Joint Spectral Radius Theory for Automated Complexity Analysis of Rewrite Systems*

Aart Middeldorp[1], Georg Moser[1], Friedrich Neurauter[1], Johannes Waldmann[2], and Harald Zankl[1]

[1] Institute of Computer Science, University of Innsbruck, Austria
[2] Fakultät Informatik, Mathematik und Naturwissenschaften, Hochschule für Technik, Wirtschaft und Kultur Leipzig, Germany

Abstract. Matrix interpretations can be used to bound the derivational complexity of term rewrite systems. In particular, triangular matrix interpretations over the natural numbers are known to induce polynomial upper bounds on the derivational complexity of (compatible) rewrite systems. Recently two different improvements were proposed, based on the theory of weighted automata and linear algebra. In this paper we strengthen and unify these improvements by using joint spectral radius theory.

Keywords: derivational complexity, matrix interpretations, weighted automata, joint spectral radius

1 Introduction

This paper is concerned with automated complexity analysis of term rewrite systems. Given a terminating rewrite system, the aim is to obtain information about the maximal length of rewrite sequences in terms of the size of the initial term. This is known as derivational complexity. Developing methods for bounding the derivational complexity of rewrite systems has become an active and competitive[3] research area in the past few years (e.g. [6, 11–15, 19, 21]).

Matrix interpretations [4] are a popular method for automatically proving termination of rewrite systems. They can readily be used to establish upper bounds on the derivational complexity of compatible rewrite systems. However, in general, matrix interpretations induce exponential (rather than polynomial) upper bounds. In order to obtain polynomial upper bounds, the matrices used in a matrix interpretation must satisfy certain (additional) restrictions, the study of which is the central concern of [14, 15, 19].

So what are the conditions for polynomial boundedness of a matrix interpretation? In the literature, two different approaches have emerged. On the one hand, there is the automata-based approach of [19], where matrices are viewed as

* This research is supported by FWF (Austrian Science Fund) project P20133. Friedrich Neurauter is supported by a grant of the University of Innsbruck.

[3] `http://www.termination-portal.org/wiki/Complexity`

weighted (word) automata computing a weight function, which is required to be polynomially bounded. The result is a complete characterization (i.e., necessary and sufficient conditions) of polynomially bounded matrix interpretations over $\mathbb{N}$. On the other hand, there is the algebraic approach pursued in [15] (originating from [14]) that can handle matrix interpretations over $\mathbb{N}$, $\mathbb{Q}$, and $\mathbb{R}$ but only provides sufficient conditions for polynomial boundedness. In what follows, we shall see, however, that these two seemingly different approaches can be unified and strengthened with the help of joint spectral radius theory [9, 10], a branch of mathematics dedicated to studying the growth rate of products of matrices taken from a set.

The remainder of this paper is organized as follows. In the next section we recall preliminaries from linear algebra and term rewriting. We give a brief account of the matrix method for proving termination of rewrite systems. In Section 3 we introduce the algebraic approach for characterizing the polynomial growth of matrix interpretations. We improve upon the results of [15] by considering the minimal polynomial associated with the component-wise maximum matrix of the interpretation. We further show that the joint spectral radius of the matrices in the interpretation provides a better characterization of polynomial growth and provide conditions for the decidability of the latter. Section 4 is devoted to automata-based methods for characterizing the polynomial growth of matrix interpretations. We revisit the characterization results of [19] and provide precise complexity statements. In Section 5 we unify the two approaches and show that, in theory at least, the joint spectral radius theory approach subsumes the automata-based approach. Automation of the results presented in earlier sections is the topic of Section 6. To this end we extend the results from [15, 19]. We also provide experimental results. We conclude with suggestions for future research in Section 7.

2 Preliminaries

As usual, we denote by $\mathbb{N}$, $\mathbb{Z}$, $\mathbb{Q}$ and $\mathbb{R}$ the sets of natural, integer, rational and real numbers. Given $D \in \{\mathbb{N}, \mathbb{Z}, \mathbb{Q}, \mathbb{R}\}$ and $m \in D$, $>_D$ denotes the standard order of the respective domain and D_m abbreviates $\{x \in D \mid x \geqslant m\}$.

Linear Algebra: Let R be a ring (e.g., $\mathbb{Z}$, $\mathbb{Q}$, $\mathbb{R}$). The ring of all *n-dimensional square matrices* over R is denoted by $R^{n \times n}$ and the *polynomial ring* in n *indeterminates* $x_1, \ldots, x_n$ by $R[x_1, \ldots, x_n]$. In the special case $n = 1$, any polynomial $p \in R[x]$ can be written as $p(x) = \sum_{k=0}^{d} a_k x^k$ for some $d \in \mathbb{N}$. For the largest k such that $a_k \neq 0$, we call $a_k x^k$ the *leading term* of p, a_k its *leading coefficient* and k its *degree*. The polynomial p is said to be *monic* if its leading coefficient is one. It is said to be *linear*, *quadratic*, *cubic* if its degree is one, two, three.

In case R is equipped with a partial order $\geqslant$, the component-wise extension of this order to $R^{n \times n}$ is also denoted as $\geqslant$. We say that a matrix A is *non-negative* if $A \geqslant 0$ and denote the set of all non-negative n-dimensional square matrices of

$\mathbb{Z}^{n\times n}$ by $\mathbb{N}^{n\times n}$. The $n \times n$ *identity matrix* is denoted by I_n and the $n \times n$ *zero matrix* is denoted by 0_n. We simply write I and 0 if n is clear from the context.

The *characteristic polynomial* $\chi_A(\lambda)$ of a square matrix $A \in R^{n\times n}$ is defined as $\det(\lambda I_n - A)$, where det denotes the determinant. It is monic and its degree is n. The equation $\chi_A(\lambda) = 0$ is called the *characteristic equation* of A. The solutions of this equation, i.e., the *roots* of $\chi_A(\lambda)$, are precisely the *eigenvalues* of A, and the *spectral radius* $\rho(A)$ of A is the maximum of the absolute values of all eigenvalues. A non-zero vector x is an *eigenvector* of A if $Ax = \lambda x$ for some eigenvalue λ of A. We say that a polynomial $p \in R[x]$ *annihilates* A if $p(A) = 0$. The Cayley-Hamilton theorem [16] states that A satisfies its own characteristic equation, i.e., χ_A annihilates A. The unique monic polynomial of minimum degree that annihilates A is called the *minimal polynomial* $m_A(x)$ of A. The *multiplicity* of a root λ of $p \in R[x]$ is denoted by $\#p(\lambda)$.

With any matrix $A \in R^{n\times n}$ we associate a directed (weighted) graph $G(A)$ on n vertices numbered from 1 to n such that there is a directed edge (of weight A_{ij}) in $G(A)$ from i to j if and only if $A_{ij} \neq 0$. In this situation, A is said to be the *adjacency matrix* of the graph $G(A)$. The *weight of a path* in $G(A)$ is the product of the weights of its edges.

With a finite set of matrices $S \subseteq R^{n\times n}$ we associate the directed (weighted) graph $G(S) := G(M)$, where M denotes the component-wise maximum of the matrices in S, i.e., $M_{ij} = \max\{A_{ij} \mid A \in S \text{ and } 1 \leqslant i,j \leqslant n\}$. Following [10], we define a directed graph $G^k(S)$ for $k \geqslant 2$ on n^k vertices representing ordered tuples of vertices of $G(S)$, such that there is an edge from vertex $(i_1,\ldots,i_k)$ to $(j_1,\ldots,j_k)$ if and only if there is a matrix $A \in S$ with $A_{i_\ell j_\ell} > 0$ for all $\ell = 1,\ldots,k$. This is akin to the k-fold Kronecker product of the matrix A.

For functions $f, g\colon \mathbb{N} \to \mathbb{N}$ we write $f(n) = O(g(n))$ if there are constants $M, N \in \mathbb{N}$ such that $f(n) \leqslant M \cdot g(n) + N$ for all $n \in \mathbb{N}$. Furthermore, $f(n) = \Theta(g(n))$ if $f(n) = O(g(n))$ and $g(n) = O(f(n))$.

Rewriting: We assume familiarity with the basics of term rewriting [1,18]. Let $\mathcal{V}$ denote a countably infinite set of variables and $\mathcal{F}$ a fixed-arity signature. The set of *terms* over $\mathcal{F}$ and $\mathcal{V}$ is denoted by $\mathcal{T}(\mathcal{F},\mathcal{V})$. The *size* $|t|$ of a term t is defined as the number of function symbols and variables occurring in it. The set of *positions* $\mathcal{P}\mathsf{os}(t)$ of a term t is defined as usual. Positions are denoted as sequences of natural numbers. A *term rewrite system* (*TRS* for short) $\mathcal{R}$ over $\mathcal{T}(\mathcal{F},\mathcal{V})$ is a *finite* set of rewrite rules $\ell \to r$ such that $\ell \notin \mathcal{V}$ and $\mathcal{V}\mathsf{ar}(\ell) \supseteq \mathcal{V}\mathsf{ar}(r)$. The smallest rewrite relation that contains $\mathcal{R}$ is denoted by $\to_{\mathcal{R}}$. The transitive (and reflexive) closure of $\to_{\mathcal{R}}$ is denoted by $\to^+_{\mathcal{R}}$ ($\to^*_{\mathcal{R}}$). Let s and t be terms. If exactly n steps are performed to rewrite s to t, we write $s \to^n t$. The *derivation height* of a term s with respect to a well-founded and finitely branching relation $\to$ is defined as $\mathsf{dh}(s,\to) = \max\{n \mid s \to^n t \text{ for some term } t\}$. The *derivational complexity function* of $\mathcal{R}$ is defined as: $\mathsf{dc}_{\mathcal{R}}(k) = \max\{\mathsf{dh}(t,\to_{\mathcal{R}}) \mid |t| \leqslant k\}$.

