A Tree-Based Approach to Joint Spectral Radius Determination

Claudia Möller and Ulrich Reif

Abstract

We suggest a novel method to determine the joint spectral radius of finite sets of matrices by validating the finiteness property. It is based on finding a certain finite tree with nodes representing sets of matrix products. Our approach accounts for cases where one or several matrix products satisfy the finiteness property. Moreover, is potentially functional even for reducible sets of matrices.

Keywords: joint spectral radius, finiteness property, exact computation
2000 MSC: 15A18, 15A60

1. Introduction

The concept of the *joint spectral radius (JSR)* has gained some attention in recent years. As a generalization of the standard spectral radius, this characteristic of a set of matrices plays an important role in various fields of modern mathematics, see e.g. the monograph [Jun09]. First introduced by Rota and Strang in 1960 ([RS60]), the JSR was almost forgotten, and then rediscovered in 1992 by Daubechies and Lagarias ([DL92]) in the context of the analysis of refineable functions. In general, the JSR is not exactly computable [BT00], and even its approximation is NP-hard in some sense [BN05]. In practice, the situation is as follows: Beyond the analysis of special cases, the few known algorithms for a potential exact evaluation are based on establishing the *finiteness property (FP)*, as introduced below, for a given problem. This approach cannot claim universality because families of matrices exist[1] that do not exhibit the FP. However, to the best of our knowledge, all problems coming from applications have permitted a treatment via FP so far.

[1]Non-constructive proofs of this fact can be found in [BM02, BTV03, Koz05], while an explicit counterexample is given in [HMST11].

In [GP13, CGSCZ10, GZ09, GZ08, Mae08, Mae00], the FP is established by the construction of a special norm with certain extremal properties. The approach suggested here also aims at the FP but is different in the way how it is verified. Some pros and cons will be discussed in the final section of this paper.

The idea of a graph-theoretical analysis has proven to be useful before. In [Gri96], a branch and bound algorithm is used for approximating the JSR, and in [HMR09], the range of C^1-parameters of the four-point scheme is determined explicitly by considering certain infinite paths in the tree of all matrix products. Though we adapted some ideas from the latter work, the trees to be considered in the following are different: Their knots represent *sets of matrices* instead of single matrix products. This crucial idea potentially reduces the analysis of infinite sets of products to a study of *finite* subtrees. In particular, this aspect facilitates automated verification of the FP by computer programs.

After introducing some notation, our main results are presented in Section 3 and then proven in the subsequent section. Section 5 illustrates our new method by some model problems but does not comprise numerical tests since the focus of this work is on theoretical aspects. Some concluding remarks can be found in Section 6.

2. Setup

We consider a finite set $\mathcal{A} = \{A_1, ..., A_m\}$ of matrices in $\mathbb{C}^{d\times d}$. To deal with products of its elements, we introduce the sets

$$\mathcal{I}_0 := \{\emptyset\}, \quad \mathcal{I}_k := \{1, \dots, m\}^k, \quad \mathcal{I} := \bigcup_{k\in\mathbb{N}_0} \mathcal{I}_k,$$

of *completely positive index vectors* of length $k \in \mathbb{N}_0$ and arbitrary length, respectively. By contrast, an *index vector* may contain also negative entries, whose special meaning will be explained in the next section.

For $k \in \mathbb{N}$, we define the matrix product

$$A_I := A_{i_k} \cdots A_{i_1}, \quad I = [i_1, \dots, i_k] \in \mathcal{I}_k. \tag{1}$$

Otherwise, if $k = 0$, let $A_\emptyset := \mathrm{Id}$ be the identity matrix. The length of the vector $I \in \mathcal{I}_k$ is denoted by $|I| := k$. That is, any index vector $I \in \mathcal{I}$ encodes a matrix product A_I with $|I|$ factors.

Let $\|\cdot\|$ denote any norm on $\mathbb{C}^d$, and also the induced matrix norm. The *joint spectral radius (JSR)* of $\mathcal{A}$ is defined as

$$\hat{\rho}(\mathcal{A}) := \limsup_{k\to\infty} \max_{I\in\mathcal{I}_k} \|A_I\|^{\frac{1}{k}}.$$

In particular, if $m = 1$, then $\hat{\rho}(\mathcal{A}) = \rho(A_1)$ is the *standard spectral radius (SSR)* of A_1. Clearly, the JSR is independent of the chosen norm. As shown in [DL92], upper and lower bounds on $\hat{\rho}(\mathcal{A})$ are given by

$$\max_{I \in \mathcal{I}_k} \rho(A_I)^{\frac{1}{k}} \leq \hat{\rho}(\mathcal{A}) \leq \max_{I \in \mathcal{I}_k} \|A_I\|^{\frac{1}{k}} \tag{2}$$

for any $k \in \mathbb{N}$. The set $\mathcal{A}$ is said to have the *finiteness property (FP)*, if there exists a completely positive index vector $J \in \mathcal{I}_k$ such that equality holds on the left hand side of the latter display,

$$\rho(A_J)^{\frac{1}{k}} = \hat{\rho}(\mathcal{A}). \tag{3}$$

The method suggested here, as well as ideas developed in [GP13, CGSCZ10, GZ09, GZ08, Mae08, Mae00], are based on verifying (3) for some (more or less) sophisticated guess J. For instance, a reasonable way is to choose a possibly large value k, determine

$$\|A_L\| = \max_{I \in \mathcal{I}_k} \|A_I\| \quad \text{or} \quad \|A_L\| = \max_{I \in \mathcal{I}_k} \rho(A_I),$$

