ELSEVIER

Systems & Control Letters 41 (2000) 135–140

SYSTEMS & CONTROL LETTERS

www.elsevier.com/locate/sysconle

The boundedness of all products of a pair of matrices is undecidable ☆

Vincent D. Blondel[a, *], John N. Tsitsiklis[b]

[a] *Division of Applied Mathematics, Center CESAME, Université Catholique de Louvain, Avenue Georges Lemaitre 4, B-1348 Louvain-la-Neuve, Belgium*

[b] *Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Cambridge, MA 02139, USA*

Received 6 February 1999; received in revised form 20 May 2000

Abstract

We show that the boundedness of the set of all products of a given pair Σ of rational matrices is undecidable. Furthermore, we show that the joint (or generalized) spectral radius $\rho(\Sigma)$ is not computable because testing whether $\rho(\Sigma) \leqslant 1$ is an undecidable problem. As a consequence, the robust stability of linear systems under time-varying perturbations is undecidable, and the same is true for the stability of a simple class of hybrid systems. We also discuss some connections with the so-called "finiteness conjecture". Our results are based on a simple reduction from the emptiness problem for probabilistic finite automata, which is known to be undecidable.

Keywords: Computability; Decidability; Matrix semigroup; Joint spectral radius; Generalised spectral radius; Finiteness conjecture; Robust stability analysis; Linear time-varying systems

1. Introduction

Let Σ be a finite set of real $n \times n$ matrices. We consider products of the form $A_t A_{t-1} \cdots A_1$, where each A_i is an arbitrary element of Σ. More specifically, we are interested in the largest possible rate of growth of such products. Issues of this type arise naturally when considering linear time-varying systems of the form $x_{t+1} = A_t x_t$, as well as in many other contexts; see [22,8].

One measure of growth of such matrix products is provided by the *joint spectral radius* $\hat{\rho}(\Sigma)$ [19], which is defined by

$$\hat{\rho}(\Sigma) = \limsup_{t \to \infty} \hat{\rho}_t(\Sigma),$$

where

$$\hat{\rho}_t(\Sigma) = \max_{A_1,\ldots,A_t \in \Sigma} \|A_t \cdots A_1\|^{1/t},$$

and $\|\cdot\|$ is some matrix norm. The value of $\hat{\rho}(\Sigma)$ turns out to be independent of the choice of the norm. Furthermore, if the matrix norm has the property $\|AB\| \leqslant \|A\| \cdot \|B\|$ (e.g., if it is an induced norm), then $\hat{\rho}_t(\Sigma)$ converges and we also have

$$\hat{\rho}(\Sigma) = \lim_{t \to \infty} \hat{\rho}_t(\Sigma) \leqslant \hat{\rho}_\tau(\Sigma), \quad \forall \tau.$$

Recall that the spectral radius of a single square matrix A is defined by

$$\rho(A) := \max\{|\lambda| \mid \lambda \text{ is an eigenvalue of } A\}.$$

☆ Supported by the NATO under Grant CRG-961115, by the AFOSR under Grant F49620-99-10320, and by the NSF under Grant ACI-9873339.

* Corresponding author. Tel.: +32-10-472-381; fax: +32-10-472-180.

E-mail addresses: blondel@inma.ucl.ac.be (V.D. Blondel), jnt@mit.edu (J.N. Tsitsiklis).

0167-6911/00/$ - see front matter
PII: S0167-6911(00)00049-9

The natural extension to a set of matrices leads to the *generalized spectral radius* $\rho(\Sigma)$, which is defined by

$$\rho(\Sigma) = \limsup_{t\to\infty} \rho_t(\Sigma),$$

where

$$\rho_t(\Sigma) = \max_{A_1,\dots,A_t\in\Sigma} \rho(A_t\cdots A_1)^{1/t}.$$

It is known that [13]

$$\rho_t(\Sigma) \leqslant \rho(\Sigma), \quad \forall t.$$

More importantly, for any finite set Σ of matrices, the generalized spectral radius is equal to the joint spectral radius [1]:

$$\rho(\Sigma) = \hat\rho(\Sigma).$$

Questions related to the computability of $\rho(\Sigma)$ have been posed in [23,13], but have largely remained open, with the exception of a negative result in [12] that refers to a restricted model of algebraic computation.