Matrix Interpretations: An $\mathcal{F}$-*algebra* $\mathcal{A}$ consists of a carrier set A and a collection of interpretations $f_{\mathcal{A}}\colon A^k \to A$ for each k-ary function symbol in $\mathcal{F}$.

By $[\alpha]_{\mathcal{A}}(\cdot)$ we denote the usual evaluation function of $\mathcal{A}$ according to an assignment α which maps variables to values in A. An $\mathcal{F}$-algebra together with a well-founded order $>$ on A is called a *monotone algebra* if every function symbol $f \in \mathcal{F}$ is monotone with respect to $>$ in all arguments. Any monotone algebra $(\mathcal{A}, >)$ (or just $\mathcal{A}$ if $>$ is clear from the context) induces a well-founded order on terms: $s >_{\mathcal{A}} t$ if and only if $[\alpha]_{\mathcal{A}}(s) > [\alpha]_{\mathcal{A}}(t)$ for all assignments α. A TRS $\mathcal{R}$ and a monotone algebra $\mathcal{A}$ are compatible if $\ell >_{\mathcal{A}} r$ for all $\ell \to r \in \mathcal{R}$.

For matrix interpretations, we fix a dimension $n \in \mathbb{N} \setminus \{0\}$ and use the set $\mathbb{R}_0^n$ as the carrier of an algebra $\mathcal{M}$, together with the order $>_\delta$ on $\mathbb{R}_0^n$ defined as $(x_1, x_2, \ldots, x_n)^{\mathrm{T}} >_\delta (y_1, y_2, \ldots, y_n)^{\mathrm{T}}$ if $x_1 >_{\mathbb{R},\delta} y_1$ and $x_i \geqslant_{\mathbb{R}} y_i$ for $2 \leqslant i \leqslant n$. Here $x >_{\mathbb{R},\delta} y$ if and only if $x \geqslant_{\mathbb{R}} y + \delta$. Each k-ary function symbol f is interpreted as a linear function of the following shape: $f_{\mathcal{M}}(\boldsymbol{v}_1, \ldots, \boldsymbol{v}_k) = F_1\boldsymbol{v}_1 + \cdots + F_k\boldsymbol{v}_k + \boldsymbol{f}$ where $\boldsymbol{v}_1, \ldots, \boldsymbol{v}_k$ are (column) vectors of variables, $F_1, \ldots, F_k \in \mathbb{R}_0^{n \times n}$ and $\boldsymbol{f}$ is a vector in $\mathbb{R}_0^n$. The $F_1, \ldots, F_k$ are called the *matrices* of the interpretation $f_{\mathcal{M}}$, while $\boldsymbol{f}$ is called the *absolute vector* of $f_{\mathcal{M}}$. We write $\mathsf{abs}(f)$ for the $\boldsymbol{f}$. To ensure monotonicity, it suffices that the top left entry $(F_i)_{11}$ of every matrix F_i is at least one. Then it is easy to see that $(\mathcal{M}, >_\delta)$ forms a monotone algebra for any $\delta > 0$. We obtain matrix interpretations over $\mathbb{Q}$ by restricting to the carrier $\mathbb{Q}_0^n$. Similarly, matrix interpretations over $\mathbb{N}$ operate on the carrier $\mathbb{N}^n$ and use $\delta = 1$.

Let α_0 denote the assignment that maps every variable to $\mathbf{0}$. Let t be a term. In the following we abbreviate $[\alpha_0]_{\mathcal{M}}(t)$ to $[t]_{\mathcal{M}}$ (or $[t]$ if $\mathcal{M}$ can be inferred from the context) and we write $[t]_j$ $(1 \leqslant j \leqslant n)$ for the j-th element of $[t]$.

Let $\mathcal{M}$ be a matrix interpretation of dimension n. Let S be the set of matrices occurring in $\mathcal{M}$ and let $S^k = \{A_1 \cdots A_k \mid A_i \in S, 1 \leqslant i \leqslant k\}$ be the set of all products of length k of matrices taken from S. Here S^0 consists of the identity matrix. Further, S^* denotes the (matrix) monoid generated by S, i.e., $S^* = \bigcup_{k=0}^{\infty} S^k$.

3 Algebraic Methods for Bounding Polynomial Growth

In this section we study an algebraic approach to characterize polynomial growth of matrix interpretations (over $\mathbb{N}$, $\mathbb{Q}$, and $\mathbb{R}$). We employ the following definition implicitly used in [14, 15].

Definition 1. *Let $\mathcal{M}$ be a matrix interpretation. We say that $\mathcal{M}$ is* polynomially bounded (with degree d) *if the growth of the entries of all matrix products in S^* is polynomial (with degree d) in the length of such products.*

The relationship between polynomially bounded matrix interpretations and the derivational complexity of compatible TRSs is as follows (cf. [15]).

Lemma 2. *Let $\mathcal{R}$ be a TRS and $\mathcal{M}$ a compatible matrix interpretation. If $\mathcal{M}$ is polynomially bounded with degree d then $\mathsf{dc}_{\mathcal{R}}(k) = O(k^{d+1})$.* □

3.1 Spectral Radius

In this subsection we over-approximate the growth of entries of matrix products of the form $A_1 \cdots A_k \in S^k$ by M^k where $M_{ij} = \max\{A_{ij} \mid A \in S\}$ for all $1 \leqslant i, j \leqslant n$. Then, by non-negativity of the matrices in S, we have

$$(A_1 \cdots A_k)_{ij} \leqslant (M^k)_{ij} \quad \text{for all } 1 \leqslant i, j \leqslant n \tag{1}$$

Thus, polynomial boundedness of the entries of $A_1 \cdots A_k$ follows from polynomial boundedness of the entries of M^k. In [15], the latter is completely characterized by the spectral radius $\rho(M)$ of M being at most one. Here we build on this result (and its proof) but establish a lemma that allows to obtain a tight bound for the degree of polynomial growth of the entries of M^k.

Lemma 3. *Let $M \in \mathbb{R}_0^{n \times n}$ and let $p \in \mathbb{R}[x]$ be a monic polynomial that annihilates M. Then $\rho(M) \leqslant 1$ if and only if all entries of M^k ($k \in \mathbb{N}$) are asymptotically bounded by a polynomial in k of degree d, where $d := \max_\lambda(0, \#p(\lambda) - 1)$ and λ are the roots of p with absolute value exactly one and multiplicity $\#p(\lambda)$.*

Proof. Straightforward adaptation of the proof of [15, Lemma 4]. □

Based on Lemma 3 (with $p(x) = \chi_M(x)$), one obtains the following theorem concerning complexity analysis via matrix interpretations.

Theorem 4 ([15, Theorem 6]**).** *Let $\mathcal{R}$ be a TRS and $\mathcal{M}$ a compatible matrix interpretation of dimension n. Further, let M denote the component-wise maximum of all matrices occurring in $\mathcal{M}$. If the spectral radius of M is at most one, then $\mathsf{dc}_\mathcal{R}(k) = O(k^{d+1})$, where $d := \max_\lambda(0, \#\chi_M(\lambda) - 1)$ and λ are the eigenvalues of M with absolute value exactly one.* □

Obviously, the set of (monic) polynomials that annihilate a matrix M is infinite. However, from linear algebra we know that this set is generated by a unique monic polynomial of minimum degree that annihilates M, namely, the *minimal polynomial* $m_M(x)$ of M. That is, if $p(x)$ is any polynomial such that $p(M) = 0$ then $m_M(x)$ divides $p(x)$. In particular, $m_M(x)$ divides the characteristic polynomial $\chi_M(x)$. Moreover, $m_M(\lambda) = 0$ if and only if λ is an eigenvalue of M, so every root of $m_M(x)$ is a root of $\chi_M(x)$. However, in case $m_M(x) \neq \chi_M(x)$, the multiplicity of a root in $m_M(x)$ may be lower than its multiplicity in $\chi_M(x)$ (cf. [8]). In light of these facts, we conclude that in Lemma 3 one should use the minimal polynomial $m_M(x)$ rather than the characteristic polynomial $\chi_M(x)$. Then the corresponding analogon of Theorem 4 is as follows.