and check the trail of $L \in \mathcal{I}_k$ for a repeating pattern, defining J. The SSR as well as the JSR are homogeneous functions, i.e.,

$$\rho(\beta A_J) = |\beta|\, \rho(A_J), \quad \hat{\rho}(\beta\mathcal{A}) = |\beta|\, \hat{\rho}(\mathcal{A}), \quad \beta \in \mathbb{C}, \tag{4}$$

where $\beta\mathcal{A} := \{\beta A_1, \ldots, \beta A_m\}$. Hence, discarding the trivial case $\rho(A_J) = 0$, we can scale the family $\mathcal{A}$ such that at least one of the dominant eigenvalues of A_J equals 1, and in particular $\rho(A_J) = 1$. Equation (3), which has to be demonstrated, then reads

$$\hat{\rho}(\mathcal{A}) = \rho(A_J) = 1. \tag{5}$$

3. Main results

In the following, we consider a matrix family $\mathcal{A}$ for which the normalized equation (5) shall be proven. By (2), this is possible only if

$$\rho(A_i) \leq 1, \quad i = 1, \ldots, m,$$

so that we assume this property, throughout. In previous work of Guglielmi and Protasov [GP13], it is requested that there is only one index vector $J \in \mathcal{I}$ (and its cyclic permutations) satisfying it, but in general, this cannot be taken for granted. To account for this, we base our analysis on the investigation of a given *family* of index vectors. This family contains index vectors with $\rho(A_J) = 1$, but, for good reasons, we allow also index vectors with $\rho(A_J) < 1$. As will be demonstrated by an example in Section 5, such index vectors may reduce significantly the complexity of the trees to be constructed.

Definition 3.1. *Given matrices $\mathcal{A} = \{A_1, \dots, A_m\}$, consider some non-empty set $\mathcal{J} = \{J_1, \dots, J_n\}$ of completely positive index vectors $J_i \in \mathcal{I}$. If*

$$\max_{J \in \mathcal{J}} \rho(A_J) = 1,$$

then $\mathcal{J}$ is called a generator set *of $\mathcal{A}$, and each element $J \in \mathcal{J}$ is called a* generator.

It is our goal to relate properties of generator sets to the equality $\hat{\rho}(\mathcal{A}) = 1$. By (2), existence of a generator set implies $\hat{\rho}(\mathcal{A}) \geq 1$ so that (5) becomes equivalent to $\hat{\rho}(\mathcal{A}) \leq 1$.

In the following, let the set $\mathcal{A} = \{A_1, \dots, A_m\}$ of matrices and the generator set $\mathcal{J} = \{J_1, \dots, J_n\}$ be fixed. To address products of matrices in $\mathcal{A}$ conveniently, we introduce the sets

$$\mathcal{K}_0 := \{\emptyset\}, \quad \mathcal{K}_\ell := \{-n, \dots, -1, 1, \dots, m\}^\ell, \quad \mathcal{K} := \bigcup_{\ell \in \mathbb{N}_0} \mathcal{K}_\ell,$$

of *index vectors* of length $\ell \in \mathbb{N}_0$ and arbitrary length, respectively. As before, the length of $K \in \mathcal{K}_\ell$ is denoted by $|K| := \ell$.

Index vectors $K \in \mathcal{K}$ encode *sets* of matrix products in the following way: While single positive indices correspond to singletons, single negative indices correspond to infinite sets containing special matrix powers,

$$\mathcal{A}_\ell := \begin{cases} \{A_\ell\} & \text{if } \ell > 0, \\ \{A_{J_{-\ell}}^k : k \in \mathbb{N}_0\} & \text{if } \ell < 0. \end{cases}$$

Defining products of sets as sets of products, i.e., $\mathcal{P} \cdot \mathcal{Q} := \{PQ : P \in \mathcal{P}, Q \in \mathcal{Q}\}$, let

$$\mathcal{A}_K := \mathcal{A}_{k_\ell} \cdots \mathcal{A}_{k_1}, \quad K = [k_1, \dots, k_\ell] \in \mathcal{K}_\ell,$$

for $\ell \in \mathbb{N}$, and $\mathcal{A}_\emptyset := \{\mathrm{Id}\}$. This definition is similar to (1), but A_I is a single matrix, while $\mathcal{A}_K$ is always a set, even if $K \in \mathcal{I}$ is completely positive. In this case, $\mathcal{A}_K = \{A_K\}$ is a singleton, while otherwise, it is typically[2] a denumerable set.

We need some more notations and definitions: *Concatenation* of vectors $P \in \mathcal{K}_i$ and $S \in \mathcal{K}_j$ is denoted by

$$[P, S] := [p_1, \dots, p_i, s_1, \dots, s_j] \in \mathcal{K}_{i+j}.$$

[2]For instance, if all eigenvalues of the matrices A_J happen to be 0, then $\mathcal{A}_K$ is finite even if K is not completely positive.

Powers indicate concatenation of an index vector with itself,

$$K^1 := K, \quad K^{\ell+1} := [K^\ell, K].$$

If $K = [P, S]$, then P is a *prefix* and S is a *suffix* of K. The sets of prefixes and suffixes of $K \in \mathcal{K}$ are denoted by

$$\mathcal{P}(K) := \{P : K = [P, S] \text{ for some } S\},$$
$$\mathcal{S}(K) := \{S : K = [P, S] \text{ for some } P\},$$

respectively.