The spectral radius $\rho(\Sigma)$ can be approximated to any desired accuracy (keep computing the upper and lower bounds $\hat\rho_t(\Sigma)$ and $\rho_t(\Sigma)$ until they get sufficiently close), but unless P = NP, there is no polynomial-time approximation algorithm [24]. From this it follows that the problems of deciding whether $\rho(\Sigma)\leqslant 1$ or whether $\rho(\Sigma) < 1$ are NP-hard; see [20,10], as well as [9] for other relevant results; see also [4] for a general discussion.

Let us also note the "finiteness conjecture" (FC) which states that

$$\forall\Sigma\ \exists t \text{ such that } \rho_t(\Sigma) = \rho(\Sigma).$$

The finiteness conjecture is discussed by Lagarias and Wang [13], who note that if the FC is true, then the problem of determining whether $\rho(\Sigma) < 1$ is decidable. This is because if $\rho(\Sigma) < 1$, then there exists t such that $\hat\rho_t(\Sigma) < 1$, whereas if $\rho(\Sigma)\geqslant 1$, the finiteness conjecture implies that there exists t such that $\rho_t(\Sigma)\geqslant 1$. By checking both conditions for increasing values of t, one of them will be eventually satisfied and a decision will be made after a finite amount of computation. Note that for a single matrix the problem is decidable, because we can use Tarski's decision procedure to test whether all roots of the characteristic polynomial have modulus less than or equal to one.

In this contribution, we show (Theorem 2 in Section 2) that the problem of determining whether $\rho(\Sigma)\leqslant 1$ is undecidable. We prove this to be the case even if Σ consists of only two matrices. We also prove that the problem of determining whether the set of all products of two matrices is bounded, is undecidable.

It is unclear whether our result has any ramifications for the problem of deciding whether $\rho(\Sigma) < 1$. But it does invalidate a stronger version of the finiteness conjecture to be discussed in Section 3.

Our result also has a number of implications for problems in systems and control. First, it proves undecidability of a certain robust stability problem under time-varying uncertainty. In that sense, it complements negative (NP-hardness) results on the robust stability of linear systems in the presence of time-invariant uncertainty [5,16,18,21]. Second, it leads to an undecidability result for a simple class of hybrid systems. These implications are discussed in Section 4.

2. Main result

Let $\|.\|$ be some submultiplicative matrix norm and consider the following matrix problems.

UNIT SPECTRAL RADIUS
Input: A finite set Σ of $n\times n$ matrices with rational entries.
Question: Is the joint spectral radius $\rho(\Sigma)\leqslant 1$?

BOUNDED MATRIX PRODUCTS
Input: A finite set Σ of $n\times n$ matrices with rational entries.
Question: Is the set of all matrix products bounded?

In Theorem 1 we show how the emptiness problem for probabilistic finite automata (PFA) can be reduced to these problems. We then use this reduction in Theorem 2 for proving that both problems are undecidable.

A PFA consists of a finite alphabet $\mathscr{A}$, a finite state space $S = \{1,\dots,n\}$, a subset F of S, a nonnegative vector $\pi\in\mathbf{R}^n$ whose entries sum to 1, and for each $\alpha\in\mathscr{A}$, an $n\times n$ stochastic matrix P^α. (Recall that a matrix is said to be stochastic if it is nonnegative and the sum of the entries in any particular row is equal to 1.) Let η be a vector in $\mathbf{R}^n$ whose ith component is equal to 1 if $i\in F$, and is equal to 0 if $i\notin F$. For any word $w = \alpha_1\alpha_2\cdots\alpha_t\in\mathscr{A}^*$, we define

$$f(w) = \pi^{\mathrm T} P^{\alpha_1}\cdots P^{\alpha_t}\eta,$$

where the superscript T stands for transposition.

For an intuitive interpretation of these definitions, consider a nonhomogeneous Markov chain $\{x(t)\mid t = 0,1,\dots\}$ on the state space S, whose initial state $x(0)$ is chosen at random according to the probability vector π, and whose transition probability matrix for the ith transition is P^{α_i}. Then, $\pi^{\mathrm T}P^{\alpha_1}\cdots P^{\alpha_t}$ is the (row) vector of state probabilities after t transitions, and the quantity $f(w)=\pi^{\mathrm{T}}P^{\alpha_1}\cdots P^{\alpha_t}\eta$ equals the probability that $x(t)$ belongs to F. Note that a word $w=\alpha_1\alpha_2\ldots\alpha_t$ amounts to an *open-loop* policy for controlling this Markov chain over t stages. The emptiness problem for PFA is as follows.

