

Solving generalised Padé approximations over polynomial rings

Johan S. R. Nielsen
Ulm University
Institute of Communications Engineering
Ulm, Germany
jsrn@jsrn.dk

ABSTRACT

We consider a type of Padé approximations over polynomial rings over fields, dubbed “2D Padé approximations”, which generalise both simultaneous and Hermitian Padé approximations. This form is related to but distinct from both matrix-Padé approximations as well as order bases. These approximations occur in e.g. decoding of algebraic codes.

We demonstrate how to solve such approximations using row reduction of polynomial matrices, and we review the achieved asymptotic complexity when using various fast, known methods for this.

We take a closer look at the Mulders–Storjohann for row reduction and show that its complexity is linked to the orthogonality defect of the input matrix, implying that it is fast when applied to 2D Padé approximations. We design a new algorithm, which essentially carries out the computations performed by Mulders–Storjohann in a demand-driven manner, yielding an algorithm which is often even faster. This algorithm could be considered as a heavy generalisation of the classical Berlekamp–Massey algorithm.

We also consider “weighted” variants of 2D Padé approximations, and it is shown how the considered algorithms can handle this with very little speed penalty.

Categories and Subject Descriptors

I.1.2 [Symbolic and Algebraic Manipulation]: Algorithms; F.2.1 [Analysis of Algorithms and Problem Complexity]: Numerical Algorithms and Problems

General Terms

Algorithms, Theory

Keywords

Padé approximation, order basis, Berlekamp–Massey

1. INTRODUCTION

Let \mathbb{F} be a field. We consider the following “2D Padé approximation” problem: given $\mathbf{S} = [S_{i,j}] \in \mathbb{F}[x]^{\rho \times \sigma}$ and $G_1, \dots, G_\sigma \in \mathbb{F}[x]$ we are seeking $(\Lambda_1, \dots, \Lambda_\rho, \Omega_1, \dots, \Omega_\sigma) \in \mathbb{F}[x]^{\rho + \sigma}$ such that

$$\sum_{i=1}^{\rho} \Lambda_i S_{i,j} \equiv \Omega_j \pmod{G_j}, \quad \text{for } j = 1, \dots, \sigma \quad (1)$$

We are specifically interested in such vectors which have low degrees. More precisely, we will be considering two *types* of these equations: the symmetric and the asymmetric. We will consider both in a general “weighted” form, but first we discuss the simpler, unweighted case.

For the *symmetric type*, we are interested in such a vector of minimal degree, i.e. minimising $\max_{i,j} \{\deg \Lambda_i, \deg \Omega_j\}$.

For the *asymmetric type*, we likewise seek the vector of minimal degree but subject to $\max_i \{\deg \Lambda_i\} > \max_j \{\deg \Omega_j\}$.

This problem directly generalises a number of well-known approximation problems over polynomial rings, see Table 1. These occur naturally in a number of settings, e.g. control theory. In algebraic coding theory, solving the so-called “key equation” has for more than 40 years been used for minimum-distance decoding of Reed–Solomon, Goppa and other algebraic codes. It is a classical Padé approximation but it is of asymmetric type and often modulo a general polynomial (and not a power of x). Classically, this equation is solved using the Berlekamp–Massey algorithm [4], or the extended Euclidean algorithm [18]. Recently, generalisations of this equation have been identified in list-decoding of Reed–Solomon codes [15, 19] or interleaved codes [16], usually solved using ad-hoc generalisations of the Berlekamp–Massey algorithm. In the PhD thesis of the author, such usage was systematised and extended to (list-)decoding of more general algebraic-geometry codes, utilising the full generality of the definition of 2D Padé approximations [12]¹.

The 2D Padé approximation is of similar expressive power as the general approximation problems matrix Padé and order bases. We compare closer with these in Section 1.2. Following work on these, we will put the 2D Padé approximation problem into a module theoretic framework and solve them using row reduction of polynomial matrices. For a given problem, the matrix to reduce will have size $(\rho + \sigma) \times (\rho + \sigma)$ and degree $\max_j \{\deg G_j\}$. We will demonstrate this in Section 2. The definition of 2D Padé approximation is here seen to naturally fit with the lattice formulation. In fact, 2D Padé approximations can be solved by any of

Submitted to the 39th International Symposium on Symbolic and Algebraic Computation, July 2014, Kobe, Japan

¹In that work, the name “2D key equation” was used for our problem, owing to the focus on coding theory.

