COMPUTING THE SPECTRAL GAP OF A FAMILY OF MATRICES

NICOLA GUGLIELMI AND VLADIMIR YU. PROTASOV

Abstract. For a single matrix (operator) it is well-known that the spectral gap is an important quantity, as well as its estimate and computation. Here we consider, for the first time in the literature, the computation of its extension to a finite family of matrices, in other words the difference between the joint spectral radius (in short JSR, which we call here the first Lyapunov exponent) and the second Lyapunov exponent (denoted as SLE). The knowledge of joint spectral characteristics and of the spectral gap of a family of matrices is important in several applications, as in the analysis of the regularity of wavelets, multiplicative matrix semigroups and the convergence speed in consensus algorithms. As far as we know the methods we propose are the first able to compute this quantity to any given accuracy.

For computation of the spectral gap one needs first to compute the JSR. A popular tool that is the Invariant Polytope algorithm, which relies on the finiteness property of the family of matrices, when this holds true.

In this paper we show that the SLE may not possess the finiteness property, although it can be efficiently approximated with an arbitrary precision. The corresponding algorithm and two effective estimates are presented. Moreover, we prove that the SLE possesses a weak finiteness property, whenever the leading eigenvalue of the dominant product is real. This allows us to find in certain situations the precise value of the SLE. Numerical results are demonstrated along with applications in the theory of multiplicative matrix semigroups and in the wavelets theory.

1. Introduction

In various contexts of the literature discrete-time linear switching systems of the form

$$\begin{cases} \boldsymbol{x}(n+1) & = & A(n)\boldsymbol{x}(n)\,, \qquad A(n) \in \mathcal{A}, \quad n \geq 0, \\ \boldsymbol{x}(0) & = & \boldsymbol{x}_0 \end{cases} \tag{1}$$

are studied extensively (see e.g. [1, 10, 18, 25] and references therein). Here $A(n)$ is an arbitrary discrete matrix-valued function called the *switching law* that takes values from a given compact family of real $d \times d$ matrices $\mathcal{A}$. We denote by $\mathcal{F}$ the set of all switching laws and by $x(\cdot)$ any trajectory, i.e., a solution of the system (1). We are interested here in trajectories having fastest asymptotic growth as $n \to \infty$.

It is well known that, under general assumptions, the fastest growth is equivalent to $\|x(n)\| \asymp \rho^n$, where $\rho = \rho(\mathcal{A})$ is the joint spectral radius (i.e. the first Lyapunov exponent) of the family of matrices $\mathcal{A}$ (whose definitions are recalled below). We focus on the the case of a finite set of real $d{\times}d$-matrices $\mathcal{A} = \{A_1, \ldots, A_m\}$. The set of products of length (degree) k of matrices from $\mathcal{A}$, with repetitions permitted, is referred to as

$$\mathcal{A}^k := \{A_{s_1} \ldots A_{s_k} \ \mid \ s_1, \ldots, s_k \in \{1, 2, \ldots, m\}\} \tag{2}$$

and $\mathcal{A}^{\mathbb{N}} := \cup_{k \in \mathbb{N}} \mathcal{A}^k$. For an arbitrary product $\Pi \in \mathcal{A}^{\mathbb{N}}$, let $|\Pi|$ denote its length and $\rho(\Pi)$ denote its spectral radius.

Date: March 26, 2023.

2020 *Mathematics Subject Classification.* Primary 15A60, 15A42; Secondary 93D20, 47D03.

Key words and phrases. Discrete time linear switching system, spectral gap, Lyapunov exponent, joint spectral radius, dominant product, trajectories, asymptotic growth, finiteness property, algorithm, branch-and-bound.

Nicola Guglielmi acknowledges that this research was supported by funds from the Italian MUR (Ministero dell'Università e della Ricerca) within the PRIN 2017 Project "Discontinuous dynamical systems: theory, numerics and applications". Nicola Guglielmi also acknowledges support by the INdAM Research group GNCS (Gruppo Nazionale di Calcolo Scientifico).

Vladimir Protasov is supported by the RFBR grant 20-01-00469.

This article is dedicated to Sara.

Based on the number

$$\bar{\rho}_k(\mathcal{A}) := \sup_{\Pi \in \mathcal{A}^k} \rho(\Pi)^{1/k}, \tag{3}$$

Daubechies and Lagarias [10] gave the following definition of generalized spectral radius,

$$\bar{\rho}(\mathcal{A}) := \limsup_{k \to \infty} \bar{\rho}_k(\mathcal{A}), \tag{4}$$

which turns out to be one of the possible definitions of joint spectral radius $\rho(\mathcal{A})$ [4].

Several methods for the numerical computation of $\rho(\mathcal{A})$ have been proposed in the literature (see e.g. [11, 12, 19, 29]). The *Invariant Polytope algorithm* - which we shall shortly review in Section 3 - presented in [12] and developed in [13, 27] computes the precise value of $\rho(\mathcal{A})$ by showing that it is equal to $\rho^{1/|\Pi|}(\Pi)$, where Π is some product of matrices (called spectrum maximizing product, in short s.m.p.) from the family $\mathcal{A}$.

Since $\rho(\mathcal{A})$ is positively homogeneous, after a natural normalization the JSR (first Lyapunov exponent) can be assumed to be equal to one, i.e. $\rho(\mathcal{A}) = 1$. Under some general assumptions that can be verified algorithmically, there exists a dominant product of matrices which has spectral radius one, while all other products (except for its cyclic permutations and their powers) have spectral radii not exceeding some number $q < 1$. The smallest such number q is referred to as the second Lyapunov exponent (SLE); it estimates the rate of decrease for those trajectories generated by periodic switching laws. Unless the period is associated to the dominant product, the trajectory decreases exponentially with the factor bounded by the SLE. We denote the SLE by $\rho^{(2)}$ and, to underline the difference with the JSR ρ, we sometimes denote the latter quantity by $\rho^{(1)}$.

One of the key differences between $\rho^{(1)}$ and $\rho^{(2)}$ is that $\rho^{(1)}$ is normalized by length, whereas $\rho^{(2)}$ is not.

It turns out that the Invariant Polytope algorithm halts if and only if the family $\mathcal{A}$ has a spectral gap: one product has the spectral radius 1, and all other products which are neither powers of Π nor powers of its cyclic permutations, have spectral radii at most $\rho^{(2)}$. According to numerical experiments, a vast majority of matrix families have a dominant product, i.e., have a positive *spectral gap* $\rho^{(1)} - \rho^{(2)}$.

The aim of this work is to analyse and estimate the spectral gap of an arbitrary matrix family and to derive methods of its computation. Before doing this, in the next section we discuss applications of the spectral gap and of the second Lyapunov exponent and give an overview of our main results.

1.1. **Motivation and applications of the research.** A main motivation is the analysis of the robustness of the joint spectral radius and of the associated spectrum maximizing product. A natural measure for this is the spectral gap, that is the difference between the joint spectral radius (first Lyapunov exponent) and the second Lyapunov exponent. In analogy to the case of a single matrix, the knowledge of the spectral gap is essential to assess the susceptivity of the joint spectral radius with respect to perturbations in the matrices forming the family $\mathcal{A}$. Due to such perturbations the spectral radii of some sensitive products of $\mathcal{A}^{\mathbb{N}}$ may change abruptly and determine significant variations in $\rho(\mathcal{A})$, even when the s.m.p. is not sensitive. A few main applications follow.

Consensus in asynchronous multiagent systems. Consensus of multi-agent systems is an important topic which receives increasing attention. Traditionally, the analysis of consensus ability for a given system with switching topology can be attacked by different techniques, which are often based on some restrictive conditions. Indeed the general consensus problem of multi-agent switching systems can be transformed into the numerical calculation of the Lyapunov exponent for a given set of matrices and resolved by using the existing numerical algorithms (see e.g. [5]). Such algorithms have been developed over recent years, and are used in the analysis of stability and stabilization of linear multiagent systems with discrete time. These methods are based on the notion of joint spectral radius of a set of matrices to analyze the rate of convergence of matrix products with factors drawn from certain sets of matrices with special properties. The knowledge of the joint spectral radius and of the associated spectrum maximizing product plays a central role and quantifies the speed of convergence of the consensus process. For this the computation of the spectral gap is important to assess its sensitivity with respect to changes/perturbations in the weights of the consensus network and also the speed of the process. For a survey of results on models of consensus in asynchronous multiagent systems we refer the reader e.g. to [23].

Local regularity of fractal curves and wavelets. Fractal curves are solutions of linear functional equations with a contraction of the argument. Among the most known examples are the classical fractal curves such as the de Rham and Koch curves, the limit functions of subdivision schemes, and the wavelet functions. Due to self-similarity, most of fractal curves possess a quite complex structure. In particular, they have a varying local regularity. The local modulus of continuity of a function f at a point x is $\omega_{f,x}(t) := \sup_{|h|\le t} |f(x+h) - f(x)|$. The local Hölder exponent at x is

$$\alpha_f(x) \ := \ \sup\Bigl\{ \alpha \ \Bigm| \ \omega_{f,x}(t) \le C\, t^{\alpha}, \ t \in (0,1) \Bigr\}.$$

Global regularity of a function is measured in a standard way by the Hölder exponent on the whole domain:

$$\alpha_f \ := \ \sup\Bigl\{ \alpha \ \Bigm| \ \omega_{f,x}(t) \ \le \ C\, t^{\alpha}\,, \ x \in \mathbb{R}, \, t \in (0,1) \ \Bigr\}.$$

If a fractal curve does not belong to $C^1(\mathbb{R})$, then the values of its local regularity at various points can fill a segment and each value from that segment is attained on an everywhere dense set of arguments x. The local regularity of fractal curves, in particular, of Daubechies wavelets and of limit functions of subdivision schemes, have been studied in detail [9, 10, 35] (see also [31] for general overview). The regularity of a fractal curve can be expressed in terms of two special matrices B_0, B_1. We have $\alpha_f = -\log_2 \rho(B_0, B_1)$, where $\rho(B_0, B_1)$ is the joint spectral radius [10, 35]. If x is a rational number, which has a period $(s_1, \dots, s_\ell)$ in its binary expansion, then $\alpha_f(x) = -\frac{1}{\ell}\,\log_2\, \rho(\Pi)$, where $\Pi = B_{s_1} \dots B_{s_\ell}$. Moreover, for the modulus of continuity at the point x, we have

$$\omega_{f,x}(t) \ \asymp \ t^{\,\alpha_f(x)}, \quad t \in (0,1),$$

(see [10, 35]). Thus, at every rational (but not dyadic rational) x, when $\ell \ge 2$, the local regularity is expressed by the spectral radius of the corresponding matrix product. If x is a dyadic number, then it has two binary expansions, one with period (0), i.e., the binary expansion ends with an infinite sequence of zeros, and one of period (1), corresponding to the sequence of ones. In this case the local regularity is equal to the maximum of numbers $-\log_2\, \rho(B_0)$ and $-\log_2\, \rho(B_1)$.

The local regularity can be estimated by the second Lyapunov exponent as follows:

If $\Pi = B_{s_1} \dots B_{s_n}$ is the dominant product, then for every rational x with the period $(s_1 \dots s_n)$, we have

$$\omega_{f,x}(t) \ \asymp \ t^{\,\alpha}, \quad t \in (0,1),,$$

where $\alpha = \alpha_f = -\log_2 \rho(B_0, B_1)$ is the Hölder exponent of f. For all other rational x, we have

$$\omega_{f,x}(t) \ \le \ C\, t^{\,\alpha}\, q^{-\frac{\log t}{\ell}}, \quad t \in (0,1).$$

where $q = \rho^{(2)}(B_0, B_2)$ and ℓ is the length of the period of x.

Thus, we spot the set of rational numbers x where the minimal regularity $\alpha_f = \alpha_f(x)$ is attained. The modulus of continuity at those points is asymptotically equivalent to $t^{\,\alpha}$, while for all other rational numbers. it does not exceed $C\, t^{\,\alpha} q^{-\frac{\log t}{\ell}}$

Periodic switching laws of the linear switching systems. One of the most important problems in the study of linear switching systems (1) is the maximal growth of their trajectories and, in particular, their stability properties [6, 18, 19, 22, 40]. Due to homogeneity this problem is reduced, with a suitable scaling, to the case when $\rho(\mathcal{A}) = 1$. In this case the system, under the usual irreducibility assumption, i.e. the matrices $A_1, A_2, \dots, A_m$ do not have a common nontrivial invariant subspace (see e.g. [12]), is weakly stable, i.e., all trajectories are bounded, and at least one trajectory is separated from zero [1]. The latter case is realized for periodic switching law with period k, when $\|\boldsymbol{x}(n)\| \asymp 1$, $n \in \mathbb{N}$. The second Lyapunov exponent allows us to estimate the growth of all other trajectories with periodic switching laws and to separate them from the fastest one. Namely: if the switching law is eventually periodic with a period Q of some length ℓ, then its trajectory $\boldsymbol{x}(n)$ converges to zero with an exponential rate as $n \to \infty$, unless Q is equal to Π (or to its cyclic permutation). Moreover, $\|\boldsymbol{x}(n)\| \le C\, \lambda^{\,n}$, where $\lambda = [\rho^{(2)}(\mathcal{A})]^{1/\ell}$.

Thus, any upper bound to the second Lyapunov exponent estimates the rate of asymptotic decay of trajectories generated by eventually periodic switching laws, whenever the period is not dominant. This result is also closely related to the concepts of *periodic stability* [8, 18].

Spectral radii of elements in matrix multiplicative semigroups.
Matrix multiplicative semigroups where all elements have the same spectral radius have been studied in various contexts in [6, 30, 32, 33, 38]. In particular, this problem is closely related to finding all finite semigroups of integer matrices [7, 20, 26]. Clearly, if all matrices of a semigroup have the same spectral radius, then it must be equal to one. The problem of full classification of such semigroups is solved only in low dimensions, for matrices of sizes $d = 2, 3$, for higher dimensions several conjectures have been proposed, although the problem of identification of such semigroups is polynomially solvable [38]. For an arbitrary matrix semigroup, the spectral gap is the distance between the maximal spectral radius of its elements (which is equal to one) and the supremum of all other spectral radii. If the semigroup is finitely generated, then this gap is strictly positive and is equal to $\rho^{(1)} - \rho^{(2)}$, provided the set of generators $\mathcal{A}$ has a dominant product. Thus, estimating the second Lyapunov exponent we determine how far is the semigroup from those having constant spectral radius.

1.2. **Overview of the main results.** We show that the second Lyapunov exponent $\rho^{(2)}$ may not possess the "finiteness property", i.e., may not be achieved on a finite product of matrices from $\mathcal{A}$. This is rather surprising since the first Lyapunov exponent $\rho^{(1)}$ does possess the finiteness property, whenever $\rho^{(2)}$ is well defined. This phenomenon suggests that the second Lyapunov exponent can hardly be found explicitly for general sets of matrices. On the other hand, we prove that under the extra assumption that the leading eigenvalue of the dominant product is real, we have a "weak finiteness property": $\rho^{(2)}$ is achieved at a finite product with one extra factor being the limit of powers of the dominant product Π. This allows us to construct an algorithm of precise computation of the second Lyapunov exponent, provided the real eigenvalue condition is satisfied. For general families, we present the branch-and-bound algorithm that computes $\rho^{(2)}$ with an arbitrary precision. Also two efficient estimates for the second Lyapunov exponent and for the spectral gap are obtained. They are both explicitly computed using the invariant polytope of the family $\mathcal{A}$.

1.3. **Structure of the paper.** In Section 2 we introduce the joint spectral radius and present its main properties; in Section 3 we describe the invariant polytope algorithm. In Section 4 we first (Section 4.1) show that the second Lyapunov exponent may not satisfy the finiteness property. Section 5 is devoted to lower and upper bounds for the second Lyapunov exponent and for the spectral gap. We derive an upper bound (which is simple) and two lower bounds (this is more complicated) that are expressed by closed formulas using the matrices of the family $\mathcal{A}$ and its invariant polytope. In Section 6 we develop the branch-and-bound algorithm and show that it is able to compute $\rho^{(2)}$ with an arbitrary precision. Section 7 analyses the real case, when the dominant product has a real leading eigenvalue. We establish the weak finiteness property and derive an algorithm for precise computation of $\rho^{(2)}$. Section 8 presents numerical examples. We analyse the efficiency of our bounds and algorithms for various families and show applications to estimate the growth of trajectories of a linear switching system and of local regularity of fractal curves and of the Daubechies wavelets.