Theorem 5. *Let $\mathcal{R}$ be a TRS and $\mathcal{M}$ a compatible matrix interpretation of dimension n. Further, let M denote the component-wise maximum of all matrices occurring in $\mathcal{M}$. If the spectral radius of M is at most one then $\mathsf{dc}_\mathcal{R}(k) = O(k^{d+1})$, where $d := \max_\lambda(0, \#m_M(\lambda) - 1)$ and λ are the eigenvalues of M with absolute value exactly one.* □

Next we illustrate the usefulness of Theorem 5 on an example.

Example 6. Consider the TRS $\mathcal{R}$ consisting of the following two rewrite rules:[4]

$$\begin{aligned} \mathsf{h}(x, \mathsf{c}(y, z)) &\to \mathsf{h}(\mathsf{c}(\mathsf{s}(y), x), z) \\ \mathsf{h}(\mathsf{c}(\mathsf{s}(x), \mathsf{c}(\mathsf{s}(\mathsf{0}), y)), z) &\to \mathsf{h}(y, \mathsf{c}(\mathsf{s}(\mathsf{0}), \mathsf{c}(x, z))) \end{aligned} \qquad M = \begin{pmatrix} 1 & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 \\ 0 & 1 & 0 & 1 \end{pmatrix}$$

There is a compatible matrix interpretation such that the component-wise maximum matrix M (given above) has characteristic polynomial $\chi_M(x) = x(x-1)^3$ and minimal polynomial $m_M(x) = x(x-1)^2$. Thus, the upper bound for the derivational complexity of $\mathcal{R}$ derived from Theorem 4 is cubic, whereas the bound obtained by Theorem 5 is quadratic.

Next we present an example that demonstrates the conceptual limitations arising from the over-approximation of matrix products by the powers of the corresponding maximum matrix.

Example 7. Consider the TRS $\mathcal{R}$ consisting of the rules

$$\mathsf{f}(\mathsf{f}(x)) \to \mathsf{f}(\mathsf{g}(\mathsf{f}(x))) \qquad \mathsf{g}(\mathsf{g}(x)) \to x \qquad \mathsf{b}(x) \to x$$

Lemma 8. *There is a matrix interpretation compatible with $\mathcal{R}$ that is polynomially bounded (in the sense of Definition 1), but there is no matrix interpretation compatible with $\mathcal{R}$ where all entries in the component-wise maximum matrix are polynomially bounded.*

Proof. For the first item consider the following matrix interpretation $\mathcal{M}$ establishing linear derivational complexity of $\mathcal{R}$ (cf. Theorem 16) from Example 7:

$$\mathsf{f}_{\mathcal{M}}(\boldsymbol{x}) = \begin{pmatrix} 1 & 1 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix} \boldsymbol{x} + \begin{pmatrix} 0 \\ 4 \\ 0 \end{pmatrix} \quad \mathsf{g}_{\mathcal{M}}(\boldsymbol{x}) = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix} \boldsymbol{x} + \begin{pmatrix} 1 \\ 0 \\ 3 \end{pmatrix} \quad \mathsf{b}_{\mathcal{M}}(\boldsymbol{x}) = \boldsymbol{x} + \begin{pmatrix} 4 \\ 4 \\ 0 \end{pmatrix}$$

For the second item assume a compatible matrix interpretation $\mathcal{M}$ of dimension n with $\mathsf{b}_{\mathcal{M}}(\boldsymbol{x}) = B\boldsymbol{x} + b$, $\mathsf{f}_{\mathcal{M}}(\boldsymbol{x}) = F\boldsymbol{x} + f$ and $\mathsf{g}_{\mathcal{M}}(\boldsymbol{x}) = G\boldsymbol{x} + g$. Assume $G \geqslant I_n$. To orient the first rule, the constraint $Ff > FGf$ must be satisfied, but by (weak) monotonicity of matrix multiplication we obtain the contradiction $Ff > FGf \geqslant FI_nf = Ff$ and hence $G \ngeqslant I_n$, i.e., there exists an index l such that $G_{ll} < 1$. Since $GG \geqslant I_n$ is needed to orient the second rule, we have $\sum_j G_{lj}G_{jl} \geqslant 1$ and consequently $\sum_{j \neq l} G_{lj}G_{jl} > 0$. The third rule demands $B \geqslant I_n$ and for the maximum matrix M we have $M \geqslant \max(I_n, G)$. From $(I_n)_{ll} \geqslant 1$ and $\sum_{j \neq l} G_{lj}G_{jl} > 0$ we conclude $(M^2)_{ll} > 1$, which gives rise to exponential growth of $(M^k)_{ll}$ since all entries in M are non-negative. □

[4] TPDB problem TRS/Endrullis_06/direct.xml

3.2 Joint Spectral Radius

Instead of using a single maximum matrix to over-approximate the growth of finite matrix products taken from a set of matrices S, in this subsection we provide a concise analysis using joint spectral radius theory [9, 10]. In particular, we shall see that the joint spectral radius of S completely characterizes polynomial growth of all such products, just like the spectral radius of a single matrix characterizes polynomial growth of the powers of this matrix (cf. Lemma 3).

Definition 9. *Let $S \subseteq \mathbb{R}^{n\times n}$ be a finite set of real square matrices, and let $\|\cdot\|$ denote a matrix norm. The growth function* growth *associated with S is defined as follows:*

$$\mathsf{growth}_S(k, \|\cdot\|) := \max\{\, \|A_1 \cdots A_k\| \mid A_i \in S, 1 \leqslant i \leqslant k \,\}$$

The asymptotic behaviour of $\mathsf{growth}_S(k, \|\cdot\|)$ can be characterized by the joint spectral radius of S.

Definition 10. *Let $S \subseteq \mathbb{R}^{n\times n}$ be finite, and let $\|\cdot\|$ denote a matrix norm. The* joint spectral radius $\rho(S)$ *of S is defined by the limit*

$$\rho(S) := \lim_{k\to\infty} \max\{\, \|A_1 \cdots A_k\|^{1/k} \mid A_i \in S, 1 \leqslant i \leqslant k \,\}$$

It is well-known that this limit always exists and that it does not depend on the chosen norm, which follows from the equivalence of all norms in $\mathbb{R}^n$; e.g., one could take the norm given by the sum of the absolute values of all matrix entries. Further, if $S = \{A\}$ is a singleton set, the joint spectral radius $\rho(S)$ of S and the spectral radius $\rho(A)$ of A coincide:

$$\rho(S) = \lim_{k\to\infty} \|A^k\|^{1/k} = \max\{\, |\lambda| \mid \lambda \text{ is an eigenvalue of } A \,\} = \rho(A)$$

Since Definition 10 is independent of the actual norm, from now on we simply write $\mathsf{growth}_S(k)$. The following theorem (due to [2]) provides a characterization of polynomial boundedness of $\mathsf{growth}_S(k)$ by the joint spectral radius of S.

Theorem 11 ([2, Theorem 1.2])**.** *Let $S \subseteq \mathbb{R}^{n\times n}$ be a finite set of matrices. Then $\mathsf{growth}_S(k) = O(k^d)$ for some $d \in \mathbb{N}$ if and only if $\rho(S) \leqslant 1$. In particular, $d \leqslant n - 1$.* □

Hence, polynomial boundedness of $\mathsf{growth}_S(k)$ is decidable if $\rho(S) \leqslant 1$ is decidable. But it is well-known that in general the latter is undecidable for arbitrary, but finite sets $S \subseteq \mathbb{R}^{n\times n}$, even if S contains only non-negative rational (real) matrices (cf. [10, Theorem 2.6]). However, if S contains only non-negative integer matrices then $\rho(S) \leqslant 1$ is decidable. In particular, there exists a polynomial-time algorithm that decides it (cf. [10, Theorem 3.1]). This algorithm is based on the following lemma.

Lemma 12 ([10, Lemma 3.3])**.** *Let $S \subseteq \mathbb{R}_0^{n\times n}$ be a finite set of non-negative, real square matrices. Then there is a product $A \in S^*$ such that $A_{ii} > 1$ for some $i \in \{1, \ldots, n\}$ if and only if $\rho(S) > 1$.* □

According to [10], for $S \subseteq \mathbb{N}^{n \times n}$, the existence of such a product can be characterized in terms of the graphs $G(S)$ and $G^2(S)$ one can associate with S. More precisely, there is a product $A \in S^*$ with $A_{ii} > 1$ if and only if

1. there is a cycle in $G(S)$ containing at least one edge of weight $w > 1$, or
2. there is a cycle in $G^2(S)$ containing at least one vertex (i, i) and at least one vertex (p, q) with $p \neq q$.

Hence, we have $\rho(S) \leqslant 1$ if and only if neither of the two conditions holds, which can be checked in polynomial time according to [10]. Furthermore, as already mentioned in [10, Chapter 3], this graph-theoretic characterization does not only hold for non-negative integer matrices, but for any set of matrices such that all matrix entries are either zero or at least one (because then all paths in $G(S)$ have weight at least one).

