Growth of Bilinear Maps

Dissertation
zur Erlangung des Grades eines
Doktors der Naturwissenschaften
(Dr. rer. nat.)

am Fachbereich Mathematik und Informatik
der Freien Universität Berlin

vorgelegt von

Vuong Bui

Berlin 2023

Erstgutachter/in: *Prof. Dr. Günter Rote*,
Institut für Informatik, Freie Universität Berlin,
Takustraße 9, 14195 Berlin, Germany

Zweitgutachter/in: *Assoc. Prof. Dr. Matthieu Rosenfeld*,
LIRMM, Université de Montpellier,
161 Rue Ada, 34095 Montpellier, France

Tag der Disputation: *18/12/2023*

Selbstständigkeitserklärung

Name: *Bui*
Vorname: *Vuong*

Ich erkläre gegenüber der Freien Universität Berlin, dass ich die vorliegende Dissertation selbstständig und ohne Benutzung anderer als der angegebenen Quellen und Hilfsmittel angefertigt habe. Die vorliegende Arbeit ist frei von Plagiaten. Alle Ausführungen, die wörtlich oder inhaltlich aus anderen Schriften entnommen sind, habe ich als solche kenntlich gemacht. Diese Dissertation wurde in gleicher oder ähnlicher Form noch in keinem früheren Promotionsverfahren eingereicht.

Mit einer Prüfung meiner Arbeit durch ein Plagiatsprüfungsprogramm erkläre ich mich einverstanden.

Datum: *24/10/2023* Unterschrift: ______________________

Vuong Bui

gửi tặng chính tôi, một ngày nào đó nhìn lại con đường dài đã qua,
và gửi tặng những ai đã đồng hành trên con đường đó...

Abstract

We study a problem that is algebraic in nature but has certain applications in graph theory. It can be seen as a generalization of the joint spectral radius.

Given a bilinear map $* : \mathbb{R}^d \times \mathbb{R}^d \to \mathbb{R}^d$ and a vector $s \in \mathbb{R}^d$, both with nonnegative coefficients and entries, among an exponential number of ways to combine n instances of s using $n-1$ applications of $*$, we are interested in the largest possible entry in a resulting vector. Let $g(n)$ denote this value, the asymptotic behaviour of $g(n)$ is investigated through the growth rate

$$\lambda = \limsup_{n\to\infty} \sqrt[n]{g(n)}.$$

It is known that checking $\lambda \leq 1$ is undecidable, as a consequence of the corresponding fact for the joint spectral radius. However, efficient algorithms are available to compute it exactly in certain cases, or approximate it to any precision in general. Furthermore, when the vector s is positive, there exists some r so that

$$\mathrm{const}\, n^{-r} \lambda^n \leq g(n) \leq \mathrm{const}\, n^r \lambda^n.$$

It means λ is actually a limit when $s > 0$. However, checking if this is the case in general is also undecidable. Some types of patterns for optimal combinations are proposed and studied as well, with some connections to the finiteness property of a set of matrices.

The techniques that are used for our problem can be applied well for the joint spectral radius, and they produce some stronger results by even simpler arguments. For example, if $\|\Sigma^n\|$ denotes the largest possible entry in a product of n matrices drawn from a finite set Σ of nonnegative matrices, whose joint spectral radius is denoted by $\rho(\Sigma)$, then there exists some r so that

$$\mathrm{const}\, n^r \rho(\Sigma)^n \leq \|\Sigma^n\| \leq \mathrm{const}\, n^r \rho(\Sigma)^n.$$

Acknowledgements

The author would express his gratitude to Günter Rote, the thesis advisor, for his various remarks during the tough revision of the thesis. Needless to say, many correspondences with him during the years have helped shaping the direction of the research. That is, the thesis may get harder to finish or may finish with some different direction without his support.

The chapter on decidability would have not existed if there were not correspondences with Matthieu Rosenfeld, for which the author is grateful. His kind sharing of a manuscript during the last phase of the preparation of the thesis has suggested several questions that the author was not even aware of at first. It would be a great pity if this last piece on the computational aspects of the problem were left out. (P.S. The author is also grateful to Matthieu Rosenfeld for being the second examiner of the thesis and his various useful remarks on the presentation of the work.)

The final thankful words are given to the Graduiertenkolleg “Facets of Complexity” for, amongst other things, the generous financial support. The author also appreciates the working group “Theoretische Informatik” of Freie Universität Berlin for pleasant mensa–coffee hours with the cheerful colleagues. There is a nice office with (a) squirrel(s) often visiting the tree near the window and a view to the small flower garden beneath.

Contents

CHAPTER 1

Introduction

We study a problem that is algebraic in nature but has certain applications in graph theory. It can be seen as a generalization of the notion of joint spectral radius.

Suppose we are given a binary operator $*$ and an operand s. For any n, there are an exponential number of ways to combine n instances of s using $n-1$ applications of $*$. For example, when $n=4$, the 5 combinations are $s*(s*(s*s))$, $s*((s*s)*s)$, $(s*s)*(s*s)$, $(s*(s*s))*s$, $((s*s)*s)*s$. In fact, the number of combinations is the $(n-1)$-th Catalan number. As the operator is not always commutative or associative, the results may vary, depending on the way we group the brackets. However, in certain situations we might still expect that the largest magnitude over all the combinations does not grow too arbitrarily, and even follows some kind of growth rate. A problem of this type was posed by Rote [**1**], where $*$ is a bilinear operator in $\mathbb{R}^d \times \mathbb{R}^d \to \mathbb{R}^d$ with nonnegative coefficients and s is a nonnegative vector in $\mathbb{R}^d$. Since every norm is equivalent up to a constant factor, we can choose any norm to be the magnitude of a resulting vector. We let the norm be the largest entry of the vectors for convenience, due to the nonnegativity of the vector s and the coefficients of the operator $*$. In other words, we are interested in the largest possible entry $g(n)$ in a vector obtained by combining n instances of s. In many cases, we have the limit

$$\lambda = \lim_{n\to\infty} \sqrt[n]{g(n)},$$

where λ is called the *growth rate* of the *bilinear system* $(*, s)$.

Let us give some examples of the limit. They are actually graph theoretical problems.

The first example: Consider a rooted binary tree T, a pruned tree of T is a tree obtained from T by removing zero or more subtrees. Let $f(n)$ be the maximum number of pruned trees of a tree T with n leaves. The function $f(n)$ can be defined recursively by $f(1) = 1$ and for $n \geq 2$,

$$f(n) = 1 + \max_{1\leq m\leq n-1} f(m)f(n-m). \tag{1.1}$$

We can express the function in a different formulation. For the vector $s = (1, 1)$ and the bilinear function $* : \mathbb{R}^2 \times \mathbb{R}^2 \to \mathbb{R}^2$ so that

$$\begin{pmatrix} x_1 \\ x_2 \end{pmatrix} * \begin{pmatrix} y_1 \\ y_2 \end{pmatrix} = \begin{pmatrix} x_1y_1 + x_2y_2 \\ x_2y_2 \end{pmatrix},$$

the function $g(n)$ of the bilinear system $(*, s)$ is identical to $f(n)$.

The function $f(n)$ was investigated for a different purpose in [**2**] where the growth rate $\lambda = \lim_{n\to\infty} \sqrt[n]{f(n)}$ was shown to be $1.502836801\ldots$. A combinatorial argument shows that the growth rate is the same as the rate $\lim_{m\to\infty}(a_m)^{1/2^m}$ of the doubly exponential sequence a_m where $a_0 = 1$ and $a_m = 1 + a_{m-1}^2$ for $m \geq 1$. Actually, the limit is expressed in a more explicit way in [**3**]:

$$\lambda = \lim_{m\to\infty} (a_m)^{1/2^m} = \exp\left(\sum_{i\geq 1} \frac{1}{2^i} \log\left(1 + \frac{1}{a_i^2}\right)\right). \tag{1.2}$$

The second example: The original motivation of the growth rate of the bilinear system was to study the maximum number of minimal dominating sets in a tree of n leaves. This number is actually of the same order as $g(n)$ for vector $s = (0, 1, 0, 0, 0, 1)$ and the operator $*$ so that

$$\begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ x_5 \\ x_6 \end{pmatrix} * \begin{pmatrix} y_1 \\ y_2 \\ y_3 \\ y_4 \\ y_5 \\ y_6 \end{pmatrix} = \begin{pmatrix} x_1y_1 + x_1y_4 + x_1y_6 + x_2y_6 + x_3y_6 \\ x_2y_4 \\ x_2y_1 + x_3y_1 + x_3y_4 \\ x_4y_1 + x_4y_2 + x_4y_4 + x_4y_5 + x_6y_1 + x_6y_2 \\ x_5y_4 + x_5y_5 + x_6y_3 \\ x_6y_4 + x_6y_5 \end{pmatrix}. \tag{1.3}$$

The relation between minimal dominating sets and the setting of $*, s$ is explained in detail in the original source [**1**] using dynamic programming. The growth rate λ curiously has some magic numbers:

$$\lambda = \lim_{n\to\infty} \sqrt[n]{g(n)} = \sqrt[13]{95}. \tag{1.4}$$

Note that while the proof of the first example, which spans some four pages, is already nontrivial enough, the proof of the second example even needs the assistance of a computer with a method we may call the "polytope method". However, the algebraic nature of the constants in (1.2) and (1.4) does not seem to suggest that the latter example is a more complicated one.

Other examples: The setting of a bilinear operator and a vector was applied by Rosenfeld to address the number of different types of dominating sets, perfect codes, different types of matchings, and maximal irredundant sets in a tree. The readers can check [**4**, Section 5] for this rich set of applications. One can find an application in graphs other than trees in [**2**] where the maximum number of cycles in an outerplanar graph is studied using the function $f(n)$ in (1.1). We believe the flexibility of the setting allows applications in more remote fields.

In general, the limit of $\sqrt[n]{g(n)}$ is not always guaranteed to exist. Rosenfeld [**5**] suggests to define the growth rate by

$$\lambda = \limsup_{n\to\infty} \sqrt[n]{g(n)}.$$

From now on, by the *growth rate* of the bilinear system we mean this limit superior λ.

Decidability. The problem of checking if $\lambda \le 1$ for a given system $(*, s)$ is shown to be undecidable in [**5**] by reducing the problem of joint spectral radius. A simpler reduction also using the joint spectral radius is given in Chapter 6.

Let JSR denote the problem of checking if the joint spectral radius $\rho \le 1$, the undecidability of JSR is actually proved in [**6**] by

$$HP \le \cdots \le PFAE \le JSR,$$

where HP denotes the halting problem and $PFAE$ denotes the problem of probabilistic finite automaton emptiness. (We denote $A \le B$ if Problem A can be reduced to Problem B.) There are several problems that can be filled into the place of the above dots.

Note that all these problems are actually Turing equivalent since we have a reduction from JSR to HP by the joint spectral radius theorem, which states that for a finite set Σ of matrices we have

$$\rho(\Sigma) = \sup_n \max_{A_1,\dots,A_n\in\Sigma} \sqrt[n]{\rho(A_1 \dots A_n)},$$

where ρ denotes both the joint spectral radius and the ordinary spectral radius, depending on the argument. Indeed, we just run the program that looks for a sequence of matrices

$A_1, \ldots, A_n$ in Σ with increasing n whose product has the spectral radius greater than 1. The program does not stop if and only if $\rho(\Sigma) \leq 1$. Note that the problem of checking $\rho(A) \leq 1$ for the ordinary spectral radius (SR) is decidable by Tarski's method [**7**].

On the other hand, a formula of λ in Chapter 5 allows a reduction from checking $\lambda \leq 1$ to the halting problem. The formula can be written in a form that looks similar to the joint spectral radius theorem:

$$\lambda = \sup_n \max_{\substack{\text{linear pattern } P \\ |P|=n}} \sqrt[n]{\rho(M(P))}.$$

We do not explain the terms in detail, which is done in Chapter 5, but we may say roughly that a linear pattern is a sequence x_n for $n = 0, 1, 2, \ldots$ so that $x_0 = s$ and x_n for $n \geq 1$ is a combination of some instances of s and precisely *one* instance of x_{n-1}. The notation $|P|$ denotes the number of instances of s and the matrix $M = M(P)$ represents the linear relation $x_n = M x_{n-1}$ for every $n \geq 1$. The reduction from checking $\lambda \leq 1$ to the halting problem is done similarly to the one for the problem of checking if $\rho(\Sigma) \leq 1$.

Let us call the problem of checking if $\lambda \leq 1$ the problem of the growth rate of a bilinear system ($GRBS$). We have established the relation of the previous problems to $GRBS$. An interesting point is that using reductions of the same kind we can show that the problem of checking the growth rate does not become harder when multiple operators and multiple starting vectors are allowed. This was first remarked by Rosenfeld [**5**], see Chapter 6 for discussion. Let us call it the problem of the *joint* growth rate of a bilinear system ($JGRBS$). In total, we have

$$SR < HP = PFAE = JSR = GRBS = JGRBS,$$

where $A < B$ means $A \leq B$ but we do not have $B \leq A$, and $A = B$ means $A \leq B$ and $B \leq A$, that is each of A, B is reducible to the other.

Note that we still do not yet have a natural reduction from $GRBS$ to JSR as the one for proving $JSR \leq GRBS$. Such a reduction is very desirable and it would have interesting consequences, as discussed in Chapter 7.

The positive setting. When the vector s is strictly positive instead of being only nonnegative, it was shown in [**8**] that the limit is always guaranteed to exist:

$$\lambda = \lim_{n \to \infty} \sqrt[n]{g(n)}.$$

In fact, Chapter 5 provides a simpler proof than the one in [**8**].

When the requirement is not met, there is chance that the limit does not exist (also see [**8**] for counterexamples, which are presented in Chapter 2). The interesting point is that the problem of checking if the limit exists when s is only nonnegative is undecidable, as in Chapter 6. It can be seen as a corollary of the undecidability of the problem of checking $\lambda \leq 1$.

Let us call the setting with nonnegative s the nonnegative setting and the setting with positive s the positive setting. In fact, it is also interesting to treat the problem when there is no condition on the signs of the entries and the coefficients, and $g(n)$ is the largest possible norm of any vector obtained from combining n instances of s for some appropriate norm. Note that the growth rate is independent of the chosen norm as two norms are in a constant factor of each other. Let us call this setting the general setting. However, we almost never treat the general setting in this work because the techniques we use depend on the nonnegativity. Moreover, the nonnegative setting seems to cover most of the applications so far, see [**4**, Section 5] for some instances.

The growth rate λ is computable for the nonnegative setting as shown in Chapter 5. By computable we mean it is possible to generate converging sequences of upper bounds and lower bounds, that is we can approximate λ to any precision. Furthermore, we have a fairly good bound for $g(n)$ in the positive setting by the following result, which is also in Chapter 5.

THEOREM. *There are some positive constants a, a' and some degrees r, r' so that they are computed from $*, s$ and for every n,*

$$an^{-r}\lambda^n \leq g(n) \leq a'n^{r'}\lambda^n.$$

A corollary is the following estimate of the growth rate λ: For any n, we have

$$\sqrt[n]{\frac{1}{a'}n^{-r'}g(n)} \leq \lambda \leq \sqrt[n]{\frac{1}{a}n^{r}g(n)}. \tag{1.5}$$

Since the ratio $\sqrt[n]{(a'/a)n^{r+r'}}$ between the upper bound and the lower bound converges to 1, we obtain a good bound when we have $g(n)$ for a large enough n.

We give examples where $g(n)$ is of order $n^r\lambda^n$ for some integer $r > 0$ in Chapter 2, where we also conjecture that r is at most 2^{d-1} for the space $\mathbb{R}^d$. However, we have not yet found any example to match the lower bound of order $n^{-r}\lambda^n$. In fact, we believe that $g(n)$ is at least a constant times λ^n. Since for a matrix A there exists a number r so that $\text{const}\, n^r\rho(A)^n \leq \|A^n\| \leq \text{const}\, n^r\rho(A)^n$, where $\rho(A)$ denotes the spectral radius of A, and a similar result also holds for the joint spectral radius of nonnegative matrices in Chapter 4, we ask the following question.

QUESTION 1.1. *Is it true that for every $*, s$ in the positive setting there always exists a number r so that*

$$\text{const}\, n^r\lambda^n \leq g(n) \leq \text{const}\, n^r\lambda^n?$$

The example of pruned trees in (1.1) was confirmed to satisfy $\lambda^{n-\frac{1}{4}} < g(n) < \lambda^n$ for $n \geq 10$ in [**2**, Theorem 2]. The example of minimal dominating sets in (1.3) also satisfies $\text{const}\, \lambda^n \leq g(n) \leq \text{const}\, \lambda^n$ by [**1**, Theorem 1.1].

By (1.5), suppose $\lambda \neq 1$, one can always decide whether $\lambda > 1$ or $\lambda < 1$ (regardless of complexity) since when n is large enough, the value 1 will be to the left or to the right of the small interval containing λ. However, when we are not guaranteed $\lambda \neq 1$, we have the following question.

QUESTION 1.2. *Is the problem of checking if $\lambda \neq 1$ for the positive setting decidable?*

Note that checking $\lambda \neq 1$ for the nonnegative setting is undecidable since one can reduce the problem of checking $\lambda \leq 1$ to it by adding one extra dimension that is always 1. The same trick also applies to the positive setting, that is checking $\lambda \leq 1$ is easier than checking $\lambda \neq 1$ up to decidability.

Another point is that the problem of checking $\lambda \neq 1$ is not harder than the halting problem since we can run a program to obtain smaller and smaller intervals containing λ. The program stops when the interval does not contain 1, that is $\lambda \neq 1$. The program never stops otherwise.

An attempt to answer this question is given in Chapter 6 when the problem of checking $\lambda \leq 1$ for the positive setting is shown to be undecidable under the assumption that it is undecidable to check if $\rho(\{A, B\}) \leq 1$ for two *positive* matrices A, B.

Organization of the thesis. Chapter 2 deals with the formal description and gives some easy examples for the readers get a feel of the problem. More complicated examples can be found in Chapters 3 and 4, which present two large classes of problems. Chapter 3 is about the growth of replacements, which is a new and interesting problem on its own. Meanwhile, Chapter 4 is about the joint spectral radius of nonnegative matrices, which is an old and quite established problem. However, we shed some lights to the latter problem under the condition of nonnegativity.

In fact, Chapter 3 and Chapter 4 can be treated as a preparation for the readers to the techniques in Chapter 5, which give a formula and some bounds for the growth rates of bilinear systems. The arguments in all three chapters are similar in one way or another, but the argument gets harder after each chapter. As a consequence, the result gets weaker after each chapter. While Chapter 5 is quite complicated, it is the core of the thesis in terms of techniques.

Chapter 6 actually confirms the fact that the joint spectral radius is an instance of the growth of bilinear maps,[1] which may be not so obvious at first. Several problems for the latter notion are shown to be undecidable, as a consequence of the undecidability of the corresponding problems for the former one. In fact, one can reduce the problem of the growth of replacements to the problem of the joint spectral radius, but it is quite meaningless since the former problem is decidable while the latter one is not. However, it is a sign that it is easier to deal with Chapter 3 than with Chapter 4. The growth of replacements can be even computed precisely in Chapter 3.

One may expect that we can obtain the growth rate by looking at the combinations that follow certain patterns. Chapter 7 discusses different kinds of patterns for which we may or may not obtain the growth rate. An interesting conjecture on the coverage of the growth rates by all the patterns is given there. We also relate the finiteness property of a set of matrices to a kind of patterns called "linear pattern", which is the key tool throughout the work.

Some related publications by the author. The thesis contains the content of several articles by the author. The starting point is [**8**], which is the first study of the problem after the introduction by Rote. We do not present most of the content of [**8**] but replace it with improved approaches. The merit of [**8**] is the introduction of the notion "linear pattern", which is so critical that its analogs can be found in almost every chapter.[2] The study on the growth of replacements is already published in [**9**]. Part of the chapter on the joint spectral radius is published in [**10**]. The remaining content of the thesis can be found in some preprints by the author, which are subjected to changes and improvements.

[1] We sometimes call the problem we are studying the *growth of bilinear maps*. In fact, the title "Growth of bilinear maps" of the thesis is also the title of the paper [**8**], which is the starting point of the research of the author in this subject.

[2] One may even title the thesis "Linear patterns and the growth of bilinear maps".

CHAPTER 2

Definitions and examples

This chapter will give the problem statement formally and present some simple examples, to prepare the readers for other chapters, in particular for Chapter 5. In fact, the problems in Chapters 3 and 4 are stated in their own languages and quite independent from other chapters. The relations to the general problem of the growth of bilinear maps are not presented until Chapter 6. However, the two problems are worth treating for their own interests.

As Fekete's lemma is used from place to place in the thesis, we give its statement in this chapter. A variant of Fekete's lemma is also introduced, which may be useful beyond the thesis.

2.1. Problem statement

We are given a *nonnegative* starting vector $s \in \mathbb{R}^d$ and a bilinear map $* : \mathbb{R}^d \times \mathbb{R}^d \to \mathbb{R}^d$ defined by *nonnegative* coefficients $c_{i,j}^{(k)}$ so that for any k and any vectors x, y, we have

$$(x * y)_k = \sum_{i,j} c_{i,j}^{(k)} x_i y_j.$$

We denote by A_n the set of all the results obtained by applying $n-1$ applications of $*$ to n instances of s, that is: $A_1 = \{s\}$ and for $n \geq 2$,

$$A_n = \bigcup_{1 \leq m \leq n-1} \{x * y : x \in A_m, y \in A_{n-m}\}.$$

The largest entry $g(n)$ over all the resulting vectors can be expressed as

$$g(n) = \max\{v_i : v \in A_n, 1 \leq i \leq d\}.$$

We denote by $g_k(n)$ the largest k-th entry over all the resulting vectors, that is

$$g_k(n) = \max\{v_k : v \in A_n\}.$$

An obvious relation between $g(n)$ and $g_k(n)$ is $g(n) = \max_k g_k(n)$.

Such pair $(*, s)$ is called a *bilinear system*. The following limit superior λ is called the *growth rate of the system*:

$$\lambda = \limsup_{n \to \infty} \sqrt[n]{g(n)}.$$

The growth rate is a well-defined number due to the boundedness of $\sqrt[n]{g(n)}$. Indeed, let s^* be the maximal entry of s and c^* be the maximal coefficient of $*$, the readers can verify by induction that

$$g(n) \leq (s^*)^n (d^2 c^*)^{n-1}.$$

In general, we do not treat the degenerate cases where the starting vector s is zero or all the coefficients of $*$ are zero. Further, we assume that there is no *degenerate dimension*, where a dimension k is said to be degenerate if $g_k(n) = 0$ for every n. In case the assumption is not met, we can safely discard the degenerate dimensions and the involved coefficients without affecting $g(n)$. The readers may take the question of how to check if a given k is a degenerate dimension as an exercise.

Note that the growth rate λ actually depends on the map and the vector of the considered system, but we do not denote it explicitly by $\lambda_{*,s}$ as they are known from context. Other notations are denoted implicitly in the same manner.

On the other hand, as the variable names d, s are so popular, they may be reused for other purposes. If they are used as defined in this section, it will be stated explicitly.

2.2. Some beginning examples

We start with a simple system: Let $s = (1, 1)$ and

$$u * v = (u_1v_2 + u_2v_1, u_2v_2),$$

we show that every combination of n instances of s using $*$ gives the same result $(n, 1)$. It holds for $n = 1$ as the vector is s then. We show that it holds for $n > 1$ provided it holds for smaller numbers than n. Let the combination be $U * V$ where U, V are the combinations of ℓ, m, respectively, instances of s, we have $U = (\ell, 1)$ and $V = (m, 1)$. The resulting vector is

$$U * V = (\ell, 1) * (m, 1) = (\ell + m, 1) = (n, 1).$$

The verification is done by induction. It follows that $g(n) = n$.

Despite being a simple example, we can see that $g(n)$ can be of the same order as a polynomial when $\lambda = 1$. In fact, Theorem 5.3 shows that $g(n)$ cannot be superpolynomial when $\lambda = 1$. We construct examples with polynomials of higher degrees. Let $s = (1, 1, 1)$ and

$$u * v = (u_1v_2 + u_2v_1, u_2v_2, u_1v_1),$$

the largest third entry should be of order n^2. We have already known by the previous example that the first two dimensions of the resulting vector are always $(n, 1)$. Let the combination for $n > 1$ be $U * V$, where U, V are the combinations of ℓ, m, respectively, instances of s. The third dimension will be ℓm. This value is at most $(\ell+m)^2/4 = n^2/4$, where the equality is attained when $\ell = m$. The situation for odd n is not very different and we obtain n^2 as the order of $g(n)$.

Let us give another example with $g(n)$ of order n^2 suggested by Rote. This is actually the puzzle "Splitting the Stacks" in the book "Mathematical Puzzles" by Peter Winkler: We are given a stack of n items and we are allowed to divide a stack into two stacks at any time and get paid the product of the two stacks. The question is what is the most money $f(n)$ we can get for some n? Of course, one should keep dividing stacks until all stacks are of size 1. But the interesting point is that no matter how we divide the stacks, the money we finally get is always $\frac{1}{2}n(n-1)$. Indeed, we can write

$$f(n) = \max_{1 \le m \le n-1} m(n-m) + f(m) + f(n-m) \tag{2.1}$$

for $n \ge 2$, while $f(1) = 0$. One can verify by induction that $f(n) = \frac{1}{2}n(n-1)$ is the answer, and we can get the optimal $f(n)$ for any m. Another way is treating the division of the stacks as a tree of n leaves, we can see that the final money we get is the total number of ways we pair two leaves, which is obviously $\binom{n}{2} = n(n-1)$ regardless of the structure of the tree. In the setting of a bilinear system, one can write $s = (0, 1, 1)$ and

$$x * y = (x_1y_3 + x_3y_1 + x_2y_2, x_2y_3 + x_3y_2, x_3x_3). \tag{2.2}$$

The third dimension is always 1, the second dimension presents the number of instances of s, and the first dimension is the money we can get. It follows that $f(n) = g(n) = O(n^2)$.

In case one may find (2.2) looks a bit hard to track, let us write it in a different way:

$$x * y = (x_1 + y_1 + x_2y_2, x_2 + y_2),$$

which resembles (2.1) better. Although it is no longer a bilinear map due to the mixture of linear and bilinear forms, one can see that there is no more expressiveness when introducing linear forms, due to the constant dimension.

As we have increased the order of $g(n)$ from $O(n)$ to $O(n^2)$ when we consider the examples in $\mathbb{R}^3$ instead of $\mathbb{R}^2$, a false impression would be that we can increase the degree of the polynomial by at most 1 when we add one dimension. The truth is that we can double the degree. Indeed, let $s = (1,1,1,1)$ and $u * v = (u_1v_2 + u_2v_1, u_2v_2, u_1v_1, u_3v_3)$, we leave the verification that n^4 is the order of $g(n)$ to the readers as it is similar to the previous verification. Inspired by the construction, we propose the conjecture that for the space of $\mathbb{R}^d$, we have

$$g(n) \leq \operatorname{const} n^{2^{d-1}} \lambda^n.$$

The corresponding bound for nonnegative matrices A is $\|A^n\| \leq \operatorname{const} n^{d-1}\lambda^n$, as one can see in Chapter 4. It is interesting that $d-1, 2^{d-1}$ are linear and exponential, in correspondence to linear and bilinear maps.

As the growth rate for the problem of minimal dominating sets is $\sqrt[13]{95}$, Rote asked in a personal communication whether it is true that if λ is algebraic then λ is a root of some number. It turns out that it is not always the case by the following example, which is related to the Fibonacci sequence. It is also the least trivial example in this chapter.

Theorem 2.1. *If $s = (1,1)$ and*

$$x * y = (x_1y_2 + x_2y_1, x_1y_2),$$

then the growth rate λ is the golden ratio $\phi = \frac{1+\sqrt{5}}{2}$. In particular, $g_1(n) = F_{n+1}$ and $g_2(n) = F_n$, where F_n is the Fibonacci sequence with $F_1 = F_2 = 1$.

It can be seen that $g_1(n) \geq F_{n+1}$ and $g_2(n) \geq F_n$ for every $n \geq 1$ since the vector (F_{n+1}, F_n) is the resulting vector of the sequence v_n with $v_1 = s$ and $v_n = v_{n-1} * s$ for $n \geq 2$ (the sequence is $s, s*s, (s*s)*s, ((s*s)*s)*s, (((s*s)*s)*s)*s, \dots$).

In order to show that they are also the upper bounds, we prove the following lemma.

Lemma 2.2. *Let $\{F_n\}_{n\geq 0}$ be the Fibonacci sequence with $F_0 = 0, F_1 = 1, F_2 = 1$, then the inequalities*

$$\begin{aligned} F_pF_{q-1} + F_{p-1}F_q &\leq F_{p+q-1}, \\ F_pF_q &\leq F_{p+q-1} \end{aligned}$$

hold for every $p, q \geq 1$.

Proof. The conclusion holds for any $(p,q) \in (\{1,2\} \times \mathbb{N}^+) \cup (\mathbb{N}^+ \times \{1,2\})$, i.e. one of the four conditions $p = 1$, $p = 2$, $q = 1$, $q = 2$ holds.

For the first inequality, if $p = 1$ (similarly for $q = 1$), then the inequality is equivalent to $F_{q-1} \leq F_q$. If $p = 2$ (similarly for $q = 2$), then it is equivalent to $F_{q-1} + F_q \leq F_{q+1}$.

For the second inequality, if $p = 1$ (similarly for $q = 1$), then the inequality is equivalent to $F_q \leq F_q$. If $p = 2$ (similarly for $q = 2$), then it is equivalent to $F_q \leq F_{q+1}$.

We prove the lemma by induction. For $p \geq 3, q \geq 3$, suppose the inequalities hold for any $(p', q') \in \{p-1, p-2\} \times \{q-1, q-2\}$. We show that they also hold for (p,q).

Indeed,

$$\begin{aligned}
F_pF_{q-1} + F_{p-1}F_q &= (F_{p-2} + F_{p-1})(F_{q-3} + F_{q-2}) + (F_{p-3} + F_{p-2})(F_{q-2} + F_{q-1}) \\
&= (F_{p-2}F_{q-3} + F_{p-3}F_{q-2}) + (F_{p-2}F_{q-2} + F_{p-3}F_{q-1}) \\
&\quad + (F_{p-1}F_{q-3} + F_{p-2}F_{q-2}) + (F_{p-1}F_{q-2} + F_{p-2}F_{q-1}) \\
&\leq F_{p+q-5} + F_{p+q-4} + F_{p+q-4} + F_{p+q-3} \\
&= F_{p+q-3} + F_{p+q-2} \\
&= F_{p+q-1}
\end{aligned}$$

and

$$\begin{aligned}
F_pF_q &= (F_{p-2} + F_{p-1})(F_{q-2} + F_{q-1}) \\
&= F_{p-2}F_{q-2} + F_{p-2}F_{q-1} + F_{p-1}F_{q-2} + F_{p-1}F_{q-1} \\
&\leq F_{p+q-5} + F_{p+q-4} + F_{p+q-4} + F_{p+q-3} \\
&= F_{p+q-3} + F_{p+q-2} \\
&= F_{p+q-1}.
\end{aligned}$$

By induction, the inequalities hold for every $p, q \geq 1$. □

Now the verification for the upper bounds of $g_1(n)$ and $g_2(n)$ becomes clear.

Proof of Theorem 2.1. The upper bounds by Fibonacci numbers hold trivially for $n = 1$. For higher n, if $g_1(n)$ is the first entry of $U * V$ where U, V are the combinations of p, q, respectively, instances of s, then we have the same bounds:

$$g_1(n) \leq g_1(p)g_2(q) + g_2(p)g_1(q) = F_{p+1}F_q + F_pF_{q+1} \leq F_{p+q+1} = F_{n+1},$$

and

$$g_2(n) \leq g_1(p)g_2(q) = F_{p+1}F_q \leq F_{p+q} = F_n.$$

Being both lower bounds and upper bounds, we have $g_1(n) = F_{n+1}$ and $g_2(n) = F_n$. □

A large class of examples is also given in Chapter 3, where we allow only one summand in the representation of the operator $*$ in the way that for each k there exist some i, j so that $(x * y)_k = x_iy_j$. We call it the problem of replacements by this condition. The growth rate is always a root of some number in this case.

Let us quickly give an example: Let $s = (1, 2, 3, 4)$ and

$$x * y = (x_2y_3, x_3y_4, x_4y_1, x_1y_2),$$

one can verify that the growth rate is $\sqrt[3]{4 \cdot 4 \cdot 3} = \sqrt[3]{48}$, by applying Theorem 3.2 in Chapter 3.

It is possible to give examples of more complicated algebraic growth rates by simulating linear recurrences. Let x_n be a sequence so that the first m elements $x_1, \ldots, x_m$ are given in advance and for $n > m$, we have

$$x_n = \sum_{i=1}^{m} a_i x_{n-i},$$

for some coefficients $a_1, \ldots, a_m$. In $\mathbb{R}^{m+1}$, we can construct an example where the first m dimensions of a combination of n instances of s, if not all zero, are $x_{n+m-1}, x_{n+m-2}, \ldots, x_n$.

Indeed, consider the bilinear system

$$s = \begin{pmatrix} x_m \\ x_{m-1} \\ x_{m-2} \\ \vdots \\ x_1 \\ 1 \end{pmatrix}, \qquad u * v = \begin{pmatrix} \sum_{i=1}^m a_i u_i v_{m+1} \\ u_1 v_{m+1} \\ u_2 v_{m+1} \\ \vdots \\ u_{m-1} v_{m+1} \\ 0 \end{pmatrix}.$$

We can see that $U * V$ is a zero vector for any combinations U, V with V containing at least 2 instances of s. Indeed, the last dimension of V is zero, which makes the first m dimensions of $U * V$ all zero as well. Therefore, if a combination has $U * V$ as a subcombination, then the vector is zero. In other words, nonzero resulting vectors can be found among the combinations $s, s * s, (s * s) * s, ((s * s) * s) * s, (((s * s) * s) * s) * s, \ldots$ only. The conclusion on the simulation of the linear recurrence easily follows.

Note that the construction still fits the requirement of the positive setting provided that $a_1, \ldots, a_m$ are nonnegative and $x_1, \ldots, x_m$ are positive. When we extend the signs of the values to the general setting, more algebraic rates can be obtained since the growth rates of linear recurrences can be characterised as the solutions of polynomials.

