

Algorithms for Simultaneous Padé Approximations*

Johan Rosenkilde, né Nielsen
Technical University of Denmark
Denmark
jsrn@jsrn.dk

Arne Storjohann
University of Waterloo
Canada
astorjoh@uwaterloo.ca

Abstract

We describe how to solve simultaneous Padé approximations over a power series ring $\mathbb{K}[[x]]$ for a field \mathbb{K} using $O^{\sim}(n^{\omega-1}d)$ operations in \mathbb{K} , where d is the sought precision and n is the number of power series to approximate. We develop two algorithms using different approaches. Both algorithms return a reduced sub-bases that generates the complete set of solutions to the input approximations problem that satisfy the given degree constraints. Our results are made possible by recent breakthroughs in fast computations of minimal approximant bases and Hermite Padé approximations.

1 Introduction

The Simultaneous Padé approximation problem concerns approximating several power series $S_1, \dots, S_n \in \mathbb{K}[[x]]$ with rational functions $\frac{\sigma_1}{\lambda}, \dots, \frac{\sigma_n}{\lambda}$, all sharing the same denominator λ . In other words, for some $d \in \mathbb{Z}_{\geq 0}$, we seek $\lambda \in \mathbb{K}[x]$ of low degree such that each of

$$\text{rem}(\lambda S_1, x^d), \text{rem}(\lambda S_2, x^d), \dots, \text{rem}(\lambda S_n, x^d)$$

has low degree. The study of Simultaneous Padé approximations traces back to Hermite's proof of the transcendence of e [18]. Solving Simultaneous Padé approximations has numerous applications, such as in coding theory, e.g. [13,28]; or in distributed, reliable computation [11]. Many algorithms have been developed for this problem, see e.g. [3,26,27,29] as well as the references therein. Usually one cares about the regime where $d \gg n$. Obtaining $O(nd^2)$ is classical through successive cancellation, see [4] or [13] for a Berlekamp–Massey-type variant. Using fast arithmetic, the previous best was $O^{\sim}(n^{\omega}d)$, where ω is the exponent for matrix multiplication, see Section 1.1. That can

*© Johan Rosenkilde, Arne Storjohann. This is the authors' version of the work. It is posted here for your personal use. Not for redistribution. The definitive version was published in ISSAG '16, <http://dx.doi.org/10.1145/2930889.2930933>.

be done by computing a minimal approximant basis with e.g. [15, 16]; this approach traces back to [2, 3]. Another possibility which achieves the same complexity is fast algorithms for solving structured linear systems, e.g. [8]; see [10] for a discussion of this approach.

A common description is to require $\deg \lambda < N_0$ for some degree bound N_0 , and similarly $\deg \text{rem}(\lambda S_i, x^d) < N_i$ for $i = 1, \dots, n$. The degree bounds could arise naturally from the application, or could be set such that a solution must exist. A natural generalisation is also to replace the x^d moduli with arbitrary $g_1, \dots, g_n \in \mathbb{K}[x]$. Formally, for any field \mathbb{K} :

Problem 1 *Given a tuple $(\mathbf{S}, \mathbf{g}, \mathbf{N})$ where*

- $\mathbf{S} = (S_1, \dots, S_n) \in \mathbb{K}[x]^n$ is a sequence of polynomials,
- $\mathbf{g} = (g_1, \dots, g_n) \in \mathbb{K}[x]^n$ is a sequence of moduli polynomials with $\deg S_i < \deg g_i$ for $i = 1, \dots, n$,
- and $\mathbf{N} = (N_0, \dots, N_n) \in \mathbb{Z}_{\geq 0}^{n+1}$ are degree bounds satisfying $1 \leq N_0 \leq \max_i \deg g_i$ and $N_i \leq \deg g_i$ for $i = 1, \dots, n$,

find, if it exists, a non-zero vector $(\lambda, \phi_1, \dots, \phi_n)$ such that

1. $\lambda S_i \equiv \phi_i \pmod{g_i}$ for $i = 1, \dots, n$, and
2. $\deg \lambda < N_0$ and $\deg \phi_i < N_i$ for $i = 1, \dots, n$.

We will call any vector $(\lambda, \phi_1, \dots, \phi_n)$ as above a *solution* to a given Simultaneous Padé approximation problem. Note that if the N_i are set too low, then it might be the case that no solution exists.

Example 2 *Consider over $\mathbb{F}_2[x]$ that $g_1 = g_2 = g_3 = x^5$, and $\mathbf{S} = (S_1, S_2, S_3) = (x^4 + x^2 + 1, x^4 + 1, x^4 + x^3 + 1)$, with degree bounds $\mathbf{N} = (5, 3, 4, 5)$. Then $\lambda_1 = x^4 + 1$ is a solution, since $\deg \lambda_1 < 5$ and*

$$\lambda_1 \mathbf{S} \equiv (x^2 + 1, 1, x^3 + 1) \pmod{x^5}.$$

$\lambda_2 = x^3 + x$ is another solution, since

$$\lambda_2 \mathbf{S} \equiv (x, x^3 + x, x^4 + x^3 + x) \pmod{x^5}.$$

These two solutions are linearly independent over $\mathbb{F}_2[x]$ and span all solutions.

Several previous algorithms for solving Problem 1 are more ambitious and produce an entire *basis* of solutions that satisfy the first output condition $\lambda S_i \equiv \phi_i \pmod{g_i}$ for $i = 1, \dots, n$, including solutions that do not satisfy the degree bounds stipulated by the second output condition. Our algorithms are slightly more restricted in that we only return the sub-basis that generates the set of solutions that satisfy both output requirements of Problem 1. Formally:

Problem 3 Given an instance of Problem 1, find a matrix $A \in \mathbb{K}[x]^{* \times (n+1)}$ such that:

- Each row of A is a solution to the instance.
- All solutions are in the $\mathbb{K}[x]$ -row space of A .
- A is $(-N)$ -row reduced¹.

The last condition ensures that A is minimal, in a sense, according to the degree bounds N , and that we can easily parametrise which linear combinations of the rows of A are solutions. We recall the relevant definitions and lemmas in Section 2.

We will call such a matrix A a *solution basis*. In the complexities we report here, we cannot afford to compute A explicitly. For example, if all $g_i = x^d$, the number of field elements required to explicitly write down all of the entries of A could be $\Omega(n^2d)$. Instead, we remark that A is completely given by the problem instance as well as the first column of A , containing the λ polynomials.² Our algorithms will therefore represent A row-wise using the following compact representation.

Definition 4 For a given instance of Problem 3, a solution specification is a tuple $(\boldsymbol{\lambda}, \boldsymbol{\delta}) \in \mathbb{K}[x]^{k \times 1} \times \mathbb{Z}_{<0}^k$ such that the completion of $\boldsymbol{\lambda}$ is a solution basis, and where $\boldsymbol{\delta}$ are the $(-N)$ -degrees of the rows of A .

The completion of $\boldsymbol{\lambda} = (\lambda_1, \dots, \lambda_k)^\top$ is the matrix

$$\begin{bmatrix} \lambda_1 & \text{rem}(\lambda_1 S_1, g_1) & \dots & \text{rem}(\lambda_1 S_n, g_n) \\ \vdots & & \ddots & \vdots \\ \lambda_k & \text{rem}(\lambda_k S_1, g_1) & \dots & \text{rem}(\lambda_k S_n, g_n) \end{bmatrix}.$$

Note that $\boldsymbol{\delta}$ will consist of only negative numbers, since any solution \boldsymbol{v} by definition has $\deg_{-N} \boldsymbol{v} < 0$.

Example 5 A solution specification for the problem in Example 2 is

$$(\boldsymbol{\lambda}, \boldsymbol{\delta}) = ([x^4 + 1, x^3 + x]^\top, (-1, -1)).$$

The completion of this is

$$A = \begin{bmatrix} x^4 + 1 & x^2 + 1 & 1 & x^3 + 1 \\ x^3 + x & x & x^3 + x & x^4 + x^3 + x \end{bmatrix}$$

One can verify that A is $(-N)$ -row reduced.

¹The notions $(-N)$ -degree, \deg_{-N} and $(-N)$ -row reduced are recalled in Section 2.