PFA Emptiness
Input: A finite number of $n\times n$ nonnegative row-stochastic matrices P^α with rational entries, a zero–one n-vector η, a nonnegative n-vector π with rational entries whose entries sum to 1, and a rational number λ with $0<\lambda<1$.
Question: Is the set $\{w\,|\,f(w)=\pi^{\mathrm{T}}P^{\alpha_1}\cdots P^{\alpha_t}\eta>\lambda\}$ empty? (i.e., do we have $\pi^{\mathrm{T}}P^{\alpha_1}\cdots P^{\alpha_t}\eta\leqslant\lambda$ for all $w=\alpha_1\alpha_2\cdots\alpha_t$?)

Theorem 1. PFA Emptiness *is reducible to* Bounded Matrix Products *and to* Unit Spectral Radius. *The reduction can be chosen so that the resulting sets of matrices Σ have only two elements.*

Proof. We first show how to reduce PFA Emptiness to Unit Spectral Radius. Given a PFA involving a set $\{P^\alpha\,|\,\alpha\in\mathscr{A}\}$ of $n\times n$ stochastic matrices, and a rational number $\lambda\in(0,1)$, we consider the finite collection of matrices

$$\Sigma=\{P^\alpha\,|\,\alpha\in\mathscr{A}\}\cup\{P^*\},$$

where

$$P^*=\frac{1}{\lambda}\eta\pi^{\mathrm{T}}.$$

Thus, the ith row of P^* is zero if $i\notin F$, and is equal to π^{T}/λ if $i\in F$.[1] The theorem will be established by showing that $\rho(\Sigma)\leqslant1$ if and only if $f(w)>\lambda$ for all words $w\in\mathscr{A}^*$.

Suppose that there exists a word $w\in\mathscr{A}^*$ such that $f(w)>\lambda$. Then, there exists a finite sequence of matrices $P_1,\ldots,P_t$, chosen from the set $\{P^\alpha\,|\,\alpha\in\mathscr{A}\}$, such that

$$\pi^{\mathrm{T}}P_1\cdots P_t\eta>\lambda.$$

By right-multiplying both sides by π^{T}/λ, and using the definition $P^*=\eta\pi^{\mathrm{T}}/\lambda$, we obtain

$$\pi^{\mathrm{T}}P_1\cdots P_tP^*>\pi^{\mathrm{T}},$$

[1] For an intuitive view, if $x(t)\notin F$ and P^* is applied, it is as if the Markov chain is terminated. If on the other hand, $x(t)\in F$ and P^* is applied, the Markov chain is restarted with the initial distribution π, but amplified by a factor of $1/\lambda$. This amplification may well result in "probabilities" that are larger than 1, but the intuition goes through if one thinks in terms of flow volumes rather than probabilities.

so

$$\pi^{\mathrm{T}}P_1\cdots P_tP^*>(1+\varepsilon)\pi^{\mathrm{T}},$$

for some $\varepsilon>0$. This implies that the largest of the absolute values of the eigenvalues of $P_1\cdots P_tP^*$ is at least $1+\varepsilon$, and $\rho(\Sigma)\geqslant\rho_{t+1}(\Sigma)\geqslant(1+\varepsilon)^{1/(t+1)}>1$.