While the previous example is a nice way to simulate a linear recurrence, it is not the most straightforward way to produce algebraic growth rates. We can use the observation that matrix multiplication is a special case of a bilinear map. Given a square matrix M in $\mathbb{R}^d$, we embed it into the space $\mathbb{R}^{d^2}$ as the starting vector and let $* : \mathbb{R}^{d^2} \times \mathbb{R}^{d^2} \to \mathbb{R}^{d^2}$ be the bilinear map corresponding to the matrix multiplication. As the multiplication for matrices is similar to the one for numbers in the sense that it is associative, the result of every combination for any n contains the entries of M^n. Therefore, we have produced a bilinear system for any algebraic root that is the spectral radius of a matrix M.

The previous examples suggest some general questions.

Question 2.3. *What is the space of all possible functions $g(n)$?*

Question 2.4. *Suppose the entries of s and the coefficients of $*$ are all integers. What is the set of all possible growth rates?*

The readers may start with whether or not we can produce a bilinear systems of growth rate e or π from integer entries and coefficients. It must be very surprising if we can obtain one.

Another issue is that the algebraic growth rates known so far are the *dominating* roots of equations. Can we obtain the non-dominating ones? For example, can we obtain the growth rate $\frac{\sqrt{5}-1}{2}$ instead of the golden ratio $\frac{\sqrt{5}+1}{2}$?

As the growth rate λ is a limit in the positive setting, it makes sense to give examples with no limit when the system is not in the positive setting. When $s = (1, 0)$ and $x * y = (x_2 y_2, x_1 y_1)$, we have $g(n) = 0$ if 3 divides n and $g(n) = 1$ otherwise. In particular, for $k \in \{1, 2\}$, we have $g_k(n) = 1$ if $n \equiv k \pmod 3$ and $g_k(n) = 0$ otherwise. One can easily verify by induction. In the general setting, let us choose an appropriate norm for $g(n)$, say the greatest absolute value of an entry (i.e. the maximum norm). If some entries of s are allowed to be negative, let us consider the system with $s = (1, -1, 1)$ and $x * y = (x_1 y_1, x_2 y_2, 3x_1 y_3 + 3x_2 y_3)$. We have $g(n) = 1$ for even n and $g(n) = 6^{\frac{n-1}{2}}$ for odd n. Indeed, the first entry of any resulting vector is 1 while the second entry is 1 for even n and is -1 for odd n. If x is the result of a combination of an odd number of instances of s, then $(x * y)_3 = 0$ for any y. When the number of instances is even, we have $(x * y)_3 = 6y_3$. It follows that $g_3(n) = 0$ for even n and $g_3(n) = 6^{\frac{n-1}{2}}$ for odd n. One can easily verify by induction. If some coefficients of $*$ are allowed to be negative, let us consider the system with $s = (1, 1, 1)$ and $x * y = (x_1y_1, -x_2y_2, 3x_1y_3 - 3x_2y_3)$. We have the same $g(n)$, that is $g(n) = 1$ for even n and $g(n) = 6^{\frac{n-1}{2}}$ for odd n. The readers can apply the same method as the previous example.

For examples of different natures, we refer the readers to Chapters 3 and 4 for two large classes of problems.

2.3. Fekete's lemma

Fekete's lemma is used in many places in the thesis. We give its statement and a variant that is also useful.

Lemma (Fekete's lemma). *Given a supperadditive sequence a_n for $n = 1, 2, \dots,$ that is $a_{m+n} \geq a_m + a_n$ for any m, n, the following limit exists and can be expressed as*

$$\lim_{n\to\infty} \frac{a_n}{n} = \sup_n \frac{a_n}{n}.$$

Note that ∞ is treated as a valid limit throughout the thesis for generality. However, we deal with bounded sequences most of the times.

We often use Fekete's lemma in the form for supermultiplicative sequences. A sequence a_n for $n = 1, 2, \dots$ is said to be supermultiplicative if $a_{m+n} \geq a_m a_n$ for any m, n. If the sequence is strictly positive, we can take the logarithm and apply Fekete's lemma to get

$$\lim_{n\to\infty} \sqrt[n]{a_n} = \sup_n \sqrt[n]{a_n}.$$

However, sometimes the sequence we treat is not strictly positive but only nonnegative. Therefore, we introduce the following variant of Fekete's lemma.

Lemma 2.5. *Given a nonnegative sequence a_n for $n = 1, 2, \dots,$ if the sequence is supermultiplicative, then the subsequence of all positive $\sqrt[n]{a_n}$, if nonempty, converges to $\sup_n \sqrt[n]{a_n}$.*

In particular, if the limit is θ, we can conclude that $a_n \leq \theta^n$ for every n.

Proof. We suppose the subsequence $\{a_n : a_n > 0\}$ is not empty. If there is any m such that $a_m > 0$, then the subsequence $\{a_n : a_n > 0\}$ is infinite ($a_{mt} > 0$ for every $t \geq 1$).

It is obvious by definition that $\limsup\{\sqrt[n]{a_n} : a_n > 0\} \leq \sup_n \sqrt[n]{a_n}$. To finish the proof, it remains to show $\liminf\{\sqrt[n]{a_n} : a_n > 0\} \geq \sup_n \sqrt[n]{a_n}$.

Consider any positive integer q. Let R be the set of integers r $(0 \leq r < q)$ such that there exists some m_r with $m_r \equiv r \pmod q$ and $a_{m_r} > 0$. For each $r \in R$, we denote by m_r the smallest such number.

For every n such that $a_n > 0$, if $n \equiv r \pmod q$, then $r \in R$. Consider the representation $n = pq + m_r$, we obtain

$$\sqrt[n]{a_n} \geq \sqrt[n]{(a_q)^p a_{m_r}}.$$

The right hand side converges to $\sqrt[q]{a_q}$ when $n \to \infty$ since m_r is bounded.

As the lower bound holds for every q, we have shown the lower bound $\sup_n \sqrt[n]{a_n}$ for the limit inferior and finished the proof. □

It should be noted that the sequence $\sqrt[n]{a_n}$ itself may not converge. For example, $a_n = 1$ for even n and $a_n = 0$ for odd n. The situation is more pleasant with submultiplicative sequences.

LEMMA 2.6. *Given a nonnegative sequence a_n for $n = 1, 2, \dots,$ if the sequence is submultiplicative, that is $a_{m+n} \le a_m a_n$ for any m, n, then the following limit exists and can be expressed as*

$$\lim_{n\to\infty} \sqrt[n]{a_n} = \inf_n \sqrt[n]{a_n}.$$

In particular, if the limit is θ, we can conclude that $a_n \ge \theta^n$ for every n.

PROOF. If $a_m = 0$ for some m, then $a_n = 0$ for every $n \ge m$, and the lemma trivially holds in this case. If the whole sequence is positive, then the supermultiplicative sequence $1/a_n$ satisfies

$$\lim_{n\to\infty} \sqrt[n]{\frac{1}{a_n}} = \sup_n \sqrt[n]{\frac{1}{a_n}}.$$

The conclusion follows. □

Throughout the work, we may deal with some forms of weak supermultiplicativity and submultiplicativity. An example is the existence of a constant K so that for every m, n we have $a_{m+n} \ge K a_m a_n$ (instead of $a_{m+n} \ge a_m a_n$ as in the ordinary supermultiplicativity). However, this does not introduce any new problem since the sequence $K a_n$, which satisfies $K a_{m+n} \ge K a_m K a_n$, is a supermultiplicative sequence. Sometimes, there is a polynomial in the place of K but the approach is still essentially the same.

As we will work with rooted binary trees, the following lemma turns out to be useful, and can be used with Fekete's lemma in some certain situations.

LEMMA 2.7. *Consider any $d \ge 1/2$. Every rooted binary tree of at least d leaves has a subtree of m leaves so that $d \le m \le 2d$.*

PROOF. Suppose otherwise that there is no such subtree. Every subtree has either less than d leaves or more than $2d$ leaves. As the tree has at least d leaves, not all subtrees have less than d leaves. Pick a minimal subtree T' with more than $2d$ leaves. This subtree cannot be a leave since $2d \ge 1$. It follows that T' has two subtrees T'_1, T'_2, each of which has less than d leaves, due to the minimality of T'. The number of leaves of T', which is the total number of leaves of T'_1, T'_2, is however less than $2d$, contradiction. □

CHAPTER 3

Growth of replacements

As checking if the growth rate $\lambda \leq 1$ is undecidable as proved in Chapter 6 and the big class of the joint spectral radius has the same situation as mentioned in Chapter 4, we present a class of problems for which fairly simple algorithms work.[1] The class consists of all the systems where the vector s is positive and the bilinear map $*$ has the following restricted structure: For every index k, there exist indices i, j so that $(x * y)_k = x_i y_j$ for every two vectors x, y. As the nature of the problem is simple, we will describe its additive version in a combinatorial language, which is quite interesting on its own, instead of the original version in terms of a bilinear system. Its simplicity allows a rather thorough study.

Suppose we have a finite number of baskets, each basket contains infinitely many balls of the same value. We start with choosing a ball from some basket to put on a table. At each subsequent step, we replace one ball on the table by two balls from some baskets with respect to a given set of rules that only involves the baskets where the balls are from. When there are n balls on the table for a given n, we stop and evaluate the sum (and the average) of the values of all the n balls. Our aim is to achieve the highest possible sum (and average) for a given n by choosing appropriately the basket of the first ball to put on the table and the ball to replace at each subsequent step. An asymptotic behavior is that when n tends to infinity, this best average converges to a constant λ, which is called the *growth rate* of the system.

Let us state the problem in an equivalent but more formal way, which will be used from now on. The formulation starts with a collection V of functions $v : \mathbb{N}^+ \to \mathbb{R}$. Denote $c_v = v(1)$ for each function $v \in V$, which will be called the *starting values* later. We have an assignment of a pair of functions $M(v) = (u, w)$ to each v ($u, v, w \in V$ not necessarily different). The value of $v(n)$ for $n \geq 2$ is given by

$$v(n) = \max_{1 \leq m \leq n-1} u(n-m) + w(m). \tag{3.1}$$

Let $g(n)$ denote the maximum of the values of the functions at n, that is

$$g(n) = \max_{v \in V} v(n).$$

We show that the sequence $\{g(n)/n\}_{n=1}^{\infty}$ converges to the so-called *growth rate λ of the system*:

$$\lambda = \lim_{n \to \infty} \frac{g(n)}{n}.$$

The equivalence between this formulation and the problem with the balls is not so hard to see. Each function v corresponds to a basket with c_v as the value of a ball in the basket. The value of $v(n)$ is the maximum sum obtained from n balls if we start with a ball from the basket corresponding to v. The value of $g(n)$ is then the maximum sum when we do not restrict which ball to start with.

On the other hand, the problem is an instance of the growth of bilinear maps where the starting vector s is positive and for each k all the coefficients $c_{i,j}^{(k)}$ are zero except that

[1] This chapter presents the content of the article [**9**].

one of them is 1. We can reduce this bilinear system to the problem of replacements by taking the logarithm. Since the problem involves replacing balls, we call it the growth of replacements, as in the chapter title.

The problem we are studying is in some sense of the same type as Fekete's lemma, which states that for a *superadditive* function $f : \mathbb{N}^+ \to \mathbb{R}$, that is $f(m+n) \geq f(m) + f(n)$ for any m, n, the limit $\lim_{n\to\infty} f(n)/n$ exists (see also Chapter 2). Our setting differs from Fekete's lemma in two points: (i) instead of the equivalent inequality $f(n) \geq \max_{1\leq m\leq n-1} f(n-m) + f(m)$, we use the equality as in (3.1), and (ii) instead of one function, a collection of functions are involved. Note that if the equality in (3.1) is replaced by the inequality, then the limit we are studying does not necessarily exist. For example, consider the functions v_0, v_1 so that $v_0(n) \geq \max_{1\leq m\leq n-1} v_0(n-m) + v_0(m)$ and $v_1(n) \geq \max_{1\leq m\leq n-1} v_0(n-m) + v_0(m)$. If $v_0(n) = n$ for every n while $v_1(n) = n$ for odd n and $v_1(n) = 2n$ for even n, then the maximum average changes between 1 and 2 as n increases. However, if the dependency graph, which will be defined later, is connected, then the limit still exists, by the proof in Section 3.6. The readers can check for themselves that the techniques there also work for the case of inequalities.

One can also formulate this problem in terms of a context free grammar (in Chomsky normal form). Let us consider the following language: There is a nonterminal symbol V associated to each function v, the production rule $V \to UW$ corresponds to the assignment of u, w to each v, and there is also a production $V \to \underline{v}$ for each nonterminal symbol V, where $\underline{v}$ is a terminal symbol for which we assign the weight c_v. We define the weight of a word to be the sum of the weights of all the symbols in that word. The function $v(n)$ is then the maximum weight of a word of length n if we start with the symbol V. Every other term is mapped accordingly.

The readers who are familiar with Petri nets [**11**] and chip-firing games [**12**] can find our problem similar to both problems in the setting but different in the object.

Four proofs for the existence of the limit λ will be given. Each is involved with different terms, for which we will give the definitions first. They are: the dependency graph, composition trees and pseudo-loops. A linear program is also related. Note that the dependency graph is defined slightly differently from chapter to chapter, to suit each particular problem. The pseudo-loops in this chapter are themselves very similar to the linear patterns in Chapter 5. However, being a special case allows the growth of replacements to have simpler approaches, techniques and proofs.

The *dependency graph* is the directed graph whose set of vertices is V and there is a directed edge from v to u if and only if one of the two functions in $M(v)$ is u (loops are allowed). As the dependency graph is directed, it can be partitioned into strongly connected components, for which we call *components* for short. A component is said to be a *single component* if it contains only one vertex and there is no loop for that vertex. In other words, the only vertex v in a single component has the outward edges vu and vw for u, w both different from v. Let us consider the *condensation* of the dependency graph, which is the acyclic graph with each vertex corresponding to a strongly connected component and there is a directed edge UV if and only if there is an edge uv with $u \in U$ and $v \in V$ in the dependency graph. The condensation defines a partial order between components where a directed edge UV means $U \geq V$. A minimal component cannot be a single component since otherwise the minimality implies that the only function in the single component depends on itself, which in turn contradicts its membership in the single component.

One can see the evaluation of a function $v(n)$ as a *composition tree* whose definition is given as follows. For any rooted binary tree of n leaves, we start with labeling the root of the tree with v, and for $M(v) = (u, w)$ we label the left child of the root with u

and the right child with w. We subsequently label all the vertices of the subtrees with the same method. If a leaf is labeled with v, then we say the leaf has value c_v. The value of a tree is the sum of the values of its leaves. Such a labeled tree with this way of evaluation is called a *composition tree*. One may see that the value of $v(n)$ is the largest evaluation over all the composition trees of n leaves. Note that the label for a vertex in a composition tree is actually a vertex of the dependency graph. Unless stated otherwise, all the trees will be regarded as composition trees.

Let us define a simple pattern for composition trees. Let T be a tree with some label for the root and a specially marked leaf that has the same label as the root. Let the sequence of trees $\{T^n\}_{n=1}^{\infty}$ be defined so that $T^1 = T$ and T^n for $n \geq 2$ is obtained from T^{n-1} by replacing the marked leaf of T^{n-1} by T. The marked leaf of T^n is defined to be the marked leaf of the instance of T. A tree T defined in this way is called a *pseudo-loop*. The *value of a pseudo-loop* is defined to be the sum of the values of all leaves excluding the marked leaf. It is not hard to see that the values of the trees $\{T^n\}_n$ grow linearly with respect to the number of leaves, which is the average of the values of all the leaves excluding the marked one. This rate will be called *the rate of the pseudo-loop*.

The following definitions on pseudo-loops will be also used later. For a subtree with one of the vertices p having the same label as the root of the subtree, the pseudo-loop obtained from the subtree by removing every descendant of p with p being the marked leaf is called an *inner pseudo-loop*. An inner pseudo-loop does not contain the marked leaf is said to be *removable*. Removing a removable inner pseudo-loop from a pseudo-loop gives another pseudo-loop. By *removing* an inner pseudo-loop we mean contracting the whole inner pseudo-loop into a vertex. Note that while an inner pseudo-loop can be in either a tree or a pseudo-loop, removability is considered only in the context of a pseudo-loop.

The relation between the growth rate λ and the notion of pseudo-loop is given in the following theorem.

Theorem 3.1. *The growth rate exists and it is the supremum of the rates of all pseudo-loops.*

Furthermore, we can find the best rate in a finite set of pseudo-loops. That is to say the supremum is always attainable.

Theorem 3.2. *There exists a pseudo-loop with the same rate as the growth rate of the system. It can be found among pseudo-loops that do not contain any removable inner pseudo-loop. In particular, such a pseudo-loop has at most $|V|2^{|V|-1}$ leaves after excluding the marked one.*

The proofs of the two above theorems can be found in Section 3.1.

The readers may relate pseudo-loops to linear patterns, for which Theorem 7.2 is a similar result to Theorem 3.1. However, the growth rate of a bilinear system is not always the rate of a linear pattern as in the specific case of the growth of replacements. In other words, there is no correspondence to Theorem 3.2, see Theorem 7.4 for a counterexample.

We can also study the growth rate by the following system[2] of $2|V|$ inequalities: For every $v \in V$,

$$\begin{aligned} z_v &\geq c_v - \theta, \\ z_v &\geq z_u + z_w, \end{aligned} \tag{3.2}$$

where $\{z_v : v \in V\}$ and θ are variables, c_v is $v(1)$ as already defined, and $M(v) = (u, w)$.

[2]This system was suggested by Günter Rote (private communication). The readers may relate it to [**1**, Proposition 5.1].

The set of the solutions is nonempty, e.g., we can set $z_v = 0$ for all v and θ is the maximum of all c_v.

Consider the linear program *minimizing* θ subject to System (3.2). We have the following representation of the growth rate.

THEOREM 3.3. *The growth rate exists and it is the optimal value of* θ *to the linear program.*

A proof is given in Section 3.2. The linear program has $|V| + 1$ variables and $2|V|$ inequalities but it is still rather simple and actually resembles the setting of the problem. In fact, given a solution of all variables, one can construct a pseudo-loop with the growth rate λ in linear time of the number of functions (variables) by a method provided during the course of the proof. A more precise order of $g(n)$ is also shown there:

$$n\lambda + O(1) \leq g(n) \leq n\lambda + O(1).$$

A corollary of this fact is a small interval bounding λ provided the value of $g(n)$ for an n large enough. The values of $O(1)$ are reasonable and can be found in the proof.

Also in Section 3.2, we consider the dual linear program. An interesting point is that given a pseudo-loop of the rate λ, we can give a solution to the dual program in linear time of the variables. Moreover, the construction is more straightforward than the other direction with the original program.

We now consider some computational aspects of the growth rate.

THEOREM 3.4. *Given any proposal* λ_0*, one can decide if* $\lambda_0 < \lambda$ *in quadratic time of the number of functions.*

Theorem 3.4 immediately gives a reasonable algorithm of approximating λ with the time complexity $O(|V|^2 \log \frac{\Delta}{\epsilon})$ for a given precision ϵ and $\Delta = \max_v c_v - \min_v c_v$ (one can observe that the growth rate must be contained in the interval $[\min_v c_v, \max_v c_v]$). In the proof in Section 3.3, the readers will find that the condition $\lambda_0 \geq \lambda$ is equivalent to whether each function v has a maximum value of $v(n)$ over all n when the considered system uses the value $c_u - \lambda_0$ instead of c_u for every u. Such maximum values are also computed as a by-product. When $\lambda_0 = \lambda$, these maximum values turn out to be a solution of the linear program for $\theta = \lambda$.

THEOREM 3.5. *When the starting values are all rational with the numerators and denominators contained in a fixed interval,*[3] *we can compute the growth rate precisely in cubic time of the number of functions.*

Theorem 3.5 is a combination of the results in Theorem 3.2 and Theorem 3.4. The idea is that the growth rate is a fraction with the denominator not too big, therefore, one can stop the binary search when the interval is small enough. Details are given in Section 3.4. How to compute the growth rate efficiently in case the starting values are not necessarily rational, such as $\pi, e, \dots$, is still open. Note that the straight algorithm by Theorem 3.2 may take double exponential time.

Although the notion of pseudo-loop is a very useful tool, we attempt to prove the existence of the limit λ without it by studying the individual functions $v(n)$.

THEOREM 3.6. *Both* $\lambda = \lim_{n\to\infty} g(n)/n$ *and* $\lambda_v = \lim_{n\to\infty} v(n)/n$ *for every* v *exist.*

The proof of Theorem 3.6 is given in Section 3.5. Although it is not as short as the other proofs, the readers can find there a nice result of the same type as Fekete's lemma.

[3]It should be noted that in the statement of [**9**, Theorem 6], the author has not assumed the condition of the boundedness of the numerators and denominators. The mistake has therefore been corrected here.

The last proof of the limit, which assumes the connectedness of the dependency graph, is given in Section 3.6. Although the fourth proof does not work without the condition, it is the shortest proof, and demonstrates a nice application of Fekete's lemma. Meanwhile, the proof that removes inner pseudo-loops in Section 3.1 is perhaps the simplest and shortest one for the general case.

3.1. Growth rate as the maximum rate over all pseudo-loops

We give two proofs of Theorem 3.1, one removes inner pseudo-loops and one extends a tree to a pseudo-loop.

At first, it is obvious that $\liminf_{n\to\infty} g(n)/n \ge \sup_T \lambda_T$, where λ_T is the rate of a pseudo-loop T. Indeed, consider a pseudo-loop T and let q be the number of leaves of T excluding the marked one. For every n, let n be expressed as $n = pq + r$ for an integer p and $1 \le r \le q$. It can be seen that $g(n) \ge pq\lambda_T + O(1)$ by considering the tree obtained from T^p by replacing the marked leaf of T^p by any tree of r leaves. The corresponding lower bound of $g(n)/n$ converges to λ_T, the conclusion follows.

Let $\bar\lambda = \limsup_{n\to\infty} g(n)/n$ and $\lambda^* = \sup_T \lambda_T$. It remains to prove that

$$\bar\lambda \le \lambda^*.$$

Proof of Theorem 3.1 that removes inner pseudo-loops. Assume the contrary that $\bar\lambda > \lambda^*$. We give a contradiction by the existence of a pseudo-loop with a higher rate than λ^*.

An inner-pseudo-loop-free tree has a bounded number of leaves. In other words, any tree of a large enough number of leaves has an inner pseudo-loop. Subsequently removing all inner pseudo-loops results in an inner-pseudo-loop-free tree. The value of the original tree is the sum of the values of all removed inner pseudo-loops and the reduced tree.

By the definition of $\bar\lambda$, for every $\epsilon > 0$ and any N_0, there exists some $N > N_0$ so that $g(N)/N > \bar\lambda - \epsilon$. Choose some ϵ small enough and consider such a large N. As the tree has the value at least $N(\bar\lambda - \epsilon)$, the sum of the values of all the removed inner pseudo-loops is $N(\bar\lambda - \epsilon) - O(1)$, since the value of the reduced tree is bounded. Since the total number of leaves of the pseudo-loops is $N - O(1)$, there must be a pseudo-loop of rate at least the average

$$\frac{N(\bar\lambda - \epsilon) - O(1)}{N - O(1)}.$$

When ϵ is small enough and N is large enough, the above average is arbitrarily close to $\bar\lambda$, hence greater than λ^*, contradiction. □

The other proof is sketched as follows.

Proof of Theorem 3.1 that extends a tree to a pseudo-loop. If there is a path from u to v then there is a composition tree $T(u, v)$ of a bounded number of leaves (and value) so that the root is labeled with u and one of the leaves is labeled with v.

If $g(N) > N(\bar\lambda - \epsilon)$ corresponds to a tree of N leaves with the root labeled with v and a leaf labeled with u so that u, v are in the same component, then replacing the leaf by $T(u, v)$, we obtain a pseudo-loop with the rate at least

$$\frac{N(\bar\lambda - \epsilon) + O(1)}{N + O(1)}, \tag{3.3}$$

which is greater than λ^* when N is large enough and ϵ is small enough.

If no leaf has the label in the same component as the label v of the root, we consider a subtree T' of T such that $|T|/3 \le |T'| \le 2|T|/3$ (which always exists by Lemma 2.7), where $|T|$ is the number of leaves of T. The value of T' is at most $|T'|(\bar\lambda + \epsilon)$ when we choose N large enough, by the definition of $\bar{\lambda}$. If the root of T' has the label in the same component as v, then we have the same situation as in (3.3). That is because the value of the tree T_0 obtained from T by contracting T' into a single leaf would have the value at least

$$N(\bar{\lambda} - \epsilon) - |T'|(\bar{\lambda} + \epsilon) + O(1),$$

while $|T_0| = N - |T'| + 1$.

If the label of the root of T' is in a lower component than the component of v, then we have the same problem for T' with the labels of the vertices being in one less components than T and the value of T' at least

$$N(\bar{\lambda} - \epsilon) - |T_0|(\bar{\lambda} + \epsilon) + O(1),$$

since the value of T_0 is at most $|T_0|(\bar{\lambda} + \epsilon)$.

Recursively treating smaller problems with N large enough and ϵ small enough would give a situation where there is a leaf having the label in the same component as the label of the root. □

Although the space of all pseudo-loops is infinite and the supremum of the rates may not belong to any particular pseudo-loop, we show that the latter is not the case by the fact that we just need to look into the set of pseudo-loops that do not contain any removable inner pseudo-loop to find one with the best rate. In other words, we prove Theorem 3.2, as follows.

Proof of Theorem 3.2. In order to prove the theorem, it suffices to show that any pseudo-loop containing a removable inner pseudo-loop does not need to be considered in the sense that there exists a pseudo-loop of fewer leaves with at least that rate. In other words, the space of pseudo-loops to be considered is finite.

Indeed, if the inner pseudo-loop has a lower or equal rate to the original one, then removing the former does not reduce the rate of the latter. If the inner one has a higher rate, then that inner one itself is a pseudo-loop with a higher rate. In both cases, we can ignore the original pseudo-loop.

It remains to show that a pseudo-loop without any removable inner pseudo-loop has at most $|V|2^{|V|-1}$ leaves after excluding the marked one. Let us call the path from the root to the marked leaf the *main path*. On the main path from the root to the marked leaf, the subpath from the vertex following the root to the marked leaf should not have two vertices of the same label, otherwise we have a removable inner pseudo-loop. That is we have at most $|V|$ vertices on the main path after excluding the marked leaf. For each vertex p on the main path other than the leaf, the subtree whose root is not on the main path is inner pseudo-loop free. Such a subtree has the depth at most $|V| - 1$ and therefore has at most $2^{|V|-1}$ leaves. In total, we have at most $|V|2^{|V|-1}$ leaves after excluding the marked one. □

Remark 3.7. *The bound $|V|2^{|V|-1}$ may not be a tight bound but we can come up with an example where a pseudo-loop of the rate λ must have at least $2^m + 1$ leaves after excluding the marked leaf for a set of $m+3$ functions $a, b, v_0, v_1, \ldots, v_m$ where $M(a) = (a, b), M(b) = (a, v_0), M(v_0) = (v_1, v_1), M(v_1) = (v_2, v_2), \ldots, M(v_{m-1}) = (v_m, v_m), M(v_m) = (a, a)$ with $c_a = c_b = c_{v_0} = \cdots = c_{v_{m-1}} = 0$ and $c_{v_m} = 1$. The verification is left to the readers as an exercise. (Hint: The growth rate is $2^m/(2^m + 1)$.)*

3.2. Growth rate as the solution of a linear program

Relation to the original program. We prove Theorem 3.3. Before that, we repeat the linear program: For every $v \in V$,

$$\begin{aligned} z_v &\geq c_v - \theta, \\ z_v &\geq z_u + z_w, \end{aligned}$$

where $\{z_v : v \in V\}$ and θ are variables, c_v is $v(1)$, and $M(v) = (u, w)$. The object is to minimize θ.

Let θ and $\{z_v\}_v$ be a solution to the linear program. We prove the following two claims.

Claim. $g(n) \leq n\theta + \max_v z_v$.

Proof. For each n, consider the composition tree corresponding to $g(n)$ and let the label of the root be v^*. Let L be the multiset of the labels of the leaves in the composition tree. Since $z_v \geq c_v - \theta$ and $z_v \geq z_u + z_w$ for any v and $M(v) = (u, w)$, we have

$$z_{v^*} \geq \sum_{u \in L} (c_u - \theta) = g(n) - n\theta \implies g(n) \leq z_{v^*} + n\theta,$$

which confirms the claim. □

Claim. $g(n) \geq n\theta + O(1)$.

Proof. We say v is decomposable if either (i) $z_v = c_v - \theta$, or (ii) $z_v = z_u + z_w$ (for $M(v) = (u, w)$) and both u, w are decomposable.

Let G be the decomposition graph, which is a directed graph with the vertices being the functions and there is an edge from v to u (resp. w) if and only if $z_v = z_u + z_w$ (for $M(v) = (u, w)$) and w (resp. u) is decomposable. (Note that the condition for a vertex to have an outward edge is weaker than the condition for a vertex to be decomposable.)

We will show that G contains a cycle. Assume otherwise, that is we have a partial order between the vertices in G with $u \leq v$ if there is an edge vu. Consider $\theta' = \theta - \epsilon$ for a small enough ϵ, we show that there is a solution with θ' (which contradicts with the minimality of θ). We first start with all decomposable functions v with $z_v = c_v - \theta$ and increase it to $z'_v = c_v - \theta'$ and gradually increase z_v for decomposable functions v with $z_v = z_u + z_w$ to $z'_v = z'_u + z'_w$. Finally, for those v with an edge vu in G whose z'_v is not established yet, we increase z_v to $z'_v = z_u + z'_w$ with z_v for smaller v in the partial order updated first. Note that we do not need to update z_v twice for any v. For the remaining functions v we keep $z'_v = z_v$ and obtain a solution $\{z'_v\}_v$ for θ'.

Now G contains a cycle, say $v_0 \to v_1 \to \cdots \to v_k \to v_0$ with $z_{v_i} = z_{v_{i+1}} + z_{w_{i+1}}$ for $M(v_i) = (v_{i+1}, w_{i+1})$ (and $z_{v_k} = z_{v_0} + z_{w_0}$). Since $z_{v_0} = \left(\sum_{i=0}^k z_{w_i}\right) + z_{v_0}$, the sum $\sum_{i=0}^k z_{w_i}$ is zero.

As each w_i is decomposable, we can construct a composition tree so that the root is labeled with w_i and z_{w_i} is the sum of $c_v - \theta$ over all the labels v of the leaves.

We now obtain a pseudo-loop whose path from the root to the marked leaf is the same as the cycle in G and the other branches are the above decomposition trees. This pseudo-loop has rate θ as the sum of z_{w_i} is zero.

Let the number of leaves excluding the marked leaf be m, then for any $n = mp + r$ $(1 \leq r \leq m)$, the claim follows from the boundedness of r and

$$g(n) \geq mp\theta + O(1).$$ □

Theorem 3.3 follows from the two claims.

Remark 3.8. *Given a solution of the program, it is possible to construct a pseudo-loop of the growth rate in linear time as in the process of the second claim. The least trivial part is to check if the functions are decomposable. We leave it as an exercise for the readers.*

Relation to the dual program. We relate the dual program[4] to pseudo-loops of the growth rate. The dual program has $2|V|$ variables $\{x_v, y_v : v \in V\}$ so that for each v we have

$$x_v + y_v = \sum_{u,w:\ M(u)=(v,w)} y_u + \sum_{u,w:\ M(u)=(w,v)} y_u,$$
$$x_v \geq 0,$$
$$y_v \geq 0,$$

and the sum of all x_v is

$$\sum_v x_v = 1.$$

The object of the program is to maximize

$$\sum_v c_v x_v.$$

The maximum value is the same solution as in the original program, which is the growth rate λ. We show that a pseudo-loop of the rate λ can give a solution to the dual program in linear time of the number of variables. In fact, the transformation is more straightforward than the other direction with the original program.

Consider a pseudo-loop with the rate λ. We let x'_v be the number of leaves labeled with v in the tree, and let y'_v be the number of non-leaf vertices labeled with v. If v is the label of the root, we reduce x'_v by 1 (not counting the marked leaf). All the variables x'_v, y'_v that have not been assigned any value will be assumed to be zero.

By the structure of the tree, we have

$$x'_v + y'_v = \sum_{u,w:\ M(u)=(v,w)} y'_u + \sum_{u,w:\ M(u)=(w,v)} y'_u.$$

Let $m = \sum_v x'_v$, we set $x_v = x'_v/m$ and $y_v = y'_v/m$ for each v. We have $\sum_v x_v = 1$, and the object $\sum_v c_v x_v$ is the rate of the pseudo-loop, which is λ. Such a solution gives the maximum value to the object.

3.3. Rate test in quadratic time

We show that it is possible to test whether a proposed rate λ_0 is smaller than the actual rate λ in quadratic time of the number of functions, which in turn immediately gives an algorithm to find an approximation to the growth rate in $O(|V|^2 \log \frac{\Delta}{\epsilon})$ for a given precision ϵ and $\Delta = \max_v c_v - \min_v c_v$. (Note that the growth rate is contained in the interval $[\min_v c_v, \max_v c_v]$.)

At first, $\lambda_0 < \lambda$ is equivalent to the existence a pseudo-loop of positive rate if we consider the system with the starting values $c_v - \lambda_0$ instead of c_v. We show that the latter fact is in turn equivalent to the existence of a function v not having a tree with the label of the root v and a maximum value z_v (regardless of the number of leaves). This equivalence will be verified after presenting the following algorithm, which gives maximum values z_v in case there are such values.

[4]An anonymous reviewer for the article form of this chapter suggested that the dual program may have some meaning.

Algorithm: For each v, initiate $z_v = c_v - \lambda_0$. We repeat the following process as long as there is a variable z_v still having the initial value and $z_v < z_u + z_w$ for $M(v) = (u, w)$:

- Update z_v by the new better value $z_u + z_w$ and mark z_v as a variable depending on z_u, z_w in the sense that any further improvement on z_u or z_w will be directly followed by an improvement on z_v.
- Make a sequence of improvements on the variables that directly or indirectly depend on z_v. If z_v is itself a variable among those variables depending on z_v, then we stop the iteration and conclude $\lambda_0 < \lambda$ right away.

If we finish without concluding $\lambda_0 < \lambda$, then we conclude otherwise $\lambda_0 \geq \lambda$.

The process can be done in $O(|V|^2)$ time since the second step in each iteration is a finite process of $O(|V|)$ time, as in the verification of the algorithm below.

We show that each z_v from our algorithm gives the largest possible value over all the compositions trees rooted by v without any inner pseudo-loop. We reason by induction on the height of trees. Consider a tree T_v^* whose root is labeled with v with the maximum value over the trees without any inner pseudo-loop. It means there is no other occurrence of v other than the root. If T_v^* is only a single vertex v, then its value is $c_v - \lambda_0$. Our algorithm gives this value in the first place and the value of T_v will never be decreased during the course. Suppose all other functions v' in the tree T_v^* than the root v have their trees $T_{v'}$ produced by the algorithm attaining their maximum values. Since T_v is the tree of two subtrees T_u, T_w, whose values are maximum due to the induction hypothesis, the value of T_v is also the maximum value for v.