²The restriction $N_i \leq \deg g_i$ in Problem 1 ensures that for a given λ , the only possibilities for the ϕ_i in a solution are $\text{rem}(\lambda S_i, g_i)$. In particular, if we allowed $N_i > \deg g_i$ then $(0, \dots, 0, g_i, 0, \dots, 0)$ would be a solution which can not be directly reconstructed from its first element.

We present two algorithms for solving Problem 3, both with complexity $O(n^{\omega-1} M(d) (\log d) (\log d/n)^2)$, where $d = \max_i \deg g_i$ and $M(d)$ is the cost of multiplying two polynomials of degree d , see Section 1.1. They both depend crucially on recent developments that allow computing minimal approximant bases of non-square matrices faster than for the square case [19, 34]. We remark that from the solution basis, one can also compute the expanded form of one or a few of the solutions in the same complexity, for instance if a single, expanded solution to the simultaneous Padé problem is needed.

Our first algorithm in Section 4 assumes $g_i = x^d$ for all i and some $d \in \mathbb{Z}_{\geq 0}$. It utilises a well-known duality between Simultaneous Padé approximations and Hermite Padé approximations, see e.g. [3]. The Hermite Padé problem is immediately solvable by fast minimal approximant basis computation. A remaining step is to efficiently compute a single row of the adjoint of a matrix in Popov form, and this is done by combining partial linearisation [16] and high-order lifting [31].

Our second algorithm in Section 5 supports arbitrary g_i . The algorithm first solves n single-sequence Padé approximations, each of S_1, \dots, S_n . The solution bases for two problem instances can be combined by computing the intersection of their row spaces; this is handled by a minimal approximant basis computation. A solution basis of the full Simultaneous Padé problem is then obtained by structuring intersections along a binary tree.

Before we describe our algorithms, we give some preliminary notation and definitions in Section 2, and in Section 3 we describe some of the computational tools that we employ.

Both our algorithms have been implemented in Sage v. 7.0 [30] (though asymptotically slower alternatives to the computational tools are used). The source code can be downloaded from <http://jsrn.dk/code-for-articles>.

1.1 Cost model

We count basic arithmetic operations in \mathbb{K} on an algebraic RAM. We will state complexity results in terms of an exponent ω for matrix multiplication, and a function $M(\cdot)$ that is a multiplication time for $\mathbb{K}[x]$ [33, Definition 8.26]. Then two $n \times n$ matrices over \mathbb{K} can be multiplied in $O(n^\omega)$ operations in \mathbb{K} , and two polynomials in $\mathbb{K}[x]$ of degree strictly less than d can be multiplied in $M(d)$ operations in \mathbb{K} . The best known algorithms allow $\omega < 2.38$ [12, 14], and we can always take $M(d) \in O(n(\log n)(\log \log n))$ [9].

In this paper we assume that $\omega > 2$, and that $M(d)$ is super-linear while $M(d) \in O(d^{\omega-1})$. The assumption $M(d) \in O(d^{\omega-1})$ simply stipulates that if fast matrix multiplication techniques are used then fast polynomial multiplication should be used also: for example, $n M(nd) \in O(n^\omega M(d))$.

2 Preliminaries

Here we gather together some definitions and results regarding row reduced bases, minimal approximant basis, and their shifted variants. For a matrix A we denote by $A_{i,j}$ the entry in row i and column j . For a matrix A over $\mathbb{K}[x]$ we denote by $\text{Row}(A)$ the $\mathbb{K}[x]$ -linear row space of A .

2.1 Degrees and shifted degrees

The degree of a nonzero vector $\mathbf{v} \in \mathbb{K}[x]^{1 \times m}$ or matrix $A \in \mathbb{K}[x]^{n \times m}$ is denoted by $\deg \mathbf{v}$ or $\deg A$, and is the maximal degree of entries of \mathbf{v} or A . If A has no zero rows the *row degrees* of A , denoted by $\text{rowdeg } A$, is the tuple (d_1, \dots, d_n) with $d_i = \deg \text{row}(A, i)$.

The (row-wise) *leading matrix* of A , denoted by $\text{LM}(A) \in \mathbb{K}^{n \times m}$, has $\text{LM}(A)_{i,j}$ equal to the coefficient of x^{d_i} of $A_{i,j}$.

Next we recall [2, 19, 34] the shifted variants of the notion of degree, row degrees, and leading matrix. For a *shift* $\mathbf{s} = (s_1, \dots, s_n) \in \mathbb{Z}^n$, define the $n \times n$ diagonal matrix $x^{\mathbf{s}}$ by

$$x^{\mathbf{s}} := \begin{bmatrix} x^{s_1} & & \\ & \ddots & \\ & & x^{s_n} \end{bmatrix}.$$

Then the \mathbf{s} -*degree* of v , the \mathbf{s} -*row degrees* of A , and the \mathbf{s} -*leading matrix* of A , are defined by $\deg_{\mathbf{s}} v := \deg vx^{\mathbf{s}}$, $\text{rowdeg}_{\mathbf{s}} A := \text{rowdeg } Ax^{\mathbf{s}}$, and $\text{LM}_{\mathbf{s}}(A) := \text{LM}(Ax^{\mathbf{s}})$. Note that we pass over the ring of Laurent polynomials only for convenience; our algorithms will only compute with polynomials. As pointed out in [19], up to negation the definition of \mathbf{s} -degree is equivalent to that used in [7] and to the notion of *defect* in [4].

For an instance $(\mathbf{S}, \mathbf{g}, \mathbf{N})$ of Problem 1, in the context of defining matrices, we will be using \mathbf{S} and \mathbf{g} as vectors, and by $\Gamma_{\mathbf{g}}$ denote the diagonal matrix with the entries of \mathbf{g} on its diagonal.

2.2 Row reduced

Although row reducedness can be defined for matrices of arbitrary shape and rank, it suffices here to consider the case of matrices of full row rank. A matrix $R \in \mathbb{K}[x]^{n \times m}$ is *row reduced* if $\text{LM}(R)$ has full row rank, and \mathbf{s} -*row reduced* if $\text{LM}_{\mathbf{s}}(R)$ has full row rank. Every $A \in \mathbb{K}[x]^{n \times m}$ of full row rank is left equivalent to a matrix $R \in \mathbb{K}[x]^{n \times m}$ that is \mathbf{s} -row reduced. The rows of R give a basis for $\text{Row}(A)$ that is minimal in the following sense: the list of \mathbf{s} -degrees of the rows of R , when sorted in non-decreasing order, will be lexicographically minimal. An important feature of row reduced matrices is the so-called “predictable degree”-property [21, Theorem 6.3-13]: for any

$\mathbf{v} \in \mathbb{K}[x]^{1 \times n}$, we have

$$\deg_{\mathbf{s}}(\mathbf{v}R) = \max_{i=1, \dots, n} (\deg_{\mathbf{s}} \text{row}(R, i) + \deg v_i) .$$

A canonical \mathbf{s} -reduced basis is provided by the \mathbf{s} -Popov form. Although an \mathbf{s} -Popov form can be defined for a matrix of arbitrary shape and rank, it suffices here to consider the case of a non-singular matrix. The following definition is equivalent to [19, Definition 1.2].

Definition 6 *A non-singular matrix $R \in \mathbb{K}[x]^{n \times n}$ is in \mathbf{s} -Popov form if $\text{LM}_{\mathbf{s}}(R)$ is unit lower triangular and the degrees of off-diagonal entries of R are strictly less than the degree of the diagonal entry in the same column.*

2.3 Adjoints of row reduced matrices

For a non-singular matrix A recall that the adjoint of A , denoted by $\text{adj}(A)$, is equal to $(\det A)A^{-1}$, and that entry $\text{adj}(A)_{i,j}^{\top}$ is equal to $(-1)^{i+j}$ times the determinant of the $(n-1) \times (n-1)$ sub-matrix that is obtained from A by deleting row i and column j .

Lemma 7 *Let $A \in \mathbb{K}[x]^{n \times n}$ be \mathbf{s} -row reduced. Then $\text{adj}(A)^{\top}$ is $(-\mathbf{s})$ -row reduced with*

$$\text{rowdeg}_{(-\mathbf{s})}\text{adj}(A)^{\top} = (\eta - s - \eta_1, \dots, \eta - s - \eta_n) ,$$

where $\boldsymbol{\eta} = \text{rowdeg}_{\mathbf{s}}A$, $\eta = \sum_i \eta_i$ and $s = \sum_i s_i$.