For the converse, let us assume that $f(w)\leqslant\lambda$ for every word w. We define the matrix

$$E=\frac{1}{\eta^{\mathrm{T}}\eta}\eta\eta^{\mathrm{T}},$$

and note that

$$EP^*=\frac{1}{\eta^{\mathrm{T}}\eta}\eta\eta^{\mathrm{T}}\frac{1}{\lambda}\eta\pi^{\mathrm{T}}=\frac{1}{\lambda}\eta\pi^{\mathrm{T}}=P^*.$$

Let us consider a product of length t of elements of Σ. Any such product is of the form $\Pi=Q_1P^*Q_2P^*Q_3\cdots P^*Q_{k+1}$, where k is the number of occurrences of P^* and each Q_i is a product of matrices in $\{P^\alpha\,|\,\alpha\in\mathscr{A}\}$ or the identity. Since $EP^*=P^*$, this product can be rewritten and grouped as

$$\Pi=(Q_1E)(P^*Q_2E)(P^*Q_3E)\cdots(P^*Q_kE)(P^*Q_{k+1}).$$

Let us consider any group other than the first and the last one. It is of the form

$$P^*QE=\frac{1}{\lambda}\eta\pi^{\mathrm{T}}\,Q\,\frac{1}{\eta^{\mathrm{T}}\eta}\eta\eta^{\mathrm{T}}.$$

Since $f(w)\leqslant\lambda$ for every word w, we have $\pi^{\mathrm{T}}Q\eta\leqslant\lambda$, and

$$P^*QE\leqslant\frac{1}{\eta^{\mathrm{T}}\eta}\,\eta\eta^{\mathrm{T}}=E,$$

where the inequality is to be interpreted componentwise. Noting also that $E^2=E$, we obtain

$$\Pi=(Q_1E)(P^*Q_2E)(P^*Q_3E)\cdots(P^*Q_kE)(P^*Q_{k+1})\leqslant Q_1EP^*Q_{k+1},$$

where we have made used of the nonnegativity of the matrices under consideration.

Consider the matrix norm $\|\cdot\|_1$, defined by

$$\|A\|_1=\max_i\sum_j|a_{ij}|.$$

Since Q_1 and Q_{k+1} are stochastic matrices, we have $\|Q_1\|_1=\|Q_{k+1}\|_1=1$, and $\|\Pi\|_1\leqslant\|EP^*\|_1$. Thus, with this choice of norm,

$$\hat\rho_t(\Sigma)\leqslant(\|EP^*\|_1)^{1/t},$$

and by taking the limit as $t\to\infty$, we obtain $\hat\rho(\Sigma)\leqslant1$.

Let us now show how to reduce the case of m matrices to the case of two matrices. The reduction is standard and is identical to the one used in [2]. Given a finite set $\Sigma = \{A_1, \ldots, A_m\}$ of matrices in $\mathbf{Q}^{n\times n}$, we define two $nm \times nm$ matrices by $A = \mathrm{diag}(A_1, \ldots, A_m)$ (i.e., A is block-diagonal with blocks $A_1, \ldots, A_m$ in that order) and

$$T = \begin{pmatrix} 0 & I_{n(m-1)} \\ I_n & 0 \end{pmatrix},$$

where I_r is the $r \times r$ identity matrix. Let $\bar{\Sigma} = \{A, T\}$. It is easily checked that $\rho(\Sigma) \leqslant 1$ if and only if $\rho(\bar{\Sigma}) \leqslant 1$, which leads to the desired result.

It remains to show how to reduce PFA Emptiness to Bounded Matrix Products. Notice therefore that for the family of matrices constructed above, we have $\rho(\Sigma) \leqslant 1$ if and only if the set of products of the matrices in the set Σ is bounded. To see this, note that if $\rho(\Sigma) > 1$, then the set of all products is clearly unbounded. On the other hand, for the case, where $\rho(\Sigma) \leqslant 1$, we have shown that the norm of any matrix product Π is bounded by $\|EP^*\|_1$. Furthermore, the technique used in the proof again allows us to restrict to the case of only two matrices. This last observation completes the proof. □

As a corollary of Theorem 1, we obtain the following:

Theorem 2. *The problems* Bounded Matrix Products *and* Unit Spectral Radius *are undecidable. They remain undecidable even in the special case where Σ consists of only two matrices.*

Proof. This is an immediate consequence of Theorem 1 and of the undecidability of PFA Emptiness.

Let us note at this point that a complete proof that PFA Emptiness is undecidable cannot be found in its entirety in the published literature. A proof (stated with a different terminology) is given in Theorem 6.17 in p. 190 of [17]. The proof given there is a few lines long and refers to a long cascade of lemma that appear at various places in the book. A full proof is hard to reconstruct.