It means if there is no pseudo-loop of positive rate, the values produced by the algorithm are also the maximum values of the trees rooted by the functions.

On the other hand, if there is any pseudo-loop of positive rate, our algorithm also detects a pseudo-loop of positive rate. In this case, $g(n)$ is unbounded. Suppose the algorithm stops without recognizing any pseudo-loop. Consider a minimal composition tree giving a value larger than any z_v given by our algorithm (minimality in the sense that no subtree has such a property). Each branch of the root should give the value at most the value given by our algorithm due to the minimality of the composition tree. Let v be the label of the root. If v is already marked as being dependent on any improvement of u, w ($M(v) = (u, w)$), then we have a contradiction as $z_v < z_u + z_w$. If the dependency has not been established, then our algorithm has not finished yet, as we still have $z_v < z_u + z_w$ and another iteration should be proceeded. In either case, we have a contradiction.

As for the matter of time complexity, we show that for the terminating condition in each iteration, we only need to check for z_v but not any other z_u whether that variable depends the improvement of z_v for the turn z_v is updated. Initially, there is no pseudo-loop in the composition trees corresponding to all z_v. Suppose we have the same situation before a given iteration. The reason for that lack is due to a missing edge of dependency. Therefore, if there is a pseudo-loop after updating z_v, it must be a pseudo-loop involving v when only two new dependencies $v \to u$ and $v \to w$ are introduced as the missing edges. Also, before reaching again v in case of a pseudo-loop, we do not have to check for other pseudo-loops when updating variables depending on v as they do not exist. The second step of the iteration can be done easily with a queue in $O(|V|)$ time. It follows that the whole algorithm takes $O(|V|^2)$ time since the outmost loop is iterated at most $|V|$ times.

We have verified the validity of the algorithm by showing that the algorithm either stops in the middle and concludes the existence of a pseudo-loop of a positive rate ($\lambda_0 < \lambda$), or finishes and gives the trees of the maximal values ($\lambda_0 \ge \lambda$).

REMARK 3.9. *The best value obtained by the algorithm is also a solution of z_v with a fixed $\theta = \lambda_0$ to System* (3.2)*. Of course, a solution only exists when $\lambda_0 \ge \lambda$.*

3.4. A cubic time algorithm to find the precise value of the growth rate

This section combines the results of Theorem 3.2 and Theorem 3.4 to give a cubic time algorithm computing the growth rate precisely provided that the starting values are rational. In other words, we settle Theorem 3.5 as follows.

At first, we can assume that the starting values are not just rational but all integers, otherwise we can scale the starting values by an appropriate factor. Note that these integers are also contained in a fixed interval. By Theorem 3.2, the growth rate of a system is the rate of a pseudo-loop without any removable inner pseudo-loop, which is of the form a/b where b is an integer at most $|V|2^{|V|-1}$. By the assumption that the starting values are integers, the numerator a is also an integer and the rates a_1/b_1 and a_2/b_2 of two pseudo-loops without any removable inner pseudo-loop are either equal or at least $1/B^2$ apart where $B = |V|2^{|V|-1}$. Therefore, if we apply the binary search to the starting interval $[\min_v c_v, \max_v c_v]$ with the quadratic time rate test algorithm in Theorem 3.4, we can stop the binary search whenever the interval is small enough, in particular less than $1/B^2$. This interval contains only one fraction whose denominator is at most B, which is the growth rate. Given the interval, we can find this precise value of the growth rate using the Farey sequence in linear time of $|V|$, which is dominated by the time finding the interval, which is $O(|V|^2 \log(B^2)) = O(|V|^3)$. In fact, instead of taking the middle value in each iteration of the binary search, one can take the mediant as in the process of the Farey sequence and avoid applying the Farey sequence in the end. However, it does not change the cubic time of the algorithm. The algorithm can be seen as a nice combination of the binary search, the Farey sequence and some insights of the problem.

REMARK 3.10. *The approach does not apply when the nature of the starting values is more complicated than rational numbers, e.g. transcendental numbers $e, \pi, \dots$. One can approximate these numbers by rationals and then recover the coefficients (the number of leaves with the corresponding label over the total number of leaves) from the estimated growth rate, however, it may take an exponential time for the recovery (and also the numerators and denominators are no longer in a fixed interval). The problem in this case seems to ask for a more direct solution than finding the value by the binary search.*

3.5. Growth rate in terms of the functions

To prove the growth rate λ and λ_v for every function v exist, we give first the following lemma, which should be of its own interest. It is kind of in the same spirit as Fekete's lemma, with more functions involved and similar proving techniques.

For convenience, in the statement of the lemma and in the proof, all the integers that are supposed to be used for indexing functions u, v will be treated as elements in $\mathbb{Z}/k\mathbb{Z}$ for the k in the statement. In particular, it is the case of the indices i, i^*, j.

LEMMA 3.11. *Given $2k$ ($k \ge 1$) functions $v_0(n), \dots, v_{k-1}(n), u_0(n), \dots, u_{k-1}(n)$: $\mathbb{N}^+ \to \mathbb{R}$ such that for every $0 \le i \le k-1$ and every $n \ge 2$,*

$$v_i(n) = \max_{1 \le m \le n-1} (v_{i+1}(n-m) + u_{i+1}(m)).$$

Then for every i,

$$\lim_{n\to\infty} \frac{v_i(n)}{n} = \sup_{m_0\ge 1,\dots,m_{k-1}\ge 1} \frac{\sum_{j=0}^{k-1} u_j(m_j)}{\sum_{j=0}^{k-1} m_j}.$$

PROOF. Denote by R the value of the supremum (note that it can be infinite). To prove the theorem, it suffices to verify the following two points for every i:

(i) $\liminf_{n\to\infty} v_i(n)/n \ge R$.

By the definition of R, for any $R' < R$, there are $m_0, \dots, m_{k-1}$ such that

$$\frac{\sum_{j=0}^{k-1} u_j(m_j)}{\sum_{j=0}^{k-1} m_j} > R'.$$

Let $m = m_0 + \cdots + m_{k-1}$. For every n, if $n = mt + p$ for some integer t and $1 \le p \le m$, we have the lower bound $v_i'(n) \le v_i(n)$ with

$$v_i'(n) = v_i(p) + t\left(\sum_{j=0}^{k-1} u_j(m_i)\right).$$

Since $v_i(p)$ is bounded, the sequence $\{v_i'(n)/n\}_n$ tends to

$$\frac{\sum_{j=0}^{k-1} u_j(m_j)}{\sum_{j=0}^{k-1} m_j} > R'.$$

It follows that $\liminf_{n\to\infty} v_i(n)/n > R'$ for any $R' < R$, which implies

$$\liminf_{n\to\infty} v_i(n)/n \ge R.$$

(ii) $\limsup_{n\to\infty} v_i(n)/n \le R$ (we assume $R \ne \infty$ otherwise it is trivial).

Assume $\limsup_{n\to\infty} v_i(n)/n = R' > R$, we will show a contradiction by giving $m_0 \ge 1, \dots, m_{k-1} \ge 1$ so that

$$\frac{\sum_{j=0}^{k-1} u_j(m_j)}{\sum_{j=0}^{k-1} m_j} > R.$$

For each i and n, due to the evaluation of $v_i(n)$, there exist a number t and tk numbers $m_j^{(s)}$ for $0 \le j \le k-1$, $1 \le s \le t$ such that: $\sum_{j,s} m_j^{(s)} = n-1$, all of them are nonzero except possibly $m_{i^*}^{(t)}, m_{i^*+1}^{(t)}, \dots, m_i^{(t)}$ for some i^* (if there is no zero, we let $i^* = i+1$), and

$$v_i(n) = v_{i^*-1}(1) + \sum_{s=1}^{t}\sum_{j=0}^{k-1} u_j(m_j^{(s)}),$$

where $u_j(0)$ is assumed to be zero for every j. (The number t can be understood as the number of rounds.)

Let ${m'}_j^{(s)} = m_j^{(s)}$, but we set ${m'}_{i^*}^{(t)} = {m'}_{i^*+1}^{(t)} = \cdots = {m'}_i^{(t)} = 1$ if there are corresponding zeroes in $\{m_j^{(s)}\}$. We have

$$\sum_{s=1}^{t}\sum_{j=0}^{k-1} u_j({m'}_j^{(s)}) = v_i(n) - v_{i^*-1}(1) + \sum_{j=i^*}^{i} u_j(1). \tag{3.4}$$

By the definition of R', for every $\epsilon > 0$, there is an arbitrarily large n such that

$$\frac{v_i(n)}{n} > R' - \epsilon.$$

Note that the right hand side of (3.4) is the sum of $v_i(n)$ and a bounded sum, and the difference between the sum of all ${m'}_j^{(s)}$ and the sum of all $m_j^{(s)}$ is also bounded. It means that for every $\epsilon' > 0$, we can choose a small enough ϵ and a large enough n such that

$$\sum_{s=1}^{t}\sum_{j=0}^{k-1} u_j({m'}_j^{(s)}) > (R' - \epsilon')\left(\sum_{s=1}^{t}\sum_{j=0}^{k-1} {m'}_j^{(s)}\right).$$

This is followed by the existence of some s^* such that

$$\frac{\sum_{j=0}^{k-1} u_j({m'}_j^{(s^*)})}{\sum_{j=0}^{k-1} {m'}_j^{(s^*)}} > R' - \epsilon'.$$

Since ϵ' can be arbitrarily small, $R' - \epsilon' > R$ for some ϵ', and since all ${m'}_j^{(s^*)} \geq 1$, we have a contradiction with the supremum R.

By (i) and (ii), the conclusion follows. $\square$

Now we can prove Theorem 3.6.

Consider the partial order between the strongly connected components of the dependency graph. The minimal component cannot be a single component. Therefore, each function in a minimal component should be in a cycle and the existence of its growth rate is confirmed by Lemma 3.11. Consider a non-minimal component with the assumption that we already have growth rates for the functions in all smaller components. If the considered component is not single, then every function has a growth rate as already reasoned. In the other case, the only function v of the component has $M(v) = (u, w)$ with u, w from smaller components, which already have growth rates by induction hypothesis. Since $v(n) = \max_m u(n-m) + w(m)$, the larger rate of u and w is the growth rate of v. By induction, all functions have growth rates. It follows from $g(n) = \max_v v(n)$ that $g(n)$ also has a growth rate, which is the largest rate over all the functions v.

Remark 3.12. *Although Lemma 3.11 also covers the case the limit is infinite, the limits in our application are obviously finite since the value $v(n)/n$ for any function $v \in V$ is always contained in the range of the minimum and maximum starting values.*

3.6. A proof of the limit for strongly connected dependency graphs using Fekete's lemma

Suppose the dependency graph is connected, this section provides a simple proof of the limit λ. It is interesting to apply Fekete's lemma here, as our problem itself can be seen as a variant of Fekete's lemma.

If there is an edge vu with $M(v) = (u, w)$, then

$$v(n) \geq u(n-1) + c_w.$$

It follows that if the distance from v to u is $d_{v,u}$, then

$$v(n) \geq u(n - d_{v,u}) + \alpha_{v,u},$$

for some constant $\alpha_{v,u}$.

Consider a function v with $M(v) = (u, w)$. For any m, n large enough, we have

$$v(m+n) \geq u(m) + w(n) \geq v(m - d_{u,v}) + \alpha_{u,v} + v(n - d_{w,v}) + \alpha_{w,v},$$

where the constants $d_{u,v}, d_{w,v}$ exist because the dependency graph is connected.

Replacing m by $m - d_{w,v}$ and n by $n - d_{u,v}$, and adding to both sides $\alpha_{u,v} + \alpha_{w,v}$, we have

$$\begin{aligned} v(m+n-d_{u,v}-d_{w,v}) + \alpha_{u,v} + \alpha_{w,v} \geq\ & v(m - d_{u,v} - d_{w,v}) + \alpha_{u,v} + \alpha_{w,v} \\ & + v(n - d_{u,v} - d_{w,v}) + \alpha_{u,v} + \alpha_{w,v}. \end{aligned}$$

Let $v'(n) = v(n - d_{u,v} - d_{w,v}) + \alpha_{u,v} + \alpha_{w,v}$, we can see that $v'(n)$ is a superadditive sequence. By Fekete's lemma, $v'(n)/n$ converges. It follows that $v(n)/n$ converges to the same limit. The convergence of $g(n)/n$ follows.

REMARK 3.13. *The approach still works when we replace the equality in* (3.1) *by the inequality* $v(n) \geq \max_{1\leq m\leq n-1}(u(n-m)+w(m))$. *However, the limit does not necessarily hold when the dependency graph is not connected, as pointed out in the introduction of the chapter.*

CHAPTER 4

Joint spectral radius

The joint spectral radius is a generalization of the spectral radius to a set of matrices, which was first introduced in [**13**]. The notion has caught a lot of attention with its theoretical interest as well as its applications in engineering fields. We advise the readers to check [**14**] for a book with a comprehensive treatment of the subject.

The reason for including a chapter[1] on the joint spectral radius is that it is actually an instance of the growth of bilinear maps, as proved in Chapter 6. Although the topic is quite established with many results, we provide some new facts, mostly bounds, in the case of nonnegative matrices. The joint spectral radius theorem for nonnegative matrices is also related.

We begin with the definition of the joint spectral radius. In this chapter, we consider only *finite* sets of matrices, unless otherwise stated. Whether the results hold for infinite sets or how they can be extended is left open.

Given a finite set Σ of square nonnegative matrices in $\mathbb{R}^{D\times D}$,[2] we denote

$$\|\Sigma^n\| = \max_{A_1,\dots,A_n\in\Sigma} \|A_1\dots A_n\|,$$

where $\|A\|$ for a matrix A is some chosen norm. In this chapter, we use the maximum norm for convenience, that is $\|A\| = \max_{i,j} |A_{i,j}|$.

As we may be interested in a specific entry, we write

$$\|\Sigma^n\|_{i,j} = \max_{A_1,\dots,A_n\in\Sigma} (A_1\dots A_n)_{i,j}.$$

A simple observation that will be frequently used is that $\|\Sigma^n\|_{i,j}$ is bounded for a bounded n (and a fixed Σ). Another one is that: For any k^*, we have

$$\|\Sigma^m\|_{i,k^*}\|\Sigma^n\|_{k^*,j} \le \|\Sigma^{m+n}\|_{i,j} \le \sum_k \|\Sigma^m\|_{i,k}\|\Sigma^n\|_{k,j},$$

where the corresponding one to the latter inequality in the case of a single matrix A is the equality $(A^{m+n})_{i,j} = \sum_k (A^m)_{i,k}(A^n)_{k,j}$.

In [**13**], the joint spectral radius $\rho(\Sigma)$ of the set Σ is defined to be the limit

$$\rho(\Sigma) = \lim_{n\to\infty} \sqrt[n]{\|\Sigma^n\|}. \tag{4.1}$$

Note that $\rho(\Sigma)$ is independent of the norm of choice, since any two norms are in a constant factor of each other.

When $|\Sigma| = 1$ with $\Sigma = \{A\}$, the joint spectral radius $\rho(\Sigma)$ becomes the ordinary spectral radius $\rho(A)$. The spectral radius $\rho(A)$ is originally defined as the largest absolute value of the eigenvalues of A, while the representation corresponding to (4.1) is known as Gelfand's formula: For every matrix A, we have

$$\rho(A) = \lim_{n\to\infty} \sqrt[n]{\|A^n\|}.$$

[1]Some part of this chapter can be found in the article [**10**].

[2]We use the capital form D instead of d to emphasize that the dimension is a constant, along with other constants U, V as defined later, and also to reserve d to denote distances and divisors.

Note that there is no natural corresponding notion to eigenvalues for a set of matrices.

The following result, which appears in most of the sources, is used to prove that the limit of $\rho(\Sigma)$ exists. We provide it again here as it also gives a bound on the radius. It is often expressed in submultiplicative norms (so-called matrix norms) and in a slightly different form.

PROPOSITION 4.1 (An adaptation of the popular statement for the maximum norm[3]). *The following limit of $\rho(\Sigma)$ exists and can be expressed as:*

$$\rho(\Sigma) = \lim_{n\to\infty} \sqrt[n]{\|\Sigma^n\|} = \inf_n \sqrt[n]{D\|\Sigma^n\|}.$$

PROOF. For any two matrices A, B, we have

$$\|AB\| = \max_{i,j}(AB)_{i,j} = \max_{i,j}\sum_k A_{i,k}B_{k,j} \le D\|A\|\|B\|.$$

For any two positive integers m, n, we have

$$\|\Sigma^{m+n}\| = \|A_1 \dots A_{m+n}\| \le D\|A_1 \dots A_m\|\|A_{m+1} \dots A_{m+n}\| \le D\|\Sigma^m\|\|\Sigma^n\|,$$

where $A_1, \dots, A_{m+n}$ are some matrices from Σ.

Writing differently, $D\|\Sigma^{m+n}\| \le (D\|\Sigma^m\|)(D\|\Sigma^n\|)$ means the sequence $\{D\|\Sigma^n\|\}_n$ is submultiplicative. By Fekete's lemma, $\sqrt[n]{D\|\Sigma^n\|}$ converges to $\inf_n \sqrt[n]{D\|\Sigma^n\|}$, which is also the limit of $\sqrt[n]{\|\Sigma^n\|}$. □

The main focus of the chapter to give a formula and some bounds for the joint spectral radius. The main tool is the dependency graph, which is actually quite similar to the dependency graph for the growth of bilinear maps in Chapter 5.

DEFINITION 4.2. *The* dependency graph *of a set of matrices Σ is a directed graph where the vertices are $1, \dots, D$, and there is an edge from i to j if and only if $A_{i,j} \neq 0$ for some matrix $A \in \Sigma$ (loops are allowed). Being a directed graph, the dependency graph can be decomposed into strongly connected components, for which we will call* components *for short. If a component contains only one vertex without loop around it, we call it a* single *component. Otherwise, we will call it a* regular *component.*

As the dependency graph is the main graph in the chapter, vertices, paths and components may be mentioned without stating explicitly the dependency graph. In a similar manner to $\|\Sigma^n\|_{i,j}$, we also denote

$$\|\Sigma^n\|_C = \max_{i,j\in C}\|\Sigma^n\|_{i,j} \tag{4.2}$$

for a component C.

When we consider $\|\Sigma^n\|_C$ for a component C instead of considering $\|\Sigma^n\|$, we are actually considering the problem reduced to C in the sense that we remove all the rows and columns not in C. This is still a problem that satisfies all the results of the original problem. Therefore, we also have the limit

$$\rho_C(\Sigma) = \lim_{n\to\infty} \sqrt[n]{\|\Sigma^n\|_C},$$

which corresponds to $\rho(\Sigma')$ with a strongly connected dependency graph for Σ'.

It is obvious that $\rho(\Sigma) \ge \max_C \rho_C(\Sigma)$, but only in later sections we will verify that the equality

$$\rho(\Sigma) = \max_C \rho_C(\Sigma)$$

holds.

[3]For a submultiplicative norm, the constant D is not necessary and the expression on the right would be $\inf_n \sqrt[n]{\|\Sigma^n\|}$.

4.1. Known bounds and estimations

This section discusses some bounds and estimations of the joint spectral radius for complex matrices, which are not necessarily nonnegative. One of the popular ways to estimate the joint spectral radius is as follows (see Proposition 1.6 and Section 2.3.3 on "Branch and Bounds Methods" in the book [**14**]). For any sequence of matrices $A_1, \dots, A_m \in \Sigma$, the m-th root of the (ordinary) spectral radius of $A_1 \dots A_m$ is a lower bound for $\rho(\Sigma)$. Denote[4]

$$P_m(\Sigma) = \max_{A_1,\dots,A_m \in \Sigma} \rho(A_1 \dots A_m).$$

Together with the bound from Proposition 4.1, we can bound $\rho(\Sigma)$ from both sides: For any m,

$$\sqrt[m]{P_m(\Sigma)} \le \rho(\Sigma) \le \sqrt[m]{D\|\Sigma^m\|}. \tag{4.3}$$

The lower bound is due to

$$\begin{aligned}\sqrt[m]{P_m(\Sigma)} &= \sqrt[m]{\max_{A_1,\dots,A_m \in \Sigma} \rho(A_1 \dots A_m)} \\ &= \sqrt[m]{\max_{A_1,\dots,A_m \in \Sigma} \lim_{t\to\infty} \sqrt[t]{\|(A_1 \dots A_m)^t\|}} \le \lim_{n\to\infty} \sqrt[n]{\|\Sigma^n\|}.\end{aligned}$$

One of the points supporting this method of bounding is that the limit superior of the sequence for the left side and the limit of the sequence for the right side (with respect to m) are equal to $\rho(\Sigma)$. In fact,

$$\check{\rho}(\Sigma) = \limsup_{n\to\infty} \sqrt[n]{P_n(\Sigma)}$$

is called the *generalized spectral radius* of Σ. That the two radii are equal for finite sets Σ is the content of the joint spectral radius theorem (see Theorem 2.3 in the book [**14**] or the article [**15**] for the first time it was proved).

THEOREM (The joint spectral radius theorem). *For every* bounded[5] *set* Σ *of matrices,*

$$\check{\rho}(\Sigma) = \rho(\Sigma).$$

Since $P_{tm}(\Sigma) \ge (P_m(\Sigma))^t$ for any t, m, we can write

$$\rho(\Sigma) = \limsup_{n\to\infty} \sqrt[n]{P_n(\Sigma)} = \sup_n \sqrt[n]{P_n(\Sigma)}. \tag{4.4}$$

We justify the latter equality by the following observation.

PROPOSITION 4.3. *If a nonnegative sequence x_n satisfies $x_{tm} \ge (x_m)^t$ for every t, m, then*

$$\limsup_{n\to\infty} \sqrt[n]{x_n} = \sup_n \sqrt[n]{x_n}. \tag{4.5}$$

PROOF. It suffices to prove $\limsup_{n\to\infty} \sqrt[n]{x_n} \ge \sup_n \sqrt[n]{x_n}$ as the other direction is obvious. For every m, $\sqrt[m]{x_m} \le \limsup_{n\to\infty} \sqrt[n]{x_n}$ due to the subsequence $\sqrt[tm]{x_{tm}} \ge \sqrt[m]{x_m}$ for $t = 1, 2, \dots$. It follows that $\sup_m \sqrt[m]{x_m} \le \limsup \sqrt[n]{x_n}$. The conclusion follows. $\square$

[4]Note that when we write $\rho(A)$ for a matrix A, we mean the classic spectral radius for A, which is actually the same as $\rho(\{A\})$. Also, the P in $P_m(\Sigma)$ is actually the capital version of ρ.

[5]The original version is stated for bounded sets (of complex matrices). In particular, every finite set is bounded. However, a bounded set may be infinite.

A set Σ is said to have the *finiteness property* if $\rho(\Sigma) = \sqrt[n]{P_n(\Sigma)}$ for some n, that is there exists a sequence $A_1, \ldots, A_n$ of matrices in Σ so that $\rho(\Sigma) = \sqrt[n]{\rho(A_1 \ldots A_n)}$. The *finiteness conjecture* [**16**], which states that every set has the finiteness property, was disproved in [**17**]. An explicit example was given in [**18**] but the nature of the entries is quite complicated. It is still open whether the conjecture holds for binary matrices (or equivalently for rational matrices), see [**19**]. Another discussion on the finiteness property can be found in Chapter 7.

An example in [**20**, Section 2] shows that the limit superior of the generalized spectral radius is not replaceable by a limit in general. This behaviour can be suggested by the fact that while $\|\Sigma^n\|$ is supermultiplicative, $P_n(\Sigma)$ is not always so (only the condition of (4.4) is guaranteed). Also, it is asked in [**20**]: Since it follows from (4.3) that for every n we have

$$\max_{1\le m\le n} \sqrt[m]{P_m(\Sigma)} \le \rho(\Sigma) \le \min_{1\le m\le n} \sqrt[m]{D\|\Sigma^m\|}, \tag{4.6}$$

what is the convergence rate of the lower bound $\max_{1\le m\le n} \sqrt[m]{P_m(\Sigma)}$ and the upper bound $\min_{1\le m\le n} \sqrt[m]{D\|\Sigma^m\|}$ (with respect to n)[6] to $\rho(\Sigma)$? This question is critical to the efficiency of the bound in (4.3). Section 4.5 will show that both sequences converge at the rate $O(\frac{\log n}{n})$ for finite sets of nonnegative matrices.

The following bound of Kozyakin [**21**] has a more explicit convergence rate than (4.3): For every n,

$$\sqrt[n]{f(n)\|\Sigma^n\|} \le \rho(\Sigma) \le \sqrt[n]{\|\Sigma^n\|}, \tag{4.7}$$

where $f(n)$ is rather complicated and it may grow very small (note that the norm in (4.7) is a submultiplicative norm). The work [**21**] describes $f(n)$ explicitly, but loosely speaking, $f(n)$ is in general roughly about C^{-n^α} for a constant C and $\alpha = \ln(D-1)/\ln D$. Although the bound in (4.7) is very interesting and $\lim_{n\to\infty} \sqrt[n]{f(n)} = 1$, it is hard to estimate $\rho(\Sigma)$ effectively as the ratio of the two bounds is large.

If one restricts the scope to *irreducible* sets of matrices,[7] we have the following bound in [**22**]: There is constant $\gamma \le 1$ so that for every n,

$$\sqrt[n]{\gamma^{1+\ln n}\|\Sigma^n\|} \le \rho(\Sigma) \le \sqrt[n]{\|\Sigma^n\|}. \tag{4.8}$$

Note that the norm here is also submultiplicative. Also note that the quantity $\gamma^{1+\ln n}$ is actually of order n^{-t} for some t. Although our bounds in the following sections have the same convergence rate as the bound (4.8), the condition in our bounds is on the signs of the entries of matrices, instead of the algebraic nature of irreducible sets in this context.

In the theme of nonnegative matrices, we mention [**23**, Theorem 16], which works for nonnegative matrices only: For the matrix S so that $S_{i,j} = \max_{A\in\Sigma} A_{i,j}$, we have

$$\frac{\rho(S)}{|\Sigma|} \le \rho(\Sigma) \le \rho(S). \tag{4.9}$$

The gap between the two bounds is obviously not so efficient when there are many matrices in Σ. However, it requires lower computational complexity due to calculating the spectral radius of a single matrix. It is asked in [**23**] whether the bound in (4.9) can be improved at a reasonable computational cost.

[6]It was actually asked in [**20**] for $\sqrt[m]{\|\Sigma^m\|}$ for submultiplicative norms. We adapt it for the maximum norm.

[7]Irreducible sets of matrices are not related to irreducible matrices: A set of matrices Σ is irreducible if the only subspaces that are invariant under all the matrices in Σ are the trivial subspaces $\{0\}$ and the whole space.

4.2. The diagonal formula for the joint spectral radius

In this section, a formula for the joint spectral radius is introduced as follows. During the course of proving, a by-product is an upper bound on $\|\Sigma^n\|$ itself.

THEOREM 4.4. *Given a finite set Σ of nonnegative matrices, we have*

$$\rho(\Sigma) = \sup_n \max_i \sqrt[n]{\|\Sigma^n\|_{i,i}}. \tag{4.10}$$

Moreover, if $\rho(\Sigma) > 0$, there exists a number r so that for every n,

$$\operatorname{const} \rho(\Sigma)^n \leq \|\Sigma^n\| \leq \operatorname{const} n^r \rho(\Sigma)^n. \tag{4.11}$$

We note here that the formula of $\rho(\Sigma)$ in Theorem 4.4 is itself not completely new, but rather a perspective on some known results.[8] We may start with the first result of the type.

THEOREM (Wimmer 1974 [**24**]). *For any complex matrix A, we have*

$$\rho(A) = \limsup_{n\to\infty} \sqrt[n]{|\operatorname{tr}(A^n)|},$$

where $\operatorname{tr}(A^n)$ *denotes the trace of* A^n.

The reason why this result is related to Theorem 4.4 will be given after discussing the following similar form for the joint spectral radius.

THEOREM (Chen and Zhou 2000 [**25**]). *For any finite set of complex matrices Σ, we have*

$$\rho(\Sigma) = \limsup_{n\to\infty} \max_{A_1,\dots,A_n} \sqrt[n]{|\operatorname{tr}(A_1 \dots A_n)|}.$$

It turns out that the formula (4.10) in Theorem 4.4 can be deduced from this theorem. Indeed, since for any nonnegative matrices $A_1, \dots, A_n$ we have

$$\max_i (A_1 \dots A_n)_{i,i} \leq \operatorname{tr}(A_1 \dots A_n) \leq D \max_i (A_1 \dots A_n)_{i,i},$$

it follows from the characterization of Chen and Zhou that

$$\begin{aligned} &\limsup_{n\to\infty} \max_{A_1,\dots,A_n\in\Sigma} \sqrt[n]{\max_i (A_1 \dots A_n)_{i,i}} \leq \rho(\Sigma) \\ &\quad \leq \limsup_{n\to\infty} \max_{A_1,\dots,A_n\in\Sigma} \sqrt[n]{D \max_i (A_1 \dots A_n)_{i,i}}. \end{aligned} \tag{4.12}$$

It means

$$\begin{aligned} \rho(\Sigma) &= \limsup_{n\to\infty} \max_{A_1,\dots,A_n\in\Sigma} \sqrt[n]{\max_i (A_1 \dots A_n)_{i,i}} \\ &= \sup_{n\to\infty} \max_{A_1,\dots,A_n\in\Sigma} \sqrt[n]{\max_i (A_1 \dots A_n)_{i,i}}. \end{aligned}$$

The former equality is due to the lower bound and the upper bound of $\rho(\Sigma)$ in (4.12) being identical since D is a constant and does not affect the asymptotic behavior when $n \to \infty$. The latter equality is due to Proposition 4.3.

By using similar arguments, Wimmer's result implies the formula (4.10) of Theorem 4.4 when Σ consists of a single matrix.

We note that the bound (4.11) on $\|\Sigma^n\|$ in Theorem 4.4 is not new either. It is well known in literature, even for bounded sets of complex matrices, e.g. see [**26**]. For finite

[8]The author was not aware of these results at first, the motivation for the formula is actually an afterthought of Theorem 5.3.

sets of nonnegative matrices, a stronger bound will be shown in Theorem 4.15 below: If $\rho(\Sigma) > 0$, there exists a nonnegative integer $r < D$ so that for every n,

$$\mathrm{const}\, n^r \rho(\Sigma)^n \le \|\Sigma^n\| \le \mathrm{const}\, n^r \rho(\Sigma)^n.$$

Note that the value of r in the proof of Theorem 4.4 is in general much larger than the value of r here.

Although the content of Theorem 4.4 is not exactly new as discussed, the point of proving it again by combining the formula and the bound is that they can be proved together in the same time and in a simple way (about two pages without relying on any result). The proof uses a kind of double induction, which may be interesting on its own.

Moreover, the formula (4.10) in Theorem 4.4 and the joint spectral radius theorem for finite sets of nonnegative matrices are closely related as follows. On one hand, suppose we already have the joint spectral radius theorem. In principle, in order to prove the formula (4.10) in Theorem 4.4, it suffices to prove the formula (4.10) for a single nonnegative matrix A, that is

$$\rho(A) = \sup_n \max_i \sqrt[n]{(A^n)_{i,i}}, \tag{4.13}$$

and then apply the joint spectral radius theorem to Σ. Indeed, $\rho(\Sigma) = \check{\rho}(\Sigma)$ while $\check{\rho}(\Sigma)$ can be written as

$$\begin{aligned}
\check{\rho}(\Sigma) &= \limsup_{n\to\infty} \sqrt[n]{P_n(\Sigma)} \\
&= \sup_n \sqrt[n]{P_n(\Sigma)} && \text{(by (4.4))} \\
&= \sup_n \max_{A_1,\dots,A_n\in\Sigma} \sqrt[n]{\rho(A_1\dots A_n)} \\
&= \sup_n \max_{A_1,\dots,A_n\in\Sigma} \sqrt[n]{\sup_t \max_i \sqrt[t]{[(A_1\dots A_n)^t]_{i,i}}} && \text{(by (4.13))} \\
&= \max_i \sup_t \sup_n \max_{A_1,\dots,A_n\in\Sigma} \sqrt[tn]{[(A_1\dots A_n)^t]_{i,i}} \\
&= \max_i \sup_n \max_{A_1,\dots,A_n\in\Sigma} \sqrt[n]{(A_1\dots A_n)_{i,i}} \\
&= \max_i \sup_n \sqrt[n]{\|\Sigma^n\|_{i,i}}.
\end{aligned} \tag{4.14}$$

Note that the sixth equality comes from the following observation: For every t, n,

$$\max_{A_1,\dots,A_n\in\Sigma} \sqrt[tn]{[(A_1\dots A_n)^t]_{i,i}} \le \max_{B_1,\dots,B_{tn}\in\Sigma} \sqrt[tn]{(B_1\dots B_{tn})_{i,i}}.$$

Although the formula (4.10) of Theorem 4.4 (for a set of matrices) can be reduced to the formula (4.13) for a single matrix, proving Theorem 4.4 for a set of matrices is not harder than proving for a single matrix. Therefore, we provide the full proof without relying on the joint spectral radius theorem.

On the other hand, suppose we have the formula (4.10) in Theorem 4.4 in the first place. This leads to a simple proof for the joint spectral radius theorem for finite sets of nonnegative matrices, since it follows from (4.14) and (4.10) that

$$\check{\rho}(\Sigma) = \max_i \sup_n \sqrt[n]{\|\Sigma^n\|_{i,i}} = \rho(\Sigma).$$

Note that (4.14) depends on (4.13), which is a special case of (4.10). However, (4.14) is independent of the joint spectral radius theorem.

Corollary 4.5. *The joint spectral radius theorem holds for finite sets of nonnegative matrices.*

In comparison to the standard proof of the joint spectral radius theorem by Elsner [**27**], which makes use of analytic–geometric tools, our combinatorial approach seems to be more elementary with basic arguments asking for less background on the subject. However, we should emphasize here that our approach works for only finite sets of nonnegative matrices.

Proof of Theorem 4.4. If $\rho(\Sigma) = 0$, then there is no cycle in the dependency graph, hence the theorem trivially holds. We assume $\rho(\Sigma) > 0$.