Proof Since A is \mathbf{s} -row reduced then $Ax^{\mathbf{s}}$ is row reduced. Note that $\text{adj}(Ax^{\mathbf{s}})^{\top}(Ax^{\mathbf{s}})^{\top} = (\det Ax^{\mathbf{s}})I_m$ with $\deg \det Ax^{\mathbf{s}} = \eta$. It follows that row i of $\text{adj}(Ax^{\mathbf{s}})^{\top}$ must have degree at least $\eta - \eta_i$ since η_i is the degree of column i of $(Ax^{\mathbf{s}})^{\top}$. However, entries in row i of $\text{adj}(Ax^{\mathbf{s}})^{\top}$ are minors of the matrix obtained from $Ax^{\mathbf{s}}$ by removing row i , hence have degree at most $\eta - \eta_i$. It follows that the (row-wise) leading coefficient matrix of $\text{adj}(Ax^{\mathbf{s}})^{\top}$ is non-singular, hence $\text{adj}(Ax^{\mathbf{s}})^{\top}$ is row reduced. Since $\text{adj}(Ax^{\mathbf{s}})^{\top} = (\det x^{\mathbf{s}})\text{adj}(A)^{\top}x^{-\mathbf{s}}$ we conclude that $\text{adj}(A)^{\top}$ is $(-\mathbf{s})$ -row reduced with $\text{rowdeg}_{(-\mathbf{s})}\text{adj}(A)^{\top} = (\eta - \eta_1 - s, \dots, \eta - \eta_n - s)$. \square

2.4 Minimal approximant bases

We recall the standard notion of minimal approximant basis, sometimes known as order basis or σ -basis [4]. For a matrix $A \in \mathbb{K}[x]^{n \times m}$ and order $d \in \mathbb{Z}_{\geq 0}$, an *order d approximant* is a vector $\mathbf{p} \in \mathbb{K}[x]^{1 \times n}$ such that $\mathbf{p}A \equiv \mathbf{0} \pmod{x^d}$.

An *approximant basis of order d* is then a matrix $F \in \mathbb{K}[x]^{n \times n}$ which is a basis of all order d approximants. Such a basis always exists and has full

rank n . For a shift $\mathbf{s} \in \mathbb{Z}^n$, F is then an \mathbf{s} -minimal approximant basis if it is \mathbf{s} -row reduced.

Let $\text{MinBasis}(d, A, \mathbf{s})$ be a function that returns $(F, \boldsymbol{\delta})$, where F is an \mathbf{s} -minimal approximant basis of A of order d , and $\boldsymbol{\delta} = \text{rowdeg}_{\mathbf{s}} F$. The next lemma recalls a well known method of constructing minimal approximant bases recursively. Although the output of MinBasis may not be unique, the lemma holds for *any* \mathbf{s} -minimal approximant basis that MinBasis might return.

Lemma 8 *Let $A = [A_1 \mid A_2]$ over $\mathbb{K}[x]$. If $(F_1, \boldsymbol{\delta}_1) = \text{MinBasis}(d, A_1, \mathbf{s})$ and $(F_2, \boldsymbol{\delta}_2) = \text{MinBasis}(d, F_1 A_2, \boldsymbol{\delta}_1)$, then $F_2 F_1$ is an \mathbf{s} -minimal approximant basis of A of order d with $\boldsymbol{\delta}_2 = \text{rowdeg}_{\mathbf{s}} F_2 F_1$.*

Sometimes only the *negative part* of an \mathbf{s} -minimal approximant bases is required, the submatrix of the approximant bases consisting of rows with negative \mathbf{s} -degree. Let function $\text{NegMinBasis}(d, A, \mathbf{s})$ have the same output as MinBasis , but with F restricted to the negative part.

Corollary 9 *Lemma 8 still holds if MinBasis is replaced by NegMinBasis , and “an \mathbf{s} -minimal” is replaced with “the negative part of an \mathbf{s} -minimal.”*

Using for example the algorithm M-Basis of [15], it is easy to show that any order d approximant basis G for an A of column dimension m has $\det G = x^D$ for some $D \in \mathbb{Z}_{\geq 0}$ with $D \leq md$.

Many problems of $\mathbb{K}[x]$ matrices or approximations reduce to the computation of (shifted) minimal approximant bases, see e.g. [4, 15], often resulting in the best known asymptotic complexities for these problems.

2.5 Direct solving of Simultaneous Padé approximations

Let $(\mathbf{S}, \mathbf{g}, \mathbf{N})$ be an instance of Problem 3 of size n . We recall some known approaches for computing a solution specification using row reduction and minimal approximant basis computation.

2.5.1 Via reduced basis

Using the predictable degree property it is easy to show that if $R \in \mathbb{K}[x]^{(n+1) \times (n+1)}$ is an $(-\mathbf{N})$ -reduced basis of

$$A = \left[\begin{array}{c|c} 1 & \mathbf{S} \\ \hline & \Gamma_{\mathbf{g}} \end{array} \right] \in \mathbb{K}[x]^{(n+1) \times (n+1)},$$

then the sub-matrix of R comprised of the rows with negative $(-\mathbf{N})$ -degree form a solution basis. A solution specification $(\boldsymbol{\lambda}, \boldsymbol{\delta})$ is then a subvector $\boldsymbol{\lambda}$ of the first column of R , with $\boldsymbol{\delta}$ the corresponding subtuple $\boldsymbol{\delta}$ of $\text{rowdeg}_{(-\mathbf{N})} R$.

Mulders and Storjohann [24] gave an iterative algorithm for performing row reduction by successive cancellation; it is similar to but faster than earlier algorithms [21, 22]. Generically on input $F \in \mathbb{K}[x]^{m \times m}$ it has complexity $O(n^3(\deg F)^2)$. Alekhovich [1] gave what is essentially a Divide & Conquer variant of Mulders and Storjohann’s algorithm, with complexity $O^{\sim}(n^{\omega+1} \deg F)$. Nielsen remarked [26] that these algorithms perform fewer iterations when applied to the matrix A above, due to its low *orthogonality defect*: $\text{OD}(F) = \sum \text{rowdeg} F - \deg \det F$, resulting in $O(n^2(\deg A)^2)$ respectively $O^{\sim}(n^{\omega} \deg A)$. Nielsen also used the special shape of A to give a variant of the Mulders–Storjohann algorithm that computes coefficients in the working matrix in a lazy manner with a resulting complexity $O(n \text{P}(\deg A))$, where $\text{P}(\deg A) = (\deg A)^2$ when the g_i are all powers of x , and $\text{P}(\deg A) = \text{M}(\deg A) \deg A$ otherwise.

Giorgi, et al. [15] gave a reduction for performing row reduction by computing a minimal approximant basis. For the special matrix A , this essentially boils down to the approach described in the following section.

When $n = 1$, the extended Euclidean algorithm on input S_1 and g_1 can solve the approximation problem by essentially computing the reduced basis of the 2×2 matrix A : each iteration corresponds to a reduced basis for a range of possible shifts [17, 20, 32]. The complexity of this is $O(\text{M}(\deg g_1) \log \deg g_1)$.

2.5.2 Via minimal approximant basis

First consider the special case when all $g_i = x^d$ for the same d . An approximant $\mathbf{v} = (\lambda, \phi_1, \dots, \phi_n)$ of order d of

$$A = \begin{bmatrix} -\mathbf{S} \\ I \end{bmatrix} \in \mathbb{K}[x]^{(n+1) \times n}$$

clearly satisfies $\lambda S_i \equiv \phi_i \pmod{x^d}$ for $i = 1, \dots, n$; conversely, any such vector \mathbf{v} satisfying these congruences must be an approximant of A of order d . So the negative part of a $(-\mathbf{N})$ -minimal approximant basis of A of order d is a solution basis.

In the general case we can reduce to a minimal approximant bases computation as shown by Algorithm 1. Correctness of the algorithm follows from the following result.