Example. Let $J_1 = [1, 2]$ be a generator for the set $\mathcal{A} = \{A_1, A_2\}$ of matrices. For $K = [1, 1, -1, 1, 2]$, we obtain

$$\mathcal{A}_K = \{A_2 A_1 (A_2 A_1)^k A_1^2 : k \in \mathbb{N}_0\}.$$

$P = [1, 1, -1] \in \mathcal{P}(K)$ is a prefix and $S = [1, 2] \in \mathcal{S}(K)$ is the complementary suffix of K. In this special case, we observe that $\mathcal{A}_K \subset \mathcal{A}_P$, what might be surprising at first sight. This phenomenon where a set of matrix products is completely covered by that of a prefix will play a prominent role below.

In a natural way, the set $\mathcal{K}$ of index vectors can be given the structure of a directed *tree*, denoted by T: The elements of $\mathcal{K}$ are the nodes, the empty vector $\emptyset$ is the root, and an edge is connecting the *parent* node P with the *child* node C if and only if $C = [P, i]$ for some index $i \in \mathcal{K}_1$.

Definition 3.2. *A node $K \in \mathcal{K}$ is called*

- positive *or* negative *if so is the suffix i when writing $K = [P, i]$.*
- 1-bounded *if $\|\mathcal{A}_K\| := \sup\{\|A\| : A \in \mathcal{A}_K\} \leq 1$.*
- covered *if there exists a prefix $P \in \mathcal{P}(K)$ such that $\mathcal{A}_K \subset \mathcal{A}_P$, and the complementary suffix S is completely positive and not empty.*

Typically, covered nodes appear in the following situation: Let $P = [P', \ell]$ be a negative node, i.e., $\ell < 0$. Then its descendant $K = [P, J_{-\ell}]$ is covered since

$$\mathcal{A}_{[\ell, J_{-\ell}]} = \mathcal{A}_{J_{-\ell}} \cdot \mathcal{A}_\ell = \{A_{J_{-\ell}}^{k+1} : k \in \mathbb{N}_0\} \subset \{A_{J_{-\ell}}^k : k \in \mathbb{N}_0\} = \mathcal{A}_\ell.$$

The example above is constructed in exactly this way.

The following theorem provides a sufficient condition for establishing the JSR of a family of matrices. This condition is based on properties of a *finite* subtree of T and thus can be verified (though not falsified) numerically or analytically in finite time.

Theorem 3.3. *Let $\mathcal{A} = \{A_1, \ldots, A_m\}$. If there exists a generator set $\mathcal{J}$ and a finite subtree $T_* \subset T$ such that*

- *the root $\emptyset$ of T_* has exactly m positive children,*
- *every leaf of T_* is either* 1*-bounded or covered,*
- *every other node of T_* has either exactly m positive children or an arbitrary number of negative children,*

then $\hat{\rho}(\mathcal{A}) = 1$.

Positive children of the root $\emptyset$ are demanded merely to simplify forthcoming arguments. Negative children, though useless in practice, would be equally fine from a theoretical point of view.

The approach suggested in [GP13] assumes irreducibility of the family $\mathcal{A}$ and is successful if and only if, possibly after scaling, this family has a spectral gap at 1, as defined below. The class of problems decidable by Theorem 3.3 is actually larger: First, Theorem 3.5 shows that a tree of the requested form exists always if $\mathcal{A}$ has these properties. Choosing an appropriate norm, this case is in fact trivial because the nodes of the tree can all be chosen to be completely positive, i.e., no infinite sets of matrices have to be employed. Second, an example in Section 5 shows that it is possible to establish the JSR of families which do *not* have a spectral gap at 1. Also irreducibility is not assumed.

The two properties mentioned above, namely irreducibility and the spectral gap at 1, mean the following: The first one concerns a possible reduction of dimensionality. The family $\mathcal{A} = \{A_1, \ldots, A_m\}$ is called *irreducible* if the matrices $A_1, \ldots, A_m$ have no common nontrivial invariant subspace. If $\mathcal{A}$ is reducible, JSR computation can be split into lower-dimensional problems until irreducibility is attained, see [Jun09], Proposition 1.5. The second one concerns separation of spectral radii in the set of matrix products. Given a single generator J, $\rho(A_J) = 1$ immediately implies $\rho(A_I) = 1$ for any index vector I which is a cyclic permutation of a power of J. A spectral gap at 1 means that the SSR of no other product matrix can come close to 1. More precisely, we define:

Definition 3.4. *The matrix family $\mathcal{A}$ has a* spectral gap at 1 *if there is exists a completely positive index vector J with $\rho(A_J) = 1$ such that*

- *there is $q < 1$ such that*
$$\rho(A_I) \leq q \tag{6}$$
for any product A_I, unless $I = \emptyset$ or $I = [S, J^r, P]$ for some $r \in \mathbb{N}_0$ and some partition $[P, S] = J$ of J,

- *the Jordan normal form Λ of A_J is*

$$\Lambda := V^{-1}A_J V = \begin{bmatrix} 1 & 0 \\ 0 & \Lambda_* \end{bmatrix}, \quad \rho(\Lambda_*) < 1. \tag{7}$$

In this case, J is called a dominant generator.

Since $\mathcal{A}$ is finite, the JSR can be characterized also by

$$\hat{\rho}(\mathcal{A}) = \limsup_{k\to\infty} \max_{I\in\mathcal{I}_k} \rho(A_I)^{\frac{1}{k}}.$$

Therefore, a spectral gap at 1 implies $\hat{\rho}(\mathcal{A}) = 1$.