Sketches of an alternative proof can be found in several recent references; see, e.g., [6,14,15]. Finally, a full proof can be found in the expanded version of [6]; see Theorem 3.2 in [7]. □

It should be noted that the problem of determining whether the number of elements in the set $\{A_1 \cdots A_t \,|\, A_i \in \Sigma,\ t = 1, 2, \ldots\}$ is *finite* is known to be decidable [11]. For matrices with integer entries, finiteness is equivalent to boundedness, and for this case, the problem Bounded Matrix Products is decidable.

3. Relation to the finiteness conjecture

As discussed in the introduction, if the finiteness conjecture is true, then the problem of determining whether $\rho(\Sigma) < 1$ is decidable. This is different than the problem we have considered, and the finiteness conjecture remains unresolved. Our results, however, disprove a somewhat stronger form of that conjecture.

Effective finiteness conjecture (EFC): *For any finite set Σ of square matrices with rational entries, there exists an* effectively computable *natural number $t(\Sigma)$ such that $\rho_{t(\Sigma)}(\Sigma) = \rho(\Sigma)$.*

Corollary 1. *The effective finiteness conjecture is false.*

Proof. Suppose that the EFC is true. Given Σ, we can first determine $t(\Sigma)$, and form all possible products of elements of Σ with $t(\Sigma)$ terms. We then have $\rho(\Sigma) \leqslant 1$ if and only if the spectral radius of all such products is bounded by 1, which can be tested in finite time. But this contradicts Theorem 2. □

4. Relation to control problems

The results of Theorem 2 have implications for the stability of time-varying and hybrid systems. A time-varying system

$$x_{t+1} = F_t(x_t),$$

is said to have bounded trajectories, if for every initial state x_0, the resulting sequence x_t is bounded.

Let Σ be a finite set of matrices and consider the family of time-varying systems

$$x_{t+1} = A_t x_t, \qquad A_t \in \Sigma, \tag{1}$$

where A_t is taken from a given finite set of matrices for each t. From the proof of Theorem 2 we deduce.

Corollary 2. *The problem of determining whether all the systems in family* (1) *have bounded trajectories, is undecidable.*

Theorem 2 also has an implication for a simple class of hybrid systems. We are given two rational $n \times n$ matrices A_+ and A_-, a rational vector $c \in \mathbf{Q}^n$, and we consider the system

$$x_{t+1} = \begin{cases} A_+ x_t & \text{if } c^T x_t \geqslant 0, \\ A_- x_t & \text{if } c^T x_t < 0. \end{cases} \quad (2)$$

Such systems were studied in [3], where it was established that deciding global convergence to the origin is NP-hard. In addition, it was shown that if the problem "$\rho(\Sigma) < 1$?" is ever shown undecidable, then global convergence to the origin will also be undecidable. A slight modification of that argument, together with Theorem 2, leads to the following result.

Corollary 3. *The problem of determining whether a given system of form* (2) *has bounded trajectories, is undecidable.*

Proof. Let A_+ and A_- be two given matrices. Consider the system described by a state vector (v_t, y_t, z_t), where v_t and y_t are scalars and z_t is a vector in $\mathbf{R}^n$, and the dynamics are of the form

$$\begin{pmatrix} v_{t+1} \\ y_{t+1} \\ z_{t+1} \end{pmatrix} = \begin{pmatrix} 1/4 & 0 & 0 \\ -1/4 & 1/2 & 0 \\ 0 & 0 & A_+ \end{pmatrix} \begin{pmatrix} v_t \\ y_t \\ z_t \end{pmatrix} \text{ when } y_t \geqslant 0,$$

and

$$\begin{pmatrix} v_{t+1} \\ y_{t+1} \\ z_{t+1} \end{pmatrix} = \begin{pmatrix} 1/4 & 0 & 0 \\ 1/4 & 1/2 & 0 \\ 0 & 0 & A_- \end{pmatrix} \begin{pmatrix} v_t \\ y_t \\ z_t \end{pmatrix} \text{ when } y_t < 0.$$

This system is of form (2), it consists of two linear systems, each of which is enabled in one of two half-spaces, as determined by the sign of y_t. Given that y_0 can be any real number, it can be verified that the sequence sign(y_t) is completely arbitrary, which then implies that the matrices A_- and A_+ can be multiplied in an arbitrary order. Since the boundedness of all products of A_- and A_+ is undecidable, it follows that the problem of determining whether the system is uniformly stable is undecidable. □

Acknowledgements

We are grateful to the reviewers for several constructive comments and to Omid Madani for sending us his work [14], thus making us aware of Paz's result.