Denoting

$$\lambda = \sup_n \max_i \sqrt[n]{\|\Sigma^n\|_{i,i}},$$

we have

$$\rho(\Sigma) = \lim_{n\to\infty} \sqrt[n]{\|\Sigma^n\|} \ge \lambda,$$

since $\|\Sigma^{tn}\| \ge \|\Sigma^{tn}\|_{i,i} \ge (\|\Sigma^n\|_{i,i})^t$ for any t, n. (The latter inequality is obtained by induction with $\|\Sigma^{tn}\|_{i,i} \ge \|\Sigma^{(t-1)n}\|_{i,i}\|\Sigma^n\|_{i,i}$.)

To finish the proof, it suffices to prove that

$$\|\Sigma^n\| \le \operatorname{const} n^r \lambda^n. \tag{4.15}$$

Indeed, suppose we have (4.15), then

$$\lambda \le \rho(\Sigma) = \lim_{n\to\infty} \sqrt[n]{\|\Sigma^n\|} \le \lim_{n\to\infty} \sqrt[n]{\operatorname{const} n^r \lambda^n} = \lambda.$$

We obtain the equality $\lambda = \rho(\Sigma)$, which is the formula (4.10) of Theorem 4.4. The bounds (4.11) of Theorem 4.4 also follow, since

$$\|\Sigma^n\| \le \operatorname{const} n^r \lambda^n = \operatorname{const} n^r \rho(\Sigma)^n,$$

while the lower bound $\|\Sigma^n\| \ge \operatorname{const} \rho(\Sigma)^n$ is due to Proposition 4.1.

The reduction to (4.15) is clarified. We now prove (4.15).

We begin with an observation: There exists some K_0 so that $\|\Sigma^n\|_{i,j} \le K_0\lambda^n$ for every n and every i, j in a regular component. Indeed, let δ be the distance from j to i in the dependency graph, we have

$$\|\Sigma^n\|_{i,j} = \frac{1}{\|\Sigma^\delta\|_{j,i}}\|\Sigma^n\|_{i,j}\|\Sigma^\delta\|_{j,i} \le \frac{1}{\|\Sigma^\delta\|_{j,i}}\|\Sigma^{n+\delta}\|_{i,i} \le \frac{1}{\|\Sigma^\delta\|_{j,i}}\lambda^{n+\delta} \le K_0\lambda^n, \tag{4.16}$$

where

$$K_0 = \max_{j',i'} \frac{\lambda^{\delta'}}{\|\Sigma^{\delta'}\|_{j',i'}}$$

with j', i' ranging over all the vertices in the same regular component so that there is a path from j' to i' and δ' is the distance from j' to i'.

Note that if the component is single, then the inequality $\|\Sigma^n\|_{i,i} \le K_0\lambda^n$ in (4.16) still trivially holds with $\|\Sigma^n\|_{i,i} = 0$ for every n.

Now we may wonder what would be the inequality when i, j are not in the same component.

The condensation of the dependency graph is the directed acyclic graph whose vertices are the components and there is an edge from C_1 to C_2 if there is an edge ij in the dependency graph with $i \in C_1$ and $j \in C_2$. We denote by $\Delta(i, j)$ the distance from the component of i to the component of j. (For i, j in the same component, we let $\Delta(i, j) = 0$, and we do not consider i, j with no path from i to j.)

For any δ, we denote
$$\|\Sigma^n\|_\delta = \max_{i,j:\ \Delta(i,j)\le\delta} \|\Sigma^n\|_{i,j}.$$
What we have shown in (4.16) is actually $\|\Sigma^n\|_0 \le K_0\lambda^n$. It is the base case of the following claim.

Claim. *For every δ, there exist a positive constant K_δ and a nonnegative number r_δ so that for every n,*
$$\|\Sigma^n\|_\delta \le K_\delta n^{r_\delta}\lambda^n.$$

Proof. By the induction method, since we have established the claim for $\delta = 0$ with K_0 in (4.16) and $r_0 = 0$, it remains to show the claim for any δ given that it holds for $\delta' = \delta - 1$ (with the corresponding numbers $K_{\delta'}, r_{\delta'}$).

Let $\alpha = \max\{\frac{\|\Sigma^1\|}{\lambda}, D(K_{\delta'})^2\}$ and $H = \max\{1, K_0D, 2^{2r_{\delta'}}\}$. It suffices to show that for every n,
$$\|\Sigma^n\|_\delta \le \alpha H^{\lceil \log n\rceil}\lambda^n, \tag{4.17}$$
since $H^{\lceil\log n\rceil} \le H^{1+\log n} = HH^{\log n} = Hn^{\log H}$ (note that $H \ge 1$), which implies
$$\|\Sigma^n\|_\delta \le \alpha H n^{\log H}\lambda^n.$$
(In other words, this is the claim for δ with $K_\delta = \alpha H$ and $r_\delta = \log H$.)

In order to prove (4.17), we again use another induction on $\lceil\log n\rceil$ as follow (i.e. double induction, first on δ, then on $\lceil \log n\rceil$). (Notation $\log n$ here denotes the logarithm of base 2.) At first, the base case trivially holds for those n with $\lceil\log n\rceil = 0$, i.e. $n = 1$. That is because $\|\Sigma^1\|_\delta \le \|\Sigma^1\| = \frac{\|\Sigma^1\|}{\lambda}\lambda \le \alpha\lambda$ while $H^{\lceil\log n\rceil} = 1$.

We assume (4.17) holds for any number n' so that $\lceil\log n'\rceil < \lceil\log n\rceil$ and proves it also holds for n. Let $n = \ell + m$ where $\ell = \lfloor n/2\rfloor$ and $m = \lceil n/2\rceil$. Suppose $\|\Sigma^n\|_\delta = \|\Sigma^n\|_{i,j}$ for some i, j. We have
$$\|\Sigma^n\|_{i,j} \le \sum_{k'} \|\Sigma^\ell\|_{i,k'}\|\Sigma^m\|_{k',j} \le D\|\Sigma^\ell\|_{i,k}\|\Sigma^m\|_{k,j}$$
for some k that maximizes $\|\Sigma^\ell\|_{i,k}\|\Sigma^m\|_{k,j}$.

We consider three cases regarding k:

- If i, k are in the same component then $\|\Sigma^\ell\|_{i,k} \le K_0\lambda^\ell$ by (4.16), and $\|\Sigma^m\|_{k,j} \le \|\Sigma^m\|_\delta \le \alpha H^{\lceil\log m\rceil}\lambda^m$ by the induction hypothesis on $\lceil\log n\rceil$, since $\lceil\log m\rceil = \lceil\log\lceil n/2\rceil\rceil = \lceil\log n\rceil - 1$. It follows that
$$\|\Sigma^n\|_{i,j} \le DK_0\lambda^\ell\alpha H^{\lceil\log m\rceil}\lambda^m \le H\alpha H^{\lceil\log m\rceil}\lambda^n = \alpha H^{\lceil\log n\rceil}\lambda^n,$$
where the latter inequality is due to $K_0D \le H$.
- If k, j are in the same component then likewise we have
$$\|\Sigma^n\|_{i,j} \le D\alpha H^{\lceil\log\ell\rceil}\lambda^\ell K_0\lambda^m \le H\alpha H^{\lceil\log\ell\rceil}\lambda^n \le \alpha H^{\lceil\log n\rceil}\lambda^n.$$
Note that $HH^{\lceil\log\ell\rceil} \le HH^{\lceil\log m\rceil} = H^{\lceil\log n\rceil}$ since $H \ge 1$ (the inequality may be strict, say when $\lceil\log\ell\rceil < \lceil\log m\rceil$, e.g. $n = 2^t + 1$, for which we need $H \ge 1$).
- If k is not in the same component with either i or j, then both $\Delta(i,k)$ and $\Delta(k,j)$ are at most $\delta' = \delta - 1$. It follows that $\|\Sigma^\ell\|_{i,k} \le \|\Sigma^\ell\|_{\delta'} \le K_{\delta'}\ell^{r_{\delta'}}\lambda^\ell$ and $\|\Sigma^m\|_{k,j} \le \|\Sigma^m\|_{\delta'} \le K_{\delta'}m^{r_{\delta'}}\lambda^m$ by the induction hypothesis on δ. We have
$$\|\Sigma^n\|_{i,j} \le DK_{\delta'}\ell^{r_{\delta'}}\lambda^\ell K_{\delta'}m^{r_{\delta'}}\lambda^m \le D(K_{\delta'})^2 n^{2r_{\delta'}}\lambda^n \le \alpha H^{\lceil\log n\rceil}\lambda^n,$$
where the last inequality is due to the condition $\alpha \ge D(K_{\delta'})^2$ and $H^{\lceil\log n\rceil} \ge H^{\log n} = n^{\log H} \ge n^{2r_{\delta'}}$ (note that we have the condition $H \ge 1$ and $H \ge 2^{2r_{\delta'}}$).

We have verified (4.17) by induction on $\lceil \log n \rceil$. The claim follows, by induction on δ. $\square$

Since every $\Delta(i,j)$ is less than D, it follows that $\|\Sigma^n\| = \|\Sigma^n\|_D$. Therefore, the claim implies (4.15), which in turn finishes the proof of Theorem 4.4.

A corollary on a limit of the trace. Following the development of Wimmer's and Chen and Zhou's results, Xu has attempted to turn the limit superior in the theorem of Chen and Zhou to a limit for nonnegative matrices with the condition that one of them is primitive. Note that a matrix A is primitive if $A^n > 0$ for some $n \geq 1$.

THEOREM (Xu 2010 [**28**]). *For any finite set of nonnegative matrices Σ with at least one primitive matrix, we have*

$$\rho(\Sigma) = \lim_{n\to\infty} \max_{A_1,\ldots,A_n \in \Sigma} \sqrt[n]{\operatorname{tr}(A_1 \ldots A_n)}.$$

Using the formula (4.10) in Theorem 4.4, we can extend Xu's result to the following theorem, by showing that the conclusion still holds with a more relaxed condition than the primitivity of at least one matrix.

THEOREM 4.6. *Given a finite set of nonnegative matrices Σ, for each i, denote*

$$\delta_i = \gcd\{n : \|\Sigma^n\|_{i,i} > 0\},$$

or we set $\delta_i = 1$ in case the set is empty. If Δ is a multiple of all δ_i, then

$$\rho(\Sigma) = \lim_{m\to\infty} \max_i \sqrt[m\Delta]{\|\Sigma^{m\Delta}\|_{i,i}} = \lim_{m\to\infty} \max_{A_1,\ldots,A_{m\Delta}} \sqrt[m\Delta]{\operatorname{tr}(A_1 \ldots A_{m\Delta})}.$$

Note that a matrix A is primitive if and only if it is irreducible, i.e. all δ_i corresponding to $\Sigma = \{A\}$ are equal to 1. If such a matrix A is an element of Σ, we still have the same $\delta_i = 1$, for which we have the conclusion for $\Delta = 1$, i.e. Xu's result.

The proof below can be seen as a corollary of Theorem 4.4, with the support of Lemma 2.5 (a variant of Fekete's lemma for nonnegative sequences).

PROOF OF THEOREM 4.6. The latter equality is quite obvious by the fact that $\frac{1}{D}\operatorname{tr}(A) \leq \max_i A_{i,i} \leq \operatorname{tr}(A)$ for any $D \times D$ nonnegative matrix A.

At first, since the sequence $\{\|\Sigma^n\|_{i,i}\}_n$ for each i is supermultiplicative, it follows from the extension of Fekete's lemma in Lemma 2.5 that the subsequence of all the positive elements is either empty or follows the growth rate

$$\rho_i = \lim \left\{ \sqrt[n]{\|\Sigma^n\|_{i,i}} : \|\Sigma^n\|_{i,i} > 0 \right\} = \sup_n \sqrt[n]{\|\Sigma^n\|_{i,i}}.$$

In case of emptiness, we still have $\rho_i = 0$ as the value of the supremum.

By Theorem 4.4, we have

$$\rho(\Sigma) = \max_i \rho_i.$$

Further, if the set $\{n : \|\Sigma^n\|_{i,i} > 0\}$ is nonempty, then $\|\Sigma^n\|_{i,i} > 0$ for every large enough multiple n of δ_i, since δ_i is the greatest common divisor. It means

$$\rho_i = \lim_{m\to\infty} \sqrt[m\delta_i]{\|\Sigma^{m\delta_i}\|_{i,i}}.$$

When the set is empty, it trivially holds with $\delta_i = 1$.

Since Δ is a multiple of all δ_i, it follows that

$$\rho(\Sigma) = \max_i \rho_i = \max_i \lim_{m\to\infty} \sqrt[m\Delta]{\|\Sigma^{m\Delta}\|_{i,i}} = \lim_{m\to\infty} \max_i \sqrt[m\Delta]{\|\Sigma^{m\Delta}\|_{i,i}}. \qquad \square$$

4.3. An explicit bound using the diagonal

Although Theorem 4.4 is not exactly new and can be seen as a perspective on known results only, the following bound is quite useful and can be seen as a nice application of the formula (4.10) in the theorem.

THEOREM 4.7. *Let m_i for each i be any number so that $\|\Sigma^{m_i}\|_{i,i} > 0$, or set $m_i = 1$ if there is no such m_i. We have*

$$\max_i \sqrt[m_i]{\|\Sigma^{m_i}\|_{i,i}} \le \rho(\Sigma) \le \max_i \sqrt[m_i]{\left(\frac{UD}{V}\right)^{3D^2} \|\Sigma^{m_i}\|_{i,i}},$$

where $D \times D$ is the dimension of the matrices, U, V are respectively the largest entry and the smallest entry over all the positive entries of the matrices in Σ, and C is taken over all components in the dependency graph.

Although U/V can be arbitrarily large, the appearance of U and V is essential to the formula. For example, let Σ contain only one matrix

$$A = \begin{bmatrix} \frac{1}{N} & 1 \\ 1 & \frac{1}{N} \end{bmatrix}$$

where N is a large number. As its square is

$$A^2 = \begin{bmatrix} 1 + \frac{1}{N^2} & \frac{2}{N} \\ \frac{2}{N} & 1 + \frac{1}{N^2} \end{bmatrix},$$

the joint spectral radius is greater than 1, since $\rho(\Sigma) \ge \sqrt{(A^2)_{1,1}} > 1$. Setting $m_1 = m_2 = 1$, we have $\sqrt[m_i]{\|\Sigma^{m_i}\|_{i,i}} = \sqrt[1]{\|\Sigma^1\|_{i,i}} = \frac{1}{N}$ for $i = 1, 2$. Therefore, the relation between U and V must present in the formula in some form.

The reason we use the quantities U, V is that they can be convenient to bound the entries of products. Indeed, given some n matrices $A_1, \dots, A_n$ from Σ, an entry $(A_1 \dots A_n)_{i,j}$ is the sum of D^{n-1} terms, each of which is the product of some n entries of the matrices from Σ. Therefore, if $(A_1 \dots A_n)_{i,j} > 0$, then

$$V^n \le (A_1 \dots A_n)_{i,j} \le D^{n-1} U^n.$$

In other words, if $\|\Sigma^n\| > 0$, then we have the same bound $V^n \le \|\Sigma^n\| \le D^{n-1}U^n$ for $\|\Sigma^n\|$. It follows that if $\rho(\Sigma) > 0$ then

$$V \le \rho(\Sigma) \le UD.$$

These simple observations will used in several arguments in this section and the following sections.

The performance of the bound in Theorem 4.7 may be not so obvious. At first, we can discard those i with no positive $\|\Sigma^{m_i}\|_{i,i}$. We also make sure that the remaining m_i have large enough values, say they are at least some m. Let us say $\max_i \sqrt[m_i]{\|\Sigma^{m_i}\|_{i,i}} = \sqrt[m_j]{\|\Sigma^{m_j}\|_{j,j}}$ and $\max_i \sqrt[m_i]{\left(\frac{UD}{V}\right)^{3D^2} \|\Sigma^{m_i}\|_{i,i}} = \sqrt[m_k]{\left(\frac{UD}{V}\right)^{3D^2} \|\Sigma^{m_k}\|_{k,k}}$. We have

$$\sqrt[m_k]{\|\Sigma^{m_k}\|_{k,k}} \le \sqrt[m_j]{\|\Sigma^{m_j}\|_{j,j}} \le \rho(\Sigma) \le \sqrt[m_k]{\left(\frac{UD}{V}\right)^{3D^2} \|\Sigma^{m_k}\|_{k,k}}.$$

It follows that the ratio between the upper bound and the lower bound in Theorem 4.7 is at most $\sqrt[m_k]{\left(\frac{UD}{V}\right)^{3D^2}} \le \sqrt[m]{\left(\frac{UD}{V}\right)^{3D^2}}$, which is the m-th root of a constant. The value of m can be taken arbitrarily large as $\|\Sigma^{tm_i}\|_{i,i} > 0$ for any $t \ge 1$.

Note that $\|\Sigma^n\|_{i,i}$ is computed based on the computation of all $|\Sigma|^n$ combinations. However, it is still reasonable since the problem of approximating the joint spectral radius is NP-hard [**29**]. The theorem applies however very well if the set contains only one matrix, i.e. the case of the ordinary spectral radius.

To prove Theorem 4.7, we need the following key lemma, which is fairly technical, and will be proved later. In fact, we use it to relate different bounds of the joint spectral radius in Section 4.6.

LEMMA 4.8. *For any index i, if m, n are two positive integers whose difference is bounded, then either $\|\Sigma^n\|_{i,i} = 0$ or the ratio $\|\Sigma^m\|_{i,i}/\|\Sigma^n\|_{i,i}$ is bounded. In particular, if $m \geq n$, then the bound can be set to*

$$(UD)^{m-n}\left(\frac{UD}{V}\right)^{3D^2-2D+1}.$$

We are interested in an explicit constant only for the case $m \geq n$ since it will be applied in the proof of Theorem 4.7 as follows. When $m < n$, we do not need the boundedness of the ratio for other results, but still prove it as an interesting fact.

PROOF OF THEOREM 4.7. The lower bound is obvious, we prove the upper bound.

Since $\|\Sigma^n\|_C$ for a component C is $\|\Sigma^n\|$ when we reduce the problem to C, it follows from Theorem 4.4 that

$$\rho_C(\Sigma) = \sup_n \max_{i\in C} \sqrt[n]{\|\Sigma^n\|_{i,i}},$$

which means

$$\rho(\Sigma) = \max_C \rho_C(\Sigma).$$

For a regular component C and any $i \in C$, by Proposition 4.1 we have

$$\rho_C(\Sigma)^{m_i} \leq D\|\Sigma^{m_i}\|_C.$$

Suppose $\|\Sigma^{m_i}\|_C = \|\Sigma^{m_i}\|_{j,k}$. Let δ_1 be the distance from i to j and δ_2 be the distance from k to i, that is $\|\Sigma^{\delta_1}\|_{i,j}$ and $\|\Sigma^{\delta_2}\|_{k,i}$ are both nonzero (when $i = j$, we have $\delta_1 = 0$, we then assume that $\|\Sigma^0\|_{i,j} = 1$, and similarly for $j = k$ with $\delta_2 = 0$ and $\|\Sigma^0\|_{k,i} = 1$). We have

$$\begin{aligned}
\|\Sigma^{m_i}\|_C &= \frac{1}{\|\Sigma^{\delta_1}\|_{i,j}\|\Sigma^{\delta_2}\|_{k,i}}\|\Sigma^{\delta_1}\|_{i,j}\|\Sigma^{m_i}\|_{j,k}\|\Sigma^{\delta_2}\|_{k,i} \\
&\leq \frac{1}{V^{\delta_1+\delta_2}}\|\Sigma^{m_i+\delta_1+\delta_2}\|_{i,i} \leq \frac{1}{V^{\delta_1+\delta_2}}(UD)^{\delta_1+\delta_2}\left(\frac{UD}{V}\right)^{3D^2-2D+1}\|\Sigma^{m_i}\|_{i,i} \\
&= \left(\frac{UD}{V}\right)^{3D^2-2D+1+\delta_1+\delta_2}\|\Sigma^{m_i}\|_{i,i} \leq \left(\frac{UD}{V}\right)^{3D^2-1}\|\Sigma^{m_i}\|_{i,i},
\end{aligned}$$

where Lemma 4.8 is used in the step bounding $\|\Sigma^{m_i+\delta_1+\delta_2}\|_{i,i}$ by a constant times $\|\Sigma^{m_i}\|_{i,i}$. (Note that $3D^2 - 2D + 1 + \delta_1 + \delta_2 \leq 3D^2 - 1$ is due to $\delta_1 \leq D - 1, \delta_2 \leq D - 1$.)

In total,

$$\rho_C(\Sigma) \leq \sqrt[m_i]{D\left(\frac{UD}{V}\right)^{3D^2-1}\|\Sigma^{m_i}\|_{i,i}} \leq \sqrt[m_i]{\left(\frac{UD}{V}\right)^{3D^2}\|\Sigma^{m_i}\|_{i,i}}.$$

When C is single, C contains a single vertex without any loop, hence $\rho_C(\Sigma) = 0$, and the above inequality trivially holds with $m_i = 1$.

As $\rho(\Sigma) = \max_C \rho_C(\Sigma)$, we obtain the conclusion

$$\rho(\Sigma) \leq \max_i \sqrt[m_i]{\left(\frac{UD}{V}\right)^{3D^2}\|\Sigma^{m_i}\|_{i,i}}. \qquad \square$$

Proof of Lemma 4.8. It suffices to consider only i for which the set $\{n : \|\Sigma^n\|_{i,i} > 0\}$ is nonempty. To prove Lemma 4.8, we need the following lemma.

LEMMA 4.9. *Let $d = \gcd\{n : (\Sigma^n)_{i,i} > 0\}$. There exists N so that $\|\Sigma^n\|_{i,i} > 0$ for every $n \geq N$ with $d|n$. In particular, one can set $N = (D-1)(2D-1)$.*

Note that Lemma 4.9 is asymptotically optimal in the worst case. For example, if the dependency graph is composed of only two disjoint cycles around i of lengths ℓ_1, ℓ_2 so that ℓ_1, ℓ_2 are not too much distant with $\gcd(\ell_1, \ell_2) = 1$, then the smallest number in the place of N would be $(\ell_1 - 1)(\ell_2 - 1)$. (We recall that the Frobenius number[9] of x, y is $(x-1)(y-1) - 1$.) An example of ℓ_1, ℓ_2 when $D = 2k$ is $\ell_1 = k+1$ and $\ell_2 = k$ (note that $\ell_1 + \ell_2 = D + 1$). When D is odd, we leave a vertex isolated and proceed with the even number of remaining vertices.

To prove Lemma 4.9, we need some preliminary results.

Two subwalks in a walk are said to be *disjoint* if the only vertex that they possibly share is a common endpoint.

OBSERVATION 4.10. *If a walk does not contain two disjoint circuits, then its length is less than $2D$.*

PROOF. Let the walk be $v_0, \ldots, v_k$. If the vertices are all distinct, then $k < D$ and we are done. Otherwise, let v_j be the first repeated vertex, that is, j is the smallest number so that there exists $i < j$ with $v_i = v_j$. It follows that $v_0, \ldots, v_{j-1}$ are distinct. Suppose the walk does not contain two disjoint circuits. This mean there is no circuit in the walk $v_j, \ldots, v_k$, that is, these vertices are distinct. It follows that $k < 2D$. □

We also need Schur's lemma[10] that gives a bound on the Frobenius number.

LEMMA 4.11 (Schur 1935 [**30**]). *Let $2 \leq p_1 < p_2 < \cdots < p_k$ be k integers such that $\gcd(p_1, \ldots, p_k) = 1$, then every integer $n \geq (p_1 - 1)(p_k - 1)$ can be expressed as a linear combination of $p_1, \ldots, p_k$ with nonnegative coefficients.*

We can now prove Lemma 4.9.

PROOF OF LEMMA 4.9. Denote $S = \{n : \|\Sigma^n\|_{i,i} > 0\}$. Let m be the smallest element of S such that the set $S^* = \{n \in S : n \leq m\}$ satisfies $\gcd S^* = d$.

We prove that $m < 2D$. Indeed, suppose $m \geq 2D$. Due to the minimality of m, we have $d^* = \gcd(S^* \setminus \{m\}) > d$ and d^* does not divide m. As $\|\Sigma^m\|_{i,i} > 0$, there is a circuit from i to i of length m. This circuit contains 2 disjoint subcircuits, by Observation 4.10. Let a and b be the lengths of the two subcircuits. Removing any of these subcircuits or both results in a circuit of length less than m, which is divisible by d^*. In other words, $d^* \mid m-a$, $d^* \mid m-b$ and $d^* \mid m-a-b$. It implies that $d^* \mid (m-a)+(m-b)-(m-a-b) = m$, contradiction.

Let $T = \{n/d : n \in S^*\}$, we have $\gcd T = 1$. As $\min S^* \leq D$ (a minimal circuit) and $\max S^* \leq m < 2D$, we have $\min T \leq D/d$ and $\max T \leq 2D/d$. It follows that $\|\Sigma^n\|_{i,i} > 0$ for every $n \geq d(\frac{D}{d} - 1)(\frac{2D}{d} - 1)$ and n divisible by d, by Schur's lemma. The conclusion follows as $d(\frac{D}{d} - 1)(\frac{2D}{d} - 1) \leq (D-1)(2D-1)$. □

Now comes the main part of the proof of Lemma 4.8.

Denote $d = \gcd\{t : \|\Sigma^t\|_{i,i} > 0\}$.

Denote $N = (D-1)(2D-1)$. By Lemma 4.9, for every $t \geq N$, if $d|t$ then $\|\Sigma^t\|_{i,i} > 0$.

[9]The Frobenius number of positive integers $p_1, \ldots, p_k$ with $\gcd(p_1, \ldots, p_k) = 1$ is the largest integer that cannot be expressed as a linear combination of $p_1, \ldots, p_k$ with nonnegative coefficients.

[10]The lemma is due to Schur in 1935 but was not published until 1942 by Brauer in [**30**].

Let $S_t = \{j : \|\Sigma^t\|_{j,i} > 0\}$ for each $t \geq 1$, that is the set of vertices from which we can reach i by a walk of length t. Let $S = \bigcup_{t:d|t} S_t$, that is the set of vertices from which we can reach i by a walk of some length divisible by d.

Let $M = N + D^2$. We have the observation.

CLAIM. *$S_t = S$ for every $t \geq M$ with $d|t$.*

PROOF. For each $j \in S$, let ud be the least multiple of d such that $\|\Sigma^{ud}\|_{j,i} > 0$. We have $u \leq D$. Indeed, suppose $u > D$. Consider the path $v_0, \dots, v_{ud}$ from $v_0 = j$ to $v_{ud} = i$. The vertices v_{kd} for $k = 0, \dots, u$ are not all distinct since $u > D$, say $v_{k'd} = v_{k''d}$. Contracting the subpath $v_{k'd}, \dots, v_{k''d}$ from the path $v_0, \dots, v_{ud}$ gives a path from j to i whose length is divisible by d but less than ud, contradiction.

Consider any $t = vd \geq M$. It follows from $M = N + D^2 \geq N + ud$ that $t - ud = vd - ud \geq N$, which implies $\|\Sigma^{vd-ud}\|_{i,i} > 0$ by Lemma 4.9. That is $\|\Sigma^t\|_{j,i} \geq \|\Sigma^{ud}\|_{j,i}\|\Sigma^{vd-ud}\|_{i,i} > 0$, i.e. the vertex j is also in S_t. □

Note that $M = N + D^2 = (D-1)(2D-1) + D^2 = 3D^2 - 3D + 1$, for which one can observe

$$M + d = 3D^2 - 3D + 1 + d \leq 3D^2 - 2D + 1.$$

We consider the case $\|\Sigma^m\|_{i,i} > 0$ and $\|\Sigma^n\|_{i,i} > 0$ only, as the conclusion is trivial otherwise. It follows that d divides both m, n. Cases regarding the magnitude of m, n are analyzed as follows:

- Suppose $n \geq M + d$ and $m > d$. It follows that there exists a greatest positive integer $t < \min\{m, n\}$ with $n - t \geq M$ and $d|t$. Such a number t exists because $t = d$ is a satisfying number. Since d divides $m - t$, we have

$$\begin{aligned}
\|\Sigma^m\|_{i,i} &\leq \sum_{j\in S} \|\Sigma^t\|_{i,j}\|\Sigma^{m-t}\|_{j,i} \\
&\leq \sum_{j\in S} \|\Sigma^t\|_{i,j}\|\Sigma^{n-t}\|_{j,i}\frac{\max_{j\in S}\|\Sigma^{m-t}\|_{j,i}}{\min_{j\in S}\|\Sigma^{n-t}\|_{j,i}} \\
&\leq \frac{\max_{j\in S}\|\Sigma^{m-t}\|_{j,i}}{\min_{j\in S}\|\Sigma^{n-t}\|_{j,i}} D\|\Sigma^n\|_{i,i}.
\end{aligned}$$

  Note that denominator is positive as $n - t \geq M$ and $d \mid n - t$.

  Since $m-t$ and $n-t$ are bounded (as t is chosen to be the greatest satisfying number), the ratio $\max_{j\in S}\|\Sigma^{m-t}\|_{j,i}/\min_{j\in S}\|\Sigma^{n-t}\|_{j,i}$ is bounded, and so is the ratio $\|\Sigma^m\|_{i,i}/\|\Sigma^n\|_{i,i}$.

  Suppose $n - m \leq M$. Now comes the explicit bound (the situation $m \geq n$ is included in this case). We have $n - t \leq M + d$. Indeed, if $n - t > M + d$, then $t + d < n - M \leq m$, $t + d < n - M < n$ and $n - (t + d) > M$, a contradiction to the maximality of t (since $t + d$ is a larger satisfying number than t). It follows that $n - t \leq M + d \leq 3D^2 - 2D + 1$. Therefore,

$$\begin{aligned}
\frac{\|\Sigma^m\|_{i,i}}{\|\Sigma^n\|_{i,i}} &\leq D\frac{\max_{j\in S}\|\Sigma^{m-t}\|_{j,i}}{\min_{j\in S}\|\Sigma^{n-t}\|_{j,i}} \leq D\frac{\frac{1}{D}(UD)^{m-t}}{V^{n-t}} \\
&= \left(\frac{UD}{V}\right)^{n-t}(UD)^{m-n} \leq \left(\frac{UD}{V}\right)^{3D^2-2D+1}(UD)^{m-n}.
\end{aligned}$$

- It remains to consider the case we do not have both $m > d$ and $n \geq M + d$.

  If $m = d$ then $\|\Sigma^m\|_{i,i}/\|\Sigma^n\|_{i,i}$ is bounded as the range of n is bounded when $m - n$ is bounded (we do not need an explicit bound here as $m < n$). If

$n < M + d \le 3D^2 - 2D + 1$, then regardless of m we still have

$$\frac{\|\Sigma^m\|_{i,i}}{\|\Sigma^n\|_{i,i}} \le \frac{\frac{1}{D}(UD)^m}{V^n} \le \left(\frac{UD}{V}\right)^n (UD)^{m-n} \le \left(\frac{UD}{V}\right)^{3D^2-2D+1} (UD)^{m-n}.$$

The conclusion follows from the verification in both cases.

4.4. A bound using the norm

The bound in the previous section gives a quite effective bound but the constant in the upper bound is not convenient to establish as we have to take care of for which n we have $\|\Sigma^n\|_{i,i} > 0$ and other issues. This section provides another approach to bounding the joint spectral radius using the "norm" $\|\Sigma^n\|$ with easier to established constants instead of the "diagonal" $\|\Sigma^n\|_{i,i}$ as follows.

Theorem 4.12. *Given a finite set Σ of nonnegative matrices. For every n,*

$$\sqrt[n]{\left(\frac{V}{UD}\right)^D \max_C \|\Sigma^n\|_C} \le \rho(\Sigma) \le \sqrt[n]{D \max_C \|\Sigma^n\|_C},$$

where $D \times D$ is the dimension of the matrices, U, V are respectively the largest entry and the smallest entry over all the positive entries of the matrices in Σ, and C is taken over all components in the dependency graph.

When D and U/V are not too large, the gap between them can be reasonably small even with a not so large n. Although U/V can be arbitrarily large, the appearance of U and V is essential to the formula. For example, let Σ contain only one matrix

$$A = \begin{bmatrix} 1 & \frac{1}{N} \\ N & 1 \end{bmatrix}$$

where N is a large number. Since

$$A^2 = \begin{bmatrix} 2 & \frac{2}{N} \\ 2N & 2 \end{bmatrix} = 2A, \qquad A^n = 2^{n-1}A,$$

the joint spectral radius is obviously 2. Meanwhile, $\max_C \|\Sigma^1\|_C = N$, therefore, the relation between U and V must present in the formula in some form. One may compare this matrix A with the matrix A in the example after Theorem 4.7, whose diagonal has the smallest values.

During the proof of Theorem 4.12, we also give the following bound on $\|\Sigma^n\|$ as in Theorem 4.15: If $\rho(\Sigma) > 0$ then there exists a nonnegative integer $r < D$ so that for every n,

$$\text{const}\, n^r \rho(\Sigma)^n \le \|\Sigma^n\| \le \text{const}\, n^r \rho(\Sigma)^n.$$

The inequalities show that $\sqrt[n]{\|\Sigma^n\|}$ converges to $\rho(\Sigma)$ at the rate $O(\frac{\log n}{n})$.

It is quite obvious that the method of bounding in Theorem 4.12 is asymptotically better than the methods in Section 4.1 except that it is not clear for the one in (4.3). The comparison to the latter one can be done only in Section 4.5 when we have all the necessary results. In short, the method in Theorem 4.12 is better than the one in (4.3) by a root of a polynomial of degree r. However, there is a modification to make the two methods asymptotically equivalent.

Comparing the bound in Theorem 4.7 and the bound in Theorem 4.12 would be a bit tricky as the former uses the diagonals while the latter uses the norms. The lower bound in the former is the trivial part while the upper bound in the latter is the trivial part. Both are asymptotically as effective as each other, while the estimation of the constants for the latter one is easier. However, the former gives a neat formula of the joint spectral radius with some interesting corollaries, e.g. a simple proof of the joint spectral radius theorem for finite sets of nonnegative matrices.

Treatment of strongly connected dependency graphs. The central argument in Proposition 4.1 is $\|\Sigma^{m+n}\| \leq \operatorname{const} \|\Sigma^m\| \|\Sigma^n\|$. The other direction of the inequality suggests an alternative approach that works in the case of nonnegative matrices. While Proposition 4.1 leads to an upper bound for $\rho(\Sigma)$, Proposition 4.13 below will lead to a lower bound.

PROPOSITION 4.13. *Given a finite set Σ of nonnegative matrices with a strongly connected dependency graph, we have the following weak form of supermultiplicativity: For every m, n,*

$$\|\Sigma^m\| \|\Sigma^n\| \leq \left(\frac{UD}{V}\right)^D \|\Sigma^{m+n}\|,$$

where U, V, D are defined as in Theorem 4.12.