Following [Els95], Lemma 4, an irreducible set $\mathcal{A}$ with $\hat{\rho}(\mathcal{A}) = 1$ is *product bounded*. That is, there exists a constant $c_{\mathcal{A}}$ such that $\|A_I\| < c_{\mathcal{A}}$ for all $I \in \mathcal{I}$. An algorithm for computing an admissible constant can be found in [Pro96]. If, in addition, $\mathcal{A}$ has a dominant generator J, we may scale the matrix V in (7) such that its columns v_j and the columns w_i of $W := V^{-\mathrm{t}}$ satisfy $\|w_1\|_2 = 1$ and $\|v_j\|_2 \le c_{\mathcal{A}}^{-1}, j \ge 2$, where the constant $c_{\mathcal{A}}$ is taken with respect to the Euclidean norm $\|\cdot\|_2$. Now, we define the matrix norm

$$\|A\|_V := \max_j \sum_i |w_i^{\mathrm{t}} A v_j|$$

as the standard 1-norm of the transformed matrix $W^{\mathrm{t}}AV$. At least with respect to this norm, the implication of Theorem 3.3 is in fact an equivalence:

Theorem 3.5. *Let $\mathcal{A}$ be irreducible with spectral gap at 1. Using the norm $\|\cdot\|_V$, there exists a subtree $T_* \subset T$ according to the specifications of Theorem 3.3. Moreover, all nodes in this tree can be requested to be completely positive.*

4. Proofs

4.1. Proof of Theorem 3.3

Below, T_* is assumed to be a fixed subtree of T according to the conditions of the theorem. Dependencies of variables on T_* will not be declared explicitly. The finite set of nodes of T_* is denoted by $\mathcal{K}_*$, and the union of all contained matrix products by $\mathcal{A}_* := \bigcup_{K\in\mathcal{K}_*} \mathcal{A}_K$. The corresponding set of completely positive index vectors is

$$\mathcal{I}_* := \{I \in \mathcal{I} : A_I \in \mathcal{A}_*\}.$$

That is, $\{A_I : I \in \mathcal{I}_*\} = \mathcal{A}_*$ and $\mathcal{K}_* \cap \mathcal{I} \subseteq \mathcal{I}_*$.

The index vectors in $\mathcal{K}$ are ordered completely by setting $K < K'$ for $|K| < |K'|$, and applying lexicographic order for index vectors of equal length. To any completely positive index vector $I \in \mathcal{I} \setminus \{\emptyset\}$, we assign the following two objects:

- The $\mathcal{I}_*$-*maximal prefix*

$$M(I) := \operatorname{argmax}(\mathcal{I}_* \cap \mathcal{P}(I))$$

  of I is defined as the longest prefix of I which is contained in $\mathcal{I}_*$.

- The $\mathcal{K}_*$-*maximal node*

$$N(I) := \max\{K \in \mathcal{K}_* : A_{M(I)} \in \mathcal{A}_K\}$$

  of I is defined as the maximal node in $\mathcal{K}_*$ with the property that the according set $\mathcal{A}_{N(I)}$ contains the product $A_{M(I)}$ corresponding to the $\mathcal{I}_*$-maximal prefix of I.

We note that, in general, $M(I)$ is not a node of T_* while $N(I)$ is by definition. On the other hand, $N(I)$ is generally not a prefix of I. Further, $M(I)$ is completely positive by definition, coding the singleton $\{A_{M(I)}\}$. In contrast, $N(I)$ may have negative entries, then coding a denumerable set of products which contains $A_{M(I)}$. If I is a node of T_*, then $M(I) = I$ if and only if I is completely positive.

Further, for any $I \in \mathcal{I} \setminus \{\emptyset\}$, the length of the $\mathcal{I}_*$-maximal prefix $M(I)$ is at least 1 because the first entry of I is necessarily a child of the root of T_*.

Lemma 4.1. *If $I \notin \mathcal{I}_*$, then its $\mathcal{K}_*$-maximal node $N(I)$ is a 1-bounded leaf of T_*.*

Proof. Let $P := M(I)$ and $K := N(I)$ denote the $\mathcal{I}_*$-maximal prefix and the $\mathcal{K}_*$-maximal node of I, respectively. First, assume that K has positive children. Writing $I = [i_1, \ldots, i_k]$ and $P = [i_1, \ldots, i_\ell]$, we know that $\ell + 1 \leq k$ because $I \notin \mathcal{I}_*$. Since, by assumption, positive children always come as a complete set of m siblings, $K' := [K, i_{\ell+1}] \in \mathcal{K}_*$ is a node of T_*. Consider the prefix $P' := [i_1, \ldots, i_{\ell+1}]$ of I. Then

$$A_{P'} = A_{i_{\ell+1}} \cdot A_P \in \mathcal{A}_{i_{\ell+1}} \cdot \mathcal{A}_K = \mathcal{A}_{K'} \subset \mathcal{A}_*.$$

This implies $P' \in \mathcal{I}_*$, contradicting maximality of the prefix P in $\mathcal{I}_*$.

Second, assume that K has a negative child $K' = [K, j]$ for some $j < 0$. Since $\mathcal{A}_j = \{A^k_{J_{-j}} : k \in \mathbb{N}_0\}$ contains the identity matrix, we have $A_P \in \mathcal{A}_K \subset \mathcal{A}_{K'}$, contradicting maximality of K.