2. Joint spectral radius

Let $\mathcal{A} = \{A_1, \ldots, A_m\}$ is a family of real $d \times d$ matrices. The set of products of length (degree) k of matrices from $\mathcal{A}$, without ordering and with repetitions permitted, is referred to as $\mathcal{A}^k$ and $\mathcal{A}^{\mathbb{N}} := \cup_{k\in\mathbb{N}}\mathcal{A}^k$.

Based on the number

$$\hat{\rho}_k(\mathcal{A}) := \max_{\Pi\in\mathcal{A}^k} \|\Pi\|^{1/k}, \tag{5}$$

Rota and Strang [39] were the first to give the definition of joint spectral radius of $\mathcal{A}$.

Definition 2.1. The number

$$\rho(\mathcal{A}) := \lim_{k\to\infty} \hat{\rho}_k(\mathcal{A}) \tag{6}$$

is said to be the joint spectral radius (JSR) of the family $\mathcal{A}$.

Note that the limit in (6) always exists and does not depend on the considered norm.

Later, Daubechies and Lagarias [10] gave Definition 4 of generalized spectral radius and shown that for all $k \geq 1$ (see (3) and (4) for the definitions of $\bar{\rho}_k(\mathcal{A})$ and of $\bar{\rho}(\mathcal{A})$)

$$\bar{\rho}_k(\mathcal{A}) \leq \bar{\rho}(\mathcal{A}) \leq \hat{\rho}(\mathcal{A}) \leq \hat{\rho}_k(\mathcal{A}). \tag{7}$$

Finally Berger and Wang [4] proved the fundamental equality

$$\hat{\rho}(\mathcal{A}) = \bar{\rho}(\mathcal{A})$$

Thus the *spectral radius* of $\mathcal{A}$ is simply denoted as $\rho(\mathcal{A})$.

An important characterization of the spectral radius $\rho(\mathcal{A})$ of a matrix family is the generalization of Gelfand's formula. In order to state this characterization, we define the *norm of the family* $\mathcal{A}$ as

$$\|\mathcal{A}\| := \hat{\rho}_1(\mathcal{A}) = \max_{i \in \mathcal{I}} \|A_i\|.$$

The following proposition provides an interesting way to either compute or approximate the joint spectral radius by constructing a suitable norm.

Proposition 2.1 (Rota and Strang [39])**.** The spectral radius of a bounded family $\mathcal{A}$ of $d \times d$-matrices is characterized by

$$\rho(\mathcal{A}) = \inf_{\|\cdot\| \in \mathrm{Op}} \|\mathcal{A}\|, \tag{8}$$

where Op denotes the set of operator norms.

Concerning the existence of a minimizer in (8), we recall the following definition.

Definition 2.2 (Extremal norm)**.** We say that a norm $\|\cdot\|$ satisfying

$$\|\mathcal{A}\| = \rho(\mathcal{A})$$

is extremal for the family $\mathcal{A}$.

The existence of an extremal norm is very important since it allows for the possibility of computing the joint spectral radius exactly by determining such an extremal norm, when the construction can be obtained by a finite procedure.

Hence a norm $\|\cdot\|$ is extremal for a family if for every $\boldsymbol{x} \in \mathbb{R}^d$, we have $\|A\boldsymbol{x}\| \leq \rho(\mathcal{A})\|\boldsymbol{x}\|,\ A \in \mathcal{A}$.

3. The invariant polytope algorithm to compute the JSR.

We report here the invariant polytope algorithm (see [12, 14, 34]), a key method to compute the joint spectral radius of a finite set of matrices by constructing an extremal norm for $\mathcal{A}$, which has a polytope unit ball. The method is sketched by Algorithm 1.

In order to illustrate Algorithm 1 we need to introduce the following definitions.

Definition 3.1 (Spectrum maximizing product)**.** If $\mathcal{A}$ is a bounded family of $d \times d$-matrices, any matrix $\Pi \in \mathcal{A}^k$ satisfying $\rho(\Pi)^{1/k} = \rho(\mathcal{A})$ for some $k \geq 1$ is called a spectrum maximizing product (s.m.p.) for $\mathcal{A}$.

We remark that, as is well-known, s.m.p.'s are not always guaranteed to exist (see, e.g., [2]).

Before stating the algorithm we need the following definition.

Definition 3.2. We say that a set $\mathcal{X} \subset \mathbb{R}^d$ is *absolutely convex* if, for all $x_1, x_2 \in \mathcal{X}$ and $\lambda_1, \lambda_2 \in \mathbb{R}$ such that $|\lambda_1| + |\lambda_2| \leq 1$, it holds that $\lambda_1 x_1 + \lambda_2 x_2 \in \mathcal{X}$. Let $\mathcal{S} \subset \mathbb{C}^d$. Then the intersection of all absolutely convex sets containing $\mathcal{S}$ is called the *absolutely convex hull* of $\mathcal{S}$ and is denoted by $\mathrm{absco}(\mathcal{S})$.

The idea implemented by Algorithm 1 is that of iteratively building a polytope which in the end is invariant for the family $\widehat{\mathcal{A}} := \mathcal{A}/\rho(\mathcal{A})$, and hence determines an extremal polytope norm for $\mathcal{A}$. The success of the algorithm would prove that the product Π guessed as a possible spectrum maximizing product in the preliminary phase is indeed an s.m.p. .

By leading eigenvalue λ of a matrix M we mean an eigenvalue of maximum modulus $\rho(M)$.

Algorithm 1: Polytope algorithm

Data: $\mathcal{A}$
Result: $\mathcal{P}$
begin
  Preliminary: find a product Π of length $k \geq 1$ such that $\rho(\Pi)^{1/k}$ is maximal among $\mathcal{A}^k$;
  Set Π as a candidate s.m.p.;
  Set $\varrho := \rho(\Pi)^{1/k}$;
  Set $\widehat{\mathcal{A}} := \varrho^{-1}\mathcal{A}$ (scaling of $\mathcal{A}$);
  Compute $\boldsymbol{v}_1, \ldots, \boldsymbol{v}_k$ leading eigenvectors of Π and of its cyclic permutations;
  Set $V_0 := \{\boldsymbol{v}_1, \ldots, \boldsymbol{v}_k\}$ initial set of vertices of the polytope $\mathcal{P}$;
  Set $W_0 = V_0$;
  Set $\ell = 0$;
  while $\mathrm{span}(V_\ell) \neq \mathbb{R}^d$ **do**
    $V_{\ell+1} := V_\ell \cup \widehat{\mathcal{A}}\, V_\ell$;
    Set $\ell = \ell + 1$;
  Set $\mathcal{P}^{(\ell)} = \mathrm{absco}(V_\ell)$;
  while $\widehat{\mathcal{A}}\, V_\ell \not\subseteq \mathcal{P}^{(\ell)}$ **do**
    Set $\ell = \ell + 1$;
    Set $W_\ell = \widehat{\mathcal{A}}\, V_{\ell-1}$;
    if $W_\ell \in \mathcal{P}^{(\ell-1)}$ **then**
      return $\mathcal{P} := \mathcal{P}^{(\ell-1)}$ (unit ball of the polytope extremal norm for $\mathcal{A}$);
      The algorithm halts;
    else
      Set $\mathcal{P}^{(\ell)} = \mathrm{absco}\,(V_{\ell-1} \cup W_\ell)$;
      Let V_ℓ the set of vertices of $\mathcal{P}^{(\ell)}$;

Algorithm 1 starts with the set W_0 of $k = |\Pi|$ leading eigenvectors of the cyclic permutations of the product Π and with the corresponding polytope $\mathcal{P}^{(0)} = \mathrm{absco}(W_0)$. Then the algorithm builds the multi-tree as follows. We set $W_1 = W_0$. For each point $\boldsymbol{v} \in W_0$, we consider all points $A_j\boldsymbol{v}$, $j = 1, \ldots, m$, remove those are inside $\mathcal{P}^{(0)}$ (the *dead vertices*) and add the others (*alive vertices*) to the set W_1. After all elements $\boldsymbol{v}$ of W_0 are exhausted we obtain the set W_1 and the corresponding polytope $\mathcal{P}^{(1)} = \mathrm{absco}(W_1)$. Then we do the next iteration, etc. The algorithm terminates after Nth iteration if $W_N = W_{N-1}$. In this case $\mathcal{P}^{(N)}$ is an invariant polytope and it is a unit ball of an extremal norm.

We address the reader to [15] as a reference with an example of application of Algorithm 1.

An implementation of Algorithm 1 would include stopping criteria to establish whether the set of new vectors W_ℓ computed at iteration ℓ lies inside or outside the current polytope $\mathcal{P}^{(\ell-1)}$.

The part of Algorithm 1 which is more computationally expensive is in most cases the preprocessing phase, with a cost which grows exponentially in k. The use of suitable criteria to exclude certain products (see, e.g., [11, 17]) is able to reduce this cost.

Definition 3.3. A product from $\mathcal{A}^{\mathbb{N}} = \cup_{k\geq 1}\mathcal{A}^k$ is called *primitive* if it is neither a power of a shorter product nor a cyclic permutation.

As an example $\Pi_1 = (A_1A_2^2)^2 = A_1A_2^2A_1A_2^2$ is not primitive because it is the square of the shorter product $\Pi = A_1A_2^2$. Similarly $\Pi_2 = A_2A_1A_2^2A_1A_2$ is not primitive because it is a cyclic permutation of Π_1 which is not primitive.

Clearly, if a product $P \in \mathcal{A}^{\mathbb{N}}$ is an s.m.p., then every its cyclic permutation is also an s.m.p. Moreover, if P is a power of a shorter product Π, then Π is an s.m.p. as well. Therefore, *every s.m.p. is either primitive or is a power of a primitive s.m.p.* Hence, in what follows we consider only primitive s.m.p.

The following definition (which is related to the joint spectral radius $\rho^{(1)}(\mathcal{A})$ plays a crucial role.

Definition 3.4 (Dominant s.m.p.)**.** A primitive s.m.p. Π of a bounded family of $d \times d$-matrices $\mathcal{A}$ is said to be dominant if there exists a constant $q < 1$ such that, for any matrix $Q \in \mathcal{A}^{\mathbb{N}}$ - other than Π, its powers and the powers of its cyclic permutations - it holds that $\rho(Q)^{1/k} \leq q \cdot \rho(\mathcal{A})$, where k is the degree of Q (i.e., $Q \in \mathcal{A}^k$).

A sufficient condition for the successful termination of Algorithm 1 within finitely many iterations is that the starting product Π be a dominant s.m.p. of the family $\mathcal{A}$ with a unique and simple leading eigenvector (see [12]).

Theorem 3.1. The Invariant polytope algorithm with the candidate s.m.p. Π halts within finite time if and only if the product Π is dominant and it has a unique and simple leading eigenvalue (i.e. of maximum modulus).

In the case of complex leading eigenvalue of Π as usual we replace in Theorem 3.1 the "unique eigenvalue" by the "unique pair of conjugate eigenvalues". All our results are proved for the general case, when the leading eigenvalue can be real or complex (see [16] for the case of complex leading conjugate eigenvalues). For the sake of simplicity we consider only the case when it is real and positive and clarify, if necessary, how the proof is modified for negative or non-real leading eigenvalues.

3.1. **Illustrative example.**

Example 3.1. We consider the following example: $\mathcal{A} = \{A_1, A_2\}$ with

$$A_1 = \alpha \begin{pmatrix} 1 & -1 \\ 1 & 0 \end{pmatrix}, \qquad A_2 = \alpha \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}, \tag{9}$$

with $\alpha = \left(\frac{1}{2}\left(3-\sqrt{5}\right)\right)^{1/5} \approx 0.824907$, which is such that $\rho(\mathcal{A}) = 1$.

The spectrum maximizing product is given by

$$\Pi = A_1 A_2 A_1^2 A_2 = \frac{1}{2}\left(3-\sqrt{5}\right)\begin{pmatrix} 2 & 1 \\ 1 & 1 \end{pmatrix} \tag{10}$$

with leading eigenvalue $\lambda = 1$ and associated eigenvector

$$v_1 = \begin{pmatrix} 1 \\ \frac{2}{1+\sqrt{5}} \end{pmatrix} = \begin{pmatrix} 1 \\ 0.618034\ldots \end{pmatrix} \tag{11}$$

In the Figure 1 we show the extremal polytope norm obtained by applying Algorithm 1.

One vertex is the vector $\boldsymbol{v}_1$, while the other vertices are given by the following vectors

$$\boldsymbol{v}_2 = A_2\boldsymbol{v}_1, \boldsymbol{v}_3 = A_1\boldsymbol{v}_2, \boldsymbol{v}_4 = A_1\boldsymbol{v}_3, \boldsymbol{v}_5 = A_2\boldsymbol{v}_4, v_6 = A_2\boldsymbol{v}_3$$

as well as all their opposites, $\{-\boldsymbol{v}_1, -\boldsymbol{v}_2, -\boldsymbol{v}_3, -\boldsymbol{v}_4, -\boldsymbol{v}_5, -\boldsymbol{v}_6\}$.

3.2. **The underlying multi-tree associated to the polytope algorithm.** Algorithm 1 constructs a multi-tree, which we denote by $\mathcal{T}$. Denote $\boldsymbol{v}_i = A_{d_{i-1}} \ldots A_{d_1}\boldsymbol{v}_1$, $i = 2, \ldots, k$ (the set of leading eigenvectors of the spectrum maximizing product Π and of its cyclic permutations). The ordered set $\{\boldsymbol{v}_1, \ldots, \boldsymbol{v}_k\}$ is called multiroot (in short *root*) and on it the algorithm builds a multi-tree $\mathcal{T}$.

Each vertex $\boldsymbol{v}_s$ of this multi-tree is associated to a vector in $\mathbb{R}^d$. Every vertex which is not in the root of the multi-tree has m children on the next level $\{A\boldsymbol{v}_s\,,\ A \in \mathcal{A}\}$.

If instead $\boldsymbol{v}_s$ belongs to the root of $\mathcal{T}$, it has one less child on the next level, i.e. $m-1$ children given by $A\boldsymbol{v}_s$ for $A \in \mathcal{A}$, $A \neq A_{d_s}$.

The root has level zero. There are k vertices on zero level and exactly $k(m-1)m^{p-1}$ vertices on the pth level, $p \geq 1$.

Each vertex is connected with $\boldsymbol{v}_1$ by a unique path (all paths are without cycles and without backtracking). We identify the vertex by the corresponding path.

A finite set of vertices $\mathcal{C}$ from positive levels of $\mathcal{T}$ is called a *minimal cut set* if every infinite path from the root intersects $\mathcal{C}$ by exactly one vertex.

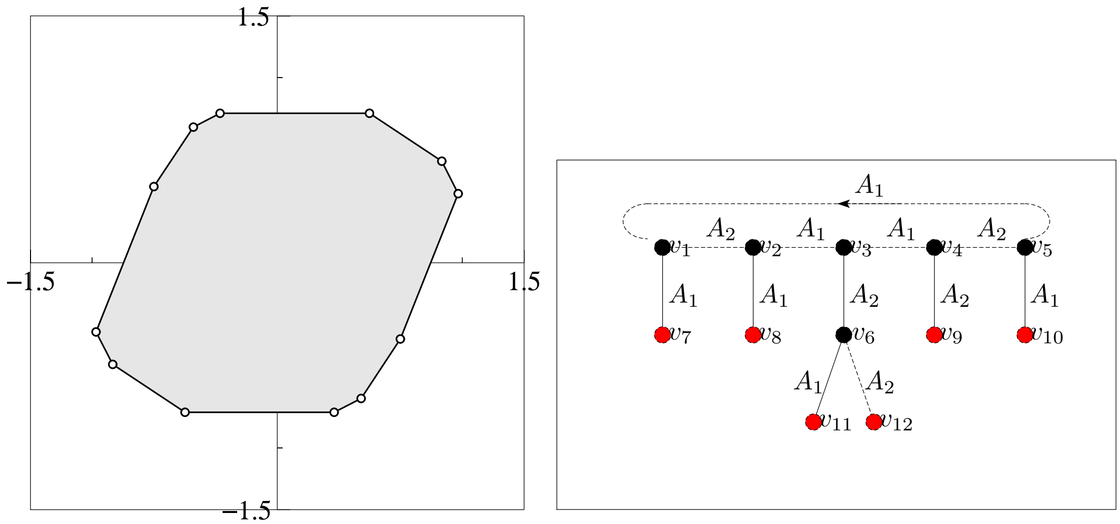

FIGURE 1. Left picture: The extremal polytope norm for Example 3.1 (see (9)). Right picture: multi-tree $\mathcal{T}$ associated to the polytope algorithm. At the top (zero) level there are two vertices, corresponding to leading eigenvectors. Dead vertices are in red.