PROOF. For any i, j, let $\delta(i,j)$ be the distance from i to j in the dependency graph. We have

$$\begin{aligned}
\|\Sigma^m\| \|\Sigma^n\| &= \|\Sigma^m\|_{i,j} \|\Sigma^n\|_{i',j'} \\
&\leq V^{-\delta(j,i')} \|\Sigma^m\|_{i,j} \|\Sigma^{\delta(j,i')}\|_{j,i'} \|\Sigma^n\|_{i',j'} \\
&\leq V^{-\delta(j,i')} \|\Sigma^{m+n+\delta(j,i')}\|_{i,j'} \\
&\leq V^{-\delta(j,i')} \|\Sigma^{m+n+\delta(j,i')}\| \\
&\leq V^{-\delta(j,i')} D \|\Sigma^{m+n}\| \|\Sigma^{\delta(j,i')}\| \\
&\leq V^{-\delta(j,i')} D \|\Sigma^{m+n}\| D^{\delta(j,i')-1} U^{\delta(j,i')} \\
&= \left(\frac{UD}{V}\right)^{\delta(j,i')} \|\Sigma^{m+n}\| \\
&\leq \left(\frac{UD}{V}\right)^D \|\Sigma^{m+n}\|,
\end{aligned}$$

where the indices i, j, i', j' are chosen to fulfill the first equation. $\square$

With $K = (\frac{V}{UD})^D$, we rewrite the inequality as

$$\|\Sigma^{m+n}\| \geq K \|\Sigma^m\| \|\Sigma^n\|.$$

It means $K\|\Sigma^{m+n}\| \geq (K\|\Sigma^m\|)(K\|\Sigma^n\|)$, i.e., the sequence $K\|\Sigma^n\|$ is supermultiplicative. As this sequence is also positive, $\sqrt[n]{K\|\Sigma^n\|}$ converges to $\sup_n \sqrt[n]{K\|\Sigma^n\|}$ by Fekete's lemma. This is also the limit of $\sqrt[n]{\|\Sigma^n\|}$.

We have now an effective two-sided bound of $\rho(\Sigma)$ for any set Σ of nonnegative matrices with a strongly connected dependency graph.

COROLLARY 4.14. *If the dependency graph is strongly connected, then for every n,*

$$\sqrt[n]{\left(\frac{V}{UD}\right)^D \|\Sigma^n\|} \leq \rho(\Sigma) \leq \sqrt[n]{D\|\Sigma^n\|},$$

where U, V, D are defined as in Theorem 4.12.

Note that the upper bound is due to Proposition 4.1.

When the concern is the value of $\|\Sigma^n\|$, we have

$$\operatorname{const} \rho(\Sigma)^n \leq \|\Sigma^n\| \leq \operatorname{const} \rho(\Sigma)^n. \tag{4.18}$$

Treatment of unconnected dependency graphs. We can extend the treatment to any set of nonnegative matrices, which do not necessarily have a strongly connected dependency graph. Consider the case when the dependency graph is not strongly connected. For each component C, by restricting Σ to the indices in C, we obtain the limit

$$\rho_C = \lim_{n\to\infty} \sqrt[n]{\|\Sigma^n\|_C},$$

where $\|\Sigma^n\|_C$ is defined as in (4.2).

The formula corresponding to (4.18) is

$$\mathrm{const}(\rho_C)^n \le \|\Sigma^n\|_C \le \mathrm{const}(\rho_C)^n.$$

Let $\lambda = \max_C \rho_C$ where the maximum is taken over all the components C. We have[11]

$$\mathrm{const}\,\lambda^n \le \max_C \|\Sigma^n\|_C \le \mathrm{const}\,\lambda^n. \tag{4.19}$$

When the concern is the maximum over all entries, we have the following theorem, which is a more precise quantitative statement than the convergence of $\sqrt[n]{\|\Sigma^n\|}$.

Theorem 4.15. *If $\lambda > 0$ then there exists a nonnegative integer $r < D$ so that for every n,*

$$\mathrm{const}\,n^r\lambda^n \le \|\Sigma^n\| \le \mathrm{const}\,n^r\lambda^n.$$

A direct corollary is $\lambda = \rho(\Sigma)$ and for every n,

$$\mathrm{const}\,n^r\rho(\Sigma)^n \le \|\Sigma^n\| \le \mathrm{const}\,n^r\rho(\Sigma)^n.$$

Note that when $\lambda = 0$, the bound still works for n large enough since there are only finitely many n so that $\|\Sigma^n\| > 0$. Only a small point in the proof uses the condition $\lambda > 0$.

The degree r is one less than the maximum number of components C with $\rho_C = \rho(\Sigma)$ that can be visited by a path in the dependency graph.

The readers may notice that this bound is stronger than the one in Theorem 4.4. In fact, it is possible to prove the same bound using the formula with the diagonal in Theorem 4.4, but it is not as convenient as using the norm. In Chapter 5, we give a bound on $g(n)$: There exists some r so that for every n,

$$\mathrm{const}\,n^{-r}\lambda^n \le g(n) \le \mathrm{const}\,n^r\lambda^n,$$

where λ is the growth rate of the bilinear system there (in case it may get confused with λ here). The lower bound and the upper bound differ by a polynomial, whose degree is not well bounded (due to the nature of the proof in Chapter 5). In contrast, the lower bound and the upper bound in Theorem 4.15 differ at most by a constant factor and we can specify that $r < D$. It is a hint that the problem for the growth of bilinear maps is harder than the one for the joint spectral radius.

Proof. For any i, j, we have

$$(A_1 \dots A_n)_{i,j} = \sum_{\substack{k_0,k_1,\dots,k_n \\ k_0=i,k_n=j}} (A_1)_{k_0,k_1}(A_2)_{k_1,k_2}\dots(A_n)_{k_{n-1},k_n}.$$

Every term of $(A_1 \dots A_n)_{i,j}$ corresponds to a path $k_0, k_1, \dots, k_n$ of length n from i to j. Partitioning the path into vertices of the same components, we obtain some ℓ components

[11]Although we will later prove that $\rho(\Sigma) = \max_C \rho_C$, we for now need another notation λ in the place of $\rho(\Sigma)$, just in case the readers may get confused of the new notation.

$C_1, \ldots, C_\ell$ with m_i edges inside each component C_i. In other words, we have

(4.20)

$$(A_1 \ldots A_n)_{i,j} = \sum_\ell \sum_{C_1,\ldots,C_\ell} \sum_{m_1,\ldots,m_\ell} \sum_{i_1,j_1,\ldots,i_\ell,j_\ell} (A_1 \ldots A_{m_1})_{i_1,j_1} (A_{m_1+1})_{j_1,i_2} (A_{m_1+2} \ldots A_{m_1+m_2+1})_{i_2,j_2} \cdots$$
$$\cdots (A_{n-m_\ell-m_{\ell-1}} \ldots A_{n-m_\ell-1})_{i_{\ell-1},j_{\ell-1}} (A_{n-m_\ell})_{j_{\ell-1},i_\ell} (A_{n-m_\ell+1} \ldots A_n)_{i_\ell,j_\ell},$$

where the sum is taken over all possible choices of: some number ℓ of components to consider, some different components $C_1, \ldots, C_\ell$, some partition of $n-\ell+1$ into nonnegative parts $m_1 + \cdots + m_\ell = n - \ell + 1$ and some indices $i_1, j_1 \in C_1, \ldots, i_\ell, j_\ell \in C_\ell$ with $i_1 = i, j_\ell = j$ (note that if $m_t = 0$ then $i_t = j_t$).

To let the summand of (4.20) be positive, the sequence $C_1, \ldots, C_\ell$ should form a chain in the sense that there is an edge uv with $u \in C_i, v \in C_{i+1}$ for any two consecutive C_i, C_{i+1}. Let r be one less than the maximal number of components C with $\rho_C = \lambda$ that lie in a common chain.

Splitting the sum (4.20) by the number m' of edges inside components C with $\rho_C < \lambda$, we have the sum (4.20) is at most

$$\begin{aligned}
&\sum_\ell \sum_{C_1,\ldots,C_\ell} \sum_{m'} \sum_{\substack{\{m_t : 0<\rho_{C_t}<\lambda\} \\ \sum m_t = m'}} \sum_{\substack{\{m_t : \rho_{C_t}=\lambda\} \\ \sum m_t = n-m'-\ell+1}} \text{const} \left(\prod_{t:0<\rho_{C_t}<\lambda} \|\Sigma^{m_t}\|_{C_t} \right) \left(\prod_{t:\rho_{C_t}=\lambda} \|\Sigma^{m_t}\|_{C_t} \right) \\
\le &\sum_\ell \sum_{C_1,\ldots,C_\ell} \sum_{m'} \text{const} (m')^{\ell-(r'+1)-1} (n-m'-\ell+1)^{r'} (\lambda')^{m'} \lambda^{n-m'-\ell+1} \\
\le &\sum_\ell \sum_{C_1,\ldots,C_\ell} \sum_{m'} \text{const} (m')^\ell n^{r'} \left(\frac{\lambda'}{\lambda}\right)^{m'} \lambda^{n-\ell+1} \\
\le &\sum_\ell \lambda^{1-\ell} \sum_{C_1,\ldots,C_\ell} \text{const}\, n^{r'} \lambda^n \sum_{m'} (m')^\ell \left(\frac{\lambda'}{\lambda}\right)^{m'} \\
\le &\sum_\ell \lambda^{1-\ell} \sum_{C_1,\ldots,C_\ell} \text{const}\, n^{r'} \lambda^n \\
\le &\,\text{const}\, n^r \lambda^n,
\end{aligned}$$

where r' is one less than the number of components C in the chain $C_1, \ldots, C_\ell$ so that $\rho_C = \lambda$, and λ' is some positive number less than λ so that if $\rho_{C_t} < \lambda$ then $\rho_{C_t} < \lambda'$. (The assumption $\lambda > 0$ is to make the products involving t not all empty.) The last inequality is due to the boundedness of ℓ and the number of choices for $C_1, \ldots, C_\ell$. In the next to last inequality, the boundedness of $\sum_{m'} (m')^\ell \left(\frac{\lambda'}{\lambda}\right)^{m'}$ can be seen by the ratio test.

In total, we have $\|\Sigma^n\| \le \text{const}\, n^r \lambda^n$. It remains to show the other direction that $\|\Sigma^n\| \ge \text{const}\, n^r \lambda^n$. Consider a chain $C_1, \ldots, C_\ell$ with $r+1$ components C attaining $\rho_C = \lambda$. We can assume that $\rho_{C_1} = \lambda$ (for a more compact representation of the sum on the right hand side of (4.22)).

Since $\|\Sigma^{n+\delta}\| \le \text{const}\, \|\Sigma^n\| \|\Sigma^\delta\| = \text{const}\, \|\Sigma^n\|$ for any fixed δ, the following inequality holds for any fixed Δ:

$$\|\Sigma^n\| \ge \text{const} \sum_{\delta=0}^{\Delta} \|\Sigma^{n+\delta}\|. \tag{4.21}$$

For two vertices i, j, consider the shortest path $k_0, k_1, \ldots, k_{\delta(i,j)}$ from i to j where $\delta(i,j)$ is the distance from i to j and $k_0 = i, k_{\delta(i,j)} = j$. There are the corresponding matrices $A_1, \dots, A_{\delta(i,j)} \in \Sigma$ so that

$$L(i,j) = (A_1)_{k_0,k_1} \dots (A_{\delta(i,j)})_{k_{\delta(i,j)-1},k_{\delta(i,j)}}$$

is positive. We prove that

$$\sum_{\delta=0}^{D^2} \|\Sigma^{n+\delta}\| \geq \sum_{\substack{m_1+\cdots+m_\ell=n-\ell+1 \\ m_t=0 \text{ for } \rho_{C_t}<\lambda \\ m_1\equiv n-\ell+1 \pmod{2D} \\ m_t\equiv 0 \pmod{2D} \text{ for } \rho_{C_t}=\lambda, t\neq 1}} \|\Sigma^{m_1}\|_{C_1} L(j_1,i_2) \|\Sigma^{m_2}\|_{C_2} \dots \|\Sigma^{m_{\ell-1}}\|_{C_{\ell-1}} L(j_{\ell-1},i_\ell) \|\Sigma^{m_\ell}\|_{C_\ell}, \tag{4.22}$$

where each pair i_t, j_t is chosen depending on m_t so that $\|\Sigma^{m_t}\|_{C_t} = \|\Sigma^{m_t}\|_{i_t,j_t}$ (for $m_t = 0$, we set $i_t = j_t = k$ for any element $k \in C_t$ and let $\|\Sigma^{m_t}\|_{C_t} = 1$).

Each summand on the right hand side of (4.22) is at most $\|\Sigma^N\|$ where $N = m_1 + \delta(j_1, i_2) + m_2 + \cdots + m_{\ell-1} + \delta(j_{\ell-1}, i_\ell) + m_\ell$. One can see that $n \leq N \leq n + D^2$, which explains why $\delta \in [0, D^2]$.

Each summand on the right hand side of (4.22) is the sum of products of the form $(A_1)_{k_0,k_1}(A_2)_{k_1,k_2} \dots (A_N)_{k_{N-1},k_N}$, which correspond to a path $k_0, k_1, \dots, k_N$ from some $k_0 \in C_1$ to some $k_N \in C_\ell$. We will show that no path appears twice in this summation. As $k_0 \in C_1$ and $k_N \in C_\ell$, the path $k_0, \dots, k_N$ goes through $C_1, \dots, C_\ell$, and the transition from C_t to C_{t+1} occurs at some edge k_{p_t}, k_{p_t+1}. Suppose two summands with different $m_1, \dots, m_\ell$ and $m'_1, \dots, m'_\ell$ have some two identical products. It follows that all the transition edges are the same at positions $p_1, \dots, p_{\ell-1}$. Let k^* be the first index so that $m_{k^*} \neq m'_{k^*}$. Due to the congruence modulo $2D$, we have $|m_{k^*} - m'_{k^*}| \geq 2D$, say $m_{k^*} - m'_{k^*} \geq 2D$. Meanwhile,

$$m_{k^*} \leq p_{k^*} - p_{k^*-1} \leq \delta(j_{k^*-1}, i_{k^*}) + m_{k^*} + \delta(j_{k^*}, i_{k^*+1})$$

and

$$m'_{k^*} \leq p_{k^*} - p_{k^*-1} \leq \delta(j'_{k^*-1}, i'_{k^*}) + m'_{k^*} + \delta(j'_{k^*}, i'_{k^*+1}).$$

It raises a contradiction by the strict inequality

$$p_{k^*} - p_{k^*-1} \leq \delta(j'_{k^*-1}, i'_{k^*}) + m'_{k^*} + \delta(j'_{k^*}, i'_{k^*+1}) < m'_{k^*} + 2D \leq m_{k^*} \leq p_{k^*} - p_{k^*-1}.$$

Since no product is accumulated twice, (4.22) follows.

As we set $m_t = 0$ for $\rho_{C_t} < \lambda$, each summand on the right hand side of (4.22) is at least $\operatorname{const} \lambda^n$. Since the number of partitions $m_1 + \cdots + m_\ell = n - \ell + 1$ in (4.22) (even with the divisibility constraints) is $\Theta(n^r)$, the right hand side of (4.22) is at least $\operatorname{const} n^r \lambda^n$. The conclusion follows from (4.21) and (4.22). □

The conclusion. Now we can see that Theorem 4.12 is a corollary of Theorem 4.15.

Proof of Theorem 4.12. If $\lambda = 0$, then $\rho(\Sigma) = 0$ and the bound in Theorem 4.12 trivially holds by (4.19). Suppose $\lambda > 0$. It follows from Theorem 4.15 that $\lambda = \rho(\Sigma)$. Since $\lambda = \max_C \rho_C$, the bound in Corollary 4.14 should now become

$$\sqrt[n]{\left(\frac{V}{UD}\right)^D \max_C \|\Sigma^n\|_C} \leq \rho(\Sigma) \leq \sqrt[n]{D \max_C \|\Sigma^n\|_C}. \qquad \square$$

4.5. On convergence rates

Although we have proved the joint spectral radius theorem for nonnegative matrices as a direct corollary of Theorem 4.4, we give another proof for the theorem using the approach with the norm. Further, we deduce the convergence rates of the sequences related to (4.6).

At first, for any regular component C and any m, we have

$$\|\Sigma^m\|_C = \|\Sigma^m\|_{i,j} \le \text{const}\,\|\Sigma^m\|_{i,j}\|\Sigma^{\delta(j,i)}\|_{j,i} \le \text{const}\,\|\Sigma^{m+\delta(j,i)}\|_{i,i} \le \text{const}\,P_{m+\delta(j,i)}(\Sigma),$$

where $i, j \in C$ are chosen to fulfill the first equality. The last inequality is due to $\|\Sigma^{m+\delta(j,i)}\|_{i,i} = (A_1 \dots A_{m+\delta(j,i)})_{i,i} \le \rho(A_1 \dots A_{m+\delta(j,i)}) \le P_{m+\delta(j,i)}(\Sigma)$ for some matrices $A_1, \dots, A_{m+\delta(j,i)}$.

Note that the inequality $\|\Sigma^m\|_C \le \text{const}\,P_{m+\delta(j,i)}(\Sigma)$ still holds trivially for a single component C.

Taking the maximum over all components, we have

$$\max_C \|\Sigma^m\|_C \le \max_{0\le\delta\le D} \text{const}\,P_{m+\delta}(\Sigma).$$

Applying Theorem 4.12 to a set of only one matrix with $n = 1$, we have

$$\begin{aligned} P_{m+\delta}(\Sigma) = \max_{A_1,\dots,A_{m+\delta}\in\Sigma} \rho(A_1 \dots A_{m+\delta}) &\le \max_{A_1,\dots,A_{m+\delta}\in\Sigma} \text{const}\max_C \|\{A_1 \dots A_{m+\delta}\}\|_C \\ &= \text{const}\max_C \|\Sigma^{m+\delta}\|_C \le \text{const}\,\lambda^{m+\delta} \le \text{const}\max_C \|\Sigma^m\|_C, \end{aligned}$$

where the two last inequalities are due to (4.19). Strictly speaking, it may be the case that several components of the dependency graph of $\{A_1 \dots A_{m+\delta}\}$ form a component of the dependency graph of Σ. However, it does not affect the results.

In total,

$$\text{const}\max_C \|\Sigma^m\|_C \le \max_{0\le\delta\le D} P_{m+\delta}(\Sigma) \le \text{const}\max_C \|\Sigma^m\|_C.$$

It follows from (4.19) that

$$\text{const}\,\lambda^m \le \max_{0\le\delta\le D} P_{m+\delta}(\Sigma) \le \text{const}\,\lambda^m. \tag{4.23}$$

Let $\tilde{P}_m(\Sigma) = \max_{0\le\delta\le D} P_{m+\delta}(\Sigma)$. Although $\sqrt[n]{P_n(\Sigma)}$ does not neccesarily converge, the sequence of $\sqrt[n]{\tilde{P}_n(\Sigma)}$ converges to λ. Together with Theorem 4.15, we have

$$\limsup_{n\to\infty} \sqrt[n]{P_n(\Sigma)} = \lim_{n\to\infty} \sqrt[n]{\|\Sigma^n\|},$$

which is the conclusion of the joint spectral radius theorem. (When Theorem 4.15 does not apply, i.e. $\lambda = 0$, the equality becomes trivial.)

Now we can see the convergence rates of some sequences that clarify the efficiency of the bound in (4.3).

At first, Theorem 4.15 gives

$$\text{const}\,m^r\lambda^m \le \|\Sigma^m\| \le \text{const}\,m^r\lambda^m,$$

which shows that the sequence $\min_{1\le m\le n} \sqrt[m]{D\|\Sigma^m\|}$ converges at the rate $O(\frac{\log n}{n})$ to λ.

On the other hand, it follows from (4.23) that

$$\lambda \min_{0\le\delta\le D} \sqrt[m+\delta]{\text{const}} \le \max_{0\le\delta\le D} \sqrt[m+\delta]{P_{m+\delta}(\Sigma)} \le \lambda \max_{0\le\delta\le D} \sqrt[m+\delta]{\text{const}}.$$

It means the same convergence rate also applies to $\max_{1\le m\le n} \sqrt[m]{P_m(\Sigma)}$.

We finish by comparing the bound in (4.3) and the bound in Theorem 4.12. When $r = 0$, the two bounds are asymptotically equivalent. When $r > 0$, the ratio between the upper bound and the lower bound in (4.3) is at least the n-th root of a polynomial, which may reduce the efficiency when r is large. One may improve the traditional method in (4.3) by considering the components separately when dealing with nonnegative matrices, which then makes the two methods as effective as each other. At any rate, Theorem 4.12 gives explicit constants. We also note that the lower bound in (4.3) is not so self-contained in the sense that it needs another method to estimate the ordinary spectral radius of a matrix. Although the bound in (4.3) in the original form is not as good as the bound in Theorem 4.12, it works for any complex matrices, not necessarily nonnegative ones. An open problem in this direction is how our results would be extended for more general matrices.

4.6. Equivalence of the bounds using diagonals and norms up to a constant

We present a way to deduce each of Theorems 4.7 and 4.12 from the other, using Lemma 4.8. However, the version of Theorem 4.12 that is deduced from Theorem 4.7 is obtained with the weaker constant $\left(\frac{V}{UD}\right)^{3D^2}$ instead of $\left(\frac{V}{UD}\right)^{D}$.

Given some i in some component C, suppose $\|\Sigma^n\|_C = \|\Sigma^n\|_{j,k}$, and let ℓ_1, ℓ_2 be the distance from i to j and from k to i, respectively. We have

(4.24)
$$\|\Sigma^n\|_C = \|\Sigma^n\|_{j,k} = \frac{1}{\|\Sigma^{\ell_1}\|_{i,j}\|\Sigma^{\ell_2}\|_{k,i}}\|\Sigma^{\ell_1}\|_{i,j}\|\Sigma^n\|_{j,k}\|\Sigma^{\ell_2}\|_{k,i} \leq \frac{1}{V^{\ell_1+\ell_2}}\|\Sigma^{n+\ell_1+\ell_2}\|_{i,i}.$$

Note that if $i = j$ (resp. $k = i$), then $\ell_1 = 0$ (resp. $\ell_2 = 0$) and we assume $\|\Sigma^0\|_{i,j} = 1$ (resp. $\|\Sigma^0\|_{k,i} = 1$).

The following proof shares some parts with the proof of Theorem 4.7.

Deduction of Theorem 4.7 from Theorem 4.12. As the lower bound of Theorem 4.7 is trivial, we prove the upper bound.

Fix a regular component C and choose any $i \in C$. Since $\|\Sigma^{m_i}\|_{i,i} > 0$, it follows from (4.24) with $n = m_i$ and Lemma 4.8 that

$$\begin{aligned}
\|\Sigma^{m_i}\|_C &\leq \frac{1}{V^{\ell_1+\ell_2}}(UD)^{\ell_1+\ell_2}\left(\frac{UD}{V}\right)^{3D^2-2D+1}\|\Sigma^{m_i}\|_{i,i} \\
&\leq \left(\frac{UD}{V}\right)^{3D^2-2D+1+\ell_1+\ell_2}\|\Sigma^{m_i}\|_{i,i} \\
&\leq \left(\frac{UD}{V}\right)^{3D^2-1}\|\Sigma^{m_i}\|_{i,i},
\end{aligned}$$

since $\ell_1 \leq D-1, \ell_2 \leq D-1$.

Multiplying by D and taking the root, we obtain

$$\sqrt[m_i]{D\|\Sigma^{m_i}\|_C} \leq \sqrt[m_i]{D\left(\frac{UD}{V}\right)^{3D^2-1}\|\Sigma^{m_i}\|_{i,i}} \leq \sqrt[m_i]{\left(\frac{UD}{V}\right)^{3D^2}\|\Sigma^{m_i}\|_{i,i}}.$$

If C is a single component, that is C contains a single vertex i with no loop, then the above inequality trivially holds with all sides being zeros.

Since $\rho(\Sigma) \leq \sqrt[n]{D\max_C\|\Sigma^n\|_C}$ by Theorem 4.12, it follows that

$$\rho(\Sigma) \leq \max_i \sqrt[m_i]{\left(\frac{UD}{V}\right)^{3D^2}\|\Sigma^{m_i}\|_{i,i}}. \qquad \square$$

Now we deduce Theorem 4.12 with a weaker constant.

Deduction of Theorem 4.12 from Theorem 4.7. As the upper bound of Theorem 4.12 is trivial, we prove the lower bound. We start with (4.24):

$$\|\Sigma^n\|_C \leq \frac{1}{V^{\ell_1+\ell_2}}\|\Sigma^{n+\ell_1+\ell_2}\|_{i,i}.$$

Suppose C is a regular component, which means both sides of the above inequality are positive. Let $\delta \leq D$ be the length of the shortest cycle from i to i. Since $\|\Sigma^{t\delta}\|_{i,i} > 0$ for any $t \geq 1$, two consecutive elements in the set $\{\ell : \|\Sigma^\ell\|_{i,i} > 0\}$ have the distance at most $\delta \leq D$. Therefore, if $n + \ell_1 + \ell_2 > D$, there exists some positive integer $m < n + \ell_1 + \ell_2$ with $n+\ell_1+\ell_2-m \leq D$ so that $\|\Sigma^m\|_{i,i} > 0$. We first consider the case $n+\ell_1+\ell_2 > D$, it follows from Lemma 4.8 that

$$\begin{aligned}
\|\Sigma^n\|_C &\leq \frac{1}{V^{\ell_1+\ell_2}}(UD)^{n+\ell_1+\ell_2-m}\left(\frac{UD}{V}\right)^{3D^2-2D+1}\|\Sigma^m\|_{i,i} \\
&\leq (UD)^{n-m}\left(\frac{UD}{V}\right)^{3D^2-2D+1+\ell_1+\ell_2}\|\Sigma^m\|_{i,i} \\
&\leq (UD)^{n-m}\left(\frac{UD}{V}\right)^{3D^2-2D+1+\ell_1+\ell_2}\rho(\Sigma)^m \\
&\leq (UD)^{n-m}\left(\frac{UD}{V}\right)^{3D^2-2D+1+\ell_1+\ell_2}\rho(\Sigma)^n\rho(\Sigma)^{m-n} \\
&\leq \left(\frac{UD}{\rho(\Sigma)}\right)^{n-m}\left(\frac{UD}{V}\right)^{3D^2-2D+1+\ell_1+\ell_2}\rho(\Sigma)^n.
\end{aligned}$$

If $n \geq m$, it follows from $\rho(\Sigma) \geq V$ that

$$\begin{aligned}
\|\Sigma^n\|_C &\leq \left(\frac{UD}{V}\right)^{n-m}\left(\frac{UD}{V}\right)^{3D^2-2D+1+\ell_1+\ell_2}\rho(\Sigma)^n \\
&= \left(\frac{UD}{V}\right)^{3D^2-2D+1+\ell_1+\ell_2+n-m}\rho(\Sigma)^n \\
&\leq \left(\frac{UD}{V}\right)^{3D^2-D+1}\rho(\Sigma)^n.
\end{aligned}$$

When $n < m$, it follows from $\rho(\Sigma) \leq UD$ that

$$\begin{aligned}
\|\Sigma^n\|_C &\leq \left(\frac{UD}{V}\right)^{3D^2-2D+1+\ell_1+\ell_2}\rho(\Sigma)^n \\
&\leq \left(\frac{UD}{V}\right)^{3D^2-1}\rho(\Sigma)^n,
\end{aligned}$$

since $\ell_1 \leq D-1, \ell_2 \leq D-1$.

Combining the two cases gives

$$\|\Sigma^n\|_C \leq \left(\frac{UD}{V}\right)^{3D^2}\rho(\Sigma)^n.$$

In the remaining case that $n + \ell_1 + \ell_2 \leq D$, we also have

$$\|\Sigma^n\|_C \leq \|\Sigma^n\| \leq \frac{1}{D}(UD)^n \leq \left(\frac{UD}{V}\right)^n V^n \leq \left(\frac{UD}{V}\right)^{3D^2}\rho(\Sigma)^n,$$

since $n \le n + \ell_1 + \ell_2 \le D \le 3D^2$.

When C is a single component, the above inequality trivially holds with $\|\Sigma^n\|_C = 0$.

Taking over all the components, we obtain the conclusion

$$\rho(\Sigma) \ge \sqrt[n]{\left(\frac{V}{UD}\right)^{3D^2} \max_C \|\Sigma^n\|_C}. \qquad \square$$

Remark. *In the above proof, only the trivial lower bound in Theorem 4.7 is used. In fact, in the deduction of Theorem 4.7 from Theorem 4.12, only the trivial upper bound of the latter is used. In other words, Lemma 4.8 is actually the key element in both proofs.*

CHAPTER 5

A formula and two bounds for the growth rate

This chapter will treat the growth of bilinear maps with a formula for the growth rate and some bounds. The key tools are the notion of linear pattern and Fekete's lemma. We remind the general assumption to avoid the degenerate cases, which is already mentioned in Chapter 2, that for every i, there exists some n so that $g_i(n) > 0$.

Before introducing linear patterns, we show a correspondence between the rooted binary trees of n leaves and the combinations of n instances of the vector s. In one direction, we let the expression associated with a tree of a single leaf be the vector s itself, and the expression associated with a tree of a higher number of leaves be $(L * R)$, where L, R are respectively the expressions associated with the left and right branches. The other direction is obvious as the previous association is a one-to-one mapping. Note that the map from the binary trees to the resulting vectors of the combinations is however not injective. Given a resulting vector v, we just pick any binary tree that gives v to be the *associated tree* with v. The arguments in the chapter are independent of the choice. The other direction is deterministic: The vector that the binary tree gives is said to be the *associated vector* with the tree. Note that from now on, all considered trees are rooted binary trees. In some places, we say the tree associated with $g_k(n)$ instead of saying the tree of n leaves associated with a vector whose k-th entry is $g_k(n)$ for short. The same manner is also applied for $g(n)$.

A *linear pattern* $P = (T, \ell)$ is a pair of a tree T and a marked leaf ℓ in T. Suppose in the expression associated with T, we put a vector variable u instead of the fixed vector s in the place associated with the leaf ℓ. The value of the expression is then a vector variable v that is a linear function of u. Let $M = M(P)$ be the matrix representing the dependency, that is $v = Mu$. The matrix M is said to be the *matrix of the linear pattern*.

For two patterns $P_1 = (T_1, \ell_1)$ and $P_2 = (T_2, \ell_2)$, the composition of the two patterns, denoted by $P_1 \oplus P_2$, is the pattern $P = (T, \ell)$ where T is obtained from T_1 by replacing ℓ_1 by T_2, and ℓ is the leaf ℓ_2 in this instance of T_2. We denote by P^q the pattern $P \oplus P \oplus \cdots \oplus P$ with q instances of P. A quick observation is $M(P_1 \oplus P_2) = M(P_1)M(P_2)$ and $M(P^q) = (M(P))^q$.

For a pattern $P = (T, \ell)$, we denote by $|P|$ the number of leaves excluding ℓ in T. We have $|P_1 \oplus P_2| = |P_1| + |P_2|$ and $|P^q| = q|P|$ for any patterns P, P_1, P_2.

When we regard "the number of leaves of pattern P", we mean $|P|$. In most of the cases, the distinction between $|P|$ and the number of leaves in T does not matter very much. However, one must be careful when taking the root, it must be the $|P|$-th root.

For convenience, we also write $P \oplus T$ for a tree T to denote a tree obtained from the tree of P by replacing its marked leaf by T.

Before presenting the formula and the bounds for the growth rate λ in the following sections, we give first some observations.

OBSERVATION 5.1. *Consider a pattern $P = (T, \ell)$. If $s_j > 0$, then $M_{i,j} \le \operatorname{const} g_i(n)$ for any i where n is the number of leaves in T. If $s_j = 0$, we have $M_{i,j} \le \operatorname{const} g_i(n + O(1))$ for any i.*

PROOF. Let v be the vector associated with T, we have

$$g_i(n) \geq v_i = \sum_k M_{i,k} s_k \geq M_{i,j} s_j.$$

It follows that $M_{i,j} \leq \operatorname{const} g_i(n)$ for the constant $\frac{1}{s_j}$ when $s_j > 0$.

For the case $s_j = 0$, we still have $g_j(m) > 0$ for some m with the associated tree T_0. Consider the tree $P \oplus T_0$ with the associated vector v, we have $v_i \geq M_{i,j} g_j(m)$. It follows that $M_{i,j} \leq \operatorname{const} v_i \leq \operatorname{const} g_i(n + O(1))$ since m is fixed. □

The *dependency graph* of a system, which expresses the dependency among the dimensions in $*$, is defined to be a directed graph that takes the dimensions as the vertices. There is an edge from k to i if and only if there exists some j so that either $c_{i,j}^{(k)} \neq 0$ or $c_{j,i}^{(k)} \neq 0$. As a directed graph, the dependency graph can be partitioned into strongly connected components, for which we call *components* for short. We define a partial order between the components: A component C_1 is said to be smaller than a component C_2 for $C_1 \neq C_2$ if there are vertices $i \in C_2$, $j \in C_1$ and a path from i to j. For any C_1, C_2, we say $C_1 \leq C_2$ if $C_1 < C_2$ or $C_1 = C_2$.

The following observation is useful in various proofs.

OBSERVATION 5.2. *If there is a path from i to j, then there exists a pattern whose matrix M satisfies $M_{i,j} > 0$. It follows that there exists some δ so that $g_i(n+\delta) \geq \operatorname{const} g_j(n)$ for every n. On the other hand, if M is the matrix of some pattern and $M_{i,j} > 0$, then there is a path from i to j.*

PROOF. Suppose there is an edge from k to i with $c_{i,j}^{(k)} > 0$. Consider the pattern (T, ℓ) for a tree T with ℓ being the left branch, the right branch is the tree associated with $g_j(m) > 0$ for some m. The matrix M of the pattern has $M_{k,i} > 0$. A similar construction is for $c_{j,i}^{(k)} > 0$ with the marked leaf ℓ being the right branch. Suppose there is a path $k_0, k_1, \dots, k_d$ from i to j, with $k_0 = i$ and $k_d = j$. The desired pattern is $P_1 \oplus \cdots \oplus P_d$ where P_t is the pattern constructed from the edge $k_{t-1}k_t$.

Given such a pattern $P = P_1 \oplus \cdots \oplus P_d$, we can see that $g_i(n + |P|) \geq \operatorname{const} g_j(n)$ by considering the tree $P \oplus T^*$ where the tree T^* is associated with $g_j(n)$.