Some approximation problems as 2D Padé		
Name	Type	Specialisations
Classical Padé [2]	S	$\rho = \sigma = 1, G_1 = x^d$
Hermite Padé [2]	S	$\sigma = 1, G_1 = x^d$
Simultaneous Padé [2]	S	$\rho = 1, G_j = x^{d_j}$
Key equation (KE) [4, 14, 18]	A	$\rho = \sigma = 1,$
Multi-LFSR [6, 16]	A	$\rho = 1, G_j = x^{d_j}$
A-G KE	A	$\rho = \sigma = 1,$
Roth-Ruckenstein KE [15]	S	$\sigma = 1, G_1 = x^d$
Roth-Ruckenstein KE Ext. [19]	S	$G_j = x^{d_j}$

Table 1: For type, “S” refers to symmetric and “A” asymmetric.

the aforementioned general approximation problems, but it would be at the cost of larger, more indirect lattices.

Row reduction can be carried out by a number of algorithms: the Mulders–Storjohann algorithm [10] is generally asymptotically good when not using fast multiplication techniques. It is subject to a divide & conquer speedup, given by Alekhovich [1]. When the matrix has small dimension compared to the degrees of its entries, the asymptotically fastest is the Las Vegas algorithm by Giorgi, Jeannerod and Villard [8] (the GJV), or the deterministic algorithm by Zhou and Labahn [20], both of which also employ fast multiplication². We will see that applying all these algorithms for solving 2D Padé approximations will have the complexities as given in Table 2. There, and in the rest of the paper, $P(n)$ is the complexity of multiplying two polynomials of degree at most n , while $M(m)$ is the complexity of multiplying two matrices in $\mathbb{F}^{m \times m}$. Also, O^\sim refers to big- O with logarithmic factors omitted, and $\#\text{supp}(G)$ for some $G \in \mathbb{F}[x]$ is the number of non-zero monomials in G .

The complexity of applying the GJV or Zhou–Labahn will follow generically. However, for the Mulders–Storjohann, we will show that it performs better than expected on matrices resulting from 2D Padé approximations. This is done by pointing out with a new analysis that the algorithm’s complexity is directly linked to the *orthogonality defect* of the input matrix. This also carries over to Alekhovich’s algorithm.

A main result of the paper is a recasting of the Mulders–Storjohann algorithm specifically for 2D Padé approximations. The intermediate results of the Mulders–Storjohann turn out to be highly redundant for such a matrix, which means that we could instead compute most of it *on demand*. The resulting “Demand–Driven” algorithm is an asymptotic improvement for many relative sizes of the approximation problem’s parameters, as well as having a very low hidden constant and low memory complexity.

In applications, e.g. in many of those from coding theory, one needs to shift the relative importance of the degrees of the Λ_i and Ω_j ; for this, we introduce *weights*: $\boldsymbol{\eta} = (\eta_1, \dots, \eta_\rho) \in \mathbb{N}_0^\rho$ and $\boldsymbol{\mu} = (\mu_1, \dots, \mu_\sigma) \in \mathbb{N}_0^\sigma$, as well as $\nu \in \mathbb{Z}_+$. For the *symmetric* type, we now wish to minimise

$$\max_{i,j} \{\nu \deg \Lambda_i + \eta_i, \nu \deg \Omega_j + \mu_j\}$$

For the *asymmetric* type, we wish to minimise the same but

²These two algorithms compute “order bases”, which can be used for row reduction. See also Section 1.2.

Complexity for solving a weighted 2D Padé approximation	
Algorithm	Field operations in big- O^\sim
Mulders–Storjohann	$\rho(\rho + \sigma)^2 \gamma^2$
Demand–Driven	$\rho^2(\rho + \sigma) \gamma P_G(\gamma)$
Alekhovich	$M(\rho + \sigma) P(\rho \gamma)$
GJV or Zhou–Labahn	$M(\rho + \sigma) P(\gamma)$

Table 2: Here, $\gamma = \max_j \{\deg G_j\}$ in the unweighted case, and $\gamma = \max\{\eta_1/\nu, \dots, \eta_\rho/\nu, \deg G_1 + \mu_1/\nu, \dots, \deg G_\sigma + \mu_\sigma/\nu\}$ in general. $P_G(\gamma)$ is the cheaper of $P(\gamma)$ and $O(\gamma \max_j \{\#\text{supp}(G_j)\})$.

subject to

$$\max_i \{\nu \deg \Lambda_i + \eta_i\} > \max_j \{\nu \deg \Omega_j + \mu_j\}$$

Notice that ν emulates allowing fractional weights in a highly controlled manner.