Let us recall that, by our assumption, the multi-tree $\mathcal{T}$ has a dominant product in its root and the root consists of leading eigenvectors of cyclic permutations of this product.

Lemma 3.1. If $\|\cdot\|$ is an extremal norm, then all vertices of the root have the same norm.

Proof. The cyclic trajectory $\boldsymbol{v}_1 \to \cdots \to \boldsymbol{v}_k \to \boldsymbol{v}_1$. has non-increasing norms of elements. If $\|\boldsymbol{v}_j\| < \|\boldsymbol{v}_1\|$ for some $j \neq k$, then $\|\boldsymbol{v}_j\| \geq \|\boldsymbol{v}_{j+1}\| \geq \cdots \geq \|\boldsymbol{v}_k\| \geq \|\boldsymbol{v}_1\|$, hence, $\|\boldsymbol{v}_j\| \geq \|\boldsymbol{v}_1\|$, which is a contradiction. Therefore, all norms of $\boldsymbol{v}_j, j = 1, \ldots, k$ are the same. $\square$

4. The second Lyapunov exponent

We also call the JSR *the first Lyapunov exponent* and use the notation $\rho = \rho^{(1)}$ to distinct it from the second Lyapunov exponent $\rho^{(2)}$, which is defined below, recalling Definition 3.4.

Definition 4.1. Let $\rho(\mathcal{A}) = 1$ ($\mathcal{A}$ is normalized) and let $\Pi \in \mathcal{A}^{\mathbb{N}}$ be a primitive dominant product, i.e. $\rho(\Pi) = 1$. The supremum q of spectral radii of all products from $\mathcal{A}^{\mathbb{N}}$ which are neither powers of Π nor of their cyclic permutations, is smaller than 1. This supremum is called *the second Lyapunov exponent* $\rho^{(2)}(\mathcal{A})$. The difference $h(\mathcal{A}) = \rho^{(1)}(\mathcal{A}) - \rho^{(2)}(\mathcal{A})$ is called the spectral gap of $\mathcal{A}$.

Remark 4.1. In view of Definition 4.1, the second Lyapunov exponent is well defined if and only if the first Lyapunov exponent is attained at a dominant product. This condition can be verified within finite time by the Invariant polytope Algorithm 1 and is related to the well known finiteness property of $\mathcal{A}$, which we are stating hereafter.

In the sequel of the paper we are making the following assumption.

Assumption 4.1. The finite family $\mathcal{A}$ has the finiteness property, that is $\mathcal{A}$ possesses a *spectrum maximising product* (s.m.p.), i.e., a product Π of some length $k \geq 1$ such that $\rho^{1/k}(\Pi) = \rho(\mathcal{A})$.

The *finiteness conjecture* set up by Lagarias and Wang [24] and, in a different form, by Gurvits [18], where it was attributed to Pyatnitskii, asserts that for every finite matrix family $\mathcal{A}$ the joint spectral radius is attained on that product. This means that the fastest growth of trajectories of the linear switching system (1) is always achieved on a periodic switching law.

The conjecture was disproved in [3, 2], although giving at least one concrete counterexample turns out to be extremely difficult [21, 22]. If a family does possess an s.m.p., then it is said to satisfy *the finiteness property*. Clearly, the dominant product is an s.m.p. Hence, the Invariant polytope algorithm actually ensures the finiteness property, that is it halts within finite time, then the family possesses the finiteness property. Invoking now Remark 4.1, we conclude that *if the second Lyapunov exponent is well-defined, then the family possesses the finiteness property.*

Assuming $\mathcal{A}$ is normalized and Π is a dominant product of $\mathcal{A}$, we have $\rho^{(1)}(\mathcal{A}) = 1$ and $\rho^{(2)}(\mathcal{A}) < 1$ and $h(\mathcal{A}) > 0$. All cyclic permutations of a dominant product are also dominant. We usually assume that we choose one of them. Their powers are not dominant since they are not primitive.

Note that if we drop the assumption $\rho(\Pi) = 1$, then the second Lyapunov exponent $\rho^{(2)}(\mathcal{A})$ can still be defined as $q\,\rho(\Pi)$, where q is the supremum of numbers such that for every $S \in \mathcal{A}^{\mathbb{N}}$, we have $\rho(S)^{1/|S|} \le q^{1/|S|}\,\rho(\Pi)^{1/|\Pi|}$ unless S is a power of Π or of its cyclic permutation.

In the generic case when the family admits a dominant product with a unique leading eigenvector, Algorithm 1 finds the invariant polytope if and only if the spectral gap is strictly positive.

From now on we address the problem of estimating the spectral gap provided it is positive.

4.1. **The second Lyapunov exponent does not possess the finiteness property.** We discuss here an important feature of the SLE under the natural convergence assumption for the invariant polytope algorithm.

Assumption 4.2. In what follows we assume that Π is a dominant product and the family $\mathcal{A}$ is normalized so that $\rho(\Pi) = 1$. Thus, $\rho^{(1)} = 1$ and $\rho^{(2)} < 1$.

Nevertheless, the finiteness property does not necessarily hold true for the second Lyapunov exponent, which may not be attained at a finite product. This is demonstrated in the following example.

Example 4.1. There are matrix families, for which the dominant product exists but the second Lyapunov exponent is not attained at a finite product. Consider - for example - the family of two matrices:

$$A_1 \;=\; \begin{pmatrix} 1 & 0 \\ 0 & \frac{1}{2} \end{pmatrix} \;; \qquad A_2 \;=\; \begin{pmatrix} \frac{1}{2} & -\frac{1}{2} \\ \frac{1}{2} & -\frac{1}{2} \end{pmatrix}$$

Let us show that $\Pi = A_1$ is a dominant product, and the second Lyapunov exponent $\rho^{(2)} = \frac{1}{2}$ is not attained at any finite product. We have $\rho(A_1) = 1$. Every product S different from a power of A_1 has a factor A_2. Taking a suitable cyclic permutation of the product S it may be assumed that it begins with A_2. Moreover, since $A_2^2 = 0$, we have $S = A_2 A_1^{n_s} \ldots A_2 A_1^{n_1}$ for some $n_1, \ldots, n_s \in \mathbb{N}$. Since the image of the operator S is contained in the image of A_2, which is the line $\{(x,x)^T \,|\, x \in \mathbb{R}\}$, it follows that the leading eigenvector $\boldsymbol{v}$ of S lies on this line, i.e., up to multiplication by a constant it may be assumed that $\boldsymbol{v} = (1,1)^T$. We have

$$A_2 A_1^n \;=\; \begin{pmatrix} 2^{-1} & -2^{-n-1} \\ 2^{-1} & -2^{-n-1} \end{pmatrix} \;,$$

hence, $A_2 A_1^n \boldsymbol{v} \;=\; (2^{-1} - 2^{-n-1})\,\boldsymbol{v}$. This yields that $S\boldsymbol{v} \;=\; \prod_{j=1}^{s}(2^{-1} - 2^{-n_j-1})\,\boldsymbol{v}$.

Thus, $\rho(S) = \prod_{j=1}^{s}(2^{-1} - 2^{-n_j-1})$. We see that the supremum of $\rho(S)$ over all admissible matrix products is equal to 2^{-1} and that it is never attained at a finite product.

Remark 4.2. In Section 7 we shall prove that in the *real case*, when the leading eigenvalue of the dominant product is real, the *weak finiteness property* always holds. This means that $\rho^{(2)}$ is reached at a finite product which may contain, apart from matrices from $\mathcal{A}$, the matrix $\Pi^{\infty} \;=\; \lim_{n\to\infty} \Pi^n$. Note that this limit is well defined since Π has the unique maximal eigenvalue 1. This result allows us to construct an algorithm that finds the precise value of $\rho^{(2)}$ (see subsection 7.2). However, this procedure in general is limited to the real case.

The non-finiteness property of the second Lyapunov exponent makes the problem of its precise computation very hard, apart from the real case (Remark 4.2). Therefore, our first goal is that to obtain estimates for $\rho^{(2)}$. We shall see that these estimates are useful in the real case as well.

The most important issue is to find a suitable upper bound for $\rho^{(2)}$ since it separates $\rho^{(2)}$ from $\rho^{(1)}$ and estimates the spectral gap from below. We develop two upper bounds: the first one is pretty rough but simple to evaluate; the second one is much sharper, however it is more technical to obtain. Then in Section 6 we develop Algorithm 2 that allows us finding the second Lyapunov exponent with an arbitrary precision, based on a branch-and-bound idea.

In all those methods we use the polytope extremal norm obtained by the invariant polytope algorithm (Algorithm 1).

Under our assumption that $\rho(\mathcal{A}) = 1$, the extremal norm $\|\boldsymbol{x}(n)\|$ is non-increasing on every trajectory of the discrete-time linear switching system $\boldsymbol{x}(n+1) = A(n)\boldsymbol{x}(n),\ A(n) \in \mathcal{A},\ n \geq 0$.

Let $\Pi = A_{d_k} \dots A_{d_1}$ be a dominant product and $\boldsymbol{v}_1$ be its leading eigenvector. Let us recall that we are consider the real case when the leading eigenvalue λ is real; hence it is not restrictive to assume it is positive, so that if the family is normalized, $\lambda = 1$. Indeed, if $\lambda = -1$, then Algorithm 1 works in the same way, with the only difference that $\Pi\,\boldsymbol{v} = -\boldsymbol{v}$ instead of $\Pi\,\boldsymbol{v} = \boldsymbol{v}$. However, we consider the absolutely convex hull of points generated by the algorithm, and hence, we work not with points $\boldsymbol{v}_k$ but with pairs $\{\boldsymbol{v}_k, -\boldsymbol{v}_k\}$. Since Π maps the pair $\{\boldsymbol{v}, -\boldsymbol{v}\}$ to the same pair, Algorithm 1 is performed in the same way in this case.

Thus, we assume $\lambda = 1$ and hence, the leading eigenvector $\boldsymbol{v}$ is a real vector.

5. Estimates for the second Lyapunov exponent

In what follows we assume that $\|\boldsymbol{v}_1\| = 1$. Consequently, if the norm is extremal, then all elements of the root have norm one (by Lemma 3.1) and all other elements of the multi-tree $\mathcal{T}$ have norms at most one. The norms of vertices do not increase along every path. Each vertex associated to a vector of unit norm one is called *alive*, otherwise it is said *dead*. Thus, all children and further descendants of dead vertices are also dead.

5.1. **Lower bounds.** Every nonempty subset $\mathcal{S} \subset \mathcal{A}^{\mathbb{N}}$ of matrix products of the family $\mathcal{A}$ gives an obvious lower bound

$$\sigma(\mathcal{S}) \;=\; \sup_{Q \in \mathcal{S}\setminus\mathcal{R}} \rho(Q),$$

where $\mathcal{R}$ is the set of powers of cyclic permutations of the dominant product Π. Clearly, $\sigma(\mathcal{S}) \leq \rho^{(2)}$ and this inequality becomes an equality for $\mathcal{S} = \mathcal{A}^{\mathbb{N}}$. Proposition 4.1 shows that this equality may not be attained for finite sets $\mathcal{S}$. Usually one can take $\mathcal{S}$ as a set of products of length bounded by a given number. Another reasonable choice of $\mathcal{S}$ is a set of products associated to paths identified by some cut set $\mathcal{C}$ of the multi-tree $\mathcal{T}$ (that is a set of nodes which separates the multitree in two parts, an upper part with respect to $\mathcal{C}$ and a bottom part, which are not mutually connected by edges, meaning that eliminating the set $\mathcal{C}$ the multi-tree is cut. This implies that every infinite path from the root intersects $\mathcal{C}$).

They are the so-called *basement sets* defined in subsection 5.3. In Section 6 we introduce a branch-and-bound method, where the set $\mathcal{S}$ is constructed step-by-step during the algorithm.

5.2. **Upper bounds.** Evaluating upper bounds for $\rho^{(2)}$ is a more difficult task. Certainly, such a bound must be smaller than one. In this case it gives a lower bound for the spectral gap. We derive two effective upper bounds for $\rho^{(2)}$. The first one is pretty rough but it is obtained "for free" during the construction of the invariant polytope. The second one is sharper but its evaluation requires additional computations.

The following theorem gives an upper bound for the second Lyapunov exponent and a lower bound for the spectral gap, respectively. Here we use the extremal polytope norm generated by the invariant polytope determined after a succesful application of Algorithm 1.

Theorem 5.1. Let $\mathcal{T}$ be a multi-tree generated by a dominant product Π with $\rho(\Pi) = 1$ and $|\Pi| = k$. Let also $\mathcal{P}$ be the corresponding invariant polytope generated by Π and $\|\cdot\| = \|\cdot\|_{\mathcal{P}}$. Suppose $\mathcal{C}$ is a minimal cut set of $\mathcal{T}$ that consists of dead nodes, ℓ is the maximal level of its elements, and $\mu = \max_{\boldsymbol{v}\in\mathcal{C}} \|\boldsymbol{v}\|$; then

$$\rho^{(2)} \;\leq\; \mu^{\,\frac{1}{k+\ell-1}}\,. \tag{12}$$

Proof. Consider an arbitrary product $S \in \mathcal{A}^{\mathbb{N}}$ of length $|S| = s$ which is not a cyclic permutation of a power of Π. Denote $r = k + \ell - 1$. We are going to prove that $S^r\mathcal{P} \subset \mu\mathcal{P}$. This will imply that $\|S^r\| \leq \mu$ and hence $\rho(S)^r \leq \mu$ which yields (12).

We need to prove that for an arbitrary vertex $\boldsymbol{v}$ of $\mathcal{P}$, we have $\|S^r\boldsymbol{v}\| \leq \mu$. Every node $\boldsymbol{u} \in \mathcal{T}$ of level at least ℓ satisfies $\|\boldsymbol{u}\| \leq \mu$. Indeed, a path going from the root through $\boldsymbol{u}$ contains one vertex from $\mathcal{C}$, denote it by $\boldsymbol{w}$. The level of $\boldsymbol{w}$ cannot exceed ℓ, hence, $\boldsymbol{u}$ is a descendent of $\boldsymbol{w}$ and therefore, $\|\boldsymbol{u}\| \leq \|\boldsymbol{w}\| \leq \max_{\boldsymbol{v}\in\mathcal{C}} \|\boldsymbol{v}\| \;=\; \mu$.

If $\boldsymbol{v}$ is not in the root, then it is located on a positive level, and so the node $S^{\ell-1}\boldsymbol{v}$ has the level at least ℓ. Therefore, $\|S^{\ell-1}\boldsymbol{v}\| \le \mu$ and since $r > \ell - 1$, we have $\|S^r\boldsymbol{v}\| \le \mu$, which is required.

If $\boldsymbol{v}$ is in the root, then denote by p the biggest integer not exceeding r such that $S^p\boldsymbol{v}$ is in the root. Then all vertices $S^j\boldsymbol{v}$, $j \in \{0,\dots,p\}$, are in the root and they are different. Otherwise, the sequence $\{S^j\boldsymbol{v}\}_{j\in\mathbb{N}}$ is periodic and lies entirely on the root.

Hence, either Π is not primitive or S is a power of some cyclic permutation of Π, which is impossible. Thus, all vertices $S^j\boldsymbol{v}$, $j \le p$, are different, therefore, $p \le k-1$, so the vertex $S^k\boldsymbol{v}$ does not belong to the root. Consequently, its level is positive and hence the level of the vertex $S^r\boldsymbol{v} = S^{r-k}S^k\boldsymbol{v}$ is at least $s(r-k)+1 = s(\ell-1)+1 \ge \ell$. This implies that $\|S^r\boldsymbol{v}\| \le \mu$, which completes the proof. ☐

The upper bound (12) is rather rough because of the power $\frac{1}{k+\ell-1}$. On the other hand, it is always less than one, i.e., less than $\rho^{(1)}$, and hence, it always separates the first and second Lyapunov exponents and provides an upper of the spectral gap. Another advantage is that the bound (12) is obtained straightforwardly, one should simply choose an arbitrary minimal cut set $\mathcal{C}$.

Every minimal cut set $\mathcal{C}$ produces its own upper bound. The simplest choice is to use the minimal cut set evaluated in the invariant polytope algorithm. During the algorithm we visit the nodes of the multi-tree $\mathcal{T}$ level-by-level and remove the dead ones with the corresponding branches, for which $\|\boldsymbol{v}\|_{\mathcal{P}_j} < 1$, where $\mathcal{P}_j$ is the polytope with vertices constructed in the previous iterations. Since $\|\boldsymbol{v}\|_{\mathcal{P}} \le \|\boldsymbol{v}\|_{\mathcal{P}_j}$, we see that $\mathcal{C}$ can be chosen as the set of dead vertices appearing during the algorithm. Since the values $\|\boldsymbol{v}\|_{\mathcal{P}_j}$ are computed during the algorithm, the estimate (12) is obtained for free as a "by-product" of the invariant polytope.