In the other direction, let P be the pattern whose matrix M satisfies $M_{i,j} > 0$. Consider the decomposition $P = P_1 \oplus \cdots \oplus P_t$ so that each P_k has the marked leaf being a child of the root. Let M_k be the matrix of P_k, we have $M = M_1 \dots M_t$, that is

$$M_{i,j} = \sum_{k_1,\dots,k_{t-1}} (M_1)_{i,k_1} (M_2)_{k_1,k_2} \dots (M_{t-1})_{k_{t-2},k_{t-1}} (M_t)_{k_{t-1},j}.$$

As $M_{i,j} > 0$, there exist $k_1, \dots, k_{t-1}$ so that all

$$(M_1)_{i,k_1}, (M_2)_{k_1,k_2}, \dots, (M_{t-1})_{k_{t-2},k_{t-1}}, (M_t)_{k_{t-1},j}$$

are positive. It follows that there are edges ik_1, k_1k_2, $\dots$, $k_{t-2}k_{t-1}$, $k_{t-1}j$, which form the path $i, k_1, k_2, \dots, k_{t-1}, j$. □

5.1. A formula for λ and a polynomial upper bound for $g(n)/\lambda^n$

We prove the following representation of the growth rate, and provide a polynomial upper bound for $g(n)/\lambda^n$ at the same time. It is done via the quantity

$$\theta = \sup_{\text{linear pattern } P} \max_i \sqrt[|P|]{M(P)_{i,i}}.$$

THEOREM 5.3. *We have*

$$\lambda = \theta.$$

Moreover, there exists r so that for every n,

$$g(n) \le \text{const}\, n^r \lambda^n.$$

OBSERVATION 5.4. $\lambda \ge \theta$.

PROOF. Consider a linear pattern P. Choose an r so that $g_i(r) > 0$ by a tree T_0, and consider the sequence of trees $P^q \oplus T_0$ for $q = 1, 2, \ldots$. Each tree $P^q \oplus T_0$ has $n = q|P| + r$ leaves and the associated vector has the i-th entry at least $g_i(r)(M(P)_{i,i})^q$. The lower bound $\lambda \ge \theta$ follows from $\lim_{q\to\infty} \sqrt[n]{g_i(r)(M(P)_{i,i})^q} = \sqrt[|P|]{M(P)_{i,i}}$. ☐

OBSERVATION 5.5. *Consider a linear pattern P. For any i, j of the same component, we have*

$$M(P)_{i,j} \le \text{const}\, \theta^{|P|}. \tag{5.1}$$

PROOF. When the component contains a single vertex without loops, the observation is trivial. In the remaining situation, let $P_{j\to i}$ be a pattern so that $M(P_{j\to i})_{j,i} > 0$. We have $M(P \oplus P_{j\to i})_{i,i} \ge M(P)_{i,j} M(P_{j\to i})_{j,i}$. Meanwhile, $M(P \oplus P_{j\to i})_{i,i} \le \theta^{|P\oplus P_{j\to i}|} \le \text{const}\, \theta^{|P|}$. The observation is established since $M(P_{j\to i})_{j,i}$ is a constant. ☐

The proof of Theorem 5.3 uses the following proposition. The strange condition in the proposition is in fact not necessary for the conclusion, by Theorem 5.3 itself. However, we need the proposition in the induction step of the proof of the theorem.

PROPOSITION 5.6. *Consider a vertex i with the condition that there exists some α so that for every j in a component lower than the component of i, we have $g_j(m) = O(m^\alpha \theta^m)$. Suppose k is a vertex in a component lower than the component of i and $g_k(m) \ne O(\theta^m)$, then there exists some r so that for every linear pattern P with $|P| = n$ and the associated matrix M, we have*

$$M_{i,k} \le \text{const}\, n^r \theta^n.$$

While $g_j(m) = O(m^r\theta^m)$ straightforwardly means the existence of some $K > 0$ so that $g_j(m) \le K m^r \theta^m$ for every m, the not very popular notation $g_k(m) \ne O(\theta^m)$ can be interpreted as: For every $K > 0$ there exists some m so that $g_k(m) > K\theta^m$.

Note that we write $f(x) = O(e(x))$ for two functions f, e in this part only, instead of writing $f(x) \le \text{const}\, e(x)$ as usual.[1] The reason is to highlight the meaning of the notation $f(x) \ne O(e(x))$.

PROOF. Decompose P into $P = P_1 \oplus \cdots \oplus P_t$ so that each pattern has the marked leaf being a child of the root. Let $M_1, \ldots, M_t$ be the associated matrices of $P_1, \ldots, P_t$, respectively. We have $M = M_1 \ldots M_t$. Let C be the component of i. The entry $M_{i,k}$ can be written as

$$M_{i,k} = \sum_{\substack{1\le s\le t\\ j_1\in C,\, j_2\notin C}} (M_1 \ldots M_{s-1})_{i,j_1} (M_s)_{j_1,j_2} (M_{s+1} \ldots M_t)_{j_2,k}, \tag{5.2}$$

where $j_1 = i$ if $s = 1$, and $j_2 = k$ if $s = t$. (Note that $j_1 j_2$ is the edge where the path leaves C.)

We proceed by considering all the nonzero summands. It means that j_2 is in a component directly lower than C, in order for $(M_s)_{j_1,j_2}$ to be nonzero. We can conclude right away that $(M_{s+1} \ldots M_t)_{j_2,k} \le \text{const}\, g_{j_2}(m + O(1)) \le \text{const}\, m^\alpha \theta^m$ for $m = |P_{s+1}| + \cdots + |P_t|$,

[1] Strictly speaking, $f(x) \le \text{const}\, e(x)$ is a stronger conclusion than $f(x) = O(e(x))$ but they are asymptotically equivalent.

by the condition of the proposition. Also, by Observation 5.5, we have $(M_1 \dots M_{s-1})_{i,j_1} \leq \operatorname{const}\theta^{|P_1|+\cdots+|P_{s-1}|}$ as i, j_1 are both in C.

It remains to consider $(M_s)_{j_1,j_2}$. Suppose the marked leaf of P_s is on the left branch, without loss of generality. Let v be the vector associated with the right branch of the tree of P_s. We have

$$(M_s)_{j_1,j_2} = \sum_j c^{(j_1)}_{j_2,j} v_j.$$

Note that in order for the summand in (5.2) to be nonzero, we need $(M_{s+1} \dots M_t)_{j_2,k} > 0$, which implies a path from j_2 to k by Observation 5.2. Another corollary of Observation 5.2 is that $g_{j_2}(m) \neq O(\theta^m)$. We can see that for any j so that $c^{(j_1)}_{j_2,j} > 0$, the vertex j is not in C. Indeed, assume otherwise, let K be fixed and Q be a tree of m leaves so that the associated vector u has $u_{j_2} = g_{j_2}(m) > K\theta^m$. Let $\hat{P}$ be the linear pattern where the left branch is Q and the right branch is the marked leaf. The associated matrix $\hat{M}$ has

$$\hat{M}_{j_1,j} = \sum_{j'} c^{(j_1)}_{j',j} u_{j'} \geq c^{(j_1)}_{j_2,j} u_{j_2} > \operatorname{const} K\theta^m,$$

contradicting Observation 5.5 as K can be arbitrarily large. Therefore, for such j we have $v_j \leq g_j(|P_s|) \leq \operatorname{const}|P_s|^\alpha \theta^{|P_s|}$ by the condition of the proposition. In other words,

$$(M_s)_{j_1,j_2} = \sum_j c^{(j_1)}_{j_2,j} v_j \leq \operatorname{const}|P_s|^\alpha \theta^{|P_s|}.$$

In total, each summand in (5.2) is at most a constant times $n^{2\alpha}\theta^n$. Meanwhile, there are only at most n options for the summation variable s and constantly many options for j_1, j_2. Therefore, for $r = 2\alpha + 1$, we have

$$M_{i,k} \leq \operatorname{const} n^r \theta^n. \qquad \square$$

We are now ready to prove Theorem 5.3.

Proof of Theorem 5.3. It suffices to prove that: For every i there exists r so that

$$g_i(n) \leq \operatorname{const} n^r \theta^n. \tag{5.3}$$

Indeed, suppose we have (5.3). Together with Observation 5.4, we have

$$\theta \leq \lambda = \limsup_{n\to\infty} \sqrt[n]{g(n)} \leq \limsup_{n\to\infty} \sqrt[n]{\operatorname{const} n^r \theta^n} = \theta,$$

which means $\lambda = \theta$ and $g(n) = \max_i g_i(n) \leq \operatorname{const} n^r \lambda^n$.

We can prove (5.3) by induction on the components. The base case that (5.3) holds for any i in a minimal component is established with the help of Observation 5.5. For such an i, let P be any pattern with the tree associated to $g_i(n)$, we have $g_i(n) = \sum_j M(P)_{i,j} s_j \leq \operatorname{const}\theta^n$, since j needs to be in the same component as i for $M(P)_{i,j}$ to be nonzero.

As for the induction step, we now consider an i that is not in a minimal component and suppose (5.3) holds for any vertex in a component lower than the component of i, we prove that it also holds for i.

Let T be the tree associated with $g_i(n)$. Pick a subtree T_0 of m leaves so that $n/3 \leq m \leq 2n/3$ by Lemma 2.7. Consider the decomposition $T = P' \oplus T_0$. Let M' be the matrix associated with P' and u be the vector associated with T_0. We have

$$g_i(n) = \sum_k M'_{i,k} u_k \leq \operatorname{const} M'_{i,j} u_j \leq \operatorname{const} M'_{i,j} g_j(m)$$

for some j that maximizes $M'_{i,j} u_j$.

Let C be the component of i. We consider the following cases regarding j:

- If j is in a lower component than C but $g_j(t) \neq O(\theta^t)$, then we have $M'_{i,j} \leq \text{const}(n-m)^\alpha \theta^{n-m}$ for some α, due to Proposition 5.6. By the induction hypothesis, we also have $g_j(m) \leq \text{const}\, m^\beta \theta^m$ for some β. In total,

$$g_i(n) \leq \text{const}(n-m)^\alpha \theta^{n-m} m^\beta \theta^m \leq \text{const}\, n^{\alpha+\beta}\theta^n = \text{const}\, n^\gamma \theta^n$$

where $\gamma = \alpha + \beta$.
- If j is in a lower component than C and $g_j(t) = O(\theta^t)$, then together with $M'_{i,j} \leq \text{const}\, g_i(n - m + O(1))$ by Observation 5.1 we have

$$g_i(n) \leq K g_i(n-m+O(1))\theta^m$$

for some large enough K.
- If j is in C, then we have $M'_{i,j} \leq \text{const}\, \theta^{n-m}$ by Observation 5.5. Therefore,

$$g_i(n) \leq K\theta^{n-m} g_j(m)$$

for some large enough K.

In any of the two latter cases, we have reduced the size n considerably by at least a fraction of n but still keep considering g_k for some k in the component of i. After repeating the process at most $O(\log n)$ times, and stopping only when the current n is small enough or we fall into the first case, we obtain

$$g_i(n) \leq \text{const}\, K^{O(\log n)} n^\gamma \theta^{n+O(\log n)}$$

for the sufficiently large constant K as specified in the two latter cases.

As $x^{\log n} = n^{\log x}$, the induction step finishes since for some r,

$$g_i(n) \leq \text{const}\, n^r \theta^n.$$

The conclusion follows by induction. □

Remark 5.7. *The value of r in the proof in principle depends on the coefficients of $*$ and the entries of s. However, it is probably due to the techniques of the proof only. We believe that the conclusion still holds if we set $r = 2^{d-1}$ where d is the dimension, regardless of the coefficients/entries of $*, s$.*

We provide a condition in the nonnegative setting so that λ is a limit.

Theorem 5.8. *Suppose there exists some n_0 so that for every $n \geq n_0$ and every i we have $g_i(n) > 0$, then λ is actually a limit.*

Proof. By Theorem 5.3, it suffices to prove that

$$\liminf_{n\to\infty} \sqrt[n]{g(n)} \geq \sup_{\text{linear pattern } P} \max_i \sqrt[|P|]{M(P)_{i,i}},$$

which can be reduced to showing that for any pattern P and any index i, we have

$$\liminf_{n\to\infty} \sqrt[n]{g(n)} \geq \sqrt[|P|]{M(P)_{i,i}}.$$

Indeed, for every n large enough, let $n = q|P| + r$ so that $n_0 \leq r < n_0 + |P|$. Consider the tree $P^q \oplus T_0$ where T_0 is the tree associated with $g_i(r)$. The i-th entry of the associated vector is at least a constant times $(M(P)_{i,i})^q$. Since r is bounded, the conclusion follows. □

Note that the condition is satisfied in the positive setting. In other words, we have obtained a simpler proof of the limit λ than the one in [**8**], whose approach is briefly mentioned in Chapter 7.

Corollary 5.9. *If $s > 0$ then the growth rate λ is a limit.*

An important corollary of the formula in Theorem 5.3 is that the growth rate is computable, in the sense that λ can be computed to any desired precision, as in Theorem 5.10 below.

THEOREM 5.10. *The growth rate in the nonnegative setting is computable.*

PROOF. It was known that the growth rate λ in the nonnegative setting is upper semi-computable [**4**], in the sense that there exists a sequence of upper bounds converging to λ. It remains to show that it is lower semi-computable, by showing a sequence of lower bounds converging to λ. As

$$\lambda = \sup_{\text{linear pattern } P} \max_i \sqrt[|P|]{M(P)_{i,i}},$$

we have

$$\lambda = \sup_n \max_{\substack{\text{linear pattern } P \\ |P|=n}} \max_i \sqrt[n]{M(P)_{i,i}}.$$

The sequence

$$a_n = \max_{\substack{\text{linear pattern } P \\ |P|\le n}} \max_i \sqrt[|P|]{M(P)_{i,i}}$$

for $n = 1, 2, \ldots$ is indeed the desired sequence since it is increasing and converges to λ. $\square$

Although the growth rate λ is computable in the sense of computability, there is no guarantee yet on the convergence rate of the approximations. The next section will give efficient bounds for λ in the positive setting.

5.2. A polynomial upper bound for $\lambda^n/g(n)$ when $s > 0$

As an upper bound for $g(n)$ is established, one might expect a matching lower bound. In this section, we attempt to give such a bound in the positive setting, where s is always positive. Note that although the growth rate is computable in the nonnegative setting, the convergence rate is still unknown and likely to be poor. On the other hand, the following matching lower bound on $g(n)$ sheds light on the convergence rate of the estimations for $s > 0$. We say it is "matching" due to the existence of the polynomial in the bound, like in Theorem 5.3.

THEOREM 5.11. *If s is a positive vector then there exists some r so that for every n,*

$$g(n) \ge \text{const}\, n^{-r} \lambda^n.$$

When s is positive, it follows from Theorem 5.11 and Theorem 5.3 that the limit λ exists. However, it is a much harder way than proving it by Corollary 5.9.

When s is not positive, the lower bound may not hold, as otherwise the limit λ exists, which we already know is not always the case as in the examples in Chapter 2. It is possible to extend the approach to the nonnegative setting, in the way that instead of $g(n)$ we give a lower bound for $\sum_{0\le\delta\le\Delta} g(n+\delta)$ for some Δ. However, we avoid treating the case as such arguments ask for a lot more care.

The leading constant of the lower bound and r in Theorem 5.11 depend on the coefficients of $*$ and the entries of the starting vector s, due to the techniques in the proof. However, we believe that it is possible to replace n^{-r} by n^t (a nonnegative degree here) so that $t \ge 0$ and $t \le 2^{d-1}$ for the space $\mathbb{R}^d$.

Before presenting the proof, we would note that we try to keep this section as independent from Section 5.1 as possible.

Preliminary lemmas.

LEMMA 5.12. *Let T be the tree associated to $g_k(n)$. If there is a subtree T_0 of m leaves, then there exists some j so that there is a path from k to j and*

$$g_k(n) \leq \text{const}\, g_k(n-m+1)g_j(m).$$

PROOF. Let v, u be the vectors associated with T, T_0, respectively. Consider the decomposition $T = P' \oplus T_0$. Let M be the matrix associated with P'. Since $v = Mu$, we have

$$g_k(n) = v_k = (Mu)_k = \sum_i M_{k,i}u_i \leq \text{const}\, M_{k,j}u_j$$

for some j that maximizes $M_{k,j}u_j$ (the constant can be chosen to be the dimension). As $M_{k,j} > 0$, there is a path from k to j.

The conclusion follows from the fact that $M_{k,j} \leq \text{const}\, g_k(n-m+1)$ (by Observation 5.1) and $u_j \leq g_j(m)$. □

COROLLARY 5.13. *For any k and any fixed d, we have $g_k(n+d) \leq \text{const}\, g_k(n)$ for every n.*

PROOF. Since every tree of at least 2 leaves has a subtree of 2 leaves, setting $m = 2$ for Lemma 5.12 gives $g_k(n+1) = g_k((n-1)+2) \leq \text{const}\, g_k(n)g_j(2) \leq \text{const}\, g_k(n)$ for every n. The conclusion follows by applying d times the inequality $g_k(n+1) \leq \text{const}\, g_k(n)$. □

The other direction also holds under a condition.

OBSERVATION 5.14. *If k is in a component that has at least one edge inside (loops are also counted), then for any fixed d, we have*

$$g_k(n+d) \geq \text{const}\, g_k(n).$$

It follows that $g_k(n)$ and $g_k(n+d)$ are in a constant factor of each other for such a component.

PROOF. By the condition of the component of k, there is a path from k to k, that is $g_k(n+d) \geq \text{const}\, g_k(n+d-\delta)$ for some δ by Observation 5.2. (We only consider n large enough, e.g. greater than δ, since smaller n can be treated by adjusting the leading constant in the lower bound.) Applying repeatedly this fact until we have $g_k(n+d) \geq \text{const}\, g_k(n_0)$ so that $n_0 < n$. Note that $n - n_0$ is bounded. By Corollary 5.13, we have $g_k(n_0) \geq \text{const}\, g_k(n)$, the conclusion follows. □

The condition in Observation 5.14 can be relaxed, but the current form is enough for later applications.

COROLLARY 5.15. *Let T be the tree associated to $g_k(n)$. If there is a subtree T_0 of m leaves, then*

$$g_k(n) \leq \text{const}\, g_k(n-m)g_k(m).$$

PROOF. By Lemma 5.12, we have

$$g_k(n) \leq \text{const}\, g_k(n-m+1)g_j(m) \leq \text{const}\, g_k(n-m)g_j(m),$$

where the latter inequality is due to Corollary 5.13.

Since there is a path from k to j, by Observation 5.2, for some d we have

$$g_j(m) \leq \text{const}\, g_k(m+d) \leq \text{const}\, g_k(m),$$

where the latter inequality is also due to Corollary 5.13.

The conclusion follows. □

Definition 5.16. *Given a component C, the C-subsystem is the system with $(*', s')$ deduced from $(*, s)$ by restricting $(*, s)$ to only the dimensions reachable from the vertices in C. In particular, $s'_i = s_i$ for $i \in C'$ and $C' \leq C$ (other dimensions i are removed). Likewise, the coefficients $c'^{(k)}_{i,j} = c^{(k)}_{i,j}$ for $k \in C'$ and $C' \leq C$.*

One can observe that the value $g_k(n)$ as well as the k-th entry of any resulting vector for a dimension k in the C-subsystem are the same as those in the original system, as they do not depend on the dimensions outside the C-subsystem.

The following proposition is [**8**, Lemma 5].

Proposition 5.17. *Given some k in a component C, in the C-subsystem we have*

$$g_k(n) \geq \text{const}\, g(n),$$

where $g(n)$ is subject to the C-subsystem.

Proof. For every dimension i in the C-subsystem, there is a path from k to i. It follows from Observation 5.2 that for some d we have

$$g_i(n) \leq \text{const}\, g_k(n + d) \leq \text{const}\, g_k(n),$$

where the latter inequality is due to Corollary 5.13.

The conclusion follows by

$$g(n) = \max_i g_i(n) \leq \text{const}\, g_k(n). \qquad \square$$

Corollary 5.18. *In the C-subsystem, the values $g_k(n)$ and $g(n)$ (subject to the C-subsystem) are in a constant factor of each other. So are the values of $g_i(n)$ and $g_j(n)$ if i, j are in the same component.*

Definition 5.19. *We denote the limit $\lambda_k = \lim_{n\to\infty} \sqrt[n]{g_k(n)}$.*

One can observe that $\lambda_i = \lambda_j$ if i, j are in the same component. Also, for a vertex k in C, the value λ_k is the growth rate of the C-subsystem. If $i \in C_1, j \in C_2$ and $C_1 \leq C_2$, we have $\lambda_i \leq \lambda_j$, by Observation 5.2. In other words, the order of the components is also the order of the growth rates.

A classification of components. We classify all the components into the three following classes, for which the growth rates for each class of components can be studied individually in a more convenient way later.

Definition 5.20. *A component C is said to be strongly self-dependent if there are three indices k, i, j in C (not necessarily different) so that $c^{(k)}_{i,j} > 0$.*

Definition 5.21. *A component C is said to be* weakly self-dependent *if for every $k \in C$ and for any $c^{(k)}_{i,j} > 0$ at least one of i, j is in a lower component than C, and for any j in a component lower than the component of k we have $\lambda_j < \lambda_k$.*

Before classifying the remaining components, we give the following observation.

Observation 5.22. *Let k be so that for each $c^{(k)}_{i,j} > 0$, both i, j are in a lower component than the component of k, we have*

$$\lambda_k = \max_{i,j:\ c^{(k)}_{i,j} > 0} \max\{\lambda_i, \lambda_j\}.$$

Note that when there is no edge from k (that is $c^{(k)}_{i,j} = 0$ for every i, j), we have $\lambda_k = 0$ with the convention that the maximum of an empty list is zero.

COROLLARY 5.23. *For any k in a weakly self-dependent component C, there exist i, j so that $c_{i,j}^{(k)} > 0$ and one of i, j is in C while the other one is in a component lower than C.*

The remaining components C (other than strongly self-dependent and weakly self-dependent) satisfy (i) for every $k \in C$ and for any $c_{i,j}^{(k)} > 0$ at least one of i, j is in a lower component than C and (ii) there exists $k \in C$ and j in a lower component than C so that $\lambda_k = \lambda_j$. It means the remaining components are included in the following class.

DEFINITION 5.24. *A component C is said to be* not self-dependent *if for each $k \in C$ we have $\lambda_k = \lambda_j$ for some j in a lower component than C.*

In total, we have the following proposition.

PROPOSITION 5.25. *The three classes of components: strongly self-dependent, weakly self-dependent and non-self-dependent components cover all the components.*

Strongly self-dependent components. We show that $g_k(n)$ is weakly supermultiplicative for k in a strongly self-dependent component.

THEOREM 5.26. *Let k be in a strongly self-dependent component. For any m, n, we have*

$$g_k(m + n) \geq \text{const}\, g_k(m)g_k(n).$$

PROOF. Let i, j be in the component of k so that $c_{i,j}^{(k)} > 0$. Using the preliminary results, we have

$$g_k(m)g_k(n) \leq \text{const}\, g_i(m + d_1)g_j(n + d_2) \leq \text{const}\, g_k(m + n + d_1 + d_2) \leq \text{const}\, g_k(m + n)$$

for some bounded d_1, d_2. The first inequality is due to Observation 5.2, while the last inequality is due to Corollary 5.13. The middle inequality is obtained by considering a tree where the left branch is associated to $g_i(m + d_1)$ and the right branch is associated to $g_j(n + d_2)$.

An alternate argument is

$$g_k(m)g_k(n) \leq \text{const}\, g_i(m)g_j(n) \leq \text{const}\, g_k(m + n),$$

where the first inequality is by Corollary 5.18. □

An instance of the condition is the strong connectedness of the dependency graph. Corollaries of the result include the limit of $\sqrt[n]{g_k(n)}$ and $g_k(n) \leq \text{const}\, \lambda_k{}^n$ by applying Fekete's lemma to the supermultiplicative sequence $\{\text{const}\, g_k(n)\}_n$. The upper bound is a case of Theorem 5.3 as we reduce the polynomial n^r to $n^0 = 1$.

REMARK 5.27. *The two arguments in the proof of Theorem 5.26 both use not so trivial propositions. We can avoid using them and obtain a slightly weaker result, which still implies the bound of $g_k(n)$ and the limit. Indeed, after obtaining the inequality $g_k(m)g_k(n) \leq K g_k(m + n + d_1 + d_2)$ for some constant K as in the first half of the first argument, we shift the sequence and multiply both sides by K to get*

$$K g_k(m - d_1 - d_2) K g_k(n - d_1 - d_2) \leq K g_k(m + n - d_1 - d_2).$$

The sequence $s_n = K g_k(n - d_1 - d_2)$ is supermultiplicative. By Fekete's lemma, we have $\sqrt[n]{s_n}$ converges to $\lambda_k = \sup_n \sqrt[n]{s_n}$. The original sequence $\sqrt[n]{g_k(n)}$ also converges to λ_k. The bound also follows.

The argument shows that the proof of the limit λ would become fairly simple when the dependency graph is strongly connected.

The upper bound const ${\lambda_k}^n$ is a nice corollary of Fekete's lemma. We naturally wonder what a lower bound would be, and whether the leading constant should be replaced by something arbitrarily small. Actually, we would conjecture that $g_k(n) \geq \text{const}\, {\lambda_k}^n$. However, what we could come up is just the following result. It is also a corollary of Fekete's lemma, but for a submultiplicative sequence. The interesting point is that the supermultiplicative form as in Theorem 5.26 is used in the proof.[2]

Theorem 5.28. *If k is in a strongly self-dependent component, then there exists some r so that*

$$g_k(n) \geq \text{const}\, n^{-r} {\lambda_k}^n.$$

We need to show a form of submultiplicativity first.

Proposition 5.29. *Let k be in a strongly self-dependent component. For any m, n we have*

$$g_k(m+n) \leq \text{const}\, K^{\log m} g_k(m) g_k(n),$$

where K is a constant.

One may recognize that the proof below is similar to the proof of Theorem 5.3.

Proof. Let T be the tree associated with $g_k(m+n)$. By Lemma 2.7, there is a subtree T_0 of m_0 leaves so that $m/2 \leq m_0 \leq m$. That means $g_k(m+n) \leq \text{const}\, g_k(m_0) g_k(n + m - m_0)$ by Corollary 5.15. We continue the process with a subtree of m_1 leaves so that $\frac{m-m_0}{2} \leq m_1 \leq m - m_0$ from the tree associated with $g_k(n + m - m_0)$, for which $g_k(n + m - m_0) \leq \text{const}\, g_k(m_1) g_k(n + m - m_0 - m_1)$. Repeating this process some $t = O(\log m)$ times, we obtain

$$g_k(m+n) \leq \text{const}\, {K_0}^{\log m} g_k(m_0) \dots g_k(m_t) g_k(n)$$

where K_0 is a constant and $m_0 + \cdots + m_t = m$.

Since k is in a strongly self-dependent component, that is $g_k(a+b) \geq \text{const}\, g_k(a) g_k(b)$ for any a, b by Theorem 5.26, we have

$$g_k(m+n) \leq \text{const}\, K^{\log m} g_k(m_0 + \cdots + m_t) g_k(n) = \text{const}\, K^{\log m} g_k(m) g_k(n)$$

where K is a constant. $\square$

We are now ready to prove Theorem 5.28.

Proof of Theorem 5.28. Consider any pair of m, n with $m \leq n$. Proposition 5.29 gives

$$g_k(m+n) \leq \text{const}\, K^{\log m} g_k(m) g_k(n).$$

We have

$$K^{\log(m+n)} g_k(m+n) \leq \alpha K^{\log m} g_k(m) K^{\log n} g_k(n)$$

for some constant α, since $m+n$ and n are in a constant factor of each other.

Writing $K^{\log n} = n^{\log K} = n^r$ for $r = \log K$, and multiplying both sides of the inequality by α, we have

$$\alpha (m+n)^r g_k(m+n) \leq \alpha m^r g_k(m) \alpha n^r g_k(n).$$

Applying Fekete's lemma for the submultiplicative sequence $\alpha n^r g_k(n)$, we have

$$\lambda_k = \inf_n \sqrt[n]{\alpha n^r g_k(n)},$$

[2]A similar technique is also applied by the author to prove a weaker lower bound on the number of polyominoes $P(n) \geq \text{const}\, n^{-\text{const} \log n} \lambda^n$, where λ is the growth rate of $P(n)$, which is also known as Klarner's constant. Bui, Vuong. "An asymptotic lower bound on the number of polyominoes." Annals of Combinatorics (2023).

which means

$$g_k(n) \geq \text{const}\, n^{-r} {\lambda_k}^n. \qquad \square$$

Weakly self-dependent components. We now treat weakly self-dependent components.

PROPOSITION 5.30. *Consider some k in a weakly self-dependent component. Let T be a tree of n leaves with the associated vector w. If both branches of T are large enough, then $g_k(n)/w_k$ is unbounded, in the sense that for any $R > 0$ there exists $L > 0$ so that if both branches of T have more than L leaves then $g_k(n)/w_k > R$.*

Before proving the proposition, we give the following lemma.

LEMMA 5.31. *Consider a vertex i so that every j in a lower component than the component of i has $\lambda_j < \lambda_i$, we have*

$$\lambda_i = \sup_{\textit{linear pattern } P} \sqrt[|P|]{M(P)_{i,i}}.$$

Furthermore, let $q_n = \max_{P:|P|=n} M(P)_{i,i}$, this supermultiplicative sequence satisfies

$$\lambda_i = \lim\{\sqrt[n]{q_n} : q_n > 0\}.$$

PROOF. Consider the C-subsystem for the component C of i, it follows from Theorem 5.3 that

$$\lambda_i = \sup_{\text{linear pattern } P} \max_j \sqrt[|P|]{M(P)_{j,j}},$$

where j is taken over all j in C or in a component lower than C.

Since $\lambda_i > \lambda_j$ for every j in a component lower than C, the maximum can be taken over all j in C only. In order to prove the first part of the lemma, it suffices to show that for two different vertices i, j in a component, we have

$$\sup_{\text{linear pattern } P} \sqrt[|P|]{M(P)_{i,i}} = \sup_{\text{linear pattern } Q} \sqrt[|Q|]{M(Q)_{j,j}}.$$

Indeed, for any pattern Q, consider the sequence of patterns $P_n = P_{i\to j} \oplus Q^n \oplus P_{j\to i}$, where $P_{i\to j}$ is a pattern of bounded number of leaves so that $M(P_{i\to j})_{i,j} > 0$ by Observation 5.2. We have $\sqrt[|P_n|]{M(P_n)_{i,i}}$ converging to $\sqrt[|Q|]{M(Q)_{j,j}}$. It means the left hand side is at least the right hand side. The other direction is obtained by exchanging the roles, hence we have an equality.

The second part of the lemma is merely a corollary of the first one, by applying the extension of Fekete's lemma for nonnegative sequences (Lemma 2.5) to the supermultiplicative sequence q_n. $\square$

Now we prove Proposition 5.30.

PROOF OF PROPOSITION 5.30. Assume that both branches of T are large enough but $g_k(n)/w_k$ is still bounded, we show contradictions.

Let u, v be the associated vectors with the left and the right branches, respectively. We can assume the number m of leaves in the right branch is smaller. Choose some ϵ small enough, we suppose m is large enough so that for every $m' \geq m$, we have $(\lambda_i - \epsilon)^{m'} < g_i(m') < (\lambda_i + \epsilon)^{m'}$ for every i.

Since $w_k = \sum_{i,j} c_{i,j}^{(k)} u_i v_j$, we have

$$w_k \leq \text{const}\, u_i v_j$$

for some i, j.

By Corollary 5.23, let i^*, j^* be a pair so that $c^{(k)}_{i^*,j^*} > 0$ with one of i^*, j^* in C, say $i^* \in C$, where C is the component of k. Consider the tree T^* where the left branch is the tree associated with $g_{i^*}(n-1)$ and the right branch is just a leaf. The associated vector w^* to T has $w^*_k \geq \text{const}\, g_{i^*}(n-1) \geq \text{const}(\lambda_k - \epsilon)^n$ (since $\lambda_{i^*} = \lambda_k$).

Back to the tree T, we have

$$u_i v_j \leq g_i(n-m) g_j(m) \leq \text{const}(\lambda_i + \epsilon)^{n-m} (\lambda_j + \epsilon)^m.$$

It follows that $i \in C$ (hence $j \notin C$) when ϵ is small enough and n is large enough, since otherwise

$$\begin{aligned} g_k(n)/w_k &\geq \text{const}\, w^*_k/(u_i v_j) \geq \text{const}\, \frac{(\lambda_k - \epsilon)^n}{(\lambda_i + \epsilon)^{n-m} (\lambda_j + \epsilon)^m} \\ &= \text{const} \left(\frac{\lambda_k - \epsilon}{\lambda_i + \epsilon} \frac{\lambda_k - \epsilon}{\lambda_j + \epsilon} \right)^m \left(\frac{\lambda_k - \epsilon}{\lambda_i + \epsilon} \right)^{n-2m} \end{aligned}$$

is unbounded. (Note that m is unbounded, $n - 2m \geq 0$ and $\frac{\lambda_k - \epsilon}{\lambda_i + \epsilon} \frac{\lambda_k - \epsilon}{\lambda_j + \epsilon} > 1$, $\frac{\lambda_k - \epsilon}{\lambda_i + \epsilon} > 1$ when ϵ is small enough.)

We can now prove that m being large enough raises a contradiction, which finishes the proof.

By Lemma 5.31, there exists a pattern of m_0 leaves so that the associated matrix M_0 satisfies $(M_0)_{i,i} \geq (\lambda_i - \epsilon)^{m_0}$. For every $m = m_0 p + r$ ($1 \leq r \leq m_0$), we have the pattern $(T_0)^p$ of $m^* = m_0 p$ leaves with the associated matrix $M^* = (M_0)^p$ satisfying $(M^*)_{i,i} \geq [(M_0)_{i,i}]^p \geq (\lambda_i - \epsilon)^{m^*}$.

