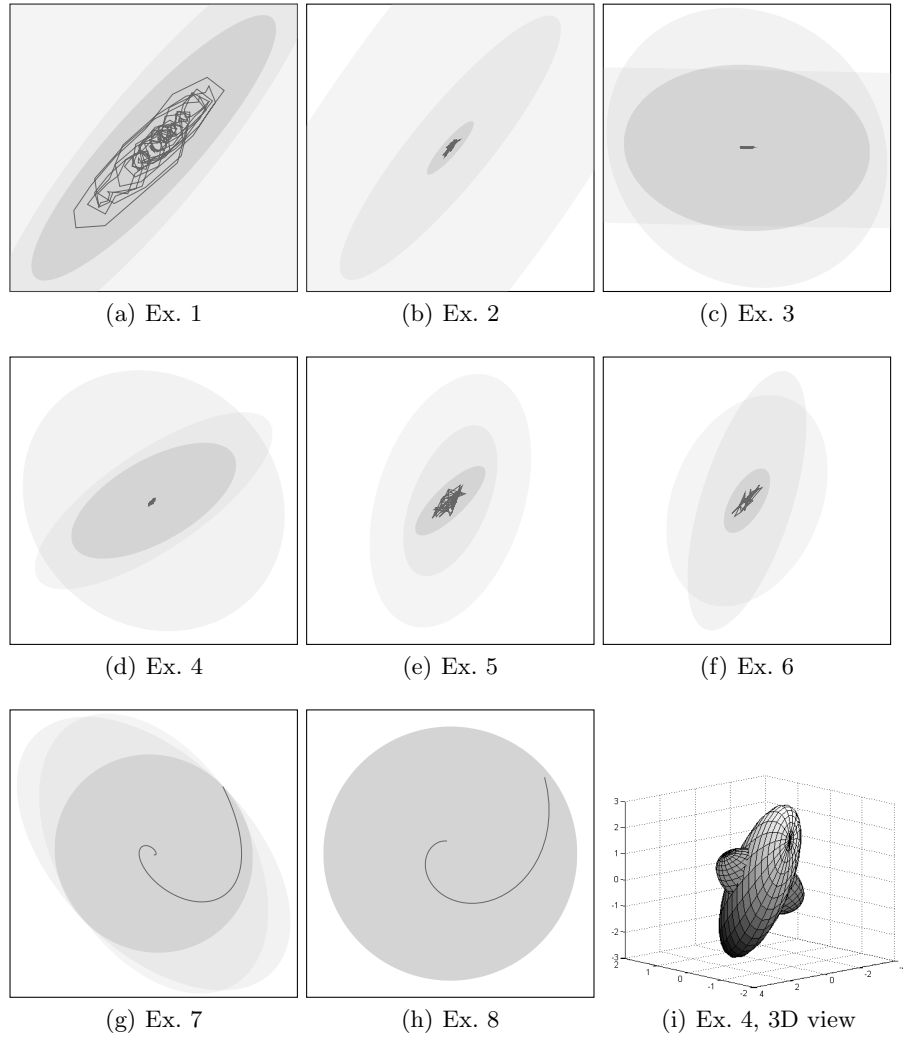


Fig. 4: Comparison of obtained ellipsoids by methods of Sections 4.1, 4.2 and 4.3 from lighter to darker, plus a random simulation trace ((b), (d), (e) and (f) are cuts along some plane).

- The core mathematical computations are done with Scilab [21], mainly with the LMI solver [18] from an OCaml front-end. This part is a set of functions that implement the algorithms presented in Sections 4, 5, 6 and 7.1, as well as projections of ellipsoids over intervals. Computation in Scilab are done using double precision floats.
- The front-end is an OCaml code using rational numbers (Num library). It loads the A and B matrices and interacts with Scilab to compute the different sequence of calls to Scilab functions.
- A last part, also in Ocaml, interfaces the obtained quadratic form with Gappa [8] to ensure its stability as explained in Sect. 7.2. Gappa is a tool intended to help verifying and formally proving properties on numerical programs dealing with floating-point or fixed-point arithmetic.