The estimate (12) cannot be improved in those terms as the following example demonstrates.

Example 5.1. *Consider the family of two matrices:*

$$A_1 \;=\; \begin{pmatrix} 0 & 1 \\ \frac{1}{2} & 0 \end{pmatrix}\;; \qquad A_2 \;=\; \begin{pmatrix} 0 & \frac{1}{2} \\ 1 & 0 \end{pmatrix}$$

The leading eigenvector of A_1A_2 is $\boldsymbol{v}_1 = (1,0)^T$ and corresponds to the eigenvalue 1. Then $\boldsymbol{v}_2 = A_2\boldsymbol{v}_1 = (0,1)^T$. Since both points $A_1\boldsymbol{v}_1$ and $A_2\boldsymbol{v}_2$ are inside the polygon $\mathcal{P} = \mathrm{co}\,\{\pm\boldsymbol{v}_1\,,\,\pm\boldsymbol{v}_2\}$ (a square), the Invariant polytope algorithm produces the polygon $\mathcal{P}$. Thus, A_1A_2 is a dominant product and $\{\boldsymbol{v}_1,\boldsymbol{v}_2\}$ is a root. Hence $\rho(\mathcal{A}) = 1$. The corresponding extremal norm is $\|\cdot\| = \|\cdot\|_{\mathcal{P}}$ is actually the L_1 norm on the plane. We have $\mathcal{C} = \{A_1\boldsymbol{v}_1, A_2\boldsymbol{v}_2\}$ and $\mu = \max_{\boldsymbol{v}\in\mathcal{C}}\|\boldsymbol{v}\| = \frac{1}{2}$. On the other hand, $k=2, \ell=1$, therefore, the right-hand side in the inequality (12) is equal to $\frac{\sqrt{2}}{2}$. This is an upper bound for the second Lyapunov exponent $\rho^{(2)}$. Let us find its precise value. Each matrix of the family $\{A_1,A_2\}$ interchanges the coordinates with multiplication of one of them by $\frac{1}{2}$. Hence, every their product of length $2s$ is a diagonal matrix with elements $2^{-a}, 2^{-b}$, where a,b are both positive integer, otherwise, this is either $(A_1A_2)^s$ or $(A_2A_1)^s$, which is prohibited. So, the spectral radius of this product is at most $\frac{1}{2} < \frac{\sqrt{2}}{2}$. Every product of length $2s+1$ is a matrix with zero diagonal and with some elements c and d on the second diagonal, where $cd = 2^{-2k-1}$, since this is the minus determinant of this product. Hence, the spectral radius is equal to $\sqrt{cd} \;=\; 2^{-s-\frac{1}{2}} \le 2^{-\frac{1}{2}}$.

Thus, for the family $\mathcal{A} = \{A_1,A_2\}$, we have $\rho^{(2)} = \frac{\sqrt{2}}{2}$ and the spectral gap is $h(\mathcal{A}) = 1 - \frac{\sqrt{2}}{2}$. This value is attained only for two products of length one: A_1 and A_2. The inequality (12) becomes an equality, hence this bound is tight.

5.3. **A tight upper bound.** The main drawback of the estimate from Theorem 5.2 is the small exponent $\frac{1}{k+\ell-1}$, which significantly narrows down the spectral gap. Can this power be possibly removed (at the cost of some extra computation)? The answer is affirmative, although it requires a certain increase of the complexity. Before formulating the results let us introduce some further notation.

Definition 5.1 (Basement)**.** For every vertex $\boldsymbol{v} \in \mathcal{T}$ of the level $r \ge 0$, we consider the backward path of length $r+k-1$ starting at $\boldsymbol{v}$. This backward path goes from $\boldsymbol{v}$ down to the root and then makes a non-closed cycle along the root visiting all its k vertices.

The inverse (usual) path will be called the *basement* of the vertex $\boldsymbol{v}$ and denoted by $\beta(\boldsymbol{v})$. This path starts at the root, makes a non-closed cycle visiting all k vertices, then jumps off the root and makes r steps to $\boldsymbol{v}$.

Each vertex of $\mathcal{T}$ has a unique basement.

Denote by $\mathcal{B}(\boldsymbol{v})$ the set of all paths contained in the basement and by $\mathcal{B}'(\boldsymbol{v}) \subset \mathcal{B}(\boldsymbol{v})$ the set of those paths which are identified by a sequence which is not a power of a cyclic permutation of the root.

Also denote by $B(\boldsymbol{v})$ the set of all paths from $\mathcal{B}(\boldsymbol{v})$ that end at the node $\boldsymbol{v}$. Similarly with $B'(\boldsymbol{v})$. Obviously, $\mathcal{B}(\boldsymbol{v}) = \bigcup_{\boldsymbol{u}\in\beta(\boldsymbol{v})} B(\boldsymbol{u})$ and $\mathcal{B}'(\boldsymbol{v}) = \bigcup_{\boldsymbol{u}\in\beta(\boldsymbol{v})} B'(\boldsymbol{u})$. We identify each path with the corresponding product of matrices. Thus, for each product $S \in \mathcal{B}'(\boldsymbol{v})$, we have $\rho(S) \le \rho^{(2)}$. For an arbitrary minimal cut set $\mathcal{C}$, we denote by $B(\mathcal{C})$, $\mathcal{B}(\mathcal{C})\,, \mathcal{B}'(\mathcal{C})$ the union of the corresponding sets over all vertices $\boldsymbol{v} \in \mathcal{C}$.

Theorem 5.2. In the notation of Theorem 5.1, we have

$$\rho^{(2)} \ \le \ \max\{\,\mu\,,\,\sigma\}, \tag{13}$$

where $\mu \ = \ \max\limits_{\boldsymbol{v}\in\mathcal{C}} \|\boldsymbol{v}\|$ and $\sigma \ = \ \max\limits_{S\in\mathcal{B}'(\mathcal{C})} \rho(S)$.

Note that $\sigma \le \rho^{(2)}$. Therefore, if $\sigma \ge \mu$, then (13) implies that $\sigma = \rho^{(2)}$ and this value is attained at the product from $\mathcal{B}'(\mathcal{C})$ with the largest spectral radius.

Thus, the estimate (13) either gives the precise value of $\rho^{(2)}$ or provides the upper bound μ, which is sharper than $\mu^{\frac{1}{k+\ell-1}}$ stated by Theorem 5.1. Therefore, Theorem 5.2 at any rate strengthens Theorem 5.1. However, its practical realization requires few extra computation of the spectral radii of all the matrix products from $\mathcal{B}'(\mathcal{C})$.

Proof of Theorem 5.2. Consider an arbitrary product $S = A_{\delta_s} \ldots A_{\delta_1}$, which is associated to the sequence of indices $\delta_s, \ldots, \delta_1$. Consider two cases:

(1) The product S does not end with the dominant product Π neither with its cyclic permutation. This means that either $s < k$ or $A_{\delta_k} \ldots A_{\delta_1}$ is not a cyclic permutation of Π. Take an arbitrary vertex $\boldsymbol{v}$ of $\mathcal{P}$ and consider the path $\boldsymbol{v} \to A_{\delta_1}\boldsymbol{v} \to A_{\delta_2}A_{\delta_1}\boldsymbol{v} \to \cdots \to S\boldsymbol{v}$ on the multi-tree $\mathcal{T}$. If this path contains a node from $\mathcal{C}$, then $\|S\boldsymbol{v}\|_{\mathcal{P}} \le \mu$. Otherwise, this path belongs to $\mathcal{B}'(S\boldsymbol{v})$. Indeed, the backward path from the vertex $S\boldsymbol{v}$ to $\boldsymbol{v}$ does not end with a cyclic permutation of Π, hence either $\boldsymbol{v}$ is not in the root or this path makes at most $k-1$ steps along the root before arriving at $\boldsymbol{v}$. This means that S is contained in the basement of the node $S\boldsymbol{v}$, that means $S \in \mathcal{B}(S\boldsymbol{v})$. Moreover, since S does not end with a cyclic permutation of Π it cannot be a power of a cyclic permutation and therefore, $S \in \mathcal{B}'(S\boldsymbol{v})$. Thus, either $\|S\boldsymbol{v}\|_{\mathcal{P}} \le \mu$ or $S \in \mathcal{B}'(S\boldsymbol{v})$. This holds for each vertex $\boldsymbol{v}$ of $\mathcal{P}$. If at least for one vertex $\boldsymbol{v}$, we have $S \in \mathcal{B}'(S\boldsymbol{v})$, then $\rho(S) \le \sigma$, which completes the proof. Otherwise, if $\|S\boldsymbol{v}\|_{\mathcal{P}} \le \mu$ for all vertices $\boldsymbol{v}$ of $\mathcal{P}$, then $S(\mathcal{P}) \subset \mu\mathcal{P}$, Hence, $\|S\|_{\mathcal{P}} \le \mu$, which implies $\rho(S) \le \mu$, and we again arrive at (13). The proof of this case is completed.

(2) The product S ends with a cyclic permutation of Π. Since S is not a power of a cyclic permutation of Π, it follows that S has the form $S = QAR\Pi^s T$, where T is a product associated to a prefix of Π (if this prefix is empty, then we set T to be the identity matrix), $s \ge 0$, $R \in \mathcal{A}^{k-1}$ is the suffix of Π of length $k-1$, A is some matrix from $\mathcal{A}$ different from the first factor of Π, and Q is some product from $\mathcal{A}^{\mathbb{N}}$. Considering a cyclic permutation of S we obtain the product $S' = \Pi^s TQAR$. Clearly, $\rho(S') = \rho(S)$. On the other hand, the product AR has length k but it is not a cyclic permutation of Π. To see this, let A be one of the factors forming Π; we observe that all cyclic permutations of Π contain the same number of factors A. Denote this number by ℓ, $\ell \ge 0$. Since the first factor of Π is different from A and R coincides with the suffix of Π of length $k-1$, it follows that the product AR contains $\ell+1$ factors A. Thus, S' does not end with a cyclic permutation of Π. Therefore, S' satisfies the assumptions of Case 1 and consequently, $\rho(S') \le \max\{\mu, \sigma\}$. Hence, $\rho(S) \le \max\{\mu, \sigma\}$, which completes the proof also of this case

5.4. **Computing the upper bounds.** Theorem 5.2 leads to several possible methods of estimating of $\rho^{(2)}$ from above depending on the choice of the minimal cut set $\mathcal{C}$. Note that this set may be chosen in many ways and each of them produces its own estimate.

The direct application of the Invariant polytope algorithm. In this method we take for $\mathcal{C}$ the set of dead vertices that appear in the Invariant polytope algorithm. This is the minimal possible set that can produce estimate (13). Hence, the direct method ensures the minimal amount of computation. However, the estimate can be rough because the Invariant polytope algorithm deal with dead vertices of the smallest levels, for which the norm can be very close to 1. For computing σ we produce at each iteration a new node $\boldsymbol{v} \in \mathcal{T}$, evaluate the spectral radii of all $r+k-1$ products from the set $B(\boldsymbol{v})$, where r is the level of $\boldsymbol{v}$, i.e., of all paths from the basement of $\boldsymbol{v}$ that end with $\boldsymbol{v}$. Thus, along with every vertex we have to compute $r+k-1$ spectral radii. To reduce the computation we can store all the performed paths in the lexicographic order and sort out the redundant products, which have been already computed in the previous iterations.

The deep search method. We run the Invariant polytope algorithm and produce the polytope $\mathcal{P}$ and the corresponding extremal norm. Thus, μ is already obtained. The parameter σ can be computed by the same iterative procedure described above. If the gap $\mu - \sigma$ is large, we can start a new routine with the set of dead vertices denoted by $\mathcal{C}_1 = \mathcal{C}$. We find the vertex with the biggest norm in $\mathcal{C}_1$ and replace it by its m children, thus we obtain $\mathcal{C}_2$, etc. Since the norm of a child does not exceed the norm of the parent, it follows that the value $\max_{\boldsymbol{v}\in\mathcal{C}_j} \|\boldsymbol{v}\|_{\mathcal{P}}$ is non-increasing in j, hence the new upper bound μ is at least non worse than the previous one. This relaxation procedure can significantly improve the upper bound for $\rho^{(2)}$.

Some conclusions on the two upper bounds. In practice the the bound (13) is significantly sharper than (12), although its evaluation requires more time. The bound (13) can be made arbitrarily close to $\rho^{(2)}$ (by convergence of σ to $\rho^{(2)}$ as the depth of $\mathcal{C}$ goes to infinity) by increasing the depth of the cut set $\mathcal{C}$. However, in practice this is not efficient since the cardinality of $\mathcal{C}$ may grow exponentially with the depth. For this reason we are going to introduce another approach based on the branch-and-bound idea.

6. The Branch-and-bound method

In this section we present an algorithm that computes $\rho^{(2)}$ with an arbitrary precision.

We are given a family $\mathcal{A}$ and a number $\varepsilon > 0$. We need to find a number σ such that $\sigma \le \rho^{(2)} \le \sigma + \varepsilon$. It is assumed that $\mathcal{A}$ possesses a dominant product Π and an invariant polytope $\mathcal{P}$ and is normalised so that $\rho(\Pi) = 1$. Thus, the family $\mathcal{A}$ is normalised: $\rho^{(1)} = 1$.

Algorithm 2.

Initialisation. We set $\sigma = 0$, $j = 1$, and $\mathcal{R}_0$ is the (multi-) root of the multi-tree $\mathcal{T}$.

The jth iteration. We have a positive number σ and a subset $\mathcal{R}_{j-1}$ of the $(j-1)$-th level of the multi-tree $\mathcal{T}$. We set $\mathcal{R}_j = \emptyset$.

If $\mathcal{R}_{j-1} = \emptyset$, then STOP, the algorithm halts. We have $\sigma \le \rho^{(2)} \le \sigma + \varepsilon$.

Otherwise, take an arbitrary node $\boldsymbol{v} \in \mathcal{R}_{j-1}$ and consider all its children $\{\boldsymbol{u} = A\boldsymbol{v} \mid A \in \mathcal{A}\}$. For each child $\boldsymbol{u}$, if $\|\boldsymbol{u}\|_{\mathcal{P}} < \sigma + \varepsilon$, then $\boldsymbol{u}$ is a dead vertex, and the whole branch growing from it is dead. Otherwise, we put $\mathcal{R}_j = \mathcal{R}_j \cup \{\boldsymbol{u}\}$ and set the new value of σ:

$$\sigma \;=\; \max\Big\{\,\sigma\,,\,\rho(S)\,,\; S \in B'(\boldsymbol{u})\Big\}\,.$$

Recall that $B'(\boldsymbol{u})$ is set of products (paths) from the basement $\beta(\boldsymbol{u})$ that end at $\boldsymbol{u}$ excluding those (if any) that are powers of cyclic permutations of Π.

We exhaust all children of $\boldsymbol{v}$. Do it for all nodes $\boldsymbol{v} \in \mathcal{R}_{j-1}$. Then the jth iteration is completed.

A pseudocode of the the algorithm follows.

Remark 6.1. We see that Algorithm 2 exploits the same idea of sorting out redundant branches as other branch-and-bound algorithms including those for computing the joint spectral radius: Gripenberg's algorithm [11], C. Möller and U. Reif algorithm [29], and the Invariant polytope algorithm [12]. The core of all those algorithms is the relaxation procedure for the lower bound, which is responsible for cutting the redundant branches. In Algorithm 2, this procedure is quite complicated. Coming to a new node we need to compute the spectral radii not only of the path leading to this node, but also of all paths of its basement. Hence, an important issue here is a proper organization of this computation. It is possible to develop the

Algorithm 2: Pseudocode of Algorithm 2

Data: Matrices $A_1, A_2, \ldots, A_m$ scaled such that $\rho(\mathcal{A}) = 1$, $\varepsilon > 0$, maxit
Result: σ, μ
begin
  Let Π the s.m.p
  Compute the invariant polytope $\mathcal{P}$ by the polytope algorithm
  Set $j = 0$
  Set $\sigma_0 = 0, \mu_0 = 2, \mathcal{R}_0 = \{\boldsymbol{v}_1, \ldots, \boldsymbol{v}_k\}$ (set of leading eigenvectors)
  while $\mu_j - \sigma_j > \varepsilon$ **do**
    Compute $\mathcal{R}_j = \mathcal{A}\mathcal{R}_{j-1}$
    Eliminate the vectors in $\mathcal{R}_j$ belonging to previous levels
    Set $\mu_j = 0$
    foreach $u \in \mathcal{R}_j$ **do**
      Compute $\mu = \|u\|$
      if $\mu < \sigma_j$ **then**
        Label u as a dead vertex of $\mathcal{R}_j$
      else
        Set $\mu_j = \max\{\mu_j, \mu\}$
        Set $\sigma_j = \max\{\sigma_j, \max\limits_{S \in B'(u)} \rho(S)\}$
      Set $j = j + 1$
    if $j >$ maxit **then**
      Stop
  Set $j = j - 1$
  Set $\mu = \mu_j, \sigma = \sigma_j$

same idea described in the end of Section 5: to store all the performed paths in the lexicographic order and compare each new path with them to avoid the double computations.