Third, assume that K is a covered node. Thus, by definition, $K = [Q, S]$ for some $S \in \mathcal{I} \setminus \{\emptyset\}$ and $A_P \in \mathcal{A}_K \subseteq \mathcal{A}_Q$. By properties of S, Q has positive children. Again, the argument used in the first part of this proof yields a contradiction to maximality of the prefix P.

The first two cases imply that the node K is one of the leaves of T_*. These are either 1-bounded or covered, the latter being excluded by the third case. □

The $\mathcal{I}_*$*-maximal prefix partition* $P_1, \dots, P_r$ of $I \in \mathcal{I} \setminus \{\emptyset\}$ is characterized by

$$\begin{aligned} I &= [P_1, \dots, P_r] \\ P_\ell &= M([P_\ell, \dots, P_r]), \quad \ell = 1, \dots, r. \end{aligned}$$

Algorithmically, the vectors P_ℓ can be determined by a recursive process, starting from $P_1 = M(I)$: Regard the complementary suffix of I respective to prefix P_ℓ. Then $P_{\ell+1}$ is its $\mathcal{I}_*$-maximal prefix. The algorithm terminates finding a $\mathcal{I}_*$-maximal prefix $P_r = M(P_r)$, which implies $P_r \in \mathcal{I}_*$. The number r cannot exceed $|I|$ because $\mathcal{I}_*$-maximal prefixes of non-empty vectors have length ≥ 1. Lemma 4.1 provides information on the $\mathcal{I}_*$-maximal prefixes $P_1, \dots, P_{r-1}$, while the suffix $S = P_r$ is covered by the following result:

Lemma 4.2. *There exists a monotone increasing polynomial* $p : \mathbb{N}_0 \to \mathbb{R}$*, depending only on* $\mathcal{A}$ *and* T_**, such that*

$$\|A_S\| \leq p(|S|)$$

for any $S \in \mathcal{I}_*$*.*

Proof. Let $\mathcal{B} := \{A_1, \dots, A_m, A_{J_1}, \dots, A_{J_n}\}$. Since $\rho(B) \leq 1$ for all $B \in \mathcal{B}$, the entries of powers B^r grow at most in a polynomial way. That is, there exists a monotone increasing polynomial $q : \mathbb{N}_0 \to [1, \infty)$ with

$$\|B^r\| \leq q(r), \quad B \in \mathcal{B}, \quad r \in \mathbb{N}_0.$$

Let $h := \max\{|K| : K \in \mathcal{K}_*\}$ denote the height of T_*. Then $p := q^h$ is also a monotone increasing polynomial on $\mathbb{N}_0$. For $S \in \mathcal{I}_*$, the product A_S is member of the set $\mathcal{A}_K$ for some $K = [k_1, \dots, k_\ell] \in \mathcal{K}_*$, and thus can be written as

$$A_S = B_\ell^{r_\ell} \cdots B_1^{r_1},$$

where $r_i \in \mathbb{N}_0, B_i \in \mathcal{B}$, and $B_i^{r_i} \in \mathcal{A}_{k_i}$. The exponents are bounded by $r_i \leq |S|$, and the number of factors by $\ell \leq h$. Hence,

$$\|A_S\| \leq \prod_{i=1}^{\ell} \|B_i^{r_i}\| \leq \prod_{i=1}^{\ell} q(r_i) \leq q(|S|)^\ell \leq p(|S|),$$

as stated. ☐

Now, we are prepared to accomplish the proof of Theorem 3.3:

Proof. Let $I \in \mathcal{I}_k$ be any completely positive index vector, and $P_1, \ldots, P_r$ its $\mathcal{I}_*$-maximal prefix partition. Then

$$A_I = A_{P_r} \cdot A_{P_{r-1}} \cdots A_{P_1}.$$

For $i = 1, \ldots, r-1$, $[P_i, \ldots, P_r] \notin \mathcal{I}_*$ and $A_{P_i} \in \mathcal{A}_{N([P_i,\ldots,P_r])}$ by construction. So,

$$\|A_{P_i}\| \leq \|\mathcal{A}_{N([P_i,\ldots,P_r])}\| \leq 1$$

due to Lemma 4.1.

A bound on the norm of P_r is given by Lemma 4.2. By monotonicity of the polynomial p and $|P_r| \leq |I| = k$, we obtain

$$\|A_I\| \leq p(|P_r|) \leq p(k).$$

Hence, by (2),

$$1 = \max\{\rho(A_J) : J \in \mathcal{J}\} \leq \hat{\rho}(\mathcal{A}) \leq \sqrt[k]{p(k)}.$$

As $k \to \infty$, the right hand side converges to 1, thus verifying the claim. ☐

4.2. Proof of Theorem 3.5

Let $I := (i_k)_{k\in\mathbb{N}}$ denote a sequence of positive indices $i_k \in \{1, \ldots, m\}$. Adapting notation in the obvious way, we denote prefixes of I by $I_k := [i_1, \ldots, i_k] \in \mathcal{P}(I)$. Further, $J^\infty := [J, J, \ldots]$ is the sequence obtained by infinite repetition of J. If $\|A_{I_k}\| > 1$ for all $k \in \mathbb{N}_0$, then I is called an *infinite path.* The following Lemma provides information about the structure of such a path:

Lemma 4.3. *Let J be a dominant generator of the irreducible family $\mathcal{A}$ with spectral gap at 1. Then any infinite path has the form $I = [P, J^\infty]$ for some prefix $P \in \mathcal{I}$.*