References

[1] M. Berger, Y. Wang, Bounded semigroups of matrices, Linear Algebra Appl. 166 (1992) 21–27.
[2] V.D. Blondel, J.N. Tsitsiklis, When is a pair of matrices mortal? Informat. Process. Lett. 63 (1997) 283–286.
[3] V.D. Blondel, J.N. Tsitsiklis, Complexity of stability and controllability of elementary hybrid systems, Automatica 35 (1999) 479–489.
[4] V.D. Blondel, J.N. Tsitsiklis, A survey of computational complexity results in systems and control, Automatica 36(9) (2000) 1249–1274
[5] R. Braatz, P. Young, J. Doyle, M. Morari, Computational complexity of μ calculation, IEEE Trans. Automat. Control 39 (1994) 1000–1002.
[6] A. Condon, R.J. Lipton, On the complexity of space bounded interactive proofs, Proceedings of the 30th Annual Symposium on Foundations of Computer Science, Research Triangle Park, NC, 1989, pp. 462–467.
[7] A. Condon, R.J. Lipton, On the complexity of space bounded interactive proofs, unpublished Technical Report, 1989.
[8] I. Daubechies, J.C. Lagarias, Sets of matrices all infinite products of which converge, Linear Algebra Appl. 162 (1992) 227–263.
[9] L. Gurvits, Stability of discrete linear inclusions, Linear Algebra Appl. 231 (1995) 47–85.
[10] L. Gurvits, Stability of linear inclusions – Part 2, Technical Report TR 96-173, NEC Research, 1996.
[11] G. Jacob, An algorithm for calculating the cardinal, finite or infinite, of the semigroups of matrices, Theoret. Comput. Sci. 5 (1977) 183–204.
[12] V.S. Kozyakin, Algebraic unsolvability of problem of absolute stability of desynchronized systems, Automat. Remote Control 51 (1990) 754–759.
[13] J.C. Lagarias, Y. Wang, The finiteness conjecture for the generalized spectral radius of a set of matrices, Linear Algebra Appl. 214 (1995) 17–42.
[14] O. Madani, On the computability of infinite-horizon partially observable Markov decision processes, AAAI98 Fall Symposium on Planning with POMDPs, Orlando, FL, 1998.
[15] O. Madani, S. Hanks, A. Condon, On the undecidability of probabilistic planning and infinite-horizon partially observable Markov decision problems, Proceedings of the Sixteenth National Conference on Artificial Intelligence, Orlando, FL, July 1999.
[16] A. Nemirovskii, Several NP-Hard problems arising in robust stability analysis, Math. Control Signals Systems 63 (1993) 99–105.
[17] A. Paz, Introduction to Probabilistic Automata, Academic Press, New York, 1971.
[18] S. Poljak, J. Rohn, Checking robust nonsingularity is NP-hard, Math. Control Signals Systems 6 (1993) 1–9.
[19] G.-C. Rota, G. Strang, A note on the joint spectral radius, Nederl. Akad. Wetensch. Indagationes Math. 22 (1960) 379–381.
[20] O. Toker, On the complexity of the robust stability problem for linear parameter varying systems, Automatica 33 (1997) 2015–2017.
[21] O. Toker, H. Ozbay, On the complexity of purely complex μ computation and related problems in multidimensional systems, IEEE Trans. Automat. Control 43 (1998) 409–414.

[22] J.N. Tsitsiklis, On the stability of asynchronous iterative processes, Math. Systems Theory 20 (1987) 137–153.

[23] J.N. Tsitsiklis, The Stability of the Products of a Finite Set of Matrices, in: T.M. Cover, B. Gopinath (Eds.), Open Problems in Communication and Computation, Springer, New York, 1987, pp. 161–163.

[24] J.N. Tsitsiklis, V.D. Blondel, The Lyapunov exponent and joint spectral radius of pairs of matrices are hard – when not impossible – to compute and to approximate, Math. Control Signals Systems 10 (1997) 31–40.