The main theorem of this section, Theorem 6.1 below, states that Algorithm 2 always halts within finite time and gives an estimate σ for $\rho^{(2)}$ satisfying $\sigma \le \rho^{(2)} \le \sigma + \varepsilon$. Before we formulate and prove this results let us clarify the key steps and the structure of Algorithm 2. We do it in the following remark.

Remark 6.2. The family $\mathcal{A}$ is normalised so that $\rho^{(1)} = 1$, which is equivalent to say that its dominant product satisfies $\rho(\Pi) = 1$. In the jth iteration of the algorithm ($j \ge 1$) we form the set $\mathcal{R}_j$ of newly appeared nodes $A\boldsymbol{v}$ $A \in \mathcal{A}$, which are the survived children of all nodes from $\boldsymbol{v} \in \mathcal{R}_j$. A child node survives if its norm (the Minkowski norm generated by the invariant polytope $\mathcal{P}$) is smaller than $\sigma + \varepsilon$. Let us remark that the parameter σ is updated each step. To determine whether the child survives we use the current (defined by this step) value of σ, and we leave it unchanged to the next step (i.e., to the new pair $(\boldsymbol{v}, A)$) if the child is dead. Otherwise, if the child survives, we update the value σ by replacing it with $\max\big\{\,\sigma\,,\,\rho(S)\,,\,S \in B'(\boldsymbol{u})\big\}$. This replacement either does non change the value of σ or increases it. Thus, in the jth iteration of the algorithm, the value of σ is updated exactly $|\mathcal{R}_j|$ times, which is the total number of survivors. Each update does not reduce σ, so, the values of σ form a non-decreasing sequence. The algorithm runs until in some iteration all children die, i.e., the set of survivors $\mathcal{R}_j$ is empty. In this case the algorithm halts and the final value of σ gives the desired estimate for $\rho^{(2)}$.

In the proof of Theorem 6.1 below we need the following auxiliary lemma which is proved in the Appendix.

Lemma 6.1. Suppose $\alpha = \boldsymbol{u}^0 \to \boldsymbol{u}^1 \to \cdots$ is an infinite path along $\mathcal{T}$ (the node $\boldsymbol{u}^h$ belongs to the hth level); then $\boldsymbol{u}^h \to 0$ as $h \to \infty$, unless α becomes periodic with the period $|\Pi| = k$ (with $\Pi \in \mathcal{A}^k$ the dominant s.m.p.).

Theorem 6.1. For every $\varepsilon > 0$, Algorithm 2 terminates within finite time. The obtained value of σ satisfies the inequality $\sigma \le \rho^{(2)} \le \sigma + \varepsilon$.

Proof. If the algorithm halts after the jth iteration, then $\mathcal{R}_j = \emptyset$ and hence, the set $\mathcal{C}$ of dead nodes obtained during all j iterations is a minimal cut set. For each $\boldsymbol{u} \in \mathcal{C}$, the norm $\|\boldsymbol{u}\|_{\mathcal{P}}$ does not exceed the current value of σ plus ε. Since σ does not decrease during the algorithm, it follows that, for the final value of σ, we have $\|\boldsymbol{u}\|_{\mathcal{P}} \le \sigma + \varepsilon$. Applying Theorem 5.2 to this minimal cut set $\mathcal{C}$ and accordingly defining μ from $\mathcal{C}$ as in the statement of the theorem, one has $\mu \le \sigma + \varepsilon$, so that we obtain $\rho^{(2)} \le \max\{\sigma, \sigma+\varepsilon\} = \sigma + \varepsilon$. Thus, $\sigma \le \rho^{(2)} \le \sigma + \varepsilon$.

It remains to prove that the algorithm always terminates within a finite number of iterations. If this is not the case, then there exists an infinite path $\alpha = \boldsymbol{v}^0 \to \boldsymbol{v}^1 \to \cdots$ (the node $\boldsymbol{v}^h$ belongs to the hth level) that does not meet any dead node. This means that $\|\boldsymbol{v}^h\|_{\mathcal{P}} \ge \sigma_h + \varepsilon$, where σ_h is the value of σ by the iteration when the node $\boldsymbol{v}^h$ was involved. The node $\boldsymbol{v}^0$ belongs to the root, hence it is a leading eigenvector of a cyclic permutation of Π. Without loss of generality it can be assumed that $\boldsymbol{v}^0$ is the leading eigenvector of Π, and so $\Pi \boldsymbol{v}^0 = \boldsymbol{v}^0$. In what follows we use a simple notation $\boldsymbol{v}^0 = \boldsymbol{v}$. The sequence $\{\sigma_h\}_{h=0}^{\infty}$ is non-decreasing, hence it has a limit σ as $h \to \infty$. Since $\sigma_h \le \rho^{(2)}$ for all h, it follows that $\sigma \le \rho^{(2)}$. The sequence $\{\boldsymbol{v}^h\}_{h=1}^{\infty}$ does not tend to zero because $\|\boldsymbol{v}^h\| > \varepsilon$ for all h. Therefore, Lemma 6.1 yields the path α is eventually periodic with the period $|\Pi|$. Consequently, α contains the nodes $\Pi^\ell S \boldsymbol{u}$ for all $\ell \in \mathbb{N}$, where $S \in \mathcal{A}^{\mathbb{N}}$ is some product. It may be assumed (possibly, replacing Π by its cyclic permutation) that the length of S is not divisible by k. Thus, for all ℓ large enough, we have

$$\|\Pi^\ell S \boldsymbol{v}\|_{\mathcal{P}} \ > \ \sigma_h + \varepsilon \ > \ \sigma + \frac{\varepsilon}{2}\,. \tag{14}$$

Let $\lambda \in \mathbb{C}$ be the leading eigenvalue of Π, $|\lambda| = 1$. Since all other eigenvalues of Π are smaller than one in absolute value, it follows that the distance between the power Π^ℓ and the rank one matrix $\lambda^\ell \boldsymbol{v}\boldsymbol{u}^T$ tends to zero as $\ell \to \infty$, where $\boldsymbol{u}$ is the left leading eigenvector of Π such that $(\boldsymbol{u}, \boldsymbol{v}) = 1$. Therefore, one can choose a subsequence $\{\ell_q\}_{q \in \mathbb{N}}$ such that Π^{ℓ_q} tends to the matrix $\boldsymbol{v}\boldsymbol{u}^T$ as $q \to \infty$, Thus, $\Pi^{\ell_q} S \boldsymbol{v} \to \boldsymbol{v}\boldsymbol{u}^T S \boldsymbol{v}$ as $q \to \infty$. Note that the path $\Pi^{\ell_q} S$ is contained in the basement of the node $\boldsymbol{v}$, hence, this path belongs to the set $B(\Pi^{\ell_q} S \boldsymbol{v})$. On the other hand, the length of $\Pi^{\ell_q} S$ is $n\ell_q + |S|$, which is not divisible by n. Therefore, $\Pi^{\ell_q} S$ is not a power of a cyclic permutation of Π, so the path $\Pi^{\ell_q} S$ belongs to $B'(\Pi^{\ell_q} S \boldsymbol{v})$. Therefore, $\rho(\Pi^{\ell_q} S) \le \sigma$. This inequality holds for all $q \in \mathbb{N}$, hence, in the limit as $q \to \infty$, we obtain $\rho(\boldsymbol{v}\boldsymbol{u}^T S) \le \sigma$. On the other hand, $(\boldsymbol{v}\boldsymbol{u}^T S)\boldsymbol{v} = \boldsymbol{v}(\boldsymbol{u}^T S \boldsymbol{v})$, so $\boldsymbol{v}$ is an eigenvector of the matrix $\boldsymbol{v}\boldsymbol{u}^T S$ with the eigenvalue $\boldsymbol{u}^T S \boldsymbol{v} = (\boldsymbol{u}, S\boldsymbol{v})$. However, this matrix has rank one, hence, this is the leading eigenvector and $\rho(\boldsymbol{v}\boldsymbol{u}^T S) = |(\boldsymbol{u}, S\boldsymbol{v})|$. Therefore, $|(\boldsymbol{u}, S\boldsymbol{v})| \le \sigma$. We see that the vector $\Pi^{\ell_q} S \boldsymbol{v}$ converges to $|(\boldsymbol{u}, S\boldsymbol{v})|\, \boldsymbol{v}$ as $q \to \infty$. On the other hand, the norm of the vector $|(\boldsymbol{u}, S\boldsymbol{v})|\, \boldsymbol{v}$ is equal to $|(\boldsymbol{u}, S\boldsymbol{v})|$ (because $\|\boldsymbol{v}\|_{\mathcal{P}} = 1$ since $\boldsymbol{v}$ belongs to the root), which does not exceed σ. Hence, for all sufficiently large ℓ_q, we have $\left\|\Pi^{\ell_q} S \boldsymbol{v}\right\|_{\mathcal{P}} < \sigma + \frac{\varepsilon}{2}$. This contradicts to (14). The proof is completed. ☐

7. The real case

We consider the case when the leading eigenvalue of the dominant product is real, in short the *real case*. In this case a weakened analogue of the finiteness property holds. This will allow us to derive an algorithm for the exact computation of the second Lyapunov exponent.

7.1. **The weak finiteness property.** The main result of this section, Theorem 7.1, asserts that if a family $\mathcal{A}$ has a dominant product Π with a real leading eigenvalue, then the second Lyapunov exponent is always achieved at a finite product of matrices from $\mathcal{A}$, possibly multiplied by the matrix $\Pi^\infty = \lim_{n\to\infty} \Pi^n$. This property will be referred to as *the weak finiteness property.* Moreover, there is an upper bound for the length of this product that can be obtained algorithmically (Corollary 7.1).

We need to introduce some further notation. We consider infinite sequences of matrices from $\mathcal{A}$. All sequences will be numbered as matrix products: from the right to the left. So, a prefix of a product is its right subword, which starts from the first matrix. The matrix sequences will be identified with the corresponding infinite path on the multi-tree starting from the root.

Every finite nonempty subword of degree τ of a periodic sequence with period τ is called a *periodic product*. For example $(A_1A_2^2)^n$ is a periodic sequence and $A_1A_2^2$ as well as $A_2A_1A_2$ are periodic products.

A product which is not periodic is referred to as *aperiodic*. The longest aperiodic prefix of an aperiodic product is the *maximal aperiodic prefix*. Clearly, each aperiodic product possesses a unique maximal aperiodic prefix, which is nonempty. In terms of the multi-tree, a product D is periodic precisely when there is an element of the root such that the path corresponding to the product and starting at that element stays at the root. Respectively, every path corresponding to an aperiodic product gets off the root, whatever node of the roots it starts.

Lemma 7.1. Let Π be the dominant s.m.p. of degree k. For an arbitrary aperiodic product, its maximal aperiodic prefix has the form $D = J_1\Pi^r J_0$ with some $r \geq 0$ and some products J_0, J_1 such that $|J_0| \leq k-1$ and $1 \leq |J_1| \leq k$.

Proof. If $|D| \leq 2k-1$, then the assertion is obvious: we set $r = 0$ and split D into two factors of lengths at most k and $k-1$ respectively. If $|D| > 2k-1$, then we remove the last factor of D and obtain a periodic product D' of length at least $2k-1$. Let r be the maximal power of Π contained in D'. Then $D' = J_2\Pi^r J_0$, where J_0, J_2 are products, possibly empty, of length at most $k-1$. Indeed, if, say, $|J_2| \geq k$, then it starts with the period Π (because D' is periodic), which contradicts to the maximality of r. Then we obtain J_1 as a concatenation of J_2 and the last factor of D. $\square$

Note that the product J_0 in Lemma 7.1 can be absent but J_1 cannot.

We use the previous notation: $\mathcal{C}$ is a minimal cut set of dead nodes produced by the Invariant polytope algorithm, ℓ is the *depth* of the multi-tree, i.e., the maximal level of elements of $\mathcal{C}$, $\mathcal{P}$ is the invariant polytope, and $\mu = \max_{\boldsymbol{v}\in\mathcal{C}} \|\boldsymbol{v}\|_{\mathcal{P}}$. As we have shown, the norm of every descendant of an element of $\mathcal{C}$ does not exceed μ. In particular, the norm of an arbitrary element of level at least ℓ does not exceed μ.

Lemma 7.2. For an arbitrary aperiodic product D and for an arbitrary product S of length at least $\ell - 1$ (with ℓ the depth of the multi-tree), we have $\|SD\|_{\mathcal{P}} \leq \mu$.

Proof. Take an arbitrary element $\boldsymbol{v}$ of the root. Since D is aperiodic, it follows that $D\boldsymbol{v}$ is out of the root and hence, $SD\boldsymbol{v}$ is an element of level at least ℓ. Therefore, $\|SD\boldsymbol{v}\|_{\mathcal{P}} \leq \mu$. This holds for every node from the root. On the other hand, if $\boldsymbol{v}$ is not in the root, then the level of $SD\boldsymbol{v}$ all the more exceeds ℓ. Thus, for every vertex of the polytope $\mathcal{P}$, we have $\|SD\boldsymbol{v}\|_{\mathcal{P}} \leq \mu$, and consequently, $\|SD\|_{\mathcal{P}} \leq \mu$. $\square$

Everything said above is true for every family of matrices that possesses a dominant product Π. Now we invoke our assumption that the leading eigenvalue of Π is real. In this case the limit $\Pi^\infty = \lim_{r\to\infty}\Pi^r$ exits and is equal to rank one matrix $\boldsymbol{v}\boldsymbol{u}^T$, where $\boldsymbol{v}, \boldsymbol{u}$ are respectively the right and the left eigenvectors of Π and $(\boldsymbol{u}, \boldsymbol{v}) = 1$. For arbitrary $n \geq 1$, $r_j \in \mathbb{N} \cup \{0, \infty\}$, and for some products S_j of matrices from the family $\mathcal{A}$, $j = 1, \ldots, n$, we consider the product

$$S_n\Pi^{r_n}\cdots S_1\Pi^{r_1} \tag{15}$$

This is a product of elements of the set of $m+1$ matrices $\mathcal{A} \cup \{\Pi^\infty\}$.

Now we can prove the weak finiteness property for the second Lyapunov exponent: it is always attained on a finite product of matrices from $\mathcal{A} \cup \{\Pi^\infty\}$. Moreover, the following theorem establishes the structure of this product.

Theorem 7.1. For every family of matrices $\mathcal{A}$ that possesses a dominant product with a real leading eigenvalue, the second Lyapunov exponent $\rho^{(2)}$ is attained at some product of the form (15) with $n < \frac{\log\sigma}{\log\mu}$, where the number σ is a certain lower bound for $\rho^{(2)}$ obtained by Algorithm 2 , the products S_j are of lengths at most $2k + \ell - 3$, the powers r_i are nonnegative integers and r_1 can take the value ∞.

Proof. See the appendix. $\square$

Thus, we see that in the real case the second Lyapunov exponent is always attained at a product of matrices from $\mathcal{A}$ with at most one factor Π^∞.

Remark 7.1. Due to Theorem 7.1, the number n of blocks in the product (15) does not exceed the number $\log\sigma/\log\mu$. All the finite powers r_i can also be estimated: they are less than a special constant ℓ that depends on the invariant polytope and can be evaluated algorithmically, see Corollary 7.1. Thus, for every family with a real leading eigenvalue of the dominant product, the second Lyapunov exponent can be found among a finite set of products.