Lemma 13. *Let $S \subseteq \mathbb{R}_0^{n \times n}$ be a finite set of matrices such that all matrix entries are either zero or at least one. Then $\rho(S) \leqslant 1$ is decidable in polynomial time.* □

So, in the situation of Lemma 13, polynomial boundedness of $\mathsf{growth}_S(k)$ is decidable in polynomial time. In addition, the exact degree of growth can be computed in polynomial time (cf. [10, Theorem 3.3]).

Theorem 14. *Let $S \subseteq \mathbb{R}_0^{n \times n}$ be a finite set of matrices such that $\rho(S) \leqslant 1$ and all matrix entries are either zero or at least one. Then*

$$\mathsf{growth}_S(k) = \Theta(k^d)$$

where the growth rate *d is the largest integer possessing the following property: there exist d different pairs of indices $(i_1, j_1), \ldots, (i_d, j_d)$ such that for every pair (i_s, j_s) the indices i_s, j_s are different and there is a product $A \in S^*$ for which $A_{i_s i_s}, A_{i_s j_s}, A_{j_s j_s} \geqslant 1$, and for each $1 \leqslant s \leqslant d-1$, there exists $B \in S^*$ with $B_{j_s i_{s+1}} \geqslant 1$. Moreover, d is computable in polynomial time.* □

Next we elaborate on the ramifications of joint spectral radius theory on complexity analysis of TRSs via polynomially bounded matrix interpretations. To begin with, we observe that the characterization of polynomially bounded matrix interpretations given in Definition 1 can be rephrased as follows: A matrix interpretation $\mathcal{M}$ is polynomially bounded if the joint spectral radius of the set of matrices occurring in $\mathcal{M}$ is at most one. This follows directly from Theorem 11 (using as $\|\cdot\|$ in $\mathsf{growth}_S(k, \|\cdot\|)$ the matrix norm given by the sum of the absolute values of all matrix entries). Due to the relationship between polynomially bounded matrix interpretations and the derivational complexity of compatible TRSs expressed in Lemma 2, we immediately obtain the following theorem, which holds for matrix interpretations over $\mathbb{N}$, $\mathbb{Q}$, and $\mathbb{R}$.

Theorem 15. *Let $\mathcal{R}$ be a TRS and $\mathcal{M}$ a compatible matrix interpretation of dimension n. Further, let $S \subseteq \mathbb{R}_0^{n \times n}$ denote the set of all matrices occurring in $\mathcal{M}$. If the joint spectral radius of S is at most one then $\mathsf{dc}_{\mathcal{R}}(k) = O(k^n)$.* □

As this theorem assumes the worst-case growth rate for $\mathsf{growth}_S(k)$, the inferred degree of the polynomial bound may generally be too high (and unnecessarily so). Yet with the help of Theorem 14, from which we obtain the exact growth rate of $\mathsf{growth}_S(k)$, Theorem 15 can be strengthened, at the expense of having to restrict the set of permissible matrices.

Theorem 16. *Let $\mathcal{R}$ be a TRS and $\mathcal{M}$ a compatible matrix interpretation of dimension n. Further, let $S \subseteq \mathbb{R}_0^{n\times n}$ denote the set of all matrices occurring in $\mathcal{M}$ and assume that all matrix entries are either zero or at least one. If the joint spectral radius of S is at most one then $\mathsf{dc}_{\mathcal{R}}(k) = O(k^{d+1})$, where d refers to the growth rate obtained from Theorem 14.* □

4 Automata Methods for Bounding Polynomial Growth

In this section we study automata-based methods to classify polynomial growth of matrix interpretations. Complementing Definition 1, we employ the following definition of polynomial growth [19].

Definition 17. *The growth function of a matrix interpretation $\mathcal{M}$ is defined as* $\mathsf{growth}_{\mathcal{M}}(k) := \max\{[t]_1 \mid t \text{ is a term and } |t| \leqslant k\}$. *We say that $\mathcal{M}$ is* polynomially bounded with degree d *if* $\mathsf{growth}_{\mathcal{M}}(k) = O(k^d)$ *for some* $d \in \mathbb{N}$.

We recall basic notions of automata theory (cf. [3,17]). A *weighted automaton (over $\mathbb{R}$)* is a quintuple $\mathcal{A} = (Q, \Sigma, \lambda, \mu, \gamma)$ where Q is a finite set of states, Σ a finite alphabet, and the mappings $\lambda\colon Q \to \mathbb{R}$, $\gamma\colon Q \to \mathbb{R}$ are weight functions for entering and leaving a state. The transition function $\mu\colon \Sigma \to \mathbb{R}^{|Q|\times|Q|}$ associates with any letter in Σ a $(|Q| \times |Q|)$-matrix over $\mathbb{R}$. For $a \in \Sigma$, $\mu(a)_{p,q}$ denotes the *weight* of the transition $p \xrightarrow{a} q$. We often view λ and γ as row and column vectors, respectively. We also write $\mathsf{weight}(p, a, q)$ for $\mu(a)_{p,q}$ and extend the transition function μ homomorphically to words. The *weight of* $x \in \Sigma^*$, denoted by $\mathsf{weight}_{\mathcal{A}}(x)$, is the sum of the weights of paths in $\mathcal{A}$ labelled with x. We define

$$\mathsf{weight}_{\mathcal{A}}(x) = \sum_{p,q\in Q} \lambda(p) \cdot \mu(x)_{p,q} \cdot \gamma(q) = \lambda \cdot \mu(x) \cdot \gamma$$

where the weight of the empty word is just $\lambda \cdot \gamma$.

A weighted automaton $\mathcal{A} = (Q, \Sigma, \lambda, \mu, \gamma)$ is called *normal* if the row vector λ is $(1, \overline{0})$ and the column vector γ contains only entries with weight 0 or 1. Here $\overline{0}$ denotes a sequence of $|Q| - 1$ zeros. Let $\mathcal{A}$ be a normal weighted automaton. The unique state q_0 such that $\lambda(q_0) \neq 0$ is called *initial* and the states in $F := \{q \mid \gamma(q) \neq 0\}$ are called *final*. We call a state $p \in Q$ *useful* if there exists a path in $\mathcal{A}$ from q_0 to $q \in F$ that contains p. An automaton $\mathcal{A}$ is *trim* if all states are useful. In the sequel all considered automata will be normal.

The main result of this section is a complete (and polytime decidable) characterization of polynomial growth of matrix interpretations over $\mathbb{N}$, thus re-stating the main result in [19]. We extend upon [19] by clarifying the polytime decidability of the properties involved. Given a matrix interpretation $\mathcal{M}$, we denote by $C^{\mathcal{M}}$ the component-wise maximum of all absolute vectors in $\mathcal{M}$.

Definition 18. *With every n-dimensional matrix interpretation $\mathcal{M}$ for a signature $\mathcal{F}$ we associate a weighted automaton $\mathcal{A} = (Q, \Sigma, \lambda, \mu, \gamma)$ as follows: $Q = \{1, \ldots, n\}$, $\Sigma = \{f_i \mid f \in \mathcal{F}$ has arity k and $1 \leqslant i \leqslant k\}$, $\lambda = (1, \overline{0})$, $\gamma \in \{0,1\}^n$ such that $\gamma(i) = 1$ if and only if $C_i^{\mathcal{M}} > 0$, and $\mu(f_i) = F_i$ where F_i denotes the i-th matrix of $f_{\mathcal{M}}$.*

Example 19. Consider the matrix interpretation $\mathcal{M}$ of dimension 3 with

$$\mathsf{a}_{\mathcal{M}}(\boldsymbol{x}) = \begin{pmatrix} 1 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{pmatrix} \boldsymbol{x} \qquad \mathsf{b}_{\mathcal{M}}(\boldsymbol{x}) = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{pmatrix} \boldsymbol{x} + \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}$$

The following automaton $\mathcal{A} = (Q, \Sigma, \lambda, \mu, \gamma)$ corresponds to $\mathcal{M}$:

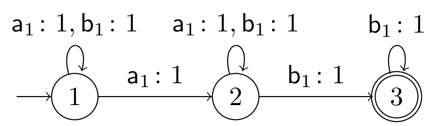

As $\mathcal{A}$ is normal, we represent λ and γ by indicating the input and output states as usual for finite automata.

In analogy to Definition 17 we define the growth function of a weighted automaton.

Definition 20. *Let $\mathcal{A} = (Q, \Sigma, \lambda, \mu, \gamma)$ be a weighted automaton. The growth function of $\mathcal{A}$ is defined as $\mathsf{growth}_{\mathcal{A}}(k) := \max\{\mathsf{weight}(x) \mid x \in \Sigma^k\}$.*

The following theorem restates [19, Theorem 3.3] in connection with the remark immediately following the theorem.