The code is released under a GPLv2 license and is available at <http://cavale.gforge.enseeiht.fr/>.

Experiments were conducted on a set of stable linear systems. These systems were extracted from [10], [1] or from basic controllers found in the literature. Table 1 illustrates the value computed using the different techniques as well as the time spent at each step. Figure 4 compares some plots of the obtained quadratic forms depending on the approach used to find the ellipsoid.

9 Related work

Many work in abstract interpretation, and its use to analyze programs, focus on linear patterns to abstract properties. However few work address non linear invariant synthesis.

Feret’s work [9,10] on the one hand is a practical approach to the problem. Its goal is to address the need by Astrée to handle the linear filters present in Airbus’ real time software. As mentioned earlier, this effort is a strict subset of ours and should be comparable on this subset.

On the other hand, there are work that target similar properties but are theory oriented and motivated. One can cite the Lagrangian Relaxation approach applied to program termination analysis as introduced by Cousot in [4] and Roozbehani, Féron and Megretski in [19], or the works of Adjé et al. [1] and Gawlitza and Seidl [12] on policy iterations and non linear forms. The latter two aim at replacing a Kleene based fixpoint computation by a symbolic reasoning based on semi-definite programming. They are more inspired by theoretical results leading to the analysis. [1,2] even cites the existence of Lyapunov based invariant as a prerequisite for the method. These works are more general than ours: they address the analysis of non linear systems, even with non convex properties. However none of them automatically finds the appropriate shape: templates need to be given, e.g by providing the Lyapunov functions, which we are automatically computing. They also do not address the floating point issues.

Our work should be considered as an in-between solution. It takes ideas from control theory results but targets the analysis of specific realistic systems.

Furthermore it addresses floating point errors as well as the validity analysis of the obtained invariants.

10 Concluding remarks and perspectives

We have presented a set of analyses allowing us to characterize quadratic invariants, i.e. ellipsoids, for a subset of linear systems: inherently stable linear systems subject to bounded inputs.

Most of the critical embedded control command systems rely on such linear systems. But intervals and linear invariants in general will not allow to precisely describe their state space.

This analysis is based on ideas from control theory. They are used to prove the stability of the system by exhibiting a proof of existence of a so-called Lyapunov quadratic form.

This work addresses the explicit computation of such a form by exploring the instantiation of multiple generic templates to find the most appropriate ellipsoids to bound the analyzed system.

Our effort also considers floating point errors and addresses the validity of the computed solution. It has been implemented and applied on several examples. The reduced product between the different templates instantiated gives extremely precise results as illustrated by the experimentations.

From here on, perspectives are easily identified: a first area of research would be to enlarge the scope of considered systems and to handle saturation operators that are often used in such critical embedded controllers; a second area would be to integrate this ellipsoid analysis in our Lustre code analyzer to strengthen its results.

References

1. Assalé Adjé, Stéphane Gaubert, and Eric Goubault. Coupling policy iteration with semi-definite relaxation to compute accurate numerical invariants in static analysis. In Andrew D. Gordon, editor, *ESOP*, volume 6012 of *Lecture Notes in Computer Science*, pages 23–42. Springer, 2010.
2. Fernando Alegre, Eric Feron, and Santosh Pande. Using ellipsoidal domains to analyze control systems software. 2009. <http://arxiv.org/abs/0909.1977>.
3. Stephen Boyd, Laurent El Ghaoui, Éric Féron, and Venkataramanan Balakrishnan. *Linear Matrix Inequalities in System and Control Theory*, volume 15 of *Studies in Applied Mathematics*. SIAM, Philadelphia, PA, June 1994.
4. Patrick Cousot. Proving program invariance and termination by parametric abstraction, lagrangian relaxation and semidefinite programming. In *Sixth International Conference on Verification, Model Checking and Abstract Interpretation (VMCAI'05)*, pages 1–24, Paris, France, LNCS 3385, January 17–19 2005. Springer.
5. Patrick Cousot and Radhia Cousot. Abstract interpretation: a unified lattice model for static analysis of programs by construction or approximation of fixpoints. In *Conference Record of the Fourth Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 238–252, Los Angeles, California, 1977. ACM Press, New York, NY.