To get an upper bound for n it suffices to take an arbitrary lower bound σ for $\rho^{(2)}$. In particular, the spectral radius of an arbitrary product (different from a power of a cyclic permutation of Π) suffices. Certainly, the bigger σ the less n, hence it makes sense to get a tight upper bound, for example, using Algorithm 2

7.2. **Precise computation of the second Lyapunov exponent.** The weak finiteness property does not mean that the finite product (15) that gives the second Lyapunov exponent can be efficiently found. Certainly, Theorem 7.1 restricts the number of factors n and the length of the products S_j in (15). Moreover, we will see that the finite powers r_i, $i \le n$, can also be bounded by an effective constant ℓ (Remark 7.1). So, the problem is merely reduced to a finite search by exhaustion. However, the number of potential candidate products can be enormous. Nevertheless, in most cases it is possible to find the optimal product using the branch-and bound approach. We present Algorithm 3 that consists of two phases. The phase 1 computes an approximate value of $\rho^{(2)}$ which is often (but not always) is an exact value. If it is not, then one applies the phase 2, which always terminates with the exact value of $\rho^{(2)}$. Note that the phase 2 can be applied immediately without the phase 1, but in practice it is more efficient to begin with the phase 1. The phase 1 is realised by the same scheme as Algorithm 2 with two modifications:

1) we set $\varepsilon = 0$;

2) in each iteration we compute the spectral radii of the path from the root to the current node multiplied by Π^{∞}.

Thus, we are given a family $\mathcal{A}$, which is assumed to possess a dominant product Π with the leading eigenvalue 1. Let $\mathcal{P}$ be the invariant polytope of $\mathcal{A}$. It is needed to find $\rho^{(2)}(\mathcal{A})$ and a product Q in which this value is attained. This is a product of the matrices from $\mathcal{A}$ and, possibly, of one extra matrix Π^{∞}.

We denote by $\mathcal{R} = \{\boldsymbol{v}_1, \ldots, \boldsymbol{v}_k\}$ the vertices of the root of the multi-tree $\mathcal{T}$, so $\Pi\boldsymbol{v}_1 = \boldsymbol{v}_1$. By $\boldsymbol{u}_1$ we denote the left eigenvector of Π such that $(\boldsymbol{u}_1, \boldsymbol{v}_1) = 1$. Thus, $\Pi^{\infty} = \boldsymbol{v}_1\boldsymbol{u}_1^T$.

Algorithm 3, phase 1.

Initialisation. We set $Q_\sigma = A_1$, $\sigma = 0$, $\varepsilon = 0, j = 1$, and $\mathcal{R}_{j-1} = \mathcal{R}$ is the root. Choose the maximal number of iterations N.

The jth iteration. We have nonnegative numbers σ, ε, and a subset $\mathcal{R}_{j-1}$ of the $(j-1)$st level of the multi-tree $\mathcal{T}$. We set $\mathcal{R}_j = \emptyset$ and $\varepsilon = 0$.

If $\mathcal{R}_{j-1} = \emptyset$, then STOP, the algorithm halts. We have $\rho^{(2)} = \sigma$ and $Q = Q_\sigma$.

Otherwise, take an arbitrary node $\boldsymbol{v} \in \mathcal{R}_{j-1}$ and consider all its children $\{\boldsymbol{u} = A\boldsymbol{v} \mid A \in \mathcal{A}\}$. For each child $\boldsymbol{u}$, if $\|\boldsymbol{u}\|_{\mathcal{P}} < \sigma$, then $\boldsymbol{u}$ is a dead vertex, and the whole branch growing from it is dead. Otherwise, we put $\mathcal{R}_j = \mathcal{R}_j \cup \{\boldsymbol{u}\}$ and set the new values of σ and ε:

$$\sigma \;=\; \max\Big\{\,\sigma\,,\;\rho(S)\,,\;S \in B'(\boldsymbol{u})\Big\}\,; \qquad \varepsilon \;=\; \max\big\{\,\varepsilon\,,\;\|\boldsymbol{u}\|_{\mathcal{P}} - \sigma\,\big\}\,.$$

If the maximum in the formula for σ is attained at some $S \in B'(\boldsymbol{u})$,

then we set $Q_\sigma = S$.

Recall that $B'(\boldsymbol{u})$ is the set of products (paths) from the basement $\beta(\boldsymbol{u})$ that end at $\boldsymbol{u}$ excluding those (if any) that are powers of the cyclic permutations of Π.

We exhaust all children of $\boldsymbol{v}$. If $|(\boldsymbol{u}_1, \boldsymbol{v})| > \sigma$, then we set $\sigma = |(\boldsymbol{u}_1, \boldsymbol{v})|$ and $Q_\sigma = \boldsymbol{v}_1\boldsymbol{u}_1^T S$, where S is the product corresponding to the path from $\boldsymbol{v}_1$ to $\boldsymbol{v}$. Then we go to the next node of $\mathcal{R}_{j-1}$.

Do it for all nodes $\boldsymbol{v} \in \mathcal{R}_{j-1}$. Then the jth iteration is completed.

If $j + 1 = N$, then STOP. The algorithm is interrupted. We have $\sigma \le \rho^{(2)} \le \sigma + \varepsilon$ and $Q = Q_\sigma$.

Remark 7.2. At each iteration of the phase 1 we compute not only spectral radii of the products from the basement of the vertex $\boldsymbol{v}$ but also of the product $\boldsymbol{v}_1\boldsymbol{u}_1^T S = \Pi^\infty S$, where S is the product corresponding to the path from $\boldsymbol{v}_1$ to $\boldsymbol{v}$. Note that $\Pi^\infty S\boldsymbol{v}_1 = \Pi^\infty \boldsymbol{v} = (\boldsymbol{u}_1, \boldsymbol{v})\boldsymbol{v}_1$, hence, the spectral radius of the matrix $\Pi^\infty S$ is equal to $|(\boldsymbol{u}_1, \boldsymbol{v})|$. If this value exceeds σ, then we update and set $\sigma = \rho(\Pi^\infty S) = |(\boldsymbol{u}_1, \boldsymbol{v})|$.

Phase 2 is applied only if the phase 1 has not found the exact value, i.e., under the following assumptions:

(a) The phase 1 has not terminated before the Nth iteration, i.e., if $\varepsilon > 0$;
(b) The product Q_σ, for which the finite value is attained, has the form $Q_\sigma = \boldsymbol{v}_1\boldsymbol{u}_1^T$.

We denote by $L = \boldsymbol{u}_1^\perp$ the hyperplane orthogonal to the left eigenvector $\boldsymbol{u}_1$ of the dominant product Π. To realise the phase 2 we add an extra *admissible polytope* $E \subset L$ such that $\Pi E \subset E$. Such a polytope exists since the spectral radius of the operator Π restricted to L is less than one. It can be iteratively constructed by the process $\boldsymbol{v}(n+1) = \Pi \boldsymbol{v}(n)$ with an arbitrary point $\boldsymbol{v}(0) \in L$ which does not belong to a smaller invariant plane of Π. Once $\boldsymbol{v}(n+1) \in \mathrm{co}\,\{\boldsymbol{v}(0), \ldots, \boldsymbol{v}(n)\}$, we stop; the polytope $E = \mathrm{co}\,\{\boldsymbol{v}(0), \ldots, \boldsymbol{v}(k)\}$ is admissible. A number $h > 0$ is an *admissible factor* if for every dead vertex $\boldsymbol{v}$ in the minimal cut set produced by the Invariant polytope algorithm, we have $\boldsymbol{v} + hSE \subset P$, where S is the product corresponding to the path from $\boldsymbol{v}_1$ to $\boldsymbol{v}$, i.e., $S\boldsymbol{v}_1 = \boldsymbol{v}$. Then for every vertex $\boldsymbol{v}$ of $\mathcal{P}$, we denote $\tilde{\boldsymbol{v}} = \boldsymbol{v} + hSE$. Thus, $\tilde{\boldsymbol{v}}$ is an $(d-1)$-dimensional polytope. The polytope

$$\tilde{\mathcal{P}} = \mathrm{co}\left\{\tilde{v} \mid v \text{ is the vertex of } \mathcal{P}\right\}$$

is an extremal polytope for $\mathcal{A}$.

Define the linear operator $B\boldsymbol{x} = \boldsymbol{x} - (\boldsymbol{u}_1, \boldsymbol{x})\,\boldsymbol{v}_1$. This is a projection onto L parallel to the vector $\boldsymbol{v}_1$. It is easily seen that B commutes with Π. For a given $t > 0$, let $\ell = \ell(t)$ be the minimal number such that $B\,\Pi^\ell \tilde{P} \subset tE$.

Algorithm 3, phase 2.

Initialisation. We denote by σ_0 the final value of σ obtained in phase 1 and respectively $Q_{\sigma_0} = Q_\sigma$. We set $j = 1$, and $\mathcal{R}_{j-1} = \mathcal{R}$ is the root. We take an admissible extra polytope $E \subset L$ and an admissible factor $h > 0$. For each node $\boldsymbol{v}$ of the multi-tree $\mathcal{T}$, we consider the corresponding polytope $\tilde{\boldsymbol{v}} = \boldsymbol{v} + hSE$, where S is the product associated to the path from $\boldsymbol{v}_1$ to $\boldsymbol{v}$. Define the extremal polytope $\tilde{\mathcal{P}}$ and the number $\ell = \ell(h\sigma_0)$.

The jth iteration. We have a positive number σ, and a subset $\mathcal{R}_{j-1}$ of the $(j-1)$st level of the multi-tree $\mathcal{T}$. We set $\mathcal{R}_j = \emptyset$.

If $\mathcal{R}_{j-1} = \emptyset$, then STOP, the algorithm halts. We have $\sigma = \rho^{(2)}$ and $Q = Q_\sigma$.

Otherwise, take an arbitrary node $\boldsymbol{v} \in \mathcal{R}_{j-1}$.

If the path from $\boldsymbol{v}_1$ to $\boldsymbol{v}$ ends with the power of $\Pi^{\,\ell}$, then $\boldsymbol{v}$ is a dead node, and the whole branch growing from $\boldsymbol{v}$ is dead.

Otherwise consider all children of $\boldsymbol{v}$. For each child $\boldsymbol{u} = A\boldsymbol{v}$, $A \in \mathcal{A}$, if for all vertices $\boldsymbol{w}$ of $\tilde{\boldsymbol{u}}$, we have $\|\boldsymbol{w}\|_{\tilde{\mathcal{P}}} < \sigma$, then $\boldsymbol{u}$ is a dead node, and the whole branch growing from it is dead.

Otherwise, we put $\mathcal{R}_j = \mathcal{R}_j \cup \{\boldsymbol{u}\}$ and set the new value of σ:

$$\sigma \;=\; \max\Big\{\,\sigma\,,\,\rho(S)\,,\;S \in B'(\boldsymbol{u})\Big\}\,.$$

If the maximum in this formula for σ is attained at some $S \in B'(\boldsymbol{u})$, then we set $Q_\sigma = S$.

We exhaust all children of $\boldsymbol{v}$. If $|(\boldsymbol{u}_1, \boldsymbol{v})| > \sigma$, then we set $\sigma = |(\boldsymbol{u}_1, \boldsymbol{v})|$ and $Q_\sigma = \boldsymbol{v}_1\boldsymbol{u}_1^T S$, where S is the product corresponding to the path from $\boldsymbol{v}_1$ to $\boldsymbol{v}$. Then we go to the next node of $\mathcal{R}_{j-1}$.

We do this for all nodes $\boldsymbol{v} \in \mathcal{R}_{j-1}$. Then the jth iteration is completed.

We continue until after some jth iteration we have $\mathcal{R}_j = \emptyset$. After this the algorithm halts. We have $\sigma = \rho^{(2)}$ and $Q = Q_\sigma$.

Theorem 7.2. Suppose the family $\mathcal{A}$ has a dominant product with a real leading eigenvalue. If Algorithm 3 halts (either in phase 1 or in phase 2), then it returns the value of $\rho^{(2)}$ and the optimal product Q. The phase 2 always halts.

Proof. The proof is in the Appendix. □

Remark 7.3. By Theorem 7.2, one could apply only phase 2 without the phase 1. Since phase 2 always halts, it returns the precise value of $\rho^{(2)}$ along with the product Q. However, in practice it is more efficient to start with phase 1 because phase 2 has a much higher complexity. The reason is that the number of vertices of the extremal polytope $\tilde{\mathcal{P}}$ can be much larger than the number of vertices of $\mathcal{P}$.

Now we can derive an upper bound for all finite powers r_i in the formula (15) for the optimal product Q.

Corollary 7.1. The phase 2 of Algorithm 3 produces the optimal product Q of the form (15) in which all finite r_i do not exceed $\ell(\sigma h)$, where h is the admissible factor for the extra polytope E and σ is an arbitrary lower bound for $\rho^{(2)}$.

Proof. Theorem 7.2 ensures that the phase 2 of Algorithm 3 always halts and returns the optimal product Q, for which $\rho(Q) = \rho^{(2)}$. From the routine of phase 2 it follows that each path terminating with Π^{ℓ} contains a dead vertex. Therefore, none of the products Q_σ can contain Π^{ℓ} and hence, neither can Q as one of those products. □

7.3. **Counterexample in the non-real case.**

Remark 7.4. In the non-real case the weak finiteness property may fail. The reason is clear: if the leading eigenvalue of the dominant product Π is non-real, then Π^{∞} is not well defined since the sequence Π^r does not have a limit as $r \to \infty$. The following simple example shows that the second Lyapunov exponent may not be attained at eventually periodic products in this case.

Example 7.1. *Choose an arbitrary real α such that α/π is irrational and set*

$$A_1 \;=\; \begin{pmatrix} \cos\alpha & -\sin\alpha \\ \sin\alpha & \cos\alpha \end{pmatrix} \;; \qquad A_2 \;=\; \begin{pmatrix} 0 & 0 \\ 0.5 & 0 \end{pmatrix} .$$

Thus, A_1 is the rotation by the angle α and A_2 is the orthogonal projection to the abscissa axis rotated by $\pi/2$ and multiplied by $\frac{1}{2}$. Since $\rho(A_1) = 1$ and all matrix product except for the powers of A_1 have the Euclidean norm less than $\frac{1}{2}$, we see that the product $\Pi = A_1$ is dominant, $\rho = 1$ and $\rho^{(2)} \le \frac{1}{2}$. If $n\alpha$ is close to $\frac{\pi}{2}$ $\pmod{2\pi}$, then the spectral radius of the product $A_1^n A_2$ is close to $\frac{1}{2}$. Hence, $\rho^{(2)} = \frac{1}{2}$. However, this value is unreachable for any eventually periodic product of matrices A_1, A_2. Indeed, if this product contains at least two matrices A_2, then its spectral radius does not exceed $\frac{1}{4}$. If it contains only one matrix A_2, then up to a cyclic permutation it coincides with $A_1^n A_2$. If a subsequence of such products is such that $\rho(A_1^{n_j} A_2) \to \frac{1}{2}$, then the sequence $n_j\alpha$ tends to $\frac{\pi}{2}$ $\pmod{2\pi}$.

8. Illustrative examples

In this section we consider two illustrative examples. The first one is given by two sign-matrices and the second one is taken from the literature on Daubechies wavelets and gives rise to finite termination of the Algorithm 2.

Example A (Algorithm 2). We consider Example 3.1 in Section 3.1: $\mathcal{A} = \{A_1, A_2\}$ with

$$A_1 = \alpha \begin{pmatrix} 1 & -1 \\ 1 & 0 \end{pmatrix}, \qquad A_2 = \alpha \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}, \tag{16}$$

with $\alpha = \left(\frac{1}{2}\left(3 - \sqrt{5}\right)\right)^{1/5} \approx 0.824907$, which is such that $\rho(\mathcal{A}) = 1$.

Flow of Algorithm 2. We indicate by $\|\cdot\| = \|\cdot\|_{\mathcal{P}}$ the polytope real norm associated to the polytope $\mathcal{P}$ constructed by the polytope algorithm (see Figure 1).

We set a tolerance $\varepsilon = 25 \cdot 10^{-4}$.

Step 1 $j = 1$, $\mathcal{R}_0 = \{\boldsymbol{v}_1, \boldsymbol{v}_2, \boldsymbol{v}_3, \boldsymbol{v}_4, \boldsymbol{v}_5\}$. We set $\sigma_1 = \rho(A_1) = \rho(A_2) = \alpha = 0.8249069\ldots$. We compute

$$v_7 = A_1 \boldsymbol{v}_1, \qquad v_8 = A_2 \boldsymbol{v}_2, \qquad v_9 = A_1 \boldsymbol{v}_4, \qquad v_{10} = A_2 \boldsymbol{v}_5$$

and get

$$\|v_7\| \approx 0.90, \|v_8\| \approx 0.68, \|v_9\| \approx 0.56, \|v_{10}\| \approx 0.68$$

which implies that only the vectors v_6 and v_7 are alive at level $j = 1$.