Theorem 21. *Let $\mathcal{M}$ be a matrix interpretation and $\mathcal{A}$ the corresponding automaton. Then $\mathsf{growth}_{\mathcal{A}}(k) = O(k^d)$ if and only if $\mathsf{growth}_{\mathcal{M}}(k) = O(k^{d+1})$.* □

As a consequence of Theorem 21 we observe that the growth of matrix interpretations and weighted automata are polynomially related. Hence if we can decide the polynomial growth rate of automata, we can decide the polynomial growth rate of matrix interpretations. In the remainder of this section we restrict to matrix interpretations over $\mathbb{N}$ (and therefore to weighted automata over $\mathbb{N}$).

The growth rate of automata has been studied in various contexts. Here we only mention two independent developments. Weber and Seidl provide in [20] a complete characterization of the degree of growth of the ambiguity of nondeterministic finite automata (NFAs). If we restrict to weights over $\{0,1\}$, weighted automata simplify to NFAs. Furthermore, if the degree of growth of the ambiguity of an NFA $\mathcal{A}$ is d then $\mathsf{growth}_{\mathcal{A}}(k) = O(k^d)$ and vice versa. Jungers [10] completely characterises polynomial growth rates of matrix products in the context of joint spectral radius theory (cf. Theorem 14). Due to the closeness of weighted automata to finite automata we follow the account of Weber and Seidl here. To make full use of the results in [20] it suffices to observe that weighted automata over $\mathbb{N}$ can be represented as finite automata with parallel edges. This will also allow to place later developments (see Section 5) into context.

Consider the following criterion for a weighted automaton $\mathcal{A} = (Q, \Sigma, \lambda, \mu, \delta)$:

$$\exists\, q \in Q\; \exists\, x \in \Sigma^* \text{ such that } \mathsf{weight}(q, x, q) \geqslant 2 \text{ and } q \text{ is useful} \qquad (\mathsf{EDA})$$

The criterion was introduced in [20] for NFAs. A similar criterion can be distilled from the decision procedures given in [10, Chapter 3]. Note that the conditions given after Lemma 12 are equivalent to EDA if we restrict our attention to trim automata. Without loss of generality let $Q = \{1, \ldots, n\}$ and let S be the set of transition matrices in $\mathcal{A}$. By definition, there exists a product $A \in S^*$ with $A_{ii} \geqslant 2$ if and only if there exists $x \in \Sigma^*$ such that $\mathsf{weight}(i, x, i) \geqslant 2$. Hence $\mathcal{A}$ fulfills EDA if and only if condition 1 or condition 2 is fulfilled. In the following we sometimes refer to EDA as the collection of these conditions.

Condition EDA is sufficient to decide the polynomial growth of $\mathbb{N}$-weighted automata. The following theorem embodies Theorem 5.2 in [19]. The polytime decision procedure stems from [20].

Theorem 22. *Let $\mathcal{A} = (Q, \Sigma, \lambda, \mu, \gamma)$ be a weighted automaton over $\mathbb{N}$. Then there exists $d \in \mathbb{N}$ such that $\mathsf{growth}_{\mathcal{A}}(k) = O(k^d)$ if and only if $\mathcal{A}$ does not admit* EDA*. Furthermore this property is decidable in time $O(|Q|^4 \cdot |\Sigma|)$.* □

To determine the degree of polynomial growth, the following criterion introduced in [20] is used:

$$\exists\, p_1, q_1, \ldots, p_d, q_d \in Q\; \exists\, v_1, u_2, v_2, \ldots, u_d, v_d \in \Sigma^* \text{ such that } \forall\, i \geqslant 1\; \forall\, j \geqslant 2$$
$$p_i \neq q_i,\ p_i \xrightarrow{v_i} p_i,\ p_i \xrightarrow{v_i} q_i,\ q_i \xrightarrow{v_i} q_i,\ q_{j-1} \xrightarrow{u_j} p_j,\ p_i \text{ and } q_i \text{ are useful} \qquad (\mathsf{IDA}_d)$$

The criterion IDA_d can be visualized as follows:

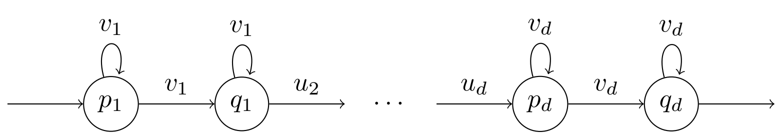

In the same vein as for EDA, one easily sees that the condition IDA_d is closely linked (on trim automata) to the conditions stated in Theorem 14. This will be detailed in Section 5.

Example 23 (continued from Example 19). It is easy to see that $\mathcal{A}$ complies with IDA_2, where $p_1 = 1$, $q_1 = p_2 = 2$, and $q_2 = 3$. Moreover, an easy calculation gives $\mathsf{weight}_{\mathcal{A}}(\mathsf{a}^n \mathsf{b}^m) = \mathsf{weight}_{\mathcal{A}}(1, \mathsf{a}^n \mathsf{b}^m, 3) = nm$ and thus the growth rate of $\mathcal{A}$ is quadratic.

The next theorem essentially follows from a close inspection of the proof of [20, Theorem 4.2] in connection with Theorem 21.

Theorem 24. *Let $\mathcal{M}$ be a matrix interpretation of dimension n and let $\mathcal{A} = (Q, \Sigma, \lambda, \mu, \gamma)$ be the corresponding weighted automaton. Then $\mathsf{growth}_{\mathcal{M}}(k) = \Theta(k^{d+1})$ if and only if $\mathcal{A}$ does not comply with* EDA *nor with* IDA_{d+1}*, but complies with* IDA_d*. Furthermore this property is decidable in time $O(|Q|^6 \cdot |\Sigma|)$.* □

As a direct consequence of the theorem together with Theorem 21 we obtain the following corollary.

Corollary 25. *Let $\mathcal{R}$ be a TRS and let $\mathcal{M}$ be a compatible matrix interpretation of dimension n. Further, let $\mathcal{A}$ be the corresponding weighted automaton such that $\mathcal{A}$ does not comply with* EDA *nor with* IDA_{d+1}. *Then* $\mathsf{dc}_{\mathcal{R}}(k) = O(k^{d+1})$. *Furthermore the conditions given are decidable in polynomial time.* □

Example 26 (continued from Example 23). The automaton $\mathcal{A}$ complies with IDA_2. It is easy to see that $\mathcal{A}$ does not comply with EDA nor with IDA_3. Hence the growth of the matrix interpretation $\mathcal{M}$ is at most cubic, i.e., $\mathsf{growth}_{\mathcal{M}}(k) = O(k^3)$.

5 Unifying Algebraic and Automata-based Methods

In this section we unify the algebraic and the automata-based approach presented in Sections 3 and 4. We start by relating the two different notions of polynomial growth of matrix interpretations that have been studied in the literature (cf. Definitions 1 and 17). The next example shows that these two definitions are not polynomially related.

Example 27. Consider the TRS $\mathsf{f}(x) \to x$ together with the following compatible matrix interpretation $\mathcal{M}$:

$$\mathsf{f}_{\mathcal{M}}(\boldsymbol{x}) = \begin{pmatrix} 1 & 1 \\ 0 & 2 \end{pmatrix} \boldsymbol{x} + \begin{pmatrix} 1 \\ 0 \end{pmatrix}$$

It is easy to see that the growth of the entries in the second column of the matrix is exponential, so $\mathcal{M}$ is not of polynomial growth with respect to Definition 1. However, $[t]_1 < |t|$ for any term t and thus $\mathcal{M}$ is of polynomial growth with respect to Definition 17.

Still, *morally* both definitions are equal. Observe that the entries that grow exponentially in the matrix in Example 27 are irrelevant when computing the value $[t]_1$. In order to prove this we exploit the connection between matrix interpretations and weighted automata emphasised in Section 4.

Example 28 (continued from Example 27). The weighted automaton $\mathcal{A}$ corresponding to $\mathcal{M}$ can be pictured as follows:

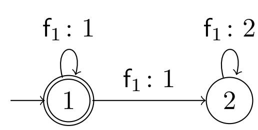

Observe that $\mathcal{A}$ is not trim because state 2 is not useful.

Lemma 29. *Let $\mathcal{R}$ be a TRS and let $\mathcal{M}$ be a compatible matrix interpretation. There exists a matrix interpretation $\mathcal{N}$ compatible with $\mathcal{R}$ such that the corresponding automaton is trim. Furthermore,* $\mathsf{growth}_{\mathcal{M}}(k) = \mathsf{growth}_{\mathcal{N}}(k)$.