We proceed with transforming the original tree. We replace the right branch by any tree of $r = m - m^*$ leaves with the associated vector v'. As r is bounded, v' is bounded. We set the left branch to be $(T_0)^p \oplus \mathcal{L}$ where $\mathcal{L}$ is the original left branch. The new tree has the k-th entry of the associated vector at least

$$c^{(k)}_{i,j} (M^*)_{i,i} u_i v'_j \geq \text{const}(\lambda_i - \epsilon)^{m^*} u_i \geq \text{const}(\lambda_i - \epsilon)^m u_i,$$

which is greater than w_k an unbounded number of times when ϵ is small enough and m is large enough. That is because

$$w_k \leq \text{const}\, u_i v_j \leq \text{const}\, u_i g_j(m) \leq \text{const}\, u_i (\lambda_j + \epsilon)^m. \qquad \square$$

Corollary 5.32. *Suppose k is in a weakly self-dependent component. Let T be the tree associated with $g_k(n)$. Then every subtree of T with at least $n/2$ leaves has a branch with a bounded number of leaves.*

Proof. Let T_0 be a subtree of $m \geq n/2$ leaves. Consider the decomposition $T = P' \oplus T_0$. Let M' be the matrix associated with P'. Let v, u be the associated vectors with T, T_0, respectively. We have $v = M'u$. It follows that $v_k = (M'u)_k = \sum_i M'_{k,i} u_i$. It means for some j,

$$v_k \leq \text{const}\, M'_{k,j} u_j.$$

Applying[3] Theorem 5.3 to the C'-subsystem with the support of Corollary 5.18 for the component C' of j, we have $u_j \leq \text{const}\, g_j(m) \leq \text{const}\, m^{r_1} {\lambda_j}^m$ for some r_1. Also, we have $M'_{k,j} \leq \text{const}\, g_k(n-m+1) \leq \text{const}(n-m+1)^{r_2} {\lambda_k}^{n-m+1}$ for some r_2. In total, $v_k \leq \text{const}\, m^{r_1} (n-m+1)^{r_2} {\lambda_j}^m {\lambda_k}^{n-m+1}$.

Suppose j is not in the component of k, that is $\lambda_j < \lambda_k$. For some ϵ small enough, we have $v_k = g_k(n) \geq (\lambda_k - \epsilon)^n$ for any n large enough. However, it follows from $m \geq n/2$ that v_k is less than $(\lambda_k - \epsilon)^n$, contradiction.

[3]Here is the only place in this section that the polynomial bound of Theorem 5.3 is applied. We avoid using the results from Section 5.1 with the intention to keep Sections 5.1 and 5.2 to be as independent as possible. Otherwise, certain arguments could be easier.

Therefore, both j, k are in the same component. By Proposition 5.30, one of the branches of T_0 has a bounded number of leaves. The conclusion follows. $\square$

We now give a bound for $g_k(n)$ for k in a weakly self-dependent component.

PROPOSITION 5.33. *Given a vertex k in a weakly self-dependent component, then the sequence $g_k(n)$ is weakly submultiplicative in the sense that $g_k(m+n) \le \operatorname{const} g_k(m)g_k(n)$ for any m, n. As a consequence, $g_k(n) \ge \operatorname{const} {\lambda_k}^n$.*

PROOF. Let T be the tree associated to $g_k(n)$. By Corollary 5.32, there exist a leaf ℓ and a decomposition $P = P_1 \oplus \cdots \oplus P_{t-1} \oplus P_t$ for $P = (T, \ell)$ so that each $|P_i|$ for $i \ne t$ is bounded and the number of leaves in P_t is at most $n/2$. The leaf ℓ is chosen by going from the root in the bigger branch at each step, and when we reach the first subtree of at most $n/2$ leaves, we assign an arbitrary leaf in the subtree to ℓ.

If $P = P' \oplus P''$, it follows from Corollary 5.13, Observation 5.14 and Corollary 5.15 that

$$g_k(|P|) \le \operatorname{const} g_k(|P'|)g_k(|P''|).$$

By the decomposition $P = P_1 \oplus \cdots \oplus P_{t-1} \oplus P_t$, for each $m \ge n/2$, there are P', P'' so that $P = P' \oplus P''$, and $|P'| - (n-m)$ and $|P''| - m$ are bounded. It follows from Corollary 5.13 and Observation 5.14 that

$$g_k(n) \le \operatorname{const} g_k(|P|) \le \operatorname{const} g_k(|P'|)g_k(|P''|) \le \operatorname{const} g_k(n-m)g_k(m).$$

Let the final constant be K. We have $Kg_k(n) \le Kg_k(n-m)Kg_k(m)$, i.e., the sequence $\{Kg_k(n)\}_n$ is submultiplicative. By Fekete's lemma, we have $\lambda_k = \inf_n \sqrt[n]{Kg_k(n)}$. The conclusion follows. $\square$

Proof of the lower bound. Now Theorem 5.11 is clear.

PROOF OF THEOREM 5.11. Consider a minimal component C so that $\lambda_k = \lambda$ for $k \in C$. The minimality means that every component C' lower than C has $\lambda_{k'} < \lambda$ for $k' \in C'$. It follows that C does not belong to the class of non-self-dependent components. By Proposition 5.25, the component C is either strongly self-dependent or weakly self-dependent. By Theorem 5.28 and Proposition 5.33, we have $g_k(n) \ge \operatorname{const} n^{-r} {\lambda_k}^n = \operatorname{const} n^{-r}\lambda^n$ for some r in both cases. The conclusion that $g(n) \ge g_k(n) \ge \operatorname{const} n^{-r}\lambda^n$ in Theorem 5.11 follows. $\square$

CHAPTER 6

Decidability and reducibility

The growth rate λ can be approximated to an arbitrary precision by Theorem 5.10. However, we sometimes need an exact solution. Rosenfeld [**5**] shows that checking $\lambda \leq 1$ is undecidable for the nonnegative setting by reducing the problem of checking $\rho \leq 1$ for the joint spectral radius ρ. Therefore, the notion of growth of bilinear maps can be seen as a generalization of the joint spectral radius.

It should be noted that the problem of checking $\lambda \leq 1$ is actually easier than the problem of checking $\lambda = 1$ in the sense that the former problem can be reduced to the latter one by adding an extra dimension that is always 1.

In this chapter, we provide another proof of the undecidability by Theorem 6.1 with a simpler reduction using the observation that matrix multiplication is also a bilinear map. The reduction is natural, and the products of the matrices can be found in an embedded form in the resulting vectors. Note that it is still left open whether the problem of checking $\lambda \leq 1$ in the positive setting is undecidable. We prove its undecidability under the assumption that it is undecidable to check $\rho \leq 1$ for the joint spectral radius ρ of a pair *positive* matrices in Section 6.6.

Suppose the coefficients of $*$ and the entries of s have no condition on the signs (they can even be complex, i.e. the general setting). Rosenfeld [**5**] asks whether the problem of checking if the system can produce a zero vector is decidable? A negative answer is given in Theorem 6.3. It uses a similar construction to the reduction from checking $\rho \leq 1$ to checking $\lambda \leq 1$, but reduces the problem of checking the mortality of a pair of matrices instead. Another application of the construction is the reduction of the problem with multiple operators and multiple starting vectors to the original problem, as in Section 6.5. This was first remarked by Rosenfeld in [**5**].

Checking if λ is actually a limit is also interesting problem, whose decidability was asked by Rosenfeld in a correspondence. Theorem 6.5 shows that it is undecidable by reducing the problem of checking $\lambda \leq 1$. During the course, there is a transformation of $(*, s)$ to a new system with the corresponding function $g'(n)$ so that for every $m \geq 1$ we have $g'(2m) = g(m)$ and $g'(2m+1) = 0$. As a related fact, we also give a transformation so that the new system has the same growth rate as the original system but with a valid limit λ, as in Section 6.4.

6.1. Checking $\lambda \leq 1$ is undecidable

The reduction in the theorem below is quite important in the sense that its variants appear throughout the chapter.

THEOREM 6.1. *The problem of checking if $\lambda \leq 1$ for the nonnegative setting is undecidable.*

PROOF. Consider the problem of checking if $\rho(\{A, B\}) \leq 1$ for the joint spectral radius $\rho(\{A, B\})$ of a pair of nonnegative matrices A, B in $\mathbb{R}^{d \times d}$, which is known to be undecidable [**6**]. We reduce this problem to the problem of checking if $\lambda \leq 1$ for the bilinear system $(*, s)$ constructed as follows.

We use some embedding of a $d \times d$ matrix C to a vector v in the space $\mathbb{R}^{d^2}$, and allow ourselves to write (C, i, j) to present a vector in $\mathbb{R}^{d^2+2}$, where C is embedded in the first d^2 dimensions and i, j are the two last dimensions.

Given a pair of matrices A, B in $\mathbb{R}^d$, we consider the system $(*, s)$ with the (d^2+2)-dimensional vector $s = (\mathbf{O}, 1, 0)$ for the zero matrix $\mathbf{O}$ and $* : \mathbb{R}^{d^2+2} \times \mathbb{R}^{d^2+2} \to \mathbb{R}^{d^2+2}$ presented by

$$\begin{pmatrix} C \\ i \\ j \end{pmatrix} * \begin{pmatrix} C' \\ i' \\ j' \end{pmatrix} = \begin{pmatrix} CC' + ij'A + ji'B \\ 0 \\ ii' \end{pmatrix}, \tag{6.1}$$

where CC' is the usual matrix multiplication. The key point here is that a matrix multiplication in $\mathbb{R}^d$ is also a bilinear map in $\mathbb{R}^{d^2} \times \mathbb{R}^{d^2} \to \mathbb{R}^{d^2}$.

Let us write down some beginning combinations (of up to 3 instances of s):

$$\begin{aligned} s &= (\mathbf{O}, 1, 0) \\ s * s &= (\mathbf{O}, 0, 1) \\ s * (s * s) &= (\mathbf{OO} + 1 \cdot 1 \cdot A + 0 \cdot 0 \cdot B, 0, 0) = (A, 0, 0) \\ (s * s) * s &= (\mathbf{OO} + 0 \cdot 0 \cdot A + 1 \cdot 1 \cdot B, 0, 0) = (B, 0, 0). \end{aligned} \tag{6.2}$$

Let n be the number of instances of s in a combination with the resulting vector v. Obviously, v_{d^2+1} (the index of i) is nonzero for only $n = 1$, and v_{d^2+2} (the index of j) is nonzero for only $n = 2$. It follows that the sum $ij'A + ji'B$ in (6.1) is nonzero only for $n = 3$. In other words, whenever $n \geq 4$, the expression for the first d^2 dimensions in (6.1) has the recursive form CC'. Together with (6.2), we have the matrix form M of the first d^2 dimensions is the product of matrices from $\{\mathbf{O}, A, B\}$. If M is not zero, then M is the product of matrices from $\{A, B\}$ where the number of instances m_A, m_B of A, B respectively correspond to the number of the occurrences of $s * (s * s)$ and $(s * s) * s$, and $3m_A + 3m_B = n$. Note that the last 2 dimensions of these combinations are always zero, due to $n \geq 3$.

On the other hand, for any sequence of matrices $M_1, \dots, M_t \in \{A, B\}$, the combination

$$(S_1 * (S_2 * (S_3 * (\cdots * (S_{t-1} * S_t) \dots),$$

where $S_k = (s * (s * s))$ if $M_k = A$ and $S_k = ((s * s) * s)$ if $M_k = B$, for $k = 1, \dots, t$, gives a vector whose first d^2 dimensions embed the matrix $M_1 \dots M_t$, and the last two dimensions are zero.

It follows from the two above directions that

$$g(3t) = \max_{M_1, \dots, M_t \in \{A, B\}} \|M_1 \dots M_t\|,$$

where $\| \cdot \|$ denotes the maximum norm.

Also, for $n \geq 3$ and n not divisible by 3, we have

$$g(n) = 0.$$

Therefore,

$$\lambda = \sqrt[3]{\rho(\{A, B\})}.$$

We have reduced the problem of the joint spectral radius to the problem of the growth rate. The conclusion on the undecidability follows. $\square$

The variant of checking $\lambda = 1$ is also undecidable due to the undecidability of the corresponding problem of checking $\rho = 1$ for the joint spectral radius. In fact, we can reduce the problem of checking $\lambda \leq 1$ to the problem of checking $\lambda = 1$ by adding an extra dimension that is always 1. However, the question for $\rho \geq 1$ still remains open (see [**14**, Section 2.2.3] for a discussion):

CONJECTURE 6.2 (Blondel and Tsitsiklis 2000 [**6**]). *It is undecidable to check if $\rho \geq 1$ for the joint spectral radius ρ.*

The conjecture has applications in the stability of dynamical systems. If it holds, then the problem of checking $\lambda \geq 1$ is also undecidable. Note that the problems of comparing ρ with 1 for a pair of matrices and a set of several matrices are equivalent, see [**6**]. However, the reduction from the problem of checking $\rho(\Sigma) \leq 1$ for a set of matrices Σ to the one for $\rho(\{A, B\}) \leq 1$ in [**6**] does not have the strong form $\rho(\Sigma) = \rho(\{A, B\})^t$ for some t as in the reduction of Theorem 6.1, where $\lambda = \sqrt[3]{\rho(\{A, B\})}$. This property can, however, be achieved for the growth of bilinear maps by an appropriate variation of the reduction in Theorem 6.1. Indeed, for any set Σ of 2 or more matrices, we can define a bilinear system with growth rate $\lambda = \sqrt[t]{\rho(\Sigma)}$ for some t. We give an example with a set Σ of 5 matrices $M_1, \ldots, M_5$ and leave the verification and the problem for higher[1] numbers of matrices to the readers: Consider the system $(*, s)$ in $\mathbb{R}^{d^2+4}$ with $s = (\mathbf{O}, 1, 0, 0, 0)$ and $*$ presented by

$$\begin{pmatrix} C \\ i \\ j \\ k \\ \ell \end{pmatrix} * \begin{pmatrix} C' \\ i' \\ j' \\ k' \\ \ell' \end{pmatrix} = \begin{pmatrix} CC' + jj'M_1 + ik'M_2 + i\ell'M_3 + ki'M_4 + \ell i'M_5 \\ 0 \\ ii' \\ ij' \\ ji' \end{pmatrix}.$$

The growth rate is $\lambda = \sqrt[4]{\rho(\Sigma)}$. (Hint for verification: $|\Sigma|$ is 5 because this is the number of combinations of 4 instances of s. The combinations yields the 5 matrices of Σ in the first d^2 dimensions.)

A corollary of the above exercise is that: Since $g(n) \leq \operatorname{const} n^r \lambda^n$ for some r by Theorem 5.3, it follows that $\|\Sigma^n\| \leq \operatorname{const} n^{r'} \rho(\Sigma)^n$ for some r'. Although the latter bound is shown in an easier way in Theorem 4.4, the point here is that the joint spectral radius can apply results from the growth of bilinear maps, since the former is an instance of the latter.

6.2. Checking the mortality is undecidable

The problems of other properties of a pair of matrices can be also reduced to the corresponding ones of a bilinear system. The following theorem is one example.

THEOREM 6.3. *When there is no condition on the signs of the coefficients and the starting entries, the problem of checking if the system can produce a zero vector is undecidable.*

PROOF. We reduce to this problem the problem of checking if a pair of matrices A, B is mortal, that is checking if there exists a sequence of matrices $M_1, \ldots, M_k$ drawn from $\{A, B\}$ for some k so that $M_1 \ldots M_k$ is a zero matrix. The problem of mortality for a pair of matrices is known to be undecidable [**31**].

For the space $\mathbb{R}^d$ of A, B, we consider the space $\mathbb{R}^{d^2+2}$ with an embedding of $d \times d$ matrices into the first d^2 dimensions. One may write (C, i, j) where C is a matrix to present a vector in $\mathbb{R}^{d^2+2}$.

---

[1]The construction covers also the cases of 3 and 4 matrices, e.g., for 3 matrices we set $M_3 = M_4 = M_5$ (or set $M_4 = M_5 = \mathbf{O}$).

Consider the system $(*, s)$ with the starting vector $s = (\mathbf{I}, 1, 0)$ where $\mathbf{I}$ is the identity matrix, and $*$ defined by

$$\begin{pmatrix} C \\ i \\ j \end{pmatrix} * \begin{pmatrix} C' \\ i' \\ j' \end{pmatrix} = \begin{pmatrix} CC' + ij'(A - \mathbf{I}) + ji'(B - \mathbf{I}) \\ 0 \\ ii' \end{pmatrix}.$$

Some begining combinations are

$$\begin{aligned} s &= (\mathbf{I}, 1, 0) \\ s * s &= (\mathbf{I}, 0, 1) \\ s * (s * s) &= (\mathbf{I} + (A - \mathbf{I}), 0, 0) = (A, 0, 0) \\ (s * s) * s &= (\mathbf{I} + (B - \mathbf{I}), 0, 0) = (B, 0, 0). \end{aligned} \tag{6.3}$$

Consider a vector v obtained by combining n instances of s. It follows from $v_{d^2+1} = 0$ for $n > 1$ that $v_{d^2+2} = 0$ for $n > 2$. The consequence is that for $n > 3$, the first d^2 dimensions of v, denoted by $\bar{v}$, are

$$CC' + ij'A + ji'B = CC'.$$

Together with (6.3), the matrix form of $\bar{v}$ for any n presents a product of matrices from $\{\mathbf{I}, A, B\}$. It follows that if $v = 0$ for some combination, then $\{A, B\}$ is mortal.

On the other hand, if $\{A, B\}$ is mortal with $M_1 \dots M_k = 0$, the combination

$$(S_1 * (S_2 * (S_3 * (\cdots * (S_{t-1} * S_t) \dots),$$

where $S_t = (s * (s * s))$ if $M_t = A$, and $S_t = ((s * s) * s)$ if $M_t = B$, for $t = 1, \dots, k$, is zero.

The equivalence means that we can reduce the problem of checking the mortality of a pair of matrices to the problem of checking if a bilinear system can produce a zero vector. The conclusion follows. $\square$

6.3. Checking if the limit exists is undecidable

Before showing that it is undecidable to check if the growth rate is a limit, we give the following nice transformation.

Proposition 6.4. *For every bilinear system $(*, s)$ with the function $g(n)$ we can construct $(*', s')$ so that for every $m \geq 1$ we have $g'(2m + 1) = 0$ and $g'(2m) = g(m)$, where $g'(n)$ is the function for $(*', s')$.*

Proof. Let $\mathbb{R}^d$ be the space of $(*, s)$. We write (x, i) for a vector $x \in \mathbb{R}^d$ and a number $i \in \mathbb{R}$ to present a vector in $\mathbb{R}^{d+1}$. Consider $(*', s')$ with the $(d+1)$-dimensional vector $s' = (\mathbf{0}, 1)$ where $\mathbf{0}$ denotes the zero vector and $*' : \mathbb{R}^{d+1} \times \mathbb{R}^{d+1} \to \mathbb{R}^{d+1}$ presented by

$$\begin{pmatrix} x \\ i \end{pmatrix} *' \begin{pmatrix} y \\ j \end{pmatrix} = \begin{pmatrix} ijs + x * y \\ 0 \end{pmatrix}. \tag{6.4}$$

Let v be the vector obtained from a combination of n instances of s' (using $*'$). For $n = 1$, we have $v = s' = (\mathbf{0}, 1)$. For $n = 2$, we have $v = s' * s' = (s, 0)$. When $n \geq 3$, the summand ijs in (6.4) is zero since either i or j is zero, for which we have the recursive form $x * y$ for the first d dimensions. It follows that the first d dimensions $\bar{v}$ of v are a combination of vectors in $\{\mathbf{0}, s\}$ (using $*$). If $\bar{v}$ is nonzero, then $\bar{v}$ is a combination of some m instances of s, with $2m = n$. Since $v_{d+1} = 0$ for any $n \geq 2$, we have $g'(2m) \leq g(m)$. Considering odd $n > 1$, we start with $s * (s * s) = (s * s) * s = (\mathbf{O}, 0)$ for $n = 3$. By induction, one can show that $g'(2m + 1) = 0$ for any $m \geq 1$. On the other hand, for any combination of m instances of s (using $*$) that is associated with $g(m)$, we also have the corresponding combination of $2m$ instances of s' (using $*'$) by replacing each instance of s by $(s' * s')$. The resulting vector of the former combination is the same as $\bar{v}$ for the resulting vector v of the latter combination. It follows that $g'(2m) = g(m)$. $\square$

THEOREM 6.5. *Checking the existence of the limit of $\sqrt[n]{g(n)}$ is undecidable.*

PROOF. We will reduce the problem of checking if $\lambda = \limsup_{n\to\infty} \sqrt[n]{g(n)} \le 1$ for a system $(*, s)$ to the problem of checking the existence of the limit of another system.

By Proposition 6.4, we can construct a system $(*', s')$ so that for every $m \ge 1$ we have $g'(2m+1) = 0$ and $g'(2m) = g(m)$. Let the space of $(*', s')$ be $\mathbb{R}^{d'}$, we construct $*'' : \mathbb{R}^{d'+1} \times \mathbb{R}^{d'+1} \to \mathbb{R}^{d'+1}$ and $s'' \in \mathbb{R}^{d'+1}$ so that the system $(*', s')$ is brought into the first d' dimensions of the new system $(*'', s'')$ and

$$s''_{d'+1} = 1, \qquad (x *'' y)_{d'+1} = x_{d'+1} y_{d'+1}.$$

The last dimension is obviously always 1. It follows that $g''(2m+1) = 1$ and $g''(2m) = \max\{g(m), 1\}$ for $m \ge 1$. It means $\liminf_{n\to\infty} \sqrt[n]{g''(n)} = 1$ since $g''(n) \ge 1$ for every n and $\liminf_{n\to\infty} \sqrt[n]{g''(n)} \le \liminf_{m\to\infty} \sqrt[2m+1]{g''(2m+1)} = 1$. Meanwhile,

$$\lambda'' = \limsup_{n\to\infty} \sqrt[n]{g''(n)} = \max\left\{ \limsup_{n\to\infty} \sqrt[n]{g'(n)}, \limsup_{n\to\infty} \sqrt[n]{g''_{d'+1}(n)} \right\} = \max\{\lambda, 1\}.$$

Therefore, the limit of $\sqrt[n]{g''(n)}$ exists if and only if $\lambda \le 1$. The reduction is finished, and the conclusion on the undecidability follows. $\square$

6.4. Transformation to make the limit valid

Beside the transformation in Proposition 6.4, we also present the following transformation, as an application of Theorem 5.3. While the former transformation makes the limit not valid, the latter ensures the opposite.

PROPOSITION 6.6. *For every bilinear system $(*, s)$ we can construct $(*', s')$ so that $(*', s')$ has the same growth rate as $(*, s)$ and the limit of $\sqrt[n]{g'(n)}$ exists, where $g'(n)$ is the function for $(*', s')$.*

PROOF. We assume $\lambda > 0$, otherwise it is trivial. (Note that $\lambda > 0$ if and only if the dependency graph has a cycle.)

For the space $\mathbb{R}^d$ of $(*, s)$, consider $*' : \mathbb{R}^{d+2} \times \mathbb{R}^{d+2} \to \mathbb{R}^{d+2}$, $s' \in \mathbb{R}^{d+2}$ so that the coefficients of $*$ and the entries of s are brought to the first d dimensions of $(*', s')$. We let $s'_{d+1} = s'_{d+2} = \alpha$ where $0 < \alpha \le \lambda$. We can take any positive lower bound of λ, e.g. by Theorem 5.3. (In fact, the value of s'_{d+1} does not matter.) The operator $*'$ is defined so that

$$(x *' y)_{d+1} = \sum_{i=1}^{d} x_i y_{d+2}$$

and

$$(x *' y)_{d+2} = x_{d+2} y_{d+2}.$$

The $(d+2)$-th entry of any vector obtained from combining n instances of s' is α^n. It follows that for any index i and any $\delta \ge 1$, we have

$$g'_{d+1}(n+\delta) \ge \alpha^\delta g_i(n)$$

by considering the composition tree where the left branch is associated with $g_i(n)$ and the right branch is any tree of δ leaves. This means that for a bounded δ, we have

$$g'(n+\delta) \ge g'_{d+1}(n+\delta) \ge \max_i \alpha^\delta g_i(n) = \alpha^\delta g(n) = \text{const}\, g(n). \tag{6.5}$$

On the other hand,

$$g'_{d+1}(n) \le d \max_{1\le i\le d} \max_{1\le \delta\le n-1} \alpha^{\delta} g_i(n-\delta),$$

which implies

$$\limsup_{n\to\infty} \sqrt[n]{g'_{d+1}(n)} \le \max\{\alpha, \limsup_{n\to\infty} \sqrt[n]{\max_i g_i(n)}\} = \max\{\alpha, \lambda\} = \lambda.$$

It follows that

$$\limsup_{n\to\infty} \sqrt[n]{g'(n)} \le \lambda.$$

For any linear pattern P with the associated matrix M and any index i, we prove that

$$\liminf_{n\to\infty} \sqrt[n]{g'(n)} \ge \sqrt[|P|]{M_{i,i}}. \tag{6.6}$$

Indeed, we pick a fixed n_0 so that $g_i(n_0) > 0$ with the associated tree T_0. For any n large enough, we write $n = q|P| + n_0 + r$ so that $0 \le r < |P|$. Consider the pattern P^q with the associated matrix M^q. Since $(M^q)_{i,i} \ge (M_{i,i})^q$, the i-th entry of the vector associated to $P^q \oplus T_0$ is at least a constant times $(M_{i,i})^q$. As r is bounded and $P^q \oplus T_0$ has $q|P| + n_0$ leaves, it follows from (6.5) that

$$g'(n) \ge \text{const}\, g(n-r) \ge \text{const}\, g_i(q|P| + n_0) \ge \text{const}(M_{i,i})^q.$$

As $n - q|P|$ is bounded, we have proved (6.6). It follows that

$$\liminf_{n\to\infty} \sqrt[n]{g'(n)} \ge \sup_{\text{linear pattern } P} \max_i \sqrt[|P|]{M(P)_{i,i}} = \lambda,$$

where the equality is due to Theorem 5.3.

In total, we have the limit

$$\lim_{n\to\infty} \sqrt[n]{g'(n)} = \liminf_{n\to\infty} \sqrt[n]{g'(n)} = \limsup_{n\to\infty} \sqrt[n]{g'(n)} = \lambda. \qquad \square$$

Assume Conjecture 6.2 holds, that is checking $\rho \ge 1$ and checking $\lambda \ge 1$ are undecidable, we give another approach to the undecidability of the problem of checking if the limit of $\sqrt[n]{g(n)}$ exists, as an application of Proposition 6.6.

Given a system $(*, s)$, let the system $(*', s')$ obtained from Proposition 6.6 be in the space $\mathbb{R}^{d'}$. Consider $*'' : \mathbb{R}^{d'+2} \times \mathbb{R}^{d'+2} \to \mathbb{R}^{d'+2}$ and $s'' \in \mathbb{R}^{d'+2}$ where the first d' dimensions are deduced from $(*', s')$. We let $s''_{d'+1} = 1$, $s''_{d'+2} = 0$, and

$$(x *'' y)_{d'+1} = x_{d'+2} y_{d'+2}, \qquad (x *'' y)_{d'+2} = x_{d'+1} y_{d'+1}.$$

We can see that the last 2 dimensions are independent of the remaining dimensions, and $\max\{g''_{d'+1}(n), g''_{d'+2}(n)\}$ is 0 if n is divisible by 3 and it is 1 otherwise, where g'' is the function for $(*'', s'')$. It follows that

$$\begin{aligned}\limsup_{n\to\infty} \sqrt[n]{g''(n)} &= \max\left\{\limsup_{n\to\infty} \sqrt[n]{\max\{g''_{d'+1}(n), g''_{d'+2}(n)\}}, \limsup_{n\to\infty} \sqrt[n]{g'(n)}\right\} \\ &= \max\{1, \lambda\}.\end{aligned}$$

Meanwhile,

$$\liminf_{n\to\infty} \sqrt[n]{g''(n)} \ge \liminf_{n\to\infty} \sqrt[n]{g'(n)} = \lim_{n\to\infty} \sqrt[n]{g'(n)} = \lambda,$$

and since $g''(3m) = g'(3m)$ for any m, we have

$$\liminf_{n\to\infty} \sqrt[n]{g''(n)} \le \liminf_{m\to\infty} \sqrt[3m]{g''(3m)} = \liminf_{m\to\infty} \sqrt[3m]{g'(3m)} = \lim_{n\to\infty} \sqrt[n]{g'(n)} = \lambda.$$

In total, $\liminf_{n\to\infty} \sqrt[n]{g''(n)} = \lambda$. It follows that we have reduced the problem of checking $\lambda \ge 1$ to the problem of checking if the limit of $\sqrt[n]{g''(n)}$ exists. Therefore, the latter problem is undecidable, under the assumption on the undecidability of $\lambda \ge 1$.

6.5. Multiple operators and multiple starting vectors

Rosenfeld [**5**] made a remark that the problem of the bilinear system does not become harder when we allow multiple operators and multiple starting vectors. We give reductions that are similar to those in Section 6.1.

The construction in Section 6.1 is well suited for reducing the problem for $(*, \{s, s'\})$ to the original problem. By the problem for $(*, \{s, s'\})$ we mean the problem where we can choose either s or s' in the place of each s instead of fixing the vector s. The two vectors s, s' play the roles of A, B in the construction. We rewrite it formally without repeating the verification.

For a bilinear map $* : \mathbb{R}^d \times \mathbb{R}^d \to \mathbb{R}^d$ and two vectors $s, s' \in \mathbb{R}^d$, consider the system $(\bullet, u)$ with the $(d+2)$-dimensional vector $u = (\mathbf{0}, 1, 0)$ where $\mathbf{0}$ is the d-dimensional zero vector and $\bullet : \mathbb{R}^{d+2} \times \mathbb{R}^{d+2} \to \mathbb{R}^{d+2}$ presented by

$$\begin{pmatrix} w \\ i \\ j \end{pmatrix} \bullet \begin{pmatrix} w' \\ i' \\ j' \end{pmatrix} = \begin{pmatrix} w * w' + ij's + ji's' \\ 0 \\ ii' \end{pmatrix}$$

By the same analysis as in Theorem 6.1, the growth rate of $(\bullet, u)$ is the cube root of the growth rate of $(*, \{s, s'\})$.

Using the idea of the previous construction, we can reduce the problem for $(\{*, *'\}, s)$ to the original problem. By the problem for $(\{*, *'\}, s)$ we mean the problem where we can choose either $*$ or $*'$ in the place of each instance of $*$ instead of fixing $*$.

For two bilinear maps $*, *' : \mathbb{R}^d \times \mathbb{R}^d \to \mathbb{R}^d$ and a vector $s \in \mathbb{R}^d$, consider the system $(\bullet, u)$ with the $(3d+2)$-dimensional vector $u = (s, s, \mathbf{0}, 1, 0)$ where $\mathbf{0}$ is the d-dimensional zero vector and $\bullet : \mathbb{R}^{3d+2} \times \mathbb{R}^{3d+2} \to \mathbb{R}^{3d+2}$ presented by

$$\begin{pmatrix} x \\ y \\ w \\ i \\ j \end{pmatrix} \bullet \begin{pmatrix} x' \\ y' \\ w' \\ i' \\ j' \end{pmatrix} = \begin{pmatrix} w * w' \\ w *' w' \\ jx' + yj' \\ 0 \\ ii' \end{pmatrix}.$$

We sketch the approach: For any vector v obtained from combining n instances of u using $\bullet$, if $v_{[2d+1,3d]} \ne 0$ then $n = 5k + 3$ for some k. Also, if $v_{[1,d]}$ or $v_{[d+1,2d]}$ is not a zero vector, then $n = 5k + 1$ for some k. The growth rate of $(\bullet, u)$ is the fifth root of the growth rate of $(\{*, *'\}, s)$. The verification is similar to that in Theorem 6.1 and we leave it to the readers.

A construction for a higher number of starting vectors or a higher number of bilinear operators, or both, can be established similarly by introducing more dimensions. We leave it to the readers as an exercise since the details would be tedious (one may consult the construction for several matrices in Section 6.1).

In conclusion, introducing more vectors and more operators does not make the problem any harder.

6.6. Conditional undecidability of checking $\lambda \le 1$ in the positive setting

As we can reduce the problem of checking $\rho \le 1$ for the joint spectral radius ρ to the problem of checking $\lambda \le 1$ for the growth of bilinear maps in the nonnegative setting, one may wonder if there is a similar reduction for the positive setting, where all the entries of s have to be positive. In this section, we give such a reduction, which implies the undecidability of checking $\lambda \le 1$ in the positive setting under the assumption that the following conjecture holds.[2]

CONJECTURE 6.7. *It is undecidable to check $\rho(\{A, B\}) \le 1$ for the joint spectral radius ρ of a pair of* positive *matrices A, B.*

The reduction is almost the same as the one in Section 6.1 but with some ideas of the reduction in Section 6.2 and a more complicated argument.

We reuse the convention of embedding a matrix into a vector in Section 6.1. For a pair of $d \times d$ *positive* matrices A, B, we consider the system $(*, s)$ with the (d^2+2)-dimensional vector $s = (\mathbf{E}, 1, \epsilon)$ where $\mathbf{E}$ denotes[3] the $d \times d$ matrix with all entries set to ϵ and $\epsilon > 0$ is small enough. The operator $* : \mathbb{R}^{d^2+2} \times \mathbb{R}^{d^2+2} \to \mathbb{R}^{d^2+2}$ is presented by

$$\begin{pmatrix} C \\ i \\ j \end{pmatrix} * \begin{pmatrix} C' \\ i' \\ j' \end{pmatrix} = \begin{pmatrix} CC' + ji'X + ij'Y \\ 0 \\ ii' \end{pmatrix}, \tag{6.7}$$

where X, Y, ϵ satisfy some requirements that are given in (6.8) below.

We denote by $\Gamma(v)$ the matrix form of the first d^2 dimensions of a vector v. Let us analyze some beginning combinations of s:

$$\begin{aligned} \Gamma(s * s) &= \mathbf{E}^2 + \epsilon X + \epsilon Y, \\ \Gamma((s * s) * s) &= (\mathbf{E}^2 + \epsilon X + \epsilon Y)\mathbf{E} + X, \\ \Gamma(s * (s * s)) &= \mathbf{E}(\mathbf{E}^2 + \epsilon X + \epsilon Y) + Y. \end{aligned}$$

We need X, Y be so that

$$\begin{aligned} \Gamma((s * s) * s) &= A, \\ \Gamma(s * (s * s)) &= B, \\ X \ge 0, \qquad & Y \ge 0. \end{aligned} \tag{6.8}$$

The requirements $X \ge 0, Y \ge 0$ are for the coefficients of $*$ to be nonnegative.