Step 2 $j = 2$, $\mathcal{R}_1 = \{v_6, v_7\}$. We get unvaried $\sigma_2 = \sigma_1$ and $\mathcal{R}_1 = \{v_6, v_7\}$. We compute

$$v_{11} = A_1 v_7, \qquad v_{12} = A_2 v_7, \qquad v_{13} = A_1 v_6, \qquad v_{14} = A_2 v_6$$

and get

$$\|v_{11}\| \approx 0.560, \|v_{12}\| \approx 0.839, \|v_{13}\| \approx 0.987, \|v_{14}\| \approx 0.680$$

which implies that only the vectors v_{12} and v_{13} are those alive at level $j = 2$.

Step 3–8 The algorithms continues similarly and only the vector v_{27} is alive at level $j = 8$.

Halt We get $\|v_{27}\| - \sigma_8 = 0.00241... < \varepsilon$. The algorithm stops with the estimate:

$$\rho^{(2)}(\mathcal{A}) = 0.9827\ldots$$

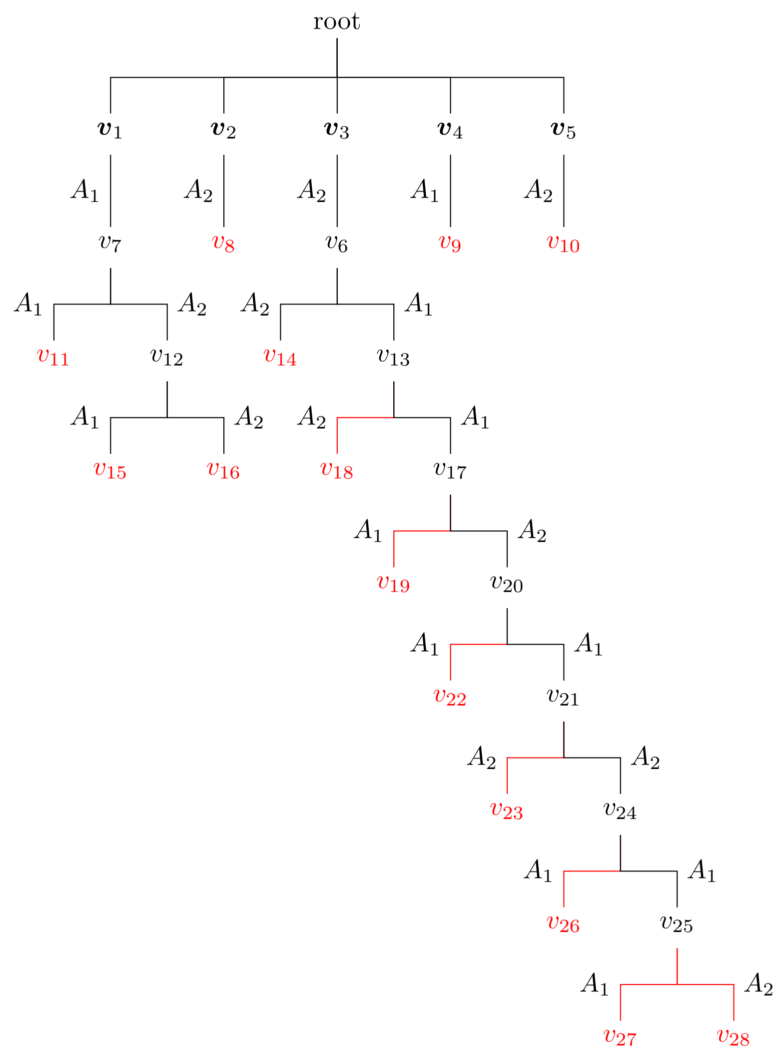

Indeed the computation suggests that the sequence $Q_\ell = A_1 A_2 \Pi^\ell$ provides the highest possible rate of growth in the product semigroup apart from the product Π and its cyclic permutations with their powers. We get in particular

$$\lim_{\ell\to\infty} \Pi^\ell = \begin{pmatrix} \frac{1}{10}\left(5+\sqrt{5}\right) & \frac{1}{\sqrt{5}} \\ \frac{1}{\sqrt{5}} & \frac{1}{10}\left(5-\sqrt{5}\right) \end{pmatrix}$$

which provides the lower bound to the second Lyapunov exponent:

$$\rho(Q_\infty) = \frac{1}{5}\left(\frac{1}{2}\left(3-\sqrt{5}\right)\right)^{2/5}\left(5+\sqrt{5}\right) \approx 0.984787.$$

This means that the spectral gap is bounded above by $0.015213\ldots$.

Example B (Algorithm 2, finite termination). We consider the wavelet function Daubechies-4, which generates the system of C^1-wavelets on the real line [31]. We denote this function by ψ. It is well-known that $\psi \in C^1(\mathbb{R})$, the Hölder exponent of its derivative ψ' is computed in [10] and is equal to $\alpha_{\psi'} = 0.61792\ldots$. The local regularity of this function was studied in [10, 35]. We are going to analyse the local regularity in rational points.

Let us recall the procedure for the construction of the matrices A_1, A_2. For every $N = 1, 2, \ldots$ we have a set of $2N$ *Daubechies filter coefficients*: $c_0, \ldots, c_{2N-1}$. They have the property $\sum_{i=0}^{2N-1} c_i = 2$, and the polynomial $m(z) = \sum_{n=0}^{2N-1} c_n z^n$ has a zero of order N at $z = -1$. We set

$$q(z) = \frac{m(z)}{\left((1+z)/2\right)^N} = \sum_{n=0}^{N-1} q_n z^n.$$

Then we are able to write the transition matrices of dimension $N-1$ as as follows:

$$(A_1)_{ij} = q_{2i-j-1}, \qquad (A_2)_{ij} = q_{2i-j}, \qquad\qquad i, j = 1, \ldots, N-1.$$

At this point we can compute $\rho(\mathcal{A})$ by the polytope algorithm and then try to estimate the spectral gap.

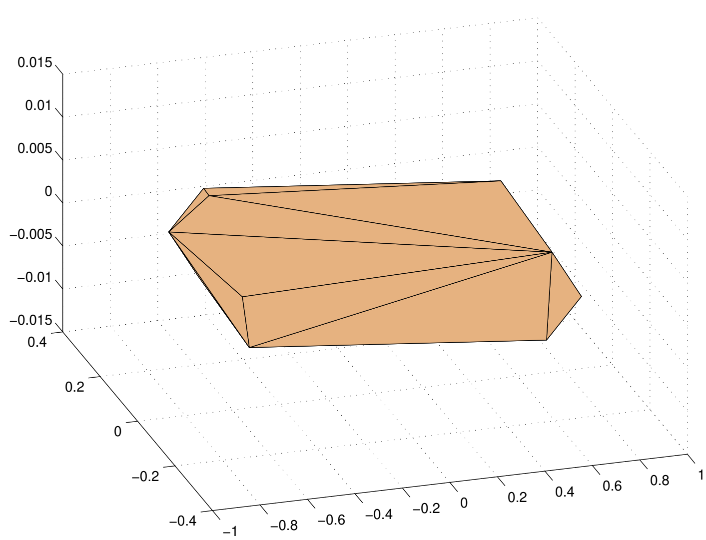

FIGURE 2. The invariant polytope for the Daubechies matrices for $N = 4$ computed by the polytope algorithm.

For $N = 4$ we get

$$T_1 = \begin{pmatrix} 5.212854848820774 & 0 & 0 \\ 1.703224934278843 & -4.676287953813834 & 5.212854848820774 \\ 0 & -0.239791829285782 & 1.703224934278843 \end{pmatrix},$$

$$T_2 = \begin{pmatrix} -4.676287953813834 & 5.212854848820774 & 0 \\ -0.239791829285782 & 1.703224934278843 & -4.676287953813834 \\ 0 & 0 & -0.239791829285782 \end{pmatrix}.$$

Applying Algorithm 1 allows to prove that the s.m.p. is T_1 having spectral radius $\rho(T_1) = 5.212854848820774\ldots$.

We normalize as

$$\mathcal{A} = \{A_1, A_2\}, \qquad A_1 = T_1/\rho(T_1), A_2 = T_2/\rho(T_1).$$

To apply the polytope algorithm we compute the normalized leading eigenvector of A_1,

$$\boldsymbol{v}_1 = \begin{pmatrix} 1.0000 \\ 0.1662 \\ -0.0113 \end{pmatrix}$$

Algorithm 1 converges and generates a centrally symmetric real polytope with $6 \cdot 2$ vertices (in 7 iterations).

The additional vertices of the polytope are the following (with their symmetric):

$$v_2 = A_2\boldsymbol{v}_1 = \begin{pmatrix} -0.7308 \\ 0.0185 \\ 5.2249 \cdot 10^{-4} \end{pmatrix}, \quad v_3 = A_1 v_2 = \begin{pmatrix} -0.7308 \\ -0.2548 \\ -6.8062 \cdot 10^{-4} \end{pmatrix}, \quad v_4 = A_1 v_3 = \begin{pmatrix} -0.7308 \\ -0.0108 \\ 0.0115 \end{pmatrix},$$

$$v_5 = A_1 v_4 = \begin{pmatrix} -0.7308 \\ -0.2175 \\ 0.0042 \end{pmatrix}, \quad v_6 = A_1 v_5 = \begin{pmatrix} -0.7308 \\ -0.0393 \\ 0.0113 \end{pmatrix}$$

and the associated polytope $\mathcal{P} = \mathrm{co}\{\pm\boldsymbol{v}_1, \pm v_2, \pm v_3, \pm v_4, \pm v_5, \pm v_6\}$ (here co denotes the convex hull) is shown in Figure 2.

Flow of Algorithm 2. We indicate by $\|\cdot\| = \|\cdot\|_{\mathcal{P}}$ the polytope real norm associated to the polytope $\mathcal{P}$ constructed by the polytope algorithm (see the Figure).

We set a tolerance $\varepsilon = 25 \cdot 10^{-4}$.

Step 1 $j = 1$, $\mathcal{R}_0 = \{\boldsymbol{v}_1\}$. We set $\sigma_1 = \rho(A_1) = \rho(A_2) = 0.8582\ldots$. We compute

$$v_2 = A_2\boldsymbol{v}_1$$

which is a vertex so that

$$\|v_2\| = 1$$

which is naturally alive at level $j = 1$.

Step 2 $j = 2$, $\mathcal{R}_1 = \{v_2\}$. We get $\sigma_2 = \sigma_1 = 0.8582\ldots$. We compute

$$v_3 = A_1 v_2, \qquad v_7 = A_2 v_2$$

and

$$\|v_3\| = 1\,, \; \|v_7\| \approx 0.9602, \tag{17}$$

which implies that v_3 and v_7 are alive at level $j = 2$.

Step 3–10 The algorithms continues similarly generating 26 vectors, and at level $j = 10$, we get $\mathcal{R}_9 = \{v_{25}\}$.

We get $\sigma_{10} = \sigma_9 = 0.8582\ldots$ and compute

$$v_{27} = A_1 v_{25}, \qquad v_{28} = A_2 v_{25}$$

and since

$$\|v_{27}\| \approx 0.8504 < \sigma_{10}, \|v_{28}\| \approx 0.7883 < \sigma_{10}, \tag{18}$$

we have that no vertices are alive at level $j = 10$.

Halt The algorithm stops with the exact value:

$$\rho^{(2)}(\mathcal{A}) = \rho(A_2) = 0.8582\ldots$$

This means that $\rho^{(2)}(\mathcal{A}) = \rho(A_2) = 0.8582\ldots$ and the spectral gap is given exactly by $0.1417\ldots$. For the Daubechies-4 wavelet function ψ, this implies that at every dyadic rational point x, we have $\omega_{\psi',x}(t) \asymp t^{\alpha}$, where $\alpha = 0.6179...$, while for all non-dyadic rationals x, we have $\omega_{\psi',x}(t) \le C\, t^{\alpha}\, q^{-\frac{\log t}{\ell}}$, where $q = 0.8582...$, where ℓ is the length of the period in the binary expansion of x.

Conclusions

In this article we have proposed novel iterative methods - based on the polytope algorithm to compute the joint spectral radius - to compute the spectral gap of a family of matrices, in other words the difference between the joint spectral radius (first Lyapunov exponent) and the second Lyapunov exponent. The knowledge of the spectral gap is important in several applications, like the behavior of a switching system towards its asymptotic state, or the Hölder regularity of Daubechies wavelets or fractal curves, where the second Lyapunov exponent plays a key role. To our knowledge there do not exist in the present literature algorithms able to compute to an arbitrary accuracy the second Lyapunov exponent and the algorithms we propose appear to be very efficient in all examples we have considered. Exploiting the special structure of the family of matrices might further improve these algorithms.

Acknowledgments

The authors are grateful to an anonymous referee for their careful reading of a preliminary version of this article and their very constructive suggestions. Nicola Guglielmi acknowledges that his research was supported by funds from the Italian MUR (Ministero dell'Università e della Ricerca) within the PRIN 2017 Project "Discontinuous dynamical systems: theory, numerics and applications" and by the INdAM Research group GNCS (Gruppo Nazionale di Calcolo Scientifico).

Appendix

Proof of Lemma 6.1. Consider the sequence of indices of the path α. Since α does not have a period Π, it follows that for every vertex $\boldsymbol{v}$ of the extremal polytope $\mathcal{P}$, there exists a prefix of this sequence such that for the corresponding matrix product S, we have $\|S\boldsymbol{v}\|_{\mathcal{P}} < 1$. Hence, $\|S\boldsymbol{v}\|_{\mathcal{P}} < \mu$, where μ is the maximum of norms of dead nodes that appear in the Invariant polytope algorithm. This inequality is true also for all longer prefixes of α, implying that there is a prefix S_1 of α such that $\|S_1\boldsymbol{v}\|_{\mathcal{P}} \le \mu$ for all vertices of $\mathcal{P}$. Hence $\|S_1\|_{\mathcal{P}} \le \mu$. Then we remove that prefix from α and repeat the same argument for the remaining sequence, whose prefix is S_2 of the such that $\|S_2\|_{\mathcal{P}} \le \mu$, etc. This way we split α to parts $S_1, S_2, \ldots$ such that $\|S_j\| \le \mu$ for all j. We have $\|S_1 \cdots S_r \boldsymbol{u}^0\| \le \mu^r$, so the sequence of numbers $\{\|\boldsymbol{u}^n\|\}_{n=0}^{\infty}$ has a subsequence that tends to zero. However, this sequence is non-increasing, so it tends to zero as well. $\Box$

Proof of Theorem 7.1. Let $q = \lceil \frac{\log \sigma}{\log \mu} \rceil - 1$. If the second Lyapunov exponent is attained at a product Q of length at most $q(2k + \ell - 3)$, then we set $r_j = 0$ and split the product into several products S_j of lengths at most $2k + \ell - 3$. Otherwise, the second Lyapunov exponent is attained at the limit of some sequence (maybe stationary) of products of lengths at least $q(2k + \ell - 2)$. Passing to a subsequence we can assume two cases:

1) All products of that sequence are periodic. In this case each of them has the form $J_1 \Pi^r J_0$ with some $r \ge 0$ and for products J_0, J_1 of lengths at most $k - 1$. Since the set of such pairs (J_0, J_1) is finite, it follows that $\rho^{(2)}$ is either achieved at one of products $J_1 \Pi^r J_0$ or at $J_1 \Pi^{\infty} J_0$, which completes the proof in this case.

2) All products of that sequence are aperiodic. By Lemma 7.1, the maximal aperiodic prefix D_1 of each product Q_1 from that sequence has the form $J_1^{(1)} \Pi^{r_1} J_0^{(1)}$ with $J_1^{(1)} \le k - 1, J_1^{(1)} \le k$. If $|Q_1| - |D_1| \le \ell - 2$ for an infinite number of elements Q_1 of that sequence, then arguing as above we conclude that $\rho^{(2)}$ is either achieved at one of products $T^{(1)} \Pi^{r_1} J_0^{(1)}$ with $r_1 \in \mathbb{N} \cup \{0, \infty\}$, where $|T^{(1)}| \le k + \ell - 2$ and $|J_0^{(1)}| \le k - 1$.

Since ρ is invariant under cyclic permutations, we see that $\rho^{(2)} = \rho\big(T^{(1)}\Pi^{r_1}J_0^{(1)}\big) = \rho\big(\Pi^{r_1}J_0^{(1)}T^{(1)}\big)$. Set $S_1 = J_0^{(1)}T^{(1)}$ and note that $|S_1| \le 2k+\ell-3$. Since $\rho(S_1\Pi^{r_1}) = \rho^{(2)}$, the proof for this case is completed.