Theorem 10 *Corresponding to an instance $(\mathbf{S}, \mathbf{g}, \mathbf{N})$ of Problem 3 of size n , define a shift \mathbf{h} and order d :*

- $\mathbf{h} := -(\mathbf{N} \mid N_0 - 1, \dots, N_0 - 1) \in \mathbb{Z}^{2n+1}$
- $d := N_0 + \max_i \deg g_i - 1$

If G is the negative part of an \mathbf{h} -minimal approximant basis of

$$H = \begin{bmatrix} -\mathbf{S} \\ I \\ \Gamma_{\mathbf{g}} \end{bmatrix} \in \mathbb{K}[x]^{(2n+1) \times n}$$

Algorithm 1 DirectSimPade

Input: $(\mathbf{S}, \mathbf{g}, \mathbf{N})$, an instance of Problem 3 of size n .

Output: $(\boldsymbol{\lambda}, \boldsymbol{\delta})$, a solution specification.

1 $\mathbf{h} \leftarrow -(\mathbf{N} \mid N_0 - 1, \dots, N_0 - 1) \in \mathbb{Z}^{2n+1}$

2 $d \leftarrow N_0 + \max_i \deg g_i - 1$

3 $H = \begin{bmatrix} -\mathbf{S} \\ I \\ \Gamma_{\mathbf{g}} \end{bmatrix}$

4 $([\boldsymbol{\lambda} \mid *], \boldsymbol{\delta}) \leftarrow \text{NegMinBasis}(d, H, \mathbf{h})$

5 **return** $(\boldsymbol{\lambda}, \boldsymbol{\delta})$

of order d , then the submatrix of G comprised of the first $n + 1$ columns is a solution basis to the problem instance.

Proof An approximant $\mathbf{v} = (\lambda, \phi_1, \dots, \phi_n, q_1, \dots, q_n)$ of order d of H clearly satisfies

$$\lambda S_i = \phi_i + q_i g_i \pmod{x^d} \quad (1)$$

for $i = 1, \dots, n$; conversely, any such vector \mathbf{v} satisfying these congruences must be an approximant of H of order d .

Now suppose \mathbf{v} is an order d approximant of H with negative \mathbf{h} -degree, so $\deg \lambda \leq N_0 - 1$, $\deg \phi_i \leq N_i - 1$, and $\deg q_i \leq N_0 - 2$. Since Problem 1 specifies that $\deg S_i < \deg g_i$ and $N_i \leq \deg g_i$, both λS_i and $q_i g_i$ will have degree bounded by $N_0 + \deg g_i - 2$. Since Problem 1 specifies that $N_0 \geq 1$, it follows that both the left and right hand sides of (1) have degree bounded by $N_0 + \deg g_i - 2$, which is strictly less than d . We conclude that

$$\lambda S_i = \phi_i + q_i g_i \quad (2)$$

for $i = 1, \dots, n$. It follows that $\mathbf{v}H = 0$ so \mathbf{v} is in the left kernel of H . Moreover, restricting \mathbf{v} to its first $n + 1$ entries gives $\bar{\mathbf{v}} := (\lambda, \phi_1, \dots, \phi_n)$, a solution to the simultaneous Padé problem with $\deg_{-\mathbf{N}} \bar{\mathbf{v}} = \deg_{\mathbf{h}} \mathbf{v}$. Conversely, if $\bar{\mathbf{v}} = (\lambda, \phi_1, \dots, \phi_n)$ is a solution to the simultaneous Padé problem, then the extension $\mathbf{v} = (\lambda, \phi_1, \dots, \phi_n, q_1, \dots, q_n)$ with $q_i = (\lambda S_i - \phi_i)/g_i \in \mathbb{K}[x]$ for $i = 1, \dots, n$ is an approximant of H of order d with $\deg_{\mathbf{h}} \mathbf{v} = \deg_{-\mathbf{N}} \bar{\mathbf{v}}$.

Finally, consider that a left kernel basis for H is given by

$$K = [K_1 \mid K_2] = \left[\begin{array}{c|c} 1 & \mathbf{S} \\ \Gamma_{\mathbf{g}} & -I \end{array} \right].$$

We must have $G = MK$ for some polynomial matrix M of full row rank. But then MK_1 also has full row rank with $\text{rowdeg}_{-\mathbf{N}} MK_1 = \text{rowdeg}_{\mathbf{h}} G$. \square

DirectSimPade can be performed in time $O^\sim(n^\omega \deg H) = O^\sim(n^\omega \max_i \deg g_i)$ using the minimal approximant basis algorithm by Jeannerod, et al. [19], see Section 3.

A closely related alternative to DirectSimPade is the recent algorithm by Neiger [25] for computing solutions to modular equations with general moduli g_i . This would give the complexity $O^\sim(n^{\omega-1} \sum_i \deg g_i) \subset O^\sim(n^\omega \max_i \deg g_i)$.

All of the above solutions ignore the sparse, simple structure of the input matrices, which is why they do not obtain the improved complexity that we do here.

3 Computational tools

The main computational tool we will use is the following very recent result from Jeannerod, Neiger, Schost and Villard [19] on minimal approximant basis computation.

Theorem 11 ([19, Special case of Theorem 1.4]) *There exists an algorithm PopovBasis(d, A, \mathbf{s}) where the input is an order $d \in \mathbb{Z}_+$, a polynomial matrix $A \in \mathbb{K}[x]^{n \times m}$ of degree at most d , and shift $\mathbf{s} \in \mathbb{Z}^n$, and which returns $(F, \boldsymbol{\delta})$, where F is an \mathbf{s} -minimal approximant basis of A of order d , F is in \mathbf{s} -Popov form, and $\boldsymbol{\delta} = \text{rowdeg}_\mathbf{s} F$. PopovBasis has complexity $O(n^{\omega-1} M(\sigma) (\log \sigma) (\log \sigma/n)^2)$ operations in \mathbb{K} , where $\sigma = md$.*

Our next result says that we can quickly compute the first row of $\text{adj}(F)$ if F is a minimal approximant basis in Popov form. In particular, since F is an approximant basis $\det F = x^D$ for some $D \leq \sigma$, where $\sigma = md$ from Theorem 11.

Theorem 12 *Let $F \in \mathbb{K}[x]^{n \times n}$ be in Popov form and with $\det F = x^D$ for some $D \in \mathbb{Z}_{\geq 0}$. Then the first row of $\text{adj}(F)$ can be computed in $O(n^{\omega-1} M(D) (\log D) (\log D/n))$ operations in \mathbb{K} .*

Proof Because F is in \mathbf{s} -Popov form, D is the sum of the column degrees of F . We consider two cases: $D \geq n$ and $D < n$.

First suppose $D \geq n$. Partial linearisation [16, Corollary 2] can produce from F , with no operations in \mathbb{K} , a new matrix $G \in \mathbb{K}[x]^{\bar{n} \times \bar{n}}$ with dimension $\bar{n} < 2n$, $\deg G \leq \lceil D/n \rceil$, $\det G = \det F$, and such that F^{-1} is equal to the principal $n \times n$ sub-matrix of G^{-1} . Let $\mathbf{v} \in \mathbb{K}[x]^{1 \times \bar{n}}$ be the first row of $x^D I_{\bar{n}}$. Then the first row of $\text{adj}(F)$ will be the first n entries of the first row of $\mathbf{v}G^{-1}$. High-order X -adic lifting [31, Algorithm 5] using the modulus $X = (x-1)^{\lceil D/n \rceil}$ will compute $\mathbf{v}G^{-1}$ in $O(n^\omega M(\lceil D/n \rceil) (\log \lceil D/n \rceil))$ operations in \mathbb{K} [31, Corollary 16]. Since $D \geq n$ this cost estimate remains valid if we replace $\lceil D/n \rceil$ with D/n . Finally, from the super-linearity assumption on $M(\cdot)$ we have $M(D/n) \leq (1/n)M(D)$, thus matching our target cost.

Now suppose $D < n$. In this case we can not directly appeal to the partial linearisation technique since the resulting $O(n^\omega \lceil D/n \rceil)$ may be asymptotically larger than our target cost. But $D < n$ means that F has — possibly many — columns of degree 0; since F is in Popov form, such columns have a 1 on the matrix's diagonal and are 0 on the remaining entries. The following describes how to essentially ignore those columns. D is then greater than or equal to the number of remaining columns, thus effectuating the gain from the partial linearisation.