PROPOSITION 6.8. *Such X, Y always exist for any ϵ small enough.*

PROOF. The first two requirements of (6.8) are equivalent to

$$\epsilon(X + Y)\mathbf{E} + X = A - \mathbf{E}^3 \tag{6.9}$$

$$\epsilon\mathbf{E}(X + Y) + Y = B - \mathbf{E}^3. \tag{6.10}$$

Let Σ_M denote the sum of all entries of a matrix M. Taking the sum of all entries in two sides of (6.9) and (6.10), we obtain

$$\begin{aligned} d\epsilon^2(\Sigma_X + \Sigma_Y) + \Sigma_X &= \Sigma_A - d^4\epsilon^3 \\ d\epsilon^2(\Sigma_X + \Sigma_Y) + \Sigma_Y &= \Sigma_B - d^4\epsilon^3. \end{aligned}$$

The solutions of Σ_X, Σ_Y are

$$\Sigma_X = \frac{1}{2}\left(\Sigma_A - \Sigma_B + \frac{1}{1 + 2d\epsilon^2}(\Sigma_A + \Sigma_B - 2d^4\epsilon^3)\right) \tag{6.11}$$

$$\Sigma_Y = \frac{1}{2}\left(\Sigma_B - \Sigma_A + \frac{1}{1 + 2d\epsilon^2}(\Sigma_A + \Sigma_B - 2d^4\epsilon^3)\right). \tag{6.12}$$

[2]During a conversation after the defense of the thesis, Rosenfeld suggested that the conjecture seems to follow from a variant of the problem Probabilistic Finite Automaton Emptiness in an upcoming note by Rote, using the technique in [**6**].

[3]$\mathbf{E}$ here is the capital version of ϵ, for the mnemonic purpose.

Taking the sum of the entries in the j-th column of both sides in (6.9) for some j, we obtain

$$\epsilon^2(\Sigma_X + \Sigma_Y) + \Sigma_{X,C_j} = \Sigma_{A,C_j} - d^3\epsilon^3,$$

where $\sum_{M,C_j}$ denotes the sum of the entries in the j-th column of a matrix M. This implies

$$\Sigma_{X,C_j} = \Sigma_{A,C_j} - d^3\epsilon^3 - \epsilon^2(\Sigma_X + \Sigma_Y). \tag{6.13}$$

The value of Σ_{Y,C_j} can be computed by taking the sum of the entries in the j-th column of both sides in (6.10):

$$d\epsilon^2(\Sigma_{X,C_j} + \Sigma_{Y,C_j}) + \Sigma_{Y,C_j} = \Sigma_{B,C_j} - d^3\epsilon^3,$$

which implies

$$\Sigma_{Y,C_j} = \frac{1}{1+d\epsilon^2}(\Sigma_{B,C_j} - d^3\epsilon^3 - d\epsilon^2\Sigma_{X,C_j}). \tag{6.14}$$

Considering the (i,j)-th entry of both sides in (6.10), we have

$$\epsilon^2(\Sigma_{X,C_j} + \Sigma_{Y,C_j}) + Y_{i,j} = B_{i,j} - d^2\epsilon^3,$$

which implies

$$Y_{i,j} = B_{i,j} - d^2\epsilon^3 - \epsilon^2(\Sigma_{X,C_j} + \Sigma_{Y,C_j}). \tag{6.15}$$

If we substitute (6.11), (6.12), (6.13) and (6.14) into (6.15), we get an explicit expression for $Y_{i,j}$ in terms of ϵ, A and B. We can check that $Y_{i,j}$ is well-defined and depends continuously on ϵ. If $\epsilon = 0$ then $Y_{i,j} = B_{i,j} > 0$. It follows that when $\epsilon > 0$ is small enough, we also have $Y_{i,j} > 0$. The entries $X_{i,j}$ are computed likewise, which are also positive, due to the symmetry of rows and columns. The readers can check for themselves that these values are indeed the solution of the system (6.9) and (6.10). □

Remark. *Note that X, Y can be shown to be positive as in the proof. However, the positivity here does not make much sense as there are already some zero coefficients in the representation of $*$ in* (6.7).

Denote $M_1 = \Gamma(s) = \mathbf{E}$ and $M_2 = \Gamma(s * s) = \mathbf{E}^2 + \epsilon X + \epsilon Y$, we have both $M_1 < \mathbf{E}'$ and $M_2 < \mathbf{E}'$ where $\mathbf{E}'$ is the matrix of all entries ϵ' that depends on ϵ. The value ϵ' can be made arbitrarily small by reducing ϵ.

We make the following observation, whose verification is simple and left to the readers.

Proposition 6.9. *The matrix form $\Gamma(v)$ for any vector v obtained by combining n instances of s is the product of some matrices from $\{A, B, M_1, M_2\}$. In particular, if m_A, m_B, m_1, m_2 are respectively the numbers of instances of A, B, M_1, M_2, then $m_1 + 2m_2 + 3(m_A + m_B) = n$. On the other hand, for any product of m matrices from $\{A, B\}$, we have a combination for $n = 3m$ so that $\Gamma(v)$ is the product.*

Since ϵ' can be made arbitrarily small, the number m_1, m_2 should be made minimal. It follows that $\lambda = \sqrt[3]{\rho(\{A, B\})}$ like in Theorem 6.1. Therefore, the problem of checking $\lambda \le 1$ is undecidable in the positive setting under the assumption that Conjecture 6.7 holds.

Remark. *In contrast to the situation in Section 6.1, the limit of $\sqrt[n]{g(n)}$ here exists, because this is always the case for a positive setting by Corollary 5.9.*

CHAPTER 7

Linear and multilinear patterns

The notion of linear pattern is a useful tool to study the growth of bilinear maps, as we have seen in Chapter 5. We revisit this matter and discuss the original motivation of linear patterns by their rates. The growth rate of a system may be recognized as the rate of a linear pattern. However, there exist cases where no linear pattern attains the growth rate. This chapter also extends the notion of linear pattern to bilinear patterns and beyond that, by allowing more leaves to be marked and replaced, in the hope that they may cover the growth rate in more cases, or even better in every case.

7.1. Linear patterns and finiteness property

The growth rate λ was expressed in [**8**] for the first time in a different form as in Theorem 7.2 follows. We need a definition first.

DEFINITION 7.1. *The rate of a linear pattern P, denoted by $\bar{\lambda}_P$, is the $|P|$-th root of the spectral radius $\lambda_P = \rho(M(P))$ of $M(P)$.*

THEOREM 7.2. *We have*

$$\lambda = \sup_{\textit{linear pattern } P} \bar{\lambda}_P.$$

We provide a quick verification using the results we have proved.

PROOF. It follows from Theorem 5.3 that

$$\begin{aligned} \lambda &= \sup_{\text{linear pattern } P} \max_i \sqrt[|P|]{M(P)_{i,i}} \\ &= \sup_{\text{linear pattern } P} \sup_n \max_i \sqrt[n|P|]{[M(P)^n]_{i,i}} \\ &= \sup_{\text{linear pattern } P} \sqrt[|P|]{\sup_n \max_i \sqrt[n]{[M(P)^n]_{i,i}}} \\ &= \sup_{\text{linear pattern } P} \sqrt[|P|]{\rho(M(P))} \\ &= \sup_{\text{linear pattern } P} \bar{\lambda}_P. \end{aligned}$$

We obtain the second equality because $M(P)^n$ is the matrix associated with P^n and $|P^n| = n|P|$. The next to last equality is due to Theorem 4.4 for a single matrix. $\square$

We would note that the formula in Theorem 4.4 for a single matrix can be deduced from Theorem 5.3 about bilinear systems as follows.

THEOREM 7.3. *For every nonnegative matrix A, the spectral radius $\rho(A)$ can be written as*

$$\rho(A) = \sup_n \max_i \sqrt[n]{(A^n)_{i,i}}.$$

PROOF. The direction that $\rho(A) \geq \sup_n \max_i \sqrt[n]{(A^n)_{i,i}}$ is trivial. We prove the other direction.

Suppose A is a $d \times d$ matrix. Consider an embedding of any $d \times d$ matrix B to a vector v in $\mathbb{R}^{d^2}$ by the function Γ so that

$$B = \Gamma(v), \qquad v = \Gamma^{-1}(B).$$

Let the system $(*, s)$ in the space $\mathbb{R}^{d^2}$ be so that $s = \Gamma^{-1}(A)$ and

$$u * v = \Gamma^{-1}(\Gamma(u)\Gamma(v)).$$

One can see that every combination of n instances of s gives $\Gamma^{-1}(A^n)$. Therefore,

$$\lambda = \rho(A).$$

On the other hand, if P is a linear pattern with $|P| = m$, then the relation between the vector at the root v and the vector at the marked leaf u is

$$\Gamma(v) = A^t \Gamma(u) A^{m-t}$$

for some $0 \leq t \leq m$. In particular, for every i, j, one can write

$$\Gamma(v)_{i,j} = \sum_{k,\ell} (A^t)_{i,k} \Gamma(u)_{k,\ell} (A^{m-t})_{\ell,j} = \sum_{k,\ell} \Gamma(u)_{k,\ell} (A^t)_{i,k} (A^{m-t})_{\ell,j}.$$

Let M be the $d^2 \times d^2$ matrix so that $v = Mu$. The diagonal $M_{(i,j),(i,j)}$ is

$$(A^t)_{i,i} (A^{m-t})_{j,j}.$$

It follows from Theorem 5.3 that

$$\begin{aligned}
\rho(A) = \lambda &= \sup_m \max_{\substack{\text{linear pattern } P \\ |P|=m}} \max_{i,j} \sqrt[m]{M(P)_{(i,j),(i,j)}} \\
&\leq \sup_m \max_{0 \leq t \leq m} \max_{i,j} \sqrt[m]{(A^t)_{i,i} (A^{m-t})_{j,j}} \\
&\leq \sup_m \max_{0 \leq t \leq m} \max_{i,j} \max \left\{ \sqrt[t]{(A^t)_{i,i}},\ \sqrt[m-t]{(A^{m-t})_{j,j}} \right\} \\
&\leq \sup_n \max_i \sqrt[n]{(A^n)_{i,i}}.
\end{aligned}$$

□

In fact, one can also deduce Theorem 4.4 about the joint spectral radius using this method with the construction in Section 6.1. It is left as an exercise for the readers.

The rate of a linear pattern is the original motivation for the proof of the limit λ in the positive setting in [**8**]. Although Theorem 7.2 is not technically more important than Theorem 5.3, the meaning of the former is worth mentioning: Consider the sequence of the trees of $P^1, P^2, \dots$, the vectors $v^{(1)}, v^{(2)}, \dots$ associated with these trees are $Ms, M^2 s, \dots$ for $M = M(P)$. As $s > 0$, the growth $\lambda_P = \lim_{n\to\infty} \sqrt[n]{\|v^{(n)}\|}$ of the norms $\|v^{(n)}\|$ is the spectral radius of M. However, a lower bound on the growth rate should be $\rho(M)$ after being normalized, by taking the $|P|$-th root, as the number of leaves in P^i grows by $|P|$ in each step, that is $\lambda \geq \bar{\lambda}_P = \sqrt[|P|]{\lambda_P}$ for any P. In other words, $\lambda \geq \sup_P \bar{\lambda}_P$. The proof in [**8**] manages to show that this is also an upper bound for λ.

Representing the growth rate in terms of the rates of linear patterns gives some new insight. While the supremum is almost never attained in the form of Theorem 5.3 (as rare as in the case of Theorem 4.4), it is quite common that some linear pattern attains the growth rate in the form of Theorem 7.2, that is $\lambda = \bar{\lambda}_P$ for some P. For example, the system in Theorem 2.1 has the growth rate attained by a linear pattern where the tree has two leaves with the marked leaf on the left. A more complicated example is the problem of the maximum number of minimal dominating sets in a tree of n leaves, as given in the introductory chapter. The underlying optimal trees are composed of beautiful snowflakes (and a linear pattern can generate these trees), see [**1**].

We now answer the following question in the negative.

> Is the growth rate always attained by a linear pattern, like Theorem 3.2?

We first relate the coverage of the rates of linear patterns to the finiteness property of a set of matrices. A set of matrices is said to have the *finiteness property* if there exists some m so that the joint spectral radius of the set is the m-th root of the spectral radius of the product of some m matrices from the set.

Given a pair of matrices A, B and the associated bilinear system that is constructed as in Section 6.1, the argument in the proof of Theorem 6.1 gives

$$\lambda = \sqrt[3]{\rho(\{A, B\})}.$$

Suppose the pair A, B has the finiteness property, i.e., there exists a sequence $M_1, \dots, M_m$ where each matrix is in $\{A, B\}$ so that $\sqrt[m]{\rho(M_1 \dots M_m)} = \rho(\{A, B\})$. We can then build a pattern $P = (T, \ell)$ so that $\bar{\lambda}_P = \sqrt[3]{\rho(\{A, B\})}$. Indeed, if T' is the tree of $3m$ leaves that is associated to $M_1 \dots M_m$ (as in Theorem 6.1), we can let T be the tree of $3m + 1$ leaves where one branch is T' and the other branch is the marked leaf ℓ. The readers can check that $\bar{\lambda}_P = \sqrt[3]{\rho(\{A, B\})}$.

On the other hand, suppose the pair A, B does not have the finiteness property, e.g. the class of pairs in [**32**], or an explicit instance in [**18**]. In this case, there is no linear pattern where $\bar{\lambda}_P = \sqrt[3]{\rho(\{A, B\})}$, since otherwise, by considering the sequence of P^t for $t = 1, 2, \dots$, we would have a periodic sequence of products of matrices whose norms follow the rate $\rho(\{A, B\})$ (with respect to the number of matrices).

It means that a pair of matrices has the finiteness property if and only if the corresponding bilinear system has the growth rate attained by a linear pattern. Therefore, we have bilinear systems where no linear pattern attains the growth rate. The matter is that the entries of the example in [**18**] have a quite complicated nature, and the verification is not trivial. Therefore, we give the following example where the entries and coefficients are binary, and the verification is not complicated. It is actually the case for the problem of pruned trees in the introductory chapter. The result first appeared in [**8**].

Theorem 7.4. *If $s = (1, 1)$ and*

$$x * y = (x_1 y_1 + x_2 y_2, x_2 y_2),$$

then $\lambda > \bar{\lambda}_P$ for every linear pattern P.

For this system, the value of $g(n)$ can be found in the vectors associated with the perfect binary trees (for n being a power of 2), see [**2**]. They cannot be generated by any linear pattern.

Proof. Consider a linear pattern $P = (T, \ell)$ with its matrix

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix}.$$

It is verifiable that $a \geq 1, b \geq 1, c = 0$ and $d = 1$ (the readers can check for themselves, e.g. by induction through the manipulations of patterns and matrices throughout the proof). The spectral radius of the matrix can be also seen to be a. Therefore, the rate is the $|P|$-th root of a.

Consider some two patterns $P_1 = (T_1, \ell_1)$ and $P_2 = (T_2, \ell_2)$ with their associated matrices respectively

$$\begin{bmatrix} a_1 & b_1 \\ 0 & 1 \end{bmatrix}, \begin{bmatrix} a_2 & b_2 \\ 0 & 1 \end{bmatrix}.$$

Their product is

$$\begin{bmatrix} a_1 a_2 & a_1 b_2 + b_1 \\ 0 & 1 \end{bmatrix},$$

which is the matrix associated with the pattern $P = P_1 \oplus P_2$.

We have

$$\bar{\lambda}_P \leq \max\{\bar{\lambda}_{P_1}, \bar{\lambda}_{P_2}\}, \tag{7.1}$$

since $|P| = |P_1| + |P_2|$ and the spectral radius of the product is $a_1 a_2$.

Suppose there is a pattern P with $\bar{\lambda}_P = \lambda$, let $P^* = (T^*, \ell^*)$ be a pattern with the minimal number of leaves among all such patterns. By (7.1), we can see that P^* is not decomposable into two patterns in that way. In other words, one child of the root of T^* is just the marked leaf ℓ^*.

Let the other branch than the branch of the marked leaf, denoted by T', have the associated vector $(a, 1)$, then the matrix associated with P^* is

$$\begin{bmatrix} a & 1 \\ 0 & 1 \end{bmatrix}.$$

We have $\bar{\lambda}_{P^*} = \sqrt[m]{a}$, where m is the number of leaves in T'.

Let T'' be a tree where each branch of the root is a copy of T'. The vector associated with T'' is $(a^2 + 1, 1)$. Since $\sqrt[2m]{a^2 + 1} > \sqrt[m]{a}$, if we replace T' in T^* by T'', we obtain another pattern with a higher rate than $\bar{\lambda}_{P^*}$, a contradiction. $\square$

The relation $JSR \leq GRBS$ seems to suggest that certain phenomena of $GRBS$ may be easier to construct than the similar ones of JSR. In fact, the finiteness conjecture, which claims that every set of matrices has the finiteness property, is still open for the case of rational (and equivalently binary) matrices, though already wrong for the general matrices, see [**19**]. Note that if we have a reduction from $GRBS$ to JSR that is as natural as the one in Section 6.1 and keeps the resulting vectors in some form in the resulting matrices, then we can obtain a set of binary matrices without the finiteness property.

When the entries of s and the coefficients of $*$ are integers, the entries of $M(P)$ for any linear pattern P are also integers, that is the spectral radius λ_P of $M(P)$ and the rate $\bar{\lambda}_P$ are algebraic. It follows that the growth rate λ is algebraic whenever there is a linear pattern attaining λ. On the other hand, the growth rate

$$\lambda = \exp\left(\sum_{i \geq 1} \frac{1}{2^i} \log\left(1 + \frac{1}{a_i^2}\right)\right) = 1.502836801\ldots$$

where $a_0 = 1$ and $a_m = 1 + a_{m-1}^2$ for $m \geq 1$ of the system in Theorem 7.4 (see [**2**]) seems to be transcendental. The fact that no linear pattern attains the growth rate in this case suggests the following question.

QUESTION 7.5. *Suppose the entries of s and the coefficients of $*$ are integers, is the following true: The growth rate λ is algebraic if and only if there exists a linear pattern attaining λ?*

Note that one direction is trivial as discussed: If there exists a linear pattern P attaining λ, then the growth rate $\lambda = \bar{\lambda}_P = \sqrt[|P|]{\rho(M(P))}$ is algebraic. Chapter 2 provides several examples of a linear pattern attaining the growth rate. However, some algebraic roots are still unknown to be constructible or not, e.g. the root $\frac{\sqrt{5}-1}{2}$, which is closely related to the golden ratio.

7.2. Bilinear and multilinear patterns

As the optimal composition trees for the system in Theorem 7.4 are more or less symmetric (and are perfect binary trees when n is a power of 2), we extend the notion of linear pattern to cover the case by allowing one more leaf to be marked as follows.

DEFINITION 7.6. *A bilinear pattern $P = (T, \ell_1, \ell_2)$ is a tree P with two marked leaves ℓ_1, ℓ_2 with the convention that the leaf ℓ_1 is on the left to ℓ_2 (they do not need to have the same father). Let $|P|$ denote the number of leaves excluding the marked leaves. We can observe that the vector at the root depends on the vectors at the two leaves by a bilinear map $\bullet$.*

Now we need to define the rate $\bar{\lambda}_P$ of a bilinear pattern P. It has the motivation from the following behaviour. Let T_0 be the tree of a single leaf, we define T_n for $n \geq 1$ be the tree obtained from T by replacing each marked leaf of T by an instance of T_{n-1}. The vector associated with T_0 is $v^{(0)} = s$ and the vector associated with T_n for $n \geq 1$ is $v^{(n)} = v^{(n-1)} \bullet v^{(n-1)}$. The norm of the sequence $v^{(n)}$ grows doubly exponentially and the rate is $\lambda_P = \lim_{n\to\infty} \|v^{(n)}\|^{\frac{1}{2^n}}$, since $\|u * w\| \leq \text{const}\, \|u\|\|w\|$ for any two vectors u, w, that is $\|v^{(n)}\| \leq \text{const}\, \|v^{(n-1)}\|^2$. However, the growth with respect to the number of leaves should be the normalized to the $(|P|+1)$-th root, that is $\bar{\lambda}_P = \sqrt[|P|+1]{\lambda_P}$. Indeed, the number of leaves x_n in T_n is actually $(|P|+1)2^n - |P|$, which satisfies the recurrence $x_n = 2x_{n-1} + |P|$ with $x_0 = 1$. As recurrences of this type will appear again, we state the following observation.

OBSERVATION 7.7. *Let the sequence $\{x_n\}_n$ be so that $x_0 = k$ and $x_n = mx_{n-1} + t$ (with $m \neq 1$) for $n \geq 1$, where k and t are some constants, we have*

$$x_n = \left(k + \frac{t}{m-1}\right) m^n - \frac{t}{m-1}.$$

The rate of a bilinear pattern plays the role of a lower bound for the growth rate. One may wonder if the growth rate can be the supremum of the rates of all bilinear patterns, as in Theorem 7.2 for linear patterns. The following theorem is an answer to such a question, however under a condition on the dependency graph.

THEOREM 7.8. *If the dependency graph is strongly connected, then we have the following representation of the growth rate:*

$$\lambda = \sup_{\text{bilinear pattern } P} \bar{\lambda}_P.$$

PROOF. Consider any bilinear pattern $P = (T, \ell_1, \ell_2)$. Let $\bullet$ be the associated bilinear map with P, that is the vector v at the root can be written as $v = s \bullet s$. It follows that there exist some i, j, k so that $\|v\| \leq \text{const}\, \hat{c}_{i,j}^{(k)}$ for the coefficients $\hat{c}$ of $\bullet$. Let $P_{i\to j}$ denote a linear pattern so that the number of leaves is bounded and the associated matrix M has $M_{i,j} > 0$. Replacing ℓ_1 by $P_{i\to k}$ and ℓ_2 by $P_{j\to k}$, we obtain from P a new bilinear pattern P' with the marked leaves being the marked leaves of $P_{i\to k}$ and $P_{j\to k}$ and the associated bilinear map $\bullet'$ having the coefficient $\hat{c}'^{(k)}_{k,k} \geq \hat{c}^{(k)}_{i,j} M^{(1)}_{i,k} M^{(2)}_{j,k} = \text{const}\, \hat{c}^{(k)}_{i,j}$ where $M^{(1)}, M^{(2)}$ are respectively the matrices associated with $P_{i\to k}$ and $P_{j\to k}$. It follows that $\lambda_{P'} \geq \text{const}\, \|v\|$. Note that $|P'| - |P|$ is bounded.

For any $\epsilon > 0$, consider a tree T of n leaves for a large enough n with the associated vector v having $\|v\| > (\lambda - \epsilon)^n$. Taking any two leaves ℓ_1, ℓ_2 of T to form a bilinear pattern $P = (T, \ell_1, \ell_2)$, we have the corresponding pattern P' with $\lambda_{P'} \geq \text{const} \|v\| \geq \text{const}(\lambda - \epsilon)^n \geq \text{const}(\lambda - \epsilon)^{|P'|+1}$, that is $\bar{\lambda}_{P'} \geq (\lambda - \epsilon) \text{const}^{\frac{1}{|P'|+1}}$. As ϵ can be arbitrarily small and n can be arbitrarily large, we have $\sup_P \bar{\lambda}_P \geq \lambda$. The other direction of the conclusion is already known, hence completing the proof. $\square$

The condition on the connectedness of the dependency graph is not artificial. Indeed, λ is greater than the supremum for the system $(*, s)$ with $s = (2, 5, 1)$ and

$$x * y = (x_1 y_2, x_3 y_3, x_3 y_3).$$

The linear pattern with the tree of two leaves where the marked leaf on the left has the rate 5. The system has the same crucial properties as the replacements, whose growth was studied in Chapter 3 (except that we use multiplication instead of addition to combine the entries: Every entry of $x * y$ is the combination of a single entry of x with a single entry of y). The rate 5 is obviously the growth rate, since the growth rate cannot be higher than the maximal entry in the starting vector. Consider any bilinear pattern $P = (T, \ell_1, \ell_2)$. We show that $\bar{\lambda}_P = 1$. Obviously the third entry of any resulting vector is 1. Also, this value 1 is the second entry of any resulting vector associated with a tree of more than one leaf. We are now basically interested in only the first entry, which is the largest entry when $n \geq 1$. We have $(v^{(n)})_1 = K (v^{(n-1)})_i (v^{(n-1)})_j$ for some constant K and some $i, j \in \{1, 2, 3\}$. The reasoning is the same as the growth of replacements. We start labeling the vertices of T with the root labeled with 1. The two children of a vertex with label k are labeled from left to right with $1, 2$ if $k = 1$, and with $3, 3$ if $k = 2, 3$. The constant K is then the product of s_k for the labels k of the leaves in T excluding ℓ_1, ℓ_2. As ℓ_1 is to the left of ℓ_2, the label j should be either 2 or 3 (the only leaf with label 1 is the leftmost leaf in T). It follows from $j \in \{2, 3\}$ that $(v^{(n)})_1 \leq K (v^{(n-1)})_1 \cdot 1 = K (v^{(n-1)})_1$ when $n \geq 2$ (i.e. $n - 1 \geq 1$). That is $(v^{(n)})_1 \leq \text{const}\, K^{n-1}$. It follows that $\limsup_{n \to \infty} [(v^{(n)})_1]^{\frac{1}{2^n}} \leq 1$. Therefore, $\lim_{n \to \infty} \|v^{(n)}\|^{\frac{1}{2^n}} = 1$ due to the other two dimensions. The rate $\bar{\lambda}_P = 1$ is smaller than the growth rate.

When the dependency graph is strongly connected, sometimes there is still no bilinear pattern attaining the growth rate λ. Consider the following system $(*, s)$ with $s = (1, 4)$ and

$$x * y = (x_1 y_2, x_1 y_2).$$

The growth rate 4 is attained by a linear pattern with the tree of two leaves and the marked leaf on the left. We show that no bilinear pattern P can attain this rate. At first $(v^{(n)})_1 = (v^{(n)})_2$ for $n \geq 1$ as the expressions for the two entries are identical. As in the previous example, when $n \geq 2$, we have $(v^{(n)})_1 = K (v^{(n-1)})_i (v^{(n-1)})_j = K((v^{(n-1)})_1)^2$ for some constant K (each of i, j is either 1 or 2). For $n = 1$, we have $(v^{(1)})_1 = K s_i s_j$ for some i and j. Let $x_n = \log(v^{(n+1)})_1$, we have $x_0 = \log(v^{(1)})_1$ and $x_n = 2x_{n-1} + \log K$. Applying Observation 7.7 to the sequence x_n, we have $(v^{(n)})_1 = (K(K s_i s_j))^{2^{n-1}} \frac{1}{K} = (K\sqrt{s_i s_j})^{2^n} \frac{1}{K}$. One can see that $K\sqrt{s_i s_j}$ is strictly less than $4^{|P|+1}$. Indeed, we assign the labels to the vertices in the tree of the pattern in the same way as for the problem of growth of replacements (like in the previous example). The constant K is the product of s_k for the $|P|$ labels k of the leaves excluding the 2 marked leaves. The labels i, j belong to the two marked leaves. Since $s_i s_j \leq 4$, it follows that $K\sqrt{s_i s_j} \leq 4^{|P|+1}$. In order to have equality, it is necessary that all leaves, including the two marked ones, are labeled with 2. (Otherwise, there is some leaf with label $k = 1$ and $s_k = 1$.) Since each dimension of the result $x * y$ depends on both dimensions of the input, the labels of the leaves are not all 2. In other words, $K\sqrt{s_i s_j} < 4^{|P|+1}$. (Note that $K s_i s_j$ can be $4^{|P|+1}$ in certain situations, however, it is not what we care, as $K s_i s_j$ is the product of $|P| + 2$ entries of

s.) The same situation also applies for the second dimension. It follows that $\bar{\lambda}_P < 4$. Therefore, no bilinear pattern attains the growth rate in this case.

Note that all the above examples are instances of the problem of the growth of replacements, where there is always a linear pattern attaining the growth rate, by Theorem 3.2.

By allowing even more leaves to be marked, we have the following extension.

DEFINITION 7.9. *A multilinear pattern $P = (T, \ell_1, \ell_2, \ldots, \ell_m)$ is a tree T with some $m \geq 2$ marked leaves $\ell_1, \ell_2, \ldots, \ell_m$. We may call such a pattern specifically an m-linear pattern (or we may call it a trilinear pattern if $m = 3$). The vector at the root depends multilinearly (m-linearly) on the vectors at the marked leaves.*

We also define the rate of a multilinear pattern in the same manner as for bilinear patterns. Let T_0 be the tree of a single leaf, we define T_n for $n \geq 1$ be the tree obtained from T by replacing each marked leaf of T by an instance of T_{n-1}. Let the multilinear map $h(u_1, \ldots, u_m) : \mathbb{R}^n \times \cdots \times \mathbb{R}^n \to \mathbb{R}^n$ represent the relation between vector at the root and the vectors at the m marked leaves. The vector associated with T_0 is $v^{(0)} = s$ and the vector associated with T_n for $n \geq 1$ is $v^{(n)} = h(v^{(n-1)}, \ldots, v^{(n-1)})$. The norm of the sequence $v^{(n)}$ has a superexponential growth rate $\lambda_P = \lim_{n\to\infty} \|v^{(n)}\|^{\frac{1}{m^n}}$, since $\|h(u_1, \ldots, u_m)\| \leq \text{const} \|u_1\| \ldots \|u_m\|$ for any vectors $u_1, \ldots, u_m$, that is $\|v^{(n)}\| \leq \text{const} \|v^{(n-1)}\|^m$.

However, the growth with respect to the number of leaves should be the normalized to the $(\frac{|P|+m-1}{m-1})$-th root, that is we have the rate $\bar{\lambda}_P = \sqrt[\frac{|P|+m-1}{m-1}]{\lambda_P}$ of the pattern P. Indeed, the number of leaves x_n in T_n is actually $\frac{|P|+m-1}{m-1} m^n - \frac{|P|}{m-1}$, which satisfies the recurrence $x_n = m x_{n-1} + |P|$ with $x_0 = 1$ of Observation 7.7.

THEOREM 7.10. *If the dependency graph is strongly connected, then for any m we have*

$$\lambda = \sup_{m\text{-linear pattern } P} \bar{\lambda}_P.$$

PROOF. The readers can apply the same method as in the proof of Theorem 7.8. $\square$

For convenience, we may call a linear pattern a 1-linear pattern in some certain cases.

The following conjecture is perhaps one of the landmarks and interesting results of the study of the growth of bilinear maps if it holds. The author would estimate that it is the hardest problem in the thesis.

CONJECTURE 7.11. *For any system $(*, s)$, there exist an integer $m \geq 1$ and an m-linear pattern P so that*

$$\lambda = \bar{\lambda}_P.$$

A direct consequence of the conjecture is that $g(n) \geq \text{const}\, \lambda^n$. There would be a simpler approach to this inequality if we could conclude that $g(n)$ is submultiplicative, but we have not been able to achieve this either.

If there is a counterexample so that no pattern of our proposed types attains the limit, it must be very interesting to see. Nevertheless, we have not yet been able to give an example where a trilinear pattern attains the limit but no linear or bilinear pattern does. We even suspect that linear and bilinear patterns are sufficient to cover all cases.

Bibliography

[1] Günter Rote. The maximum number of minimal dominating sets in a tree. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1201–1214. SIAM, 2019.

[2] Anna de Mier and Marc Noy. On the maximum number of cycles in outerplanar and series–parallel graphs. *Graphs and Combinatorics*, 28(2):265–275, 2012.

[3] Alfred V. Aho and Neil J.A. Sloane. Some doubly exponential sequences. *Fibonacci Quart*, 11(4):429–437, 1973.

[4] Matthieu Rosenfeld. The growth rate over trees of any family of sets defined by a monadic second order formula is semi-computable. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 776–795. SIAM, 2021.

[5] Matthieu Rosenfeld. It is undecidable whether the growth rate of a given bilinear system is 1. *Linear Algebra and its Applications*, 651:131–143, 2022.

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

[7] Alfred Tarski. A decision method for elementary algebra and geometry. In *Quantifier elimination and cylindrical algebraic decomposition*, pages 24–84. Springer, 1998.

[8] Vuong Bui. Growth of bilinear maps. *Linear Algebra and its Applications*, 624:198–213, 2021.

[9] Vuong Bui. Growth of Replacements. *The Electronic Journal of Combinatorics*, 29:P4.15, 2022.

[10] Vuong Bui. On the joint spectral radius of nonnegative matrices. *Linear Algebra and its Applications*, 654:89–101, 2022.

[11] James L. Peterson. Petri nets. *ACM Computing Surveys*, 9(3):223–252, 1977.

[12] Anders Björner and László Lovász. Chip-firing games on directed graphs. *Journal of Algebraic Combinatorics*, 1(4):305–328, 1992.

[13] Gian-Carlo Rota and Gilbert Strang. A note on the joint spectral radius. *Indag. Math*, 22(4):379–381, 1960.

[14] Raphaël Jungers. *The joint spectral radius: theory and applications*, volume 385. Springer Science & Business Media, 2009.

[15] Marc A. Berger and Yang Wang. Bounded semigroups of matrices. *Linear Algebra and its Applications*, 166:21–27, 1992.

[16] Jeffrey C. Lagarias and Yang Wang. The finiteness conjecture for the generalized spectral radius of a set of matrices. *Linear Algebra and its Applications*, 214:17–42, 1995.

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

[18] Kevin G. Hare, Ian D. Morris, Nikita Sidorov, and Jacques Theys. An explicit counterexample to the Lagarias–Wang finiteness conjecture. *Advances in Mathematics*, 226(6):4667–4701, 2011.

[19] Raphaël M. Jungers and Vincent D. Blondel. On the finiteness property for rational matrices. *Linear Algebra and its Applications*, 428(10):2283–2295, 2008.

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

[21] Victor Kozyakin. On accuracy of approximation of the spectral radius by the Gelfand formula. *Linear Algebra and its Applications*, 431(11):2134–2141, 2009.

[22] Fabian Wirth. On the calculation of time-varying stability radii. *International Journal of Robust and Nonlinear Control*, 8(12):1043–1058, 1998.

[23] Vincent D. Blondel, Yurii Nesterov, and Jacques Theys. On the accuracy of the ellipsoid norm approximation of the joint spectral radius. *Linear Algebra and its Applications*, 394:91–107, 2005.

[24] H. K. Wimmer. Spectral radius and radius of convergence. *The American Mathematical Monthly*, 81(6):625–627, 1974.

[25] Quande Chen and Xinlong Zhou. Characterization of joint spectral radius via trace. *Linear Algebra and its Applications*, 315(1-3):175–188, 2000.

[26] Jason P. Bell. A gap result for the norms of semigroups of matrices. *Linear Algebra and its Applications*, 402:101–110, 2005.

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

[28] Jianhong Xu. On the trace characterization of the joint spectral radius. *The Electronic Journal of Linear Algebra*, 20:367–375, 2010.

[29] John N. Tsitsiklis and Vincent D. Blondel. The Lyapunov exponent and joint spectral radius of pairs of matrices are hard—when not impossible—to compute and to approximate. *Mathematics of Control, Signals and Systems*, 10(1):31–40, 1997.

[30] Alfred Brauer. On a problem of partitions. *American Journal of Mathematics*, 64(1):299–312, 1942.

[31] Vincent D. Blondel and John N. Tsitsiklis. When is a pair of matrices mortal? *Information Processing Letters*, 63(5):283–286, 1997.

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