Otherwise, if $|Q_1| - |D_1| \ge \ell - 1$ for an infinite number of elements Q_1, then we denote by $T^{(1)}$ the concatenation of $J_1^{(1)}$ with the subword of Q_1 of length $\ell - 1$ following $J_1^{(1)}$. Then we invoke Lemma 7.2 and conclude that $\|T^{(1)}\Pi^{r_1}J_0^{(1)}\|_{\mathcal{P}} \le \mu$. The suffix of Q_1 following this product will be denoted by Q_2. Since the $\mathcal{P}$-norm of every matrix from $\mathcal{A}$ does not exceed one, it follows that $\|Q_1\|_{\mathcal{P}} \le \mu$. Then we apply the same line of reasoning to Q_2 instead of Q_1. We obtain that either the second Lyapunov exponent is attained on the product (15) with $n=2$ or Q_2 has a prefix $T^{(2)}\Pi^{r_2}J_0^{(2)}$ of the norm at most μ. In the latter case $\|Q_2\| \le \mu$ and therefore $\|Q_1\|_{\mathcal{P}} \le \mu\,\|Q_2\| \le \mu^2$. Running q iterations we obtain: either $\rho^{(2)}$ is reached for a product (15) with $n<q$ or $\|Q_1\|_{\mathcal{P}} \le \mu^q$. The latter is impossible since $\mu^q < \sigma \le \rho^{(2)}$ and therefore the sequence of such products Q_1 cannot tend to $\rho^{(2)}$.

Thus, we have shown that the second Lyapunov exponent is attained at a product (15) with some powers r_j, possibly infinite. Among those products, take one with the smallest number n and show that at most one of the powers r_i is infinite. In this case we can put it at the first position using the invariance of the spectral radius under cyclic permutation. Assume the contrary: let at least two of them be infinite. Without loss of generality it can be assumed that $r_1 = r_j = \infty$, where $j \ge 2$. After a cyclic permutation setting r_1 to the last position we write (15) in the form $Q = \Pi^\infty A \Pi^\infty B$, where A and B are some nonempty products. Since $\Pi^\infty = \boldsymbol{v}\,\boldsymbol{u}^T$, the range of this operator is the line spanned by $\boldsymbol{v}$. Hence, so is the range of Q and of $\Pi^\infty B$. Therefore, $\boldsymbol{v}$ is their common leading eigenvector: $Q\boldsymbol{v} = \lambda\boldsymbol{v}$ and $\Pi^\infty Bv = \lambda\boldsymbol{v}$, where $|\lambda| = \rho^{(2)}$ and $|\mu| \le 1$ (because the $\mathcal{P}$-norm of $\Pi^\infty B$ is at most one).

Thus, we have $\lambda\boldsymbol{v} = Q\boldsymbol{v} = \Pi^\infty A\mu\boldsymbol{v}$, and so $\Pi^\infty A\boldsymbol{v} = \frac{\lambda}{\mu}\,\boldsymbol{v}$. However, $\left|\frac{\lambda}{\mu}\right| \ge |\lambda| = \rho^{(2)}$. On the other hand, the product A is not periodic with period Π (otherwise it can be concatenated with Π^∞), hence, the equality $\Pi^\infty A\boldsymbol{v} = \pm\boldsymbol{v}$ is impossible, and consequently, the spectral radius of $\Pi^\infty A$ is less than one, i.e., it does not exceed $\rho^{(2)}$. On the other hand, we have shown that it is not smaller than $\rho^{(2)}$, hence, it is equal to this number. Thus, the second Lyapunov exponent is attained at the product $\Pi^\infty A$ with the smaller number of factors than Q. This contradicts to the choice of Q and therefore, concludes the proof. $\square$

Proof of Theorem 7.2. If Algorithm 3 halts in phase 1, then $\sigma = \rho^{(2)}$ and $Q = Q_\sigma$. The proof is literally the same as in Theorem 6.1 after setting $\varepsilon = 0$.

Suppose Algorithm 3 halts in phase 2 and returns the value σ. In what follows we use the short notation $\|\cdot\|_{\tilde{\mathcal{P}}} = \|\cdot\|$. Let us show that for every dead node $\boldsymbol{v}$, we have $\|\boldsymbol{w}\| \le \sigma$ for all vertices $\boldsymbol{w}$ of $\tilde{v}$. In this case $\sigma = \rho^{(2)}$. Indeed, the union of nodes from $\tilde{\boldsymbol{v}}$ over all dead nodes $\boldsymbol{v}$ form a minimal cut set $\mathcal{C}$. Applying now Theorem 5.2 for the case case $\mu = \sigma$, we obtain $\rho^{(2)} \le \max\{\sigma, \sigma\} = \sigma$, and hence, $\rho^{(2)} = \sigma$. It remain to show that $\|\boldsymbol{w}\| \le \sigma$ for all vertices $\boldsymbol{w}$ of $\tilde{\boldsymbol{v}}$. In the phase 2 a node $\boldsymbol{v}$ is dead if either $\|\boldsymbol{w}\| \le \sigma_n$ for all vertices $\boldsymbol{w}$ of $\tilde{v}$, where σ_n is the value of σ when the algorithm arrives at the node v, or if $\boldsymbol{v} = \Pi^\ell S\boldsymbol{v}_1$. In the former case there is nothing to prove because $\sigma_n \le \sigma$ (the values of σ_n do not decrease during the algorithm), and hence, $\|\boldsymbol{w}\| \le \sigma_n \le \sigma$. Let us consider the latter case: $\boldsymbol{v} = \Pi^\ell S\boldsymbol{v}_1$. Note that the value of σ by the moment when the node $\boldsymbol{v}$ is involved is not less that by the moment when $S\boldsymbol{v}_1$ was involved, which is at least $|(\boldsymbol{u}_1, S\boldsymbol{v}_1)|$. Without loss of generality we assume that $(\boldsymbol{u}_1, S\boldsymbol{v}_1) \ge 0$ and denote this value by α. Thus, $\sigma \ge (\boldsymbol{u}_1, S\boldsymbol{v}_1) = \alpha$. We have $\Pi^\ell B\tilde{\mathcal{P}} \subset \sigma_0 hE$. Every vertex $\boldsymbol{v}$ form the set $S\tilde{\boldsymbol{v}}_1$ belongs to $\tilde{\mathcal{P}}$, therefore, $\Pi^\ell B\boldsymbol{w} \in \sigma_0 hE$. On the other hand, $\boldsymbol{w} = \alpha\boldsymbol{v} + B\boldsymbol{w}$. Hence,

$$\Pi^\ell\boldsymbol{w} \;=\; \Pi^\ell\big(\alpha\boldsymbol{v} + B\boldsymbol{w}\big) \;=\; \alpha\,\Pi^\ell\boldsymbol{v} + \Pi^\ell B\boldsymbol{w} \;=\; \alpha\,\boldsymbol{v} + \Pi^\ell B\boldsymbol{w} \tag{19}$$

Since $\Pi^\ell B\boldsymbol{w} \in \sigma_0 hE$, we can denote $\Pi^\ell\boldsymbol{w}B = \sigma_0\boldsymbol{e}$, where $\boldsymbol{e} \in hE$. Taking into account that $\boldsymbol{v} + hE \subset \tilde{\mathcal{P}}$, we have $\boldsymbol{v} + h\boldsymbol{e} \in \tilde{\mathcal{P}}$ and so, $\|\boldsymbol{v} + \boldsymbol{e}\| \le 1$. This yields that the norm of the last expression in (19) can be estimated as follows:

$$\Big\|\alpha\,\boldsymbol{v} + \Pi^r B\boldsymbol{w}\Big\| \;=\; \Big\|\alpha\,\boldsymbol{v} + \sigma_0\boldsymbol{w}\Big\| \;=\; \Big\|\sigma_0\,\boldsymbol{v} + \sigma_0\boldsymbol{e} + (\alpha-\sigma_0)\,\boldsymbol{v}\Big\| \;\le$$

$$\sigma_0\,\Big\|\boldsymbol{v} + \boldsymbol{e}\Big\| + (\alpha-\sigma_0)\,\Big\|\boldsymbol{v}\Big\| \;\le\; \sigma_0 + (\alpha-\sigma_0) \;=\; \alpha$$

Combining with (19) we obtain $\|\Pi^{\ell}\boldsymbol{w}\| \le \alpha$. Since $\alpha \le \sigma$, we conclude that $\|\Pi^{\ell}\boldsymbol{w}\| \le \sigma$, the proof of the first part of the theorem is completed.

It remains to prove that the phase 2 always halts. If it does not, then there is an infinite path $\boldsymbol{v}_1 \to \boldsymbol{v}_2 \to \cdots$ which starts with the leading eigenvector $\boldsymbol{v}_1$ of Π, leaves the root and does not meet any dead node. This means that $\|\boldsymbol{v}_n\| \ge \sigma_n$, where σ_n is the value of σ by the iteration when the node $\boldsymbol{v}_n$ is involved. Hence, $\|\boldsymbol{v}_n\| \ge \sigma_0$ for all n, and thus, $\boldsymbol{v}_n$ does not tend to zero as $n \to \infty$. Invoking Lemma 6.1 we conclude that the path α is eventually periodic with the period Π. However, after ℓ periods it arrives at a dead node, which is a contradiction. $\square$

References

[1] N. E. Barabanov, *Lyapunov indicator for discrete inclusions, I-III*, Autom. Remote Control, 49 (1988), No 2, 152–157.

[2] V.D. Blondel, J. Theys and A.A. Vladimirov, *An elementary counterexample to the finiteness conjecture*, SIAM J. Matrix Anal. Appl. 24 (2003), 963-–970.

[3] T. Bousch and J. Mairesse, Asymptotic height optimization for topical IFS, Tetris heaps, and the finiteness conjecture, J. Amer. Math. Soc., 15 (2002), pp. 77-–111.

[4] M. A. Berger and Y. Wang, *Bounded semigroups of matrices*, Lin. Algebra Appl. 166 (1992), 21–27.

[5] Y. Chen, J. Lü, H. Dong and X. Yu, *On the Lyapunov exponent of consensus algorithm*, Proceedings of the 10th World Congress on Intelligent Control and Automation, (2012), 931–936.

[6] Y. Chitour, G. Mazanti, M. Sigalotti, *On the gap between deterministic and probabilistic joint spectral radii for discrete-time linear systems*, Linear Alg. Appl., 613 (2021), 24-–45.

[7] V. Crespi, G. Cybenko, and G. Jiang, *The theory of trackability with applications to sensor networks*, ACM Tran. Sensor Networks 4 (2008), no 3, 1–42.

[8] X. Dai, Y. Huang, and M. Xiao. *Periodically switched stability induces exponential stability of discrete-time linear switched systems in the sense of Markovian probabilities*, Automatica J. IFAC, 47 (2011), 1512-–1519.

[9] G. Deslauriers and S. Dubuc, *Symmetric iterative interpolation processes*, Constr. Approx., 5 (1989), 49–68.

[10] I. Daubechies and J. Lagarias, *Two-scale difference equations. II. Local regularity, infinite products of matrices and fractals*, SIAM. J. Math. Anal., 23 (1992), 1031–1079.

[11] G. Gripenberg, *Computing the joint spectral radius*, Lin. Alg. Appl., 234 (1996), 43–60.

[12] N. Guglielmi and V.Yu. Protasov, *Exact computation of joint spectral characteristics of linear operators*, Found. Comput. Math., 13 (2013), no 1, 37–97.

[13] N. Guglielmi, V. Yu. Protasov, *Invariant polytopes of sets of matrices with applications to regularity of wavelets and subdivisions*, SIAM J. Matr. Anal. Appl., 37 (2016) 1, 18–52.

[14] N. Guglielmi, F. Wirth, and M. Zennaro, *Complex polytope extremality results for families of matrices*, SIAM J. Matrix Anal. Appl. 27 (2005), 721–743.

[15] N. Guglielmi and M. Zennaro, *An algorithm for finding extremal polytope norms of matrix families*, Lin. Alg. Appl. 428 (2008), 2265–2282.

[16] N. Guglielmi and M. Zennaro, *Finding extremal complex polytope norms for families of real matrices*, SIAM J. Matrix Anal. Appl. 31 (2009), 602–620.

[17] N. Guglielmi and M. Zennaro, *Stability of linear problems: joint spectral radius of sets of matrices,* in "Current Challenges in Stability Issues for Numerical Differential Equations", Lecture Notes in Mathematics 2082, CIME Foundation Subseries, L. Dieci and N. Guglielmi Eds., Springer, 2014.

[18] L. Gurvits, *Stability of discrete linear inclusions*, Lin. Alg. Appl., 231 (1995), 47–85.

[19] R. Jungers, *The joint spectral radius. Theory and applications*, Lecture Notes in Control and Information Sciences (2009), Springer.

[20] R. M. Jungers, V. Yu. Protasov, and V. D. Blondel, *Efficient algorithms for deciding the type of growth of products of integer matrices*, Linear Alg. Appl., 428 (2008), No 10, 2296–2312.

[21] K. Hare, I.D. Morris, N. Sidorov, and J. Theys, *An explicit counterexample to the Lagarias-Wang finiteness conjecture*, Adv. Math. 226 (2011), no. 6, 4667-–4701.

[22] V. S. Kozyakin, *Structure of extremal trajectories of discrete linear systems and the finiteness conjecture*, Automat. Remote Control, 68 (2007), 174–209.

[23] V. S. Kozyakin, N. ,A. Kuznetsov, and P. Yu Chebotarev. *Consensus in asynchronous multiagent systems. II. The joint spectral radius method*, Avtomat. i Telemekh., 5 (2019), 3–31.

[24] J.C. Lagarias and Y. Wang, *The finiteness conjecture for the generalized spectral radius of a set of matrices*, Linear Algebra Appl., 214 (1995), pp. 17-–42.

[25] D. Liberzon, *Switching in systems and control*, Birkhauser, Boston, MA, 2003.

[26] A. Mandel and I. Simon, *On finite semigroups of matrices*, Theor. Comput. Sci. 5 (1977), 101–111.

[27] T. Mejstrik, *Improved invariant polytope algorithm and applications*, ACM Trans. Math. Softw., 46 (2020) 3 (29), 1–26.

[28] T. Mejstrik and V. Yu. Protasov, *Elliptic polytopes and invariant norms of linear operators*, (2021), arXiv:2109.12159

[29] C. Möller and U. Reif, *A tree-based approach to joint spectral radius determination*, Linear Alg. Appl., 563 (2014), 154-170.

[30] R. Mc Naughton and Y. Zalcstein, *The Burnside problem for semigroups*, J. Algebra 34 (1975), 202–299.

[31] I. Y. Novikov, V. Yu. Protasov, and M. A. Skopina, *Wavelets theory*, AMS, Translations Mathematical Monographs, 239 (2011), 506 pp.

[32] M. Omladič and H. Radjavi, *Irreducible semigroups with multiplicative spectral radius*, Linear Alg. Appl., 251 (1997), 59–72.

[33] A. I. Popov, *On matrix semigroups bounded above and below*, Linear Alg. Appl., 438 (2013), no 11, 4439–4447.

[34] V. Yu. Protasov, *The generalized spectral radius. A geometric approach*, Izvestiya Math. 61 (1997), 995–1030.

[35] V. Yu. Protasov, *Fractal curves and wavelets*, Izvestiya Math., 70 (2006), 123–162.

[36] V. Yu. Protasov, *The Euler binary partition function and subdivision schemes*, Math. Comp. 86 (2017), 1499–1524.

[37] V. Yu. Protasov, *The Barabanov norm is generically unique, simple, and easily computed*, SIAM J. Contr. Optimiz., 60 (2022), no), 2246—2267.

[38] V. Yu. Protasov, A. Voynov, *Matrix semigroups with constant spectral radius*, Linear Alg. Appl., 513 (2017), 376–408.

[39] G. C. Rota and G. Strang, *A note on the joint spectral radius*, Kon. Nederl. Acad. Wet. Proc. Vol. 63 (1960), 379–381.

[40] F. Wirth, *The generalized spectral radius and extremal norms*, Linear Algebra Appl., 342 (2002), 17-–40.

Gran Sasso Science Institute (GSSI), via Crispi 7, I-67010 L' Aquila, Italy.
Email address: `nicola.guglielmi@gssi.it`

DISIM, University of L'Aquila, via Vetoio 1, I-67010, L'Aquila, Italy
Email address: `vladimir.protasov@univaq.it`