Proof. Following [19], with every position $p \in \mathcal{P}\mathsf{os}(t)$ we associate a word t_p over Σ: $t_\epsilon = \epsilon$ and if $p = i \cdot q$ and $t = f(t_1, \ldots, t_n)$, then $t_p = f_i \cdot t_q$. The set $\mathsf{Path}(t)$ consists of all these words. We write $t(p)$ for the root symbol of the subterm of t at position p. Let $\mathcal{M}$ be a matrix interpretation of dimension n and let $\mathcal{A} = (Q, \Sigma, \lambda, \mu, \gamma)$ be the corresponding normal weighted automaton. A key observation is the following identity, which holds for all assignments α of vectors in $\mathbb{N}^n$ to the variables in t:

$$[\alpha]_{\mathcal{M}}(t) = \sum_{p \in \mathcal{P}\mathsf{os}(t)} \mu(t_p) \cdot \mathsf{abs}(\alpha, t(p))$$

where $\mathsf{abs}(\alpha, t(p)) = \mathsf{abs}(t(p))$ if $t(p) \in \mathcal{F}$ and $\mathsf{abs}(\alpha, t(p)) = \alpha(t(p))$ otherwise. Without loss of generality we assume that $\mathcal{R}$ is non-empty. As $\mathcal{M}$ is compatible with $\mathcal{R}$, there must be a word in Σ^* of positive weight. This implies in particular that state 1 is useful. Now suppose that $\mathcal{A}$ is not trim. Let $\mathcal{B} = (Q', \Sigma, \lambda', \mu', \gamma')$ denote an equivalent trim automaton. We have $\lambda' = \lambda{\restriction_{Q'}}$, $\gamma' = \gamma{\restriction_{Q'}}$, and for all $a \in \Sigma$, $\mu'(a) = \mu(a){\restriction_{Q' \times Q'}}$. To simplify notation we assume (without loss of generality; recall that state 1 is useful) that $Q' = \{1, \ldots, m\}$ for some $m < n$.

Let $\mathcal{N}$ be the matrix interpretation corresponding to $\mathcal{B}$, where the absolute vector of every $f_\mathcal{N}$ equals $\mathsf{abs}(f){\restriction_{Q'}}$. We prove that $\mathcal{N}$ is compatible with $\mathcal{R}$. Let $\ell \to r$ be a rule in $\mathcal{R}$ and let $\beta\colon \mathcal{V} \to \mathbb{N}^m$ be an arbitrary assignment. Let $\alpha\colon \mathcal{V} \to \mathbb{N}^n$ be the assignment that is obtained from β by zero padding, i.e., $\beta(x)_i = \alpha(x)_i$ for $1 \leqslant i \leqslant m$ and $\beta(x)_i = 0$ for $m < i \leqslant n$. From the compatibility of $\mathcal{M}$ and $\mathcal{R}$ we obtain $[\alpha]_\mathcal{M}(\ell)_1 > [\alpha]_\mathcal{M}(r)_1$ and $[\alpha]_\mathcal{M}(\ell)_j \geqslant [\alpha]_\mathcal{M}(r)_j$ for all $1 < j \leqslant n$. We claim that $[\alpha]_\mathcal{M}(t)_i = [\beta]_\mathcal{N}(t)_i$ for all $1 \leqslant i \leqslant m$. From the claim we immediately obtain $[\beta]_\mathcal{N}(\ell)_1 > [\beta]_\mathcal{N}(r)_1$ and $[\beta]_\mathcal{N}(\ell)_j \geqslant [\beta]_\mathcal{N}(r)_j$ for all $1 < j \leqslant m$. It follows that $\mathcal{N}$ is compatible with $\mathcal{R}$ and by Definition 17, $\mathsf{growth}_\mathcal{M}(k) = \mathsf{growth}_\mathcal{N}(k)$.

To prove the claim, fix $i \in \{1, \ldots, m\}$ and let $\lambda_i = (\overline{0}, 1, \overline{0})$ be the row vector of dimension n where the 1 is at position i. Let $p \in \mathcal{P}\mathsf{os}(t)$. We have

$$\begin{aligned}
\lambda_i \cdot \mu(t_p) \cdot \mathsf{abs}(\alpha, t(p)) &= (\mu(t_p) \cdot \mathsf{abs}(\alpha, t(p)))_i \\
&= (\mu'(t_p) \cdot \mathsf{abs}(\alpha, t(p)){\restriction_{Q'}})_i \\
&= (\mu'(t_p) \cdot \mathsf{abs}(\beta, t(p)))_i = \lambda'_i \cdot \mu'(t_p) \cdot \mathsf{abs}(\beta, t(p))
\end{aligned}$$

where $\lambda'_i = \lambda_i{\restriction_{Q'}}$. It follows that $[\alpha]_\mathcal{M}(t)_i = [\beta]_\mathcal{N}(t)_i$ for all $1 \leqslant i \leqslant m$. ☐

Example 30 (continued from Example 28). Because state 2 in the weighted automaton $\mathcal{A}$ is useless, it is removed to obtain an equivalent trim automaton $\mathcal{B}$. This one-state automaton gives rise to the 1-dimensional matrix interpretation (i.e., linear polynomial interpretation) $\mathcal{N}$ with $\mathsf{f}_\mathcal{N}(x) = x + 1$.

An immediate consequence of Theorem 21 and Lemma 29 is that Definitions 1 and 17 are equivalent in the following sense.

Corollary 31. *Let $\mathcal{R}$ be a TRS. The following statements are equivalent.*

1. *There exists a matrix interpretation $\mathcal{M}$ compatible with $\mathcal{R}$ that is polynomially bounded according to Definition 1.*

2. *There exists a matrix interpretation $\mathcal{N}$ compatible with $\mathcal{R}$ that is polynomially bounded according to Definition 17.*

Furthermore the entries of all products of matrices in $\mathcal{M}$ are polynomial with degree d in the length of such products if and only if $\mathsf{growth}_{\mathcal{N}}(k) = O(k^{d+1})$.

Proof. First, assume that $\mathcal{M}$ is a compatible matrix interpretation that is polynomially bounded according to Definition 1. Further assume the entries of all matrix products of matrices in $\mathcal{M}$ are polynomial with degree d in the length of such products. Since the matrices occurring in $\mathcal{M}$ form the transitions of the automaton $\mathcal{A}$ corresponding to $\mathcal{M}$, $\mathsf{growth}_{\mathcal{A}}(k) = O(k^d)$. Hence, by Theorem 21, $\mathsf{growth}_{\mathcal{M}}(k) = O(k^{d+1})$, which means that $\mathcal{M}$ is polynomially bounded according to Definition 17.

As to the converse statement, assume that $\mathcal{N}$ is a compatible matrix interpretation that is polynomially bounded according to Definition 17 and let $\mathsf{growth}_{\mathcal{N}}(k) = O(k^{d+1})$. By Lemma 29 there exists a compatible matrix interpretation $\mathcal{M}$ such that $\mathsf{growth}_{\mathcal{N}}(k) = \mathsf{growth}_{\mathcal{M}}(k)$ and the corresponding automaton $\mathcal{A}$ is trim. By Theorem 21, we have $\mathsf{growth}_{\mathcal{A}}(k) = \mathrm{O}(k^d)$. This entails that all entries of all matrix products of matrices in $\mathcal{M}$ are polynomial with degree d, as $\mathcal{M}$ is trim, which means that $\mathcal{M}$ is polynomially bounded according to Definition 1. ☐

In what follows we show that the algebraic approach presented in Section 3 readily subsumes the automata theory based approach of Section 4. To be precise, we show that the restriction of Theorem 16 to matrix interpretations over $\mathbb{N}$ applies in any situation where Corollary 25, the main result of Section 4, applies.

Theorem 32. *Let $\mathcal{R}$ be a TRS. If one can establish polynomial derivational complexity of some degree via Corollary 25, then one can also establish polynomial derivational complexity of the same degree via Theorem 16.*

Proof. Let $\mathcal{M}$ be a compatible matrix interpretation over $\mathbb{N}$ of dimension n and let $S \subseteq \mathbb{N}^{n\times n}$ denote the set of all matrices occurring in $\mathcal{M}$. By Lemma 29 we may assume that the automaton $\mathcal{A} = (Q, \Sigma, \lambda, \mu, \gamma)$ corresponding to $\mathcal{M}$ is trim. By assumption $\mathcal{A}$ does not comply with EDA nor with IDA_{d+1} and $\mathsf{dc}_{\mathcal{R}}(k) = O(k^{d+1})$. As one obtains the tightest bound for $\mathsf{dc}_{\mathcal{R}}(k)$ if $d+1$ is the least integer such that $\neg\mathsf{IDA}_{d+1}$ holds in $\mathcal{A}$, or, equivalently, if d is the largest integer such that IDA_d holds in $\mathcal{A}$, this will be assumed in what follows.

First we show that $\neg\mathsf{EDA}$ implies $\rho(S) \leqslant 1$. By Lemma 12 we have $\rho(S) > 1$ if and only if there is a product $A \in S^*$ such that $A_{ii} > 1$ for some $i \in \{1, \ldots, n\}$. By construction of $\mathcal{A}$ the latter is equivalent to the existence of a word $x \in \Sigma^*$ (corresponding to the product $A \in S^*$) such that $\mu(x)_{i,i} > 1$ for some state $i \in Q = \{1, \ldots, n\}$, or, equivalently, $\mu(x)_{i,i} = \mathsf{weight}(i, x, i) \geqslant 2$ since $\mathcal{A}$ is $\mathbb{N}$-weighted. This means that $\mathcal{A}$ complies with EDA because state i is useful.