6. Patrick Cousot and Radhia Cousot. Systematic design of program analysis frameworks. In *Conference Record of the Sixth Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 269–282, San Antonio, Texas, 1979. ACM Press, New York, NY.
7. Patrick Cousot, Radhia Cousot, Jérôme Feret, Laurent Mauborgne, Antoine Miné, David Monniaux, and Xavier Rival. Combination of abstractions in the ASTRÉE static analyzer. In M. Okada and I. Satoh, editors, *Eleventh Annual Asian Computing Science Conference (ASIAN'06)*, pages 272–300, Tokyo, Japan, LNCS 4435, 2007. Springer.
8. Florent de Dinechin, Christoph Quirin Lauter, and Guillaume Melquiond. Certifying the floating-point implementation of an elementary function using gappa. *IEEE Trans. Computers*, 60(2):242–253, 2011.
9. Jérôme Feret. Static analysis of digital filters. In *European Symposium on Programming (ESOP'04)*, number 2986 in LNCS. Springer-Verlag, 2004.
10. Jérôme Feret. Numerical abstract domains for digital filters. In *International workshop on Numerical and Symbolic Abstract Domains (NSAD 2005)*, 2005.
11. Jürgen Garloff. Pivot tightening for the interval cholesky method. In *Proc. in Applied Mathematics and Mechanics*, pages 549–550, 2010.
12. Thomas Gawlitza and Helmut Seidl. Precise fixpoint computation through strategy iteration. In Rocco De Nicola, editor, *ESOP*, volume 4421 of *Lecture Notes in Computer Science*, pages 300–315. Springer, 2007.
13. Thomas Martin Gawlitza and Helmut Seidl. Computing relaxed abstract semantics w.r.t. quadratic zones precisely. In Radhia Cousot and Matthieu Martel, editors, *SAS*, volume 6337 of *Lecture Notes in Computer Science*, pages 271–286. Springer, 2010.
14. Wassim M. Haddad and Vijay S. Chellaboina. *Nonlinear Dynamical Systems and Control: A Lyapunov-Based Approach*. Princeton University Press, 2008.
15. Ulf T. Jönsson. A lecture on the S-Procedure, 2001.
16. Antoine Miné. The octagon abstract domain. In *AST 2001 in WCRE 2001*, IEEE, pages 310–319. IEEE CS Press, October 2001.
17. David Monniaux. The pitfalls of verifying floating-point computations. *ACM Trans. Program. Lang. Syst.*, 30(3), 2008.
18. Ramine Nikoukhah, François Delebecque, and Laurent El Ghaoui. LMITOOL: a Package for LMI Optimization in Scilab User's Guide. Research Report RT-0170, INRIA, February 1995. Projet META2.
19. Mardavij Roozbehani, Eric Feron, and Alexandre Megretski. Modeling, optimization and computation for software verification. In Manfred Morari and Lothar Thiele, editors, *HSCC*, volume 3414 of *Lecture Notes in Computer Science*, pages 606–622. Springer, 2005.
20. Sriram Sankaranarayanan, Michael Colón, Henny B. Sipma, and Zohar Manna. Efficient strongly relational polyhedral analysis. In E. Allen Emerson and Kedar S. Namjoshi, editors, *VMCAI*, volume 3855 of *Lecture Notes in Computer Science*, pages 111–125. Springer, 2006.
21. Scilab Team. Scilab. <http://www.scilab.org>.
22. Qinqing Yang. *Minimum Decay Rate of a Family of Dynamical Systems*. PhD thesis, Stanford University, 1992.