If $n - k$ is the number of such columns in F that means we can find a permutation matrix P such that

$$\hat{F} := PFP^\top = \left[\begin{array}{c|c} F_1 & \\ \hline F_2 & I_{n-k} \end{array} \right],$$

with each column of F_1 having degree strictly greater than zero. Let i be the row index of the single 1 in the first column of P^\top . Since $F^{-1} = P^\top \hat{F}^{-1} P$, we have

$$\text{row}(\text{adj}(F), 1)P^{-1} = x^D \text{row}(\hat{F}^{-1}, i). \quad (3)$$

Considering that

$$\hat{F}^{-1} = \left[\begin{array}{c|c} F_1^{-1} & \\ \hline -F_2 F_1^{-1} & I_{n-k} \end{array} \right],$$

it will suffice to compute the first k entries of the vector on the right hand side of (3). If $i \leq k$ then let $\mathbf{v} \in \mathbb{K}[x]^{1 \times k}$ be row i of $x^D I_k$. Otherwise, if $i > k$ then let \mathbf{v} be row $i - k$ of $-x^D F_2$. Then in both cases, $\mathbf{v} F_1^{-1}$ will be equal to the first k entries of the vector on the right hand side of (3). Like before, high-order lifting combined with partial linearisation will compute this vector in $O(k^\omega \mathbb{M}(\lceil D/k \rceil) (\log \lceil D/k \rceil))$ operations in \mathbb{K} . Since $D \geq k$ the cost estimate remains valid if $\lceil D/k \rceil$ is replaced with D/k . \square

4 Reduction to Hermite Padé

In this section we present an algorithm for solving Problem 3 when $g_1 = \dots = g_n = x^d$ for some $d \in \mathbb{Z}_{\geq 0}$. The algorithm is based on the well-known duality between the Simultaneous Padé problem and the Hermite Padé problem, see for example [3]. This duality, first observed in a special case [23], and then later in the general case [5], was exploited in [6] to develop algorithms for the fraction free computation of Simultaneous Padé approximation. We begin with a technical lemma that is at the heart of this duality.

Lemma 13 *Let $\hat{A}, \hat{B} \in \mathbb{K}[x]^{(n+1) \times (n+1)}$ be as follows.*

$$\hat{A} = \left[\begin{array}{c|c} x^d & -\mathbf{S} \\ \hline & I \end{array} \right] \quad \hat{B} = \left[\begin{array}{c|c} 1 & \\ \hline \mathbf{S}^\top & x^d I \end{array} \right]$$

Then \hat{B} is the adjoint of \hat{A}^\top . Furthermore, \hat{A}^\top is an approximant basis for \hat{B} of order d , and \hat{B}^\top is an approximant basis of \hat{A} of order d .

Proof Direct computation shows that $\hat{A}^\top \hat{B} = x^d I_m = \det \hat{A}^\top I_m$, so \hat{B} is the adjoint of \hat{A}^\top .

Let now G be an approximant basis of \hat{B} . By the above computation the row space of \hat{A}^\top must be a subset of the row space of G . But since $G\hat{B} = (x^d I_m)R$ for some $R \in \mathbb{K}[x]^{(n+1) \times (n+1)}$, then $\det G = x^d \det R$. Thus $x^d \mid \det G$. But $\det \hat{A}^\top = x^d$, so the row space of \hat{A}^\top can not be smaller than the row space of G . That is, \hat{A}^\top is an approximant basis for B of order d . Taking the transpose through the argument shows that \hat{B}^\top is an approximant basis of \hat{B} of order d . \square

Theorem 14 *Let A and B be as follows.*

$$A = \begin{bmatrix} -\mathbf{S} \\ I \end{bmatrix} \in \mathbb{K}[x]^{(n+1) \times (n+1)} \quad B = \begin{bmatrix} 1 \\ \mathbf{S} \end{bmatrix} \in \mathbb{K}[x]^{(n+1) \times 1}$$

If G is an \mathbf{N} -minimal approximant basis of B of order d with shift $\mathbf{N} \in \mathbb{Z}_{>0}^{n+1}$, then $\text{adj}(G^\top)$ is a $(-\mathbf{N})$ -minimal approximant basis of A of order d . Moreover, if $\boldsymbol{\eta} = \text{rowdeg}_{\mathbf{N}} G$, then $\text{rowdeg}_{-\mathbf{N}} \text{adj}(G) = (\eta - N - \eta_1, \dots, \eta - N - \eta_{n+1})$, where $\eta = \sum_i \eta_i$ and $N = \sum_i N_i$.

Proof Introduce \hat{A} and \hat{B} as in Lemma 13. Clearly G is also an \mathbf{N} -minimal approximant basis of \hat{B} of order d . Likewise, \hat{A} and A have the same minimal approximant bases for given order and shift.

Assume, without loss of generality, that we have scaled G such that $\det G$ is monic. Since \hat{A}^\top is also an approximant basis for \hat{B} of order d , then $\det G = \det \hat{A}^\top = x^d$. By definition $G\hat{B} = x^d R$ for some matrix $R \in \mathbb{K}[x]^{(n+1) \times (n+1)}$. That means

$$\begin{aligned} x^{2d}((G\hat{B})^\top)^{-1} &= x^{2d}((x^d R)^\top)^{-1}, & \text{so} \\ (x^d(G^\top)^{-1})(x^d(\hat{B}^\top)^{-1}) &= x^d(R^\top)^{-1}, & \text{that is} \\ \text{adj}(G^\top)\hat{A} &= x^d(R^\top)^{-1}. \end{aligned}$$

Now $\det R = 1$ since $(x^d)^{n+1} \det R = \det(G\hat{B}) = x^{d+nd}$, so $(R^\top)^{-1} = \text{adj}(R^\top) \in \mathbb{K}[x]^{(n+1) \times (n+1)}$. Therefore $\text{adj}(G^\top)$ is an approximant basis of \hat{A} of order d . The theorem now follows from Lemma 7 by noting that G is \mathbf{N} -row reduced. \square

Example 15 *We apply Theorem 14 to the problem of Example 2 with shifts*

Algorithm 2 DualitySimPade

Input: $(\mathcal{S}, (x^d, \dots, x^d), \mathbf{N})$, an instance of Problem 3 of size n .

Output: $(\boldsymbol{\lambda}, \boldsymbol{\delta})$, solution specification.

- 1 $B \leftarrow [1, S_1, \dots, S_n]^T \in \mathbb{K}[x]^{(n+1) \times 1}$
 - 2 $G \leftarrow \text{PopovBasis}(d, B, \mathbf{N})$
 - 3 $\boldsymbol{\eta} \leftarrow \text{rowdeg}_{\mathbf{N}} G$
 - 4 $\hat{\boldsymbol{\lambda}} \leftarrow \text{first column of } \text{adj}(G^\top)$
 - 5 $\hat{\boldsymbol{\delta}} \leftarrow (\eta - N - \eta_1, \dots, \eta - N - \eta_{n+1})$, where $\eta = \sum_i \eta_i$ and $N = \sum_i N_i$
 - 6 $I \leftarrow \{i \mid \hat{\delta}_i < 0\}$, and $k \leftarrow |I|$
 - 7 $(\boldsymbol{\lambda}, \boldsymbol{\delta}) \leftarrow (\hat{\boldsymbol{\lambda}}_{i \in I}, (\hat{\delta}_i)_{i \in I}) \in \mathbb{K}[x]^{k \times 1} \times \mathbb{Z}^k$
 - 8 **return** $(\boldsymbol{\lambda}, \boldsymbol{\delta})$
-

$\mathbf{N} = (5, 3, 4, 5)$. We have

$$A = \begin{bmatrix} x^4 + x^2 + 1 & x^4 + 1 & x^4 + x^3 + 1 \\ & 1 & \\ & & 1 \\ & & & 1 \end{bmatrix}$$

$$B = \begin{bmatrix} 1 \\ x^4 + x^2 + 1 \\ x^4 + 1 \\ x^4 + x^3 + 1 \end{bmatrix}$$

An \mathbf{N} -minimal approximant basis to order $d = 5$ of B is

$$G = \begin{bmatrix} x & 0 & x & 0 \\ 1 & x^2 + 1 & 0 & 0 \\ 0 & 1 & x^2 + 1 & 0 \\ 0 & x & x + 1 & 1 \end{bmatrix}, \text{ and}$$

$$\text{adj}(G)^\top = \begin{bmatrix} x^4 + 1 & x^2 + 1 & 1 & x^3 + 1 \\ x & x^3 + x & x & x^4 + x \\ x^3 + x & x & x^3 + x & x^4 + x^3 + x \\ 0 & 0 & 0 & x^5 \end{bmatrix}.$$

$\text{adj}(G)^\top$ can be confirmed to be an $(-\mathbf{N})$ -minimal approximant basis of A , since $\text{adj}(G)^\top A \equiv 0 \pmod{x^d}$, and since the $(-\mathbf{N})$ -leading coefficient matrix of $\text{adj}(G)^\top$ has full rank.