Proof. Since $\mathcal{A}$ is product bounded, there exists a convergent subsequence $B_\ell := A_{I_{k(\ell)}}, \ell \in \mathbb{N}_0$, with limit $B^* := \lim_\ell B_\ell$. Let $\lambda \in \mathbb{N}$ and $L_{\ell,\lambda} \in \mathcal{I}$ such that $\left[I_{k(\ell)}, L_{\ell,\lambda}\right] = I_{k(\ell+\lambda)}$ for $\ell \in \mathbb{N}_0$. Then

$$B_{\ell+\lambda} = C_{\ell,\lambda} B_\ell, \quad C_{\ell,\lambda} := A_{L_{\ell,\lambda}}, \quad \ell \in \mathbb{N}_0.$$

The right hand side of

$$\|(C_{\ell,\lambda} - \mathrm{Id})B^*\| = \|C_{\ell,\lambda}(B^* - B_\ell) + B_{\ell+\lambda} - B^*\| \leq c_{\mathcal{A}}\|B^* - B_\ell\| + \|B^* - B_{\ell+\lambda}\|$$

tends to 0. Thus,

$$\lim_{\ell\to\infty} C_{\ell,\lambda}B^* = B^*.$$

By assumption, $\|B_\ell\| > 1$ for all ℓ and hence $\|B^*\| \neq 0$. So, recalling (6), the upper equality shows that there exists $\ell_0 \in \mathbb{N}$ such that $\rho(C_{\ell,\lambda}) > q$ for all $\ell \geq \ell_0$. Since J is dominant, we obtain by definition

$$L_{\ell,\lambda} = S_{\ell,\lambda} J^{r_{\ell,\lambda}} P_{\ell,\lambda}, \quad \ell \geq \ell_0,$$

for partitions $[P_{\ell,\lambda}, S_{\ell,\lambda}] = J$. Substituting this representation into $[L_{\ell,1}, L_{\ell+1,1}] = L_{\ell,2}$ yields

$$S_{\ell,1} J^{r_{\ell,1}} P_{\ell,1} S_{\ell+1,1} J^{r_{\ell+1,1}} P_{\ell+1,1} = S_{\ell,2} J^{r_{\ell,2}} P_{\ell,2}.$$

This is possible only if $P_{\ell,1}S_{\ell+1,1} = J$, implying $P_{\ell,1} = P_{\ell+1,1}$ and $S_{\ell,1} = S_{\ell+1,1}$. That is, $P_{\ell,1} = P_{\ell_0,1}$ and $S_{\ell,1} = S_{\ell_0,1}$ for all $\ell \geq \ell_0$. We recall $B_0 := A_{I_{k(0)}}$ and abbreviate $r_i := r_{\ell_0+i,1}$, $\tilde{P} := P_{\ell_0,1}$, $\tilde{S} := S_{\ell_0,1}$ to find

$$I = [I_{k(0)}, \tilde{S}, J^{r_0}, \tilde{P}, \tilde{S}, J^{r_1}, \tilde{P}, \dots] = [I_{k(0)}, \tilde{S}, J^\infty],$$

and the claim follows with $P := [I_{k(0)}, \tilde{S}]$. □

While the lemma above makes no assumptions concerning the underlying norm, the next one shows that the V-norm $\|\cdot\|_V$ is special.

Lemma 4.4. *If $\mathcal{A}$ is irreducible with spectral gap at 1, then there is no infinite path with respect to the V-norm.*

Proof. Assume that there exists an infinite path. According to the previous lemma, it is given by $I = [P, J^\infty]$ with J being a dominant generator. The corresponding limit of matrix products $S := \lim_k A_{[P,J^k]}$ exists because $\Lambda^\infty := \lim_k \Lambda^k = \operatorname{diag}([1, 0, \dots, 0])$ exists. Since $\|A_{J^k}\|_V = \|\Lambda^k\|_1 = 1$ for k sufficiently large, P cannot be empty. In this case, we know that $\|A_{[P,J^k]}\|_V > 1$ and $\rho(A_{[P,J^k]}) \leq q < 1$ for all $k \in \mathbb{N}$. Hence, $\|S\|_V \geq 1$ and $\rho(S) \leq q$. With $T := A_P$, the matrix $\tilde{S} := W^{\mathrm{t}} S V$ is given by

$$\tilde{S} = \Lambda^\infty W^{\mathrm{t}} T V = \begin{bmatrix} w_1^{\mathrm{t}} T v_1 & w_1^{\mathrm{t}} T v_2 & \cdots & w_1^{\mathrm{t}} T v_d \\ 0 & 0 & \cdots & 0 \\ \vdots & \vdots & & \vdots \\ 0 & 0 & \cdots & 0 \end{bmatrix}.$$

Since S and $\tilde{S}$ are similar, we have $|w_1^{\mathrm{t}} T v_1| = \rho(\tilde{S}) = \rho(S) < 1$. Further, recalling $\|w\|_2 = 1$ and $\|v_j\|_2 \leq c_{\mathcal{A}}^{-1}$, we find $|w_1^{\mathrm{t}} T v_j| \leq \|w_1^{\mathrm{t}}\|_2 \, \|T\|_2 \, \|v_j\|_2 < 1$ for $j \geq 2$. Thus, we obtain the contradiction $1 \leq \|S\|_V = \|\tilde{S}\|_1 < 1$. □

The lemma enables us to prove Theorem 3.5:

Proof. Let T_* be the largest subtree of T with the following properties: $\emptyset$ is contained as a node, all nodes are completely positive and each 1-bounded node is a leaf. Assuming that the set $\mathcal{K}_*$ of nodes is infinite, we define $\mathcal{K}' \subset \mathcal{K}_*$ as the set of nodes which are prefix of infinitely many other nodes in $\mathcal{K}_*$. Then $\mathcal{K}'$ is not empty because the root $\emptyset$ belongs to it. Further, if $K \in \mathcal{K}'$, then there exists at least one child $[K, k] \in \mathcal{K}'$. That is, the recursion

$$i_k := \min\{i \in \mathcal{I}_1 : [i_1, \ldots, i_{k-1}, i] \in \mathcal{K}'\}$$

defines an infinite path $I = (i_k)_{k \in \mathbb{N}}$, contradicting Lemma 4.4. □

5. Examples

In this section, we illustrate some aspects of our method by considering a simple model problem. Let $\mathcal{A} = \{A_1, A_2\}$ with

$$A_1 = \begin{bmatrix} \frac{10}{9} & \frac{1}{3} \\ -\frac{1}{3} & 0 \end{bmatrix}, \quad A_2 = \begin{bmatrix} 0 & \frac{1}{5}\sqrt{1-\varepsilon} \\ -\frac{1}{5}\sqrt{1-\varepsilon} & \frac{26}{25} - \varepsilon \end{bmatrix}.$$

Then $\rho(A_1) = 1$ and $\rho(A_2) = 1 - \varepsilon$. All trees will be constructed with respect to the maximum absolute row sum norm $\|\cdot\|_\infty$.

For $\varepsilon = \frac{1}{8}$, the family $\mathcal{A}$ has a dominant generator $J = [1]$, leading to a finite tree that satisfies the conditions of Theorem 3.3, i.e., $\hat{\rho}(\mathcal{A}) = 1$. The tree is visualized in Figure 1 *(left)*.

For $\varepsilon = 0$, both matrices A_1, A_2 have spectral radius 1. Hence, $\mathcal{A}$ does *not* have a spectral gap at 1. While this case is explicitly excluded in [GP13], Figure 1 *(right)* shows the corresponding tree according to Theorem 3.3 with generators $J_1 = [1], J_2 = [2]$, thus proving $\hat{\rho}(\mathcal{A}) = 1$ also in this case.

The asserted benefits of generators J with $\rho(A_J) < 1$ become apparent when considering small values of ε. For $\varepsilon = 0.01$, the spectral radius of A_2 is 0.99. In principle, it is sufficient to use only the generator $J_1 = [1]$, see Figure 2 *(left)*. However, Figure 2 *(right)* shows that the depth of the resulting tree is reduced significantly when using the additional generator $J_2 = [2]$. This is because the slow decay of norms of matrix powers $\|A_2^k\|$ is subsumed in the single negative node marked by a triangle. For smaller values of ε, the effect becomes even more drastic. For instance, $\varepsilon = 0.001$ leads to depth 224 when using a single generator, while the two generators still yield the same trim pattern shown in Figure 2 *(right)*.

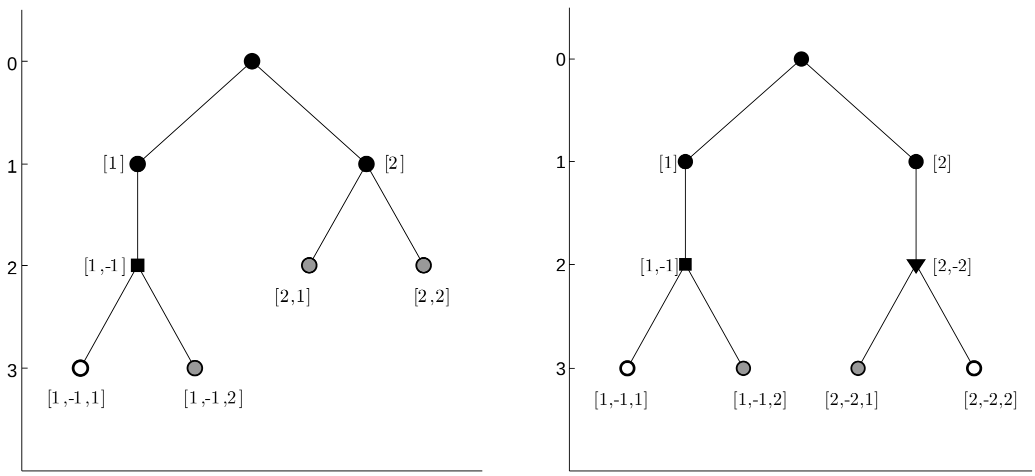

Figure 1: Tree for $\varepsilon = 1/8$ *(left)* and $\varepsilon = 0$ *(right)* according to Theorem 3.3. Colors and shapes of markers indicate properties of nodes: gray $\triangleq$ 1-bounded, white $\triangleq$ covered, square $\triangleq$ negative child wrt. generator $J_1 = [1]$, triangle $\triangleq$ negative child wrt. generator $J_2 = [2]$, black $\triangleq$ other.