Next we show that d is exactly the growth rate mentioned in Theorem 16, that is, the growth rate inferred from Theorem 14. As $\mathcal{A}$ complies with IDA_d, there are useful states $p_1, q_1, \ldots, p_d, q_d \in Q$ and words $v_1, u_2, v_2, \ldots, u_d, v_d \in \Sigma^*$

such that $p_s \neq q_s$, $p_s \xrightarrow{v_s} p_s$, $p_s \xrightarrow{v_s} q_s$, and $q_s \xrightarrow{v_s} q_s$ for all $s = 1, \ldots, d$ and $q_{s-1} \xrightarrow{u_s} p_s$ for all $s = 2, \ldots, d$. Letting $A_s = \mu(v_s) \in S^*$ for $s = 1, \ldots, d$ and $B_s = \mu(u_s) \in S^*$ for $s = 2, \ldots, d$ (by the one-to-one correspondence between letters in Σ and matrices in S), this is equivalent to the existence of d pairs of indices $(p_1, q_1), \ldots, (p_d, q_d)$ and matrix products $A_1, B_2, A_2, \ldots, B_d, A_d \in S^*$ such that $p_s \neq q_s$ and $(A_s)_{p_s p_s}, (A_s)_{p_s q_s}, (A_s)_{q_s q_s} \geqslant 1$ for all $s = 1, \ldots, d$ and $(B_s)_{q_{s-1} p_s} \geqslant 1$ for all $s = 2, \ldots, d$. Note that all pairs of indices (states) must be different because otherwise $\mathcal{A}$ would comply with EDA, contradicting our assumption. Hence, all conditions concerning the growth rate in Theorem 14 are satisfied and Theorem 16 yields $\mathsf{dc}_{\mathcal{R}}(k) = O(k^{d+1})$. □

6 Automation and Experimental Results

6.1 Automation

In this section we describe certificates for polynomial boundedness of matrix interpretations. These certificates will be described by finite-domain constraint systems and are solved by machine. The constraint domain consists of relations on the state set of an automaton. In our implementations, the constraint system is solved via translation to a constraint system in propositional logic. This is motivated by the intended application in conjunction with a constraint system for the entries in the matrices that describes its compatibility with a given TRS [4]. So if there is a solution, it fulfills both properties (compatibility and polynomial growth).

First we focus on how to ensure $\rho(A) \leqslant 1$ for a single n-dimensional matrix with non-negative real entries, which is needed to implement Theorem 5. We base our encoding of the minimal polynomial m_A on the factorization approach **(C)** from [15] but instead of demanding certain properties of the characteristic polynomial we encode a monic polynomial p that annihilates A such that $|\lambda| \leqslant 1$ for every root λ of p. One candidate for p is the characteristic polynomial of A, so the degree of p can be chosen n. To cancel some factors we introduce variables $C, C_j \in \{0, 1\}$ such that

$$p(\lambda) = (C\lambda - Cr + 1 - C)^b \cdot \prod_j (C_j\lambda^2 + C_j p_j \lambda + C_j q_j + 1 - C_j)$$

Here $r, p_j, q_j \in \mathbb{R}$ and $b = 0$ if n is even and $b = 1$ otherwise. Note that if C is zero then this factor simplifies to one and hence does not affect the product, while if C is one, the factor contributes to p. The same property holds for the C_j. We add a constraint $p(A) = 0$ to the encoding which ensures that p annihilates A. By the shape of the factorization, p is automatically monic. The condition $|\lambda| \leqslant 1$ is encoded based on Cr, $C_j p_j$, $C_j q_j$ as in [15] if $C_j = 1$, while $C_j = 0$ does not require additional constraints. Finding the minimal polynomial m_A is then an optimization problem, i.e., it amounts to minimize the maximum multiplicity of roots with absolute value exactly one.

Next we sketch how to count the number of roots equal to one. We abbreviate $(C_j p_j)^2 - 4C_j q_j$ by D_j. Obviously $Cr = 1$ corresponds to an occurrence of root one and $Cr = -1$ to an occurrence of root minus one. For the other factors, inspecting the quadratic formula gives root one if $C_j p_j + C_j q_j = -1$. Then $D_j = 0$ corresponds to a double occurrence and $D_j > 0$ to a single occurrence. The reasoning for root minus one is similar and based on $C_j p_j = C_j q_j + 1$. To keep the encoding simple, we over-approximate the multiplicity of a complex root with absolute value one (i.e., we possibly identify different complex roots for counting). This counter is incremented by one if $D_j < 0$ and $C_j q_j = 1$.

The following example shows that computing the minimal polynomial after establishing a compatible matrix interpretation may not result in optimal bounds.

Example 33. Consider the TRS consisting of the single rule $\mathsf{f}(x) \to x$ and assume a compatible matrix interpretation of dimension n having maximum matrix A with $A_{ij} = 0$ if $i > j$ and $A_{ij} = 1$ otherwise. Then $\chi_A(x) = m_A(x)$ and hence Theorem 5 establishes a polynomial bound of degree n. However, if the maximum matrix equals I_n then Theorem 5 establishes a linear upper bound on the derivational complexity.

Next we aim to describe polynomial growth of arbitrary matrix products with entries from $\mathbb{N}$. Section 4 gives polynomial-time algorithms for checking polynomial boundedness and computing the degree of growth for weighted automata. We will express these algorithms as constraint systems. By Cook's theorem, we know that any P (even NP) computation can be simulated by a polynomially-sized SAT formula. The size of the formula is the product of space and time of the computation. Thus a straightforward translation of the decision procedure for polynomial boundedness ($\neg\mathsf{EDA}$) results in a formula size of $O(n^6)$, and for polynomial boundedness of a fixed degree ($\neg\mathsf{IDA}_{d+1}$), in a formula size of $O(n^9)$, where n is the dimension of the matrices.

Note that Cook's translation produces a formula where the number of variables is of the same order as the number of clauses. We will strive to reduce theses numbers, and in particular, the number of variables. E.g., we implement the criterion $\neg\mathsf{IDA}_{d+1}$ by $O(n^8)$ clauses but only $O(n^5)$ variables. The fewer variables, the less choice points there are for the SAT solver, and for a fixed set of variables, more constraints allow for more unit propagations, hopefully speeding up the solver even more.

The formula size is reduced by replacing parts of the original algorithms by approximations, but using them in such a way as to still ensure correctness and completeness. E.g., for reachability with respect to a set of matrices S, we specify a relation $R \subseteq Q^2$ such that R is reflexive and transitive and, for all $p, q \in Q$, $G(S)_{p,q} > 0$ implies $(p, q) \in R$,

While the presence of patterns (EDA, IDA_d) is easily described by a constraint system in existential logic, the challenge is to specify the *absence* of those patterns.

$\neg$EDA*:* By the algorithm given right after Lemma 12, we use the following constraints, for a given set S of matrices:

- Condition 1 follows from the constraint $G(S)_{i,j} > 1 \Rightarrow (j,i) \notin R$, for all i and j, where R is the over-approximation of reachability discussed above.
- For condition 2, denote by G_2 the adjacency relation of $G^2(S)$. A cycle from (i,i) through some (p,q) for $p \neq q$ means that (p,q) is strongly G_2-connected to (i,i). Therefore, we set $D = \{(q,q) \mid q \in Q\}$ and specify an over-approximation D_1 of the index pairs reachable from the main diagonal by $D \subseteq D_1$ and $G_2(D_1) \subseteq D_1$, and an over-approximation D_2 of the index pairs that can reach the main diagonal by $D \subseteq D_2$ and $G_2^-(D_2) \subseteq D_2$, and then require $D_1 \cap D_2 \subseteq D$.

The resulting constraint system is satisfiable if and only if $\mathcal{A}$ does not comply with EDA. So we have a complete characterization of polynomial boundedness, equivalently, of $\rho(S) \leqslant 1$ for a set of non-negative integer matrices (cf. Lemma 13 and Theorem 22).

The constraint system has $O(n^2)$ variables and $O(n^4)$ constraints, which is small in comparison to the size of the constraint system that describes compatibility of the interpretation with the rewrite system [4].

$\neg$IDA$_{d+1}$*:* We use the relation $I \subseteq Q^2$ given by

$$I = \{(p,q) \mid p \neq q \text{ and } p \xrightarrow{x} p,\ p \xrightarrow{x} q,\ q \xrightarrow{x} q \text{ for some } x \in \Sigma^+\}$$

and let $J = I \circ R$, where R is an over-approximation of reachability as defined previously. Then IDA$_d$ implies that there is a J-chain of length d, and conversely, if there is no J-chain of length $d+1$ then $\neg$IDA$_{d+1}$ holds. To specify I, we use the graph $G^3(S)$, and denote its adjacency relation by G_3. For each (p,q) with $p \neq q$ we specify a set $T \subseteq Q^3$ by $(p,p,q) \in T$ and $G_3(T) \subseteq T$, Then $(p,q,q) \in T$ implies $I(p,q)$. We bound the length of J-chains by encoding a height function $h\colon Q \to \{0,1,\ldots,d\}$ and a monotonicity property $J(p,q) \Rightarrow h(p) > h(q)$. Since only comparison but no arithmetic is needed here, it is convenient to represent numbers in unary notation.