Algorithm 2 uses Theorem 14 to solve a Simultaneous Padé approximation by computing a minimal approximant basis of B in Popov form.

Theorem 16 *Algorithm 2 is correct. The cost of the algorithm is $O(n^{\omega-1} \mathbf{M}(d)(\log d)(\log d/n)^2)$ operations in \mathbb{K} .*

Proof Correctness follows from Theorem 14. The complexity estimate is achieved if the algorithms supporting Theorem 11 and Theorem 12 are used for the computation in lines 2 and 4, respectively. \square

5 A Divide & Conquer algorithm

Our second algorithm can handle the full generality of Problem 3. It works by first solving n single Padé approximations, one for each of the S_i individually, and then intersecting these solutions to form approximations of multiple S_i simultaneously. The intersection is structured in a Divide & Conquer tree, and performed by computing minimal approximant bases. Let $(\mathbf{S}, \mathbf{g}, \mathbf{N})$ be an instance of Problem 3 of size n .

The idea of the intersection algorithm is the following: consider that we have solution specifications for two different Simultaneous Padé problems, $(\boldsymbol{\lambda}_1, \boldsymbol{\delta}_1)$ and $(\boldsymbol{\lambda}_2, \boldsymbol{\delta}_2)$. We then compute an approximant basis G of the following matrix:

$$R = \left[\begin{array}{c|c} 1 & 1 \\ \hline -\boldsymbol{\lambda}_1 & \\ \hline & -\boldsymbol{\lambda}_2 \end{array} \right] \quad (4)$$

G then encodes the *intersection* of the $\mathbb{K}[x]$ -linear combinations of the $\boldsymbol{\lambda}_1$ with the $\mathbb{K}[x]$ -linear combinations of the $\boldsymbol{\lambda}_2$: any $\lambda \in \mathbb{K}[x]$ residing in both sets of polynomials will appear as the first entry of a vector in the row space of G . We compute G as an \mathbf{r} -minimal approximant basis to high enough order, where \mathbf{r} is selected carefully such that the \mathbf{r} -degree of any $(\lambda \mid \dots) \in \text{Row}(G)$ will equal the $(-\mathbf{N})$ -degree of the completion of λ according to the combined Simultaneous Padé problem, whenever this degree is negative. From those rows of G with negative \mathbf{r} -degree we then get a solution specification for the combined problem.

Example 17 Consider again Example 2. We divide the problem into two sub-problems $\mathbf{S}_1 = (S_1, S_2)$, $\mathbf{N}_1 = (5, 3, 4)$, and $\mathbf{S}_2 = (S_3)$ and $\mathbf{N}_2 = (5, 5)$. Note that $N_{1,0} = N_{2,0} = 5$, since this is the degree bound on the sought λ for the combined problem. The sub-problems have the following solution specifications and their completions:

$$\begin{aligned} (\boldsymbol{\lambda}_1, \boldsymbol{\delta}_1) &= ([x^4 + 1, x^3 + x]^\top, (-1, -1)) \\ A_1 &= \begin{pmatrix} x^4 + 1 & x^2 + 1 & 1 \\ x^3 + x & x & x^3 + x \end{pmatrix} \\ (\boldsymbol{\lambda}_2, \boldsymbol{\delta}_2) &= ([x^2, x^3 + x + 1]^\top, (-3, -2)) \\ A_2 &= \begin{pmatrix} & x^2 & x^2 \\ x^3 + x + 1 & x + 1 & \end{pmatrix} \end{aligned}$$

We construct R as in (4), and compute G , a minimal approximant basis of R of order 7 and with shifts $\mathbf{r} = (-5 \mid -1, -1 \mid -3, -2)$ (the G below is actually

Algorithm 3 RecursiveSimPade

Input: $(\mathbf{S}, \mathbf{g}, \mathbf{N})$, an instance of Problem 3 of size n .

Output: $(\boldsymbol{\lambda}, \boldsymbol{\delta})$, a solution specification.

```

1 if  $n = 1$  then
2   return DirectSimPade( $\mathbf{S}, \mathbf{g}, \mathbf{N}$ )
3 else
4    $\mathbf{S}_1, \mathbf{g}_1 \leftarrow$  the first  $\lfloor n/2 \rfloor$  elements of  $\mathbf{S}, \mathbf{g}$ 
5    $\mathbf{S}_2, \mathbf{g}_2 \leftarrow$  the last  $\lfloor n/2 \rfloor$  elements of  $\mathbf{S}, \mathbf{g}$ 
6    $\mathbf{N}_1 \leftarrow (N_0, N_1, \dots, N_{\lfloor n/2 \rfloor})$ 
7    $\mathbf{N}_2 \leftarrow (N_0, N_{\lfloor n/2 \rfloor + 1}, \dots, N_n)$ 
8    $(\boldsymbol{\lambda}_1, \boldsymbol{\delta}_1) \leftarrow$  RecursiveSimPade( $\mathbf{S}_1, \mathbf{g}_1, \mathbf{N}_1$ )
9    $(\boldsymbol{\lambda}_2, \boldsymbol{\delta}_2) \leftarrow$  RecursiveSimPade( $\mathbf{S}_2, \mathbf{g}_2, \mathbf{N}_2$ )
10   $\mathbf{r} \leftarrow (-N_0 \mid \boldsymbol{\delta}_1 \mid \boldsymbol{\delta}_2)$ 
11   $d \leftarrow N_0 + \max_i \deg g_i - 1$ 
12   $R \leftarrow \left[ \begin{array}{c|c} 1 & 1 \\ \hline -\boldsymbol{\lambda}_1 & \\ \hline & -\boldsymbol{\lambda}_2 \end{array} \right]$ 
13   $([\boldsymbol{\lambda} \mid *], \boldsymbol{\delta}) \leftarrow$  NegMinBasis( $d, R, \mathbf{r}$ )
14  return  $(\boldsymbol{\lambda}, \boldsymbol{\delta})$ 
15 end if

```

in \mathbf{r} -Popov form):

$$G = \begin{pmatrix} & x^8 & 0 & 0 & 0 & 0 \\ & x^3 + x + 1 & x^4 + 1 & 1 & 0 & 1 \\ x^3 + x^2 + x + 1 & & 1 & x + 1 & 1 & 1 \\ x^4 + x^3 + x + 1 & & 1 & 1 & x^2 & 1 \\ & x^4 + 1 & 1 & 0 & x + 1 & x + 1 \end{pmatrix}$$

G has \mathbf{r} -row degree $(3, 3, 0, -1, -1)$. Only rows 4 and 5 have negative \mathbf{r} -degree, and their first entries are the linearly independent solutions $x^4 + x^3 + x + 1$ and $x^4 + 1$. Both solutions complete into vectors with $(-\mathbf{N})$ -degree -1.

To prove the correctness of the above intuition, we will use Algorithm 1 (DirectSimPade). The following lemma says that to solve two simultaneous Padé approximations, one can compute a minimal approximant basis of one big matrix A constructed essentially from two of the matrices employed in DirectSimPade. Afterwards, Lemma 19 uses this to show that a minimal approximant basis of R in (4) provides the crucial information in a minimal approximant basis of A .