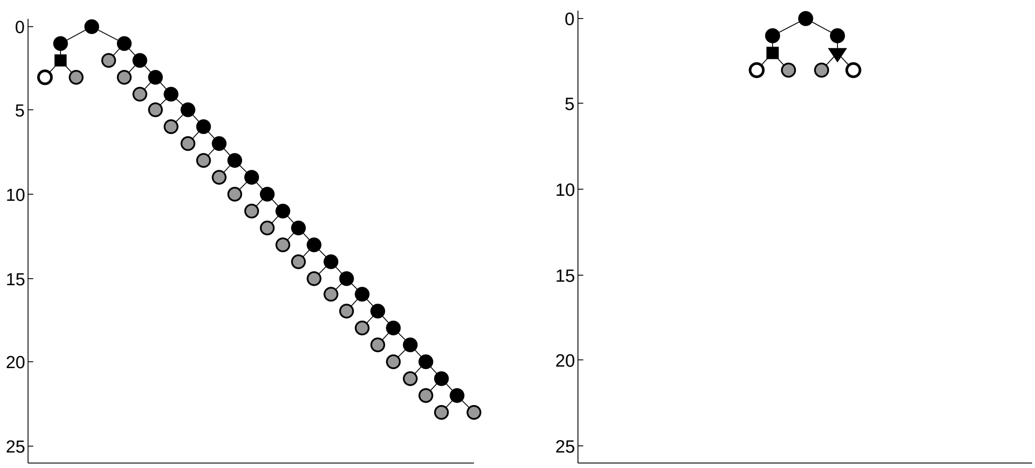

Figure 2: Tree for $\varepsilon = 0.01$ using a single generator $J_1 = [1]$ *(left)* and two generators $J_1 = [1], J_2 = [2]$ *(right)*.

6. Conclusion

We have presented an alternative to known algorithms for rigorous certification of the JSR of a finite set of matrices. It is based on the observation that existence of a certain finite tree with nodes representing sets of matrix products implies validity of the FP for a given set of matrices. Such trees can be searched automatically by computer programs by a variant on depth-first search. Details of such an algorithm will be presented in a forthcoming report. We have to acknowledge that the run-time of our current implementation cannot compete with the impressive results of Protasov, Guglielmi and others. However, the main advantage of our method is the fact that it can cope with situations where more than one generator is satisfying the FP. Even reducibility of the set of matrices is not mandatory, though favorable in applications. Problems revealing a spectral gap at 1 are guaranteed to be decidable by our algorithm when choosing the norm appropriately.

References

[BM02] T. Bousch and J. Mairesse. Asymptotic height optimization for topical IFS, tetris heaps, and the finiteness conjecture. *Journal of the American Mathematical Society*, 15(1):77–111, 2002.

[BN05] V.D. Blondel and Y. Nesterov. Computationally efficient approximations of the joint spectral radius. *SIAM J. Matrix Anal. Appl.*, 27(1):256–272, 2005.

[BT00] V.D. Blondel and J.N. Tsitsiklis. The boundedness of all products of a pair of matrices is undecidable. *Systems and Control Letters*, 41(2):135–140, 2000.

[BTV03] V.D. Blondel, J. Theys, and A.A. Vladimirov. An elementary counterexample to the finiteness conjecture. *SIAM Journal on Matrix Analysis and Applications*, 24(4):963–970, 2003.

[CGSCZ10] A. Cicone, N. Guglielmi, S. Serra-Capizzano, and M. Zennaro. Finiteness property of pairs of 2×2 sign-matrices via real extremal polytope norms. *Linear Algebra and its Applications*, 432(2-3):796 – 816, 2010.

[DL92] I. Daubechies and J.C. Lagarias. Sets of matrices all infinite products of which converge. *Linear Algebra and its Applications*, 161:227–263, 1992.

[Els95] L. Elsner. The generalized spectral-radius theorem: An analytic-geometric proof. *Linear Algebra and its Applications*, 220:151 – 159, 1995.

[GP13] N. Guglielmi and V.Y. Protasov. Exact computation of joint spectral characteristics of linear operators. *Foundations of Computational Mathematics*, 13(1):37–97, 2013.

[Gri96] G. Gripenberg. Computing the joint spectral radius. *Linear Algebra and its Applications*, 234:43–60, 1996.

[GZ08] N. Guglielmi and M. Zennaro. An algorithm for finding extremal polytope norms of matrix families. *Linear Algebra and its Applications*, 428(10):2265 – 2282, 2008.

[GZ09] N. Guglielmi and M. Zennaro. Finding extremal complex polytope norms for families of real matrices. *SIAM J. Matrix Anal. Appl.*, 31(2):602–620, 2009.

[HMR09] J. Hechler, B. Mößner, and U. Reif. C^1-Continuity of the generalized four-point scheme. *Linear Algebra and its applications*, 430(11-12):3019–3029, 2009.

[HMST11] K.G. Hare, I.D. Morris, N. Sidorov, and J. Theys. An explicit counterexample to the Lagarias-Wang finiteness conjecture. *Advances in Mathematics*, 226(6):4667–4701, 2011.

[Jun09] R.M. Jungers. *The joint spectral radius, theory and applications.* Springer, 2009.

[Koz05] V. Kozyakin. A dynamical systems construction of a counterexample to the finiteness conjecture. In *Proceedings of the 44th IEEE Conference on Decision and Control and ECC 2005*, pages 2338–2343, 2005.

[Mae00] M. Maesumi. Joint spectral radius and Hölder regularity of wavelets. *Computers and Mathematics with Applications*, 40(1):145 – 155, 2000.

[Mae08] M. Maesumi. Optimal norms and the computation of joint spectral radius of matrices. *Linear Algebra and its Applications*, 428(10):2324 – 2338, 2008.

[Pro96] V.Y. Protasov. The joint spectral radius and invariant sets of the several linear operators. *Fundamentalnaya i prikladnaya matematika*, 2(1):205–231, 1996.

[RS60] G.C. Rota and W.G. Strang. A note on the joint spectral radius. *Indag. Math.*, 22(4):379–381, 1960.