The resulting constraint system is satisfiable if and only if $\neg$IDA$_{d+1}$, so we obtain a characterization of growth $O(n^d)$. The system has $O(n^5)$ variables and $O(n^8)$ constraints. This is comparable in size to the compatibility constraints.

6.2 Experimental Results

The criteria proposed in this paper have been implemented in the complexity tools CaT [21] and matchbox [19]. All tests have been performed on a server equipped with 64 GB of main memory and eight dual-core AMD Opteron® 885 processors running at a clock rate of 2.6 GHz with a time limit of 60 seconds per system. For experiments[5] the 295 non-duplicating TRSs in TPDB 8.0 for which

[5] For full details see http://colo6-c703.uibk.ac.at/hzankl/11cai.

Table 1. Polynomial bounds for 295 systems

	$O(k)$	$O(k^2)$	$O(k^3)$	$O(k^n)$
triangular	92	194	208	210
Theorem 4	66	180	194	196
Theorem 5	72	187	193	194
$\sum$	92	201	215	216
Theorem 15 ($\neg$EDA)	73	173	188	213
Theorem 16 ($\neg$EDA, $\neg\mathsf{IDA}_{d+1}$)	107	214	219	224
$\sum$	110	225	231	236
$\sum$	112	227	235	239

GT could find a compatible matrix interpretation (not necessarily polynomially bounded) have been considered.

We searched for matrix interpretations of dimension $d \in \{1, \dots, 5\}$ by encoding the constraints as an SMT problem (quantifier-free non-linear arithmetic), which is solved by "bit-blasting". That means the problem is transformed into a finite domain problem by prescribing discrete and finite domains for numerical unknowns. Then we encode these discrete domains by propositional values, to obtain a propositional satisfiability problem, for which efficient solvers are available. For finding matrix interpretations, this is a successful method, even with surprisingly small domains that can be represented by at most 5 bits. In the experiments we only considered matrix interpretations over the natural numbers.

Table 1 indicates the number of systems where polynomial upper bounds on the derivational complexity could be established. The rows in the table correspond to different approaches to polynomially bound the derivational complexity of rewrite systems and the columns give the degree of these polynomials. Triangular matrix interpretations [14, 15, 19] serve as a reference (here the degree of the polynomial is determined by the number of ones in the main diagonal).

The first three rows (and the accumulated results in row four) operate on the component-wise maximum matrix. While in theory Theorem 4 allows to establish strictly more polynomial bounds than triangular matrices, in practice the constraints are harder to solve, resulting in a worse overall score. A similar argument holds when comparing Theorems 4 and 5, i.e., explaining the better performance of the latter for low bounds but worse result for polynomial bounds. The accumulative score of the first three rows justifies each of the approaches.

The criteria in the next two rows are not restricted to the growth of the maximum matrix. In the former row no explicit bounds on the degree are added to the constraints (so the dimension of the matrices dominates the degree of polynomial growth) while in the latter row upper bounds are explicitly added to the search. Somehow surprisingly, the accumulated score of these two rows is noticeable larger than the score for each row on its own.

The final row in Table 1 shows that the methods from the first block do not significantly increase the power of the second block (in our experiments). However, the criteria from the first block are also suitable for matrix interpretations over $\mathbb{R}$, while the methods involved in the second block become undecidable over these domains.

7 Concluding Remarks

In this paper we employed joint spectral radius theory to unify as well as clarify the different approaches to obtain upper bounds on the derivational complexity of rewrite systems from compatible matrix interpretations. Our results are not limited to the study of *derivational* complexity, but can also be employed in the context of *runtime* complexity of rewrite systems [7].

It remains to be seen whether joint spectral radius theory will advance implementations (for matrix interpretations over $\mathbb{Q}$ and $\mathbb{R}$). For matrices with rational or real entries the joint spectral radius is not computable in general. In the future we will investigate whether good approximations can be obtained to improve the complexity bounds obtained from matrix interpretations over $\mathbb{Q}$ and $\mathbb{R}$ [5, 22].

We conclude the paper by observing that matrix interpretations are incomplete when it comes to establishing polynomial derivational complexity. This shows that new ideas are necessary to obtain a complete characterization of TRSs with polynomial derivational complexity (cf. RTA open problem #107).[6]

Lemma 34. *There exists a TRS with linear derivational complexity that is compatible with a matrix interpretation but not with a polynomially bounded one.*

Proof. To prove this result we extend the TRS $\mathcal{R}$ from Example 7 with the rule $\mathsf{b}(x) \to \mathsf{g}(x)$. The resulting TRS $\mathcal{S}$ has linear derivational complexity since it is match-bounded by 3 [6]. To obtain a matrix interpretation compatible with $\mathcal{S}$, we change the interpretation of b in the proof of Lemma 8 to

$$\mathsf{b}_{\mathcal{M}}(\boldsymbol{x}) = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 1 \\ 0 & 1 & 1 \end{pmatrix} \boldsymbol{x} + \begin{pmatrix} 4 \\ 4 \\ 3 \end{pmatrix}$$

That there cannot exist a polynomially bounded matrix interpretation for $\mathcal{S}$ follows from Corollary 31 and the proof of Lemma 8 since $B \geqslant \max(I_n, G)$ and hence entries in B^k grow exponentially. □

Since the TRS in the proof of the above lemma is a string rewrite system and match-bounded, this lemma also answers a question by Waldmann, cf. [19, Section 10]: *Do all match-bounded string rewrite systems have a polynomially (even linearly) bounded matrix interpretation?* Here linearly bounded refers to Definition 17.

[6] `http://rtaloop.mancoosi.univ-paris-diderot.fr/problems/107.html`

References

1. Baader, F., Nipkow, T.: Term Rewriting and All That. Cambridge University Press (1998)
2. Bell, J.: A gap result for the norms of semigroups of matrices. LAA 402, 101–110 (2005)
3. Droste, M., Kuich, W., Vogler, H. (eds.): Handbook of Weighted Automata. Springer Verlag (2009)
4. Endrullis, J., Waldmann, J., Zantema, H.: Matrix interpretations for proving termination of term rewriting. JAR 40(2–3), 195–220 (2008)
5. Gebhardt, A., Hofbauer, D., Waldmann, J.: Matrix evolutions. In: WST 2007. pp. 4–8 (2007)
6. Geser, A., Hofbauer, D., Waldmann, J., Zantema, H.: On tree automata that certify termination of left-linear term rewriting systems. I&C 205(4), 512–534 (2007)
7. Hirokawa, N., Moser, G.: Automated complexity analysis based on the dependency pair method. In: IJCAR 2008. LNCS, vol. 5195, pp. 364–379 (2008)
8. Horn, R., Johnson, C.: Matrix Analysis. Cambridge University Press (1990)
9. Jungers, R.M., Protasov, V., Blondel, V.D.: Efficient algorithms for deciding the type of growth of products of integer matrices. LAA 428(10), 2296–2311 (2008)
10. Jungers, R.: The Joint Spectral Radius: Theory and Applications. Springer Verlag (2009)
11. Koprowski, A., Waldmann, J.: Arctic termination ... below zero. In: RTA 2008. LNCS, vol. 5117, pp. 202–216 (2008)
12. Moser, G.: Proof Theory at Work: Complexity Analysis of Term Rewrite Systems. Habilitation thesis, University of Innsbruck (2009)
13. Moser, G., Schnabl, A.: The derivational complexity induced by the dependency pair method. LMCS (2011), to appear. Available at `http://arxiv.org/abs/0904.0570`
14. Moser, G., Schnabl, A., Waldmann, J.: Complexity analysis of term rewriting based on matrix and context dependent interpretations. In: FSTTCS 2008. LIPIcs, vol. 2, pp. 304–315 (2008)
15. Neurauter, F., Zankl, H., Middeldorp, A.: Revisiting matrix interpretations for polynomial derivational complexity of term rewriting. In: LPAR-17. LNCS, vol. 6397, pp. 550–564 (2010)
16. Rose, H.E.: Linear Algebra: A Pure Mathematical Approach. Birkhäuser (2002)
17. Sakarovitch, J.: Elements of Automata Theory. Cambridge University Press (2009)
18. Terese: Term Rewriting Systems, Cambridge Tracts in Theoretical Computer Science, vol. 55. Cambridge University Press (2003)
19. Waldmann, J.: Polynomially bounded matrix interpretations. In: RTA 2010. LIPIcs, vol. 6, pp. 357–372 (2010)
20. Weber, A., Seidl, H.: On the degree of ambiguity of finite automata. TCS 88(2), 325–349 (1991)
21. Zankl, H., Korp, M.: Modular complexity analysis via relative complexity. In: RTA 2010. LIPIcs, vol. 6, pp. 385–400 (2010)
22. Zankl, H., Middeldorp, A.: Satisfiability of non-linear (ir)rational arithmetic. In: LPAR-16. LNCS, vol. 6355, pp. 481–500 (2010)