Lemma 18 *Let $(\mathbf{S}_1, \mathbf{g}_1, \mathbf{N}_1)$ and $(\mathbf{S}_2, \mathbf{g}_2, \mathbf{N}_2)$ be two instances of Problem 3 of lengths n_1, n_2 respectively, and where $\mathbf{N}_1 = (N_0 \mid \dot{\mathbf{N}}_1)$ and $\mathbf{N}_2 = (N_0 \mid \dot{\mathbf{N}}_2)$. Let $\mathbf{S} = (\mathbf{S}_1 \mid \mathbf{S}_2)$, $\mathbf{g} = (\mathbf{g}_1 \mid \mathbf{g}_2)$ and $\mathbf{N} = (N_0 \mid \dot{\mathbf{N}}_1 \mid \dot{\mathbf{N}}_2)$ be the combined problem having length $n = n_1 + n_2$.*

Let $\mathbf{h}_i = (-\mathbf{N}_i \mid N_0 - 1 \dots N_0 - 1) \in \mathbb{Z}^{2n_i+1}$ for $i = 1, 2$. Let $(F, \boldsymbol{\delta}) = \text{NegMinBasis}(d, A, \mathbf{a})$, where A of dimension $(2n + 3) \times (n + 2)$ is given as:

$$A = [A_1 \mid A_2] = \left[\begin{array}{cc|cc} & & 1 & 1 \\ & & -1 & \\ -\mathbf{S}_1 & & & \\ I & & & \\ \Gamma_{g_1} & & & \\ & -\mathbf{S}_2 & & -1 \\ & I & & \\ & \Gamma_{g_2} & & \end{array} \right],$$

with $\mathbf{a} = (-N_0 \mid \mathbf{h}_1 \mid \mathbf{h}_2)$ and $d = N_0 + \max_i \deg g_i - 1$. Then $(\boldsymbol{\lambda}, \boldsymbol{\delta})$ is a solution specification to $(\mathbf{S}, \mathbf{g}, \mathbf{N})$, where $\boldsymbol{\lambda}$ is the first column of F .

Proof Note that the matrix A is right equivalent to the following matrix B :

$$B := A \left[\begin{array}{ccc|ccc} & & I & & & \\ & & & I & & \\ 1 & & \mathbf{S}_1 & & & \\ & 1 & & \mathbf{S}_2 & & \end{array} \right] = \left[\begin{array}{cc|cc} 1 & 1 & -\mathbf{S}_1 & -\mathbf{S}_2 \\ -1 & & & \\ & & I & \\ & & \Gamma_{g_1} & \\ & -1 & & \\ & & & I \\ & & & \Gamma_{g_2} \end{array} \right].$$

Since F is an \mathbf{a} -minimal approximant of A of order d , then it will also be one for B . Let P be the permutation matrix that produces the following matrix $C := PB$:

$$C = PB = \left[\begin{array}{cc|cc} 1 & 1 & -\mathbf{S}_1 & -\mathbf{S}_2 \\ & & I & \\ & & & I \\ & & \Gamma_{g_1} & \\ \hline -1 & & & \Gamma_{g_2} \\ -1 & & & \end{array} \right] = \left[\begin{array}{cc|c} 1 & 1 & -\mathbf{S} \\ & & I \\ & & \Gamma_{\mathbf{g}} \\ \hline -1 & & \\ & -1 & \end{array} \right].$$

Define $\mathbf{c} := \mathbf{a}P^{-1}$, and note that $\mathbf{c} = (\mathbf{h} \mid -N_0, -N_0)$. Since $F = \text{NegMinBasis}(d, A, \mathbf{a})$, then $(FP^{-1}, \boldsymbol{\delta})$ is a valid output of $\text{NegMinBasis}(d, C, \mathbf{c})$. Furthermore, since the first column of P is $(1, 0, \dots, 0)$, the first column of F will be equal to the first column of FP^{-1} .

We are therefore finished if we can show that if $(F', \boldsymbol{\delta}')$ is any valid output of $\text{NegMinBasis}(d, C, \mathbf{c})$, then the first column of F' together with $\boldsymbol{\delta}'$ form a solution specification to $(\mathbf{S}, \mathbf{g}, \mathbf{N})$.

Consider therefore such an $(F', \boldsymbol{\delta}')$. By the first two columns of C , we must have $F'_{*,1} \equiv F'_{*,2n+2} \equiv F'_{*,2n+3} \pmod{x^d}$, where $F'_{*,i}$ denotes the i 'th column of F' . Since each row of F' have negative \mathbf{c} -degree, and since

$N_0 < d$, then the congruences must lift to equalities. We can therefore write $F = [G \mid F'_{*,1} \mid F'_{*,1}]$ for some $G \in \mathbb{K}[x]^{k \times (2n+1)}$ for some k , and we have $\text{rowdeg}_{\mathbf{h}} G = \text{rowdeg}_{\mathbf{c}} F' = \boldsymbol{\delta}'$.

By the last n columns of C , we have $GH \equiv 0 \pmod{x^d}$, where

$$H = \begin{bmatrix} -\mathbf{S} \\ I \\ \Gamma_{\mathbf{g}} \end{bmatrix}.$$

In fact, $(G, \boldsymbol{\delta}')$ is a valid output for $\text{NegMinBasis}(d, H, \mathbf{h})$: for G has full row rank since F' does; G is \mathbf{h} -row reduced since F' is \mathbf{c} -row reduced; and any negative \mathbf{h} -order d approximant of H must clearly be in the span of G since F' is a negative \mathbf{c} -minimal approximant basis of C .

By the choice of d , then Theorem 10 therefore implies that the first column of G together with $\boldsymbol{\delta}'$ form a solution specification to the problem $(\mathbf{S}, \mathbf{g}, \mathbf{N})$. Since the first column of G is also the first column of F' , this finishes the proof. \square

Lemma 19 *In the context of Lemma 18, let $(\boldsymbol{\lambda}_1, \boldsymbol{\delta}_1)$ and $(\boldsymbol{\lambda}_2, \boldsymbol{\delta}_2)$ be solution specifications to the two sub-problems, and let $\mathbf{r} = (-N_0 \mid \boldsymbol{\delta}_1 \mid \boldsymbol{\delta}_2)$. If $([\boldsymbol{\lambda} \mid *], \boldsymbol{\delta}) = \text{NegMinBasis}(d, R, \mathbf{r})$, where $\boldsymbol{\lambda}$ is a column vector and*

$$R = \left[\begin{array}{c|c} 1 & 1 \\ \hline -\boldsymbol{\lambda}_1 & \\ \hline & -\boldsymbol{\lambda}_2 \end{array} \right],$$

then $(\boldsymbol{\lambda}, \boldsymbol{\delta})$ is a solution specification for the combined problem.

Proof We will prove the lemma by using Lemma 9 to relate valid outputs of $\text{NegMinBasis}(d, R, \mathbf{r})$ with valid outputs of $\text{NegMinBasis}(d, A, \mathbf{a})$ from Lemma 18.

For $i = 1, 2$, since $(\boldsymbol{\lambda}_i, \boldsymbol{\delta}_i)$ is a solution specification to the i 'th problem, then by Theorem 10 there is some $G_i \in \mathbb{K}[x]^{k_i \times 2n_i + 1}$ whose first column is $\boldsymbol{\lambda}_i$ and such that G_i is a valid output of $\text{NegMinBasis}(d, H_i, \mathbf{h}_i)$, where

$$H_i = \begin{bmatrix} -\mathbf{S}_i \\ I \\ \Gamma_{\mathbf{g}_i} \end{bmatrix} \in \mathbb{K}[x]^{(2n_i+1) \times n_i},$$

and \mathbf{h}_i is as in Lemma 18. Note now that if

$$F_1 := \begin{bmatrix} 1 & & \\ & G_1 & \\ & & G_2 \end{bmatrix} \in \mathbb{K}[x]^{(k_1+k_2+1) \times (2n_1+2n_2+3)},$$

then (F_1, \mathbf{r}) is a valid output of $\text{NegMinBasis}(d, A_1, \mathbf{a})$: for $\text{rowdeg}_{\mathbf{a}} F_1$ is clearly \mathbf{r} ; F_1 has full row rank and is \mathbf{r} -row reduced; and the rows of F_1 must

span all \mathbf{a} -order d approximants of A_1 , since the three column “parts” of F_1 correspond to the three row parts of A_1 .

Note now that $F_1 A_2 = R$. Thus by Lemma 9, if $(F_2, \boldsymbol{\delta}) = \text{NegMinBasis}(d, R, \mathbf{r})$, then $(F_2 F_1, \boldsymbol{\delta})$ is a valid output of $\text{NegMinBasis}(d, A, \mathbf{a})$. Note that by the shape of F_1 then the first column $\boldsymbol{\lambda}$ of $F_2 F_1$ is the first column of F_2 . Thus $\boldsymbol{\lambda}, \boldsymbol{\delta}$ are exactly as stated in the lemma, and by Lemma 18 they must be a solution specification to the combined problem. \square

Theorem 20 *Algorithm 3 is correct. The cost of the algorithm is $O(n^{\omega-1} M(d)(\log d)(\log d/n)^2)$, $d = \max_i \deg g_i$.*

Proof Correctness follows from Lemma 19. For complexity, note that the choice of order in Line 11 is bounded by $2 \max_i \deg g_i$, i.e. twice the value of d of this theorem. So if $T(n)$ is the cost Algorithm 3 for given n and where the order will be bounded by $O(d)$, then we have the following recursion:

$$T(n) = \begin{cases} 2T(n/2) + P(n) & \text{if } n > 1 \\ O(M(d) \log d) & \text{if } n = 1 \text{ (see Section 2.5.1)} \end{cases} ,$$

where $P(n)$ is the cost of line 13. Using algorithm `PopovBasis` for the computation of the negative part of the minimal approximant bases we can set $P(n)$ to the target cost. The recursion then implies $T(n) \in O(P(n))$. \square

Acknowledgements. The authors would like to thank George Labahn for valuable discussions, and for making us aware of the Hermite–Simultaneous Padé duality. We would also like to thank Vincent Neiger for making preprints of [19] available to us. The first author would like to thank the Digiteo Foundation for funding the research visit at Waterloo, during which most of the ideas of this paper were developed.

References

- [1] M. Alekhovich. Linear Diophantine Equations Over Polynomials and Soft Decoding of Reed–Solomon Codes. *IEEE Trans. Inf. Th.*, 51(7):2257–2265, 2005.
- [2] M. V. Barel and A. Bultheel. A general module theoretic framework for vector M-*Padé* and matrix rational interpolation. *Numerical Algorithms*, 3(1):451–461, Dec. 1992.
- [3] B. Beckermann and G. Labahn. A uniform approach for Hermite *Padé* and simultaneous *Padé* approximants and their matrix-type generalizations. *Numerical Algorithms*, 3(1):45–54, 1992.
- [4] B. Beckermann and G. Labahn. A Uniform Approach for the Fast Computation of Matrix-Type *Padé* Approximants. *SIAM J. Matr. Anal. Appl.*, 15(3):804–823, July 1994.

- [5] B. Beckermann and G. Labahn. Recursiveness in matrix rational interpolation problems. *J. Comp. App. Math.*, 77(1–2):5–34, Jan. 1997.
- [6] B. Beckermann and G. Labahn. Fraction-Free Computation of Simultaneous Padé Approximants. In *Proc. of ISSAC*, pages 15–22, 2009.
- [7] B. Beckermann, G. Labahn, and G. Villard. Normal forms for general polynomial matrices. *J. Symb. Comp.*, 41(6):708–737, 2006.
- [8] A. Bostan, C.-P. Jeannerod, and E. Schost. Solving structured linear systems with large displacement rank. *Th. Comp. Sc.*, 407(1–3):155–181, Nov. 2008.
- [9] D. G. Cantor and E. Kaltofen. On fast multiplication of polynomials over arbitrary algebras. *Acta Informatica*, 28(7):693–701, 1991.
- [10] M. Chowdhury, C.-P. Jeannerod, V. Neiger, E. Schost, and G. Villard. Faster Algorithms for Multivariate Interpolation With Multiplicities and Simultaneous Polynomial Approximations. *IEEE Trans. Inf. Theory*, 61(5):2370–2387, May 2015.
- [11] Clément Pernet. *High Performance and Reliable Algebraic Computing*. Nov. 2014. Habilitation.
- [12] D. Coppersmith and S. Winograd. Matrix Multiplication via Arithmetic Progressions. *J. Symb. Comp.*, 9(3):251–280, 1990.
- [13] G.-L. Feng and K. K. Tzeng. A Generalization of the Berlekamp-Massey Algorithm for Multisequence Shift-Register Synthesis with Applications to Decoding Cyclic Codes. *IEEE Trans. Inf. Theory*, 37(5):1274–1287, 1991.
- [14] F. L. Gall. Powers of tensors and fast matrix multiplication. In *Proc. of ISSAC*, pages 296–303, 2014.
- [15] P. Giorgi, C. Jeannerod, and G. Villard. On the Complexity of Polynomial Matrix Computations. In *Proc. of ISSAC*, pages 135–142, 2003.
- [16] S. Gupta, S. Sarkar, A. Storjohann, and J. Valeriote. Triangular x -basis decompositions and derandomization of linear algebra algorithms over $K[x]$. *J. Symb. Comp.*, 47(4):422–453, 2012.
- [17] F. Gustavson and D. Yun. Fast algorithms for rational Hermite approximation and solution of Toeplitz systems. *IEEE Trans. Circ. Sys.*, 26(9):750–755, 1979.
- [18] C. Hermite. Sur la Formule d’Interpolation de Lagrange. *J. Reine und Angewandte Math.*, 84(1):70–79, 1878.

- [19] C.-P. Jeannerod, V. Neiger, E. Schost, and G. Villard. Fast computation of minimal interpolation bases in Popov form for arbitrary shifts. Submitted to ISSAC'16.
- [20] J. Justesen. On the complexity of decoding Reed-Solomon codes (Corresp.). *IEEE Trans. Inf. Theory*, 22(2):237–238, Mar. 1976.
- [21] T. Kailath. *Linear Systems*. Prentice-Hall, 1980.
- [22] A. Lenstra. Factoring Multivariate Polynomials over Finite Fields. *J. Comp. Syst. Sc.*, 30(2):235–248, 1985.
- [23] K. Mahler. Perfect systems. *Compos. Math*, 19:95–168, 1968.
- [24] T. Mulders and A. Storjohann. On Lattice Reduction for Polynomial Matrices. *J. Symb. Comp.*, 35(4):377–401, 2003.
- [25] V. Neiger. Fast computation of shifted Popov forms of polynomial matrices via systems of modular polynomial equations. Submitted to ISSAC'16.
- [26] J. S. R. Nielsen. Generalised Multi-sequence Shift-Register Synthesis using Module Minimisation. In *Proc. of IEEE ISIT*, 2013.
- [27] Z. Olesh and A. Storjohann. The vector rational function reconstruction problem. In *Proc. of WWCA*, pages 137–149, 2006.
- [28] G. Schmidt, V. Sidorenko, and M. Bossert. Collaborative Decoding of Interleaved Reed–Solomon Codes and Concatenated Code Designs. *IEEE Trans. Inf. Theory*, 55(7):2991–3012, 2009.
- [29] V. Sidorenko and G. Schmidt. A Linear Algebraic Approach to Multisequence Shift-Register Synthesis. *Prob. Inf. Trans.*, 47(2):149–165, 2011.
- [30] W. A. Stein et al. SageMath Software. <http://www.sagemath.org>.
- [31] A. Storjohann. High-order lifting and integrality certification. *J. Symb. Comp.*, 36(3):613–648, 2003.
- [32] Y. Sugiyama, M. Kasahara, S. Hirasawa, and T. Namekawa. Further Results on Goppa Codes and their Applications to Constructing Efficient Binary Codes. *IEEE Trans. Inf. Theory*, 22(5):518–526, 1976.
- [33] J. von zur Gathen and J. Gerhard. *Modern Computer Algebra*. Cambridge Univ. Press, 3rd edition, 2012.
- [34] W. Zhou and G. Labahn. Efficient algorithms for order basis computation. *J. Symb. Comp.*, 47(7):793–819, 2012.