We will show in Section 5 two mappings which embed these weights into the matrix to be row reduced. The first is simple but will generically incur an overhead of at least a factor ν in the complexity. We therefore derive from it a slightly more involved mapping which embeds the weights with practically no overhead. Finally, in Section 5.1 we show how the weights can be handled by the new demand–driven algorithm.

The special case of weighted 2D Padé approximations with $\rho = 1$ were investigated by the author in [11] where also a simpler version of the Demand–Driven algorithm was presented.

1.1 Notation

For $\mathbf{v} = (v_1, \dots, v_m) \in \mathbb{F}[x]^m$, let $\deg \mathbf{v} = \max_i \{\deg v_i\}$, and similarly for $\mathbf{V} = [v_{i,j}] \in \mathbb{F}[x]^{n \times m}$ then $\deg \mathbf{V} = \max_{i,j} \{\deg v_{i,j}\}$. We will also deal with row-degrees of matrices: $\text{rowdeg } \mathbf{V} = \sum_{i=1}^n \deg \mathbf{v}_i$ where the \mathbf{v}_i are the rows of \mathbf{V} .

By the leading position $\text{LP}(\mathbf{v})$, we will mean the *rightmost* element of \mathbf{v} which has degree $\deg \mathbf{v}$. This breaking of ties should be kept in mind; it has subtle effects in e.g. Section 5. From this we let $\text{LT}(\mathbf{v}) = v_{\text{LP}(\mathbf{v})}$. For a polynomial $p = \sum_{h=0}^{\deg p} p_h x^h \in \mathbb{F}[x]$, then $\text{LC}(p) = p_{\deg p}$.

We will be working with free $\mathbb{F}[x]$ -modules which are subsets of $\mathbb{F}[x]^m$. When we say “module” we will always mean such a module, and we will be representing bases of them by matrices whose rows are basis elements. We will refer to such matrices directly as “bases”.

1.2 Relation to other approximations

There are other, already studied approximations of high generality. We briefly compare our new definition with two of these: solving 2D Padé approximations with either of these are possible but less direct as the manner we propose in Section 2.

Slightly paraphrased, a matrix Padé [3] is, given $\mathbf{T} \in \mathbb{F}[x]^{m \times s}$ and $N, M, \tau \in \mathbb{Z}_+$, the task of finding $\mathbf{A} \in \mathbb{F}[x]^{n \times m}$ and $\mathbf{B} \in \mathbb{F}[x]^{n \times s}$ such that $\mathbf{A}\mathbf{T} \equiv \mathbf{B} \pmod{x^\tau}$ while $\deg \mathbf{A} \leq N$ and $\deg \mathbf{B} \leq M$, and \mathbf{A} has full rank. It is not embedded in the definition that the rows of \mathbf{A} should in any way represent the *smallest* solutions, or equivalently, that \mathbf{A} is row reduced. Therefore, though the congruence part of 2D Padé approximations can be emulated with matrix Padé, we can not simply choose any off-the-shelf algorithm for matrix

Padé and apply.

The related “Order bases” [3] embeds this kind of minimality, but does so by requiring a complete basis of solutions: given $\mathbf{T} \in \mathbb{F}[x]^{m \times s}$ and an order $\tau \in \mathbb{Z}_+$, the order basis is a matrix $\mathbf{A} \in \mathbb{F}[x]^{n \times m}$ whose rows span the $\mathbb{F}[x]$ -module of vectors $\mathbf{a} \in \mathbb{F}[x]^m$ satisfying $\mathbf{a}\mathbf{T} \equiv \mathbf{0} \pmod{x^\tau}$. Furthermore, \mathbf{A} should have minimal row degree, or rather, it should be in weak Popov form, see Section 2. Unweighted, symmetric 2D Padé approximations where all G_j are the same power of x can be solved as an order basis: set $\mathbf{T} = [\mathbf{S}^\top \mid -\mathbf{I}]^\top$, and then $(\Lambda_1, \dots, \Lambda_\rho, \dots, \Omega_1, \dots, \Omega_\sigma)$ will be a minimal-degree row in the order basis. For asymmetric 2D Padé, we instead find the minimal-degree row with leading position at most ρ . When the G_j are different powers of x , we can multiply the columns of \mathbf{T} by x -powers in an appropriate way such that one value of $\tau = \max_j \{\deg G_j\}$ suffice for the congruence.

For general G_j , one would need to convert to a larger order basis problem. One approach is to set $\mathbf{T} = [\mathbf{S}^\top \mid -\mathbf{I} \mid \text{diag}(G_1, \dots, G_\sigma)]^\top$ and set τ large enough so that the required congruence actually lifts to an equality for all rows in the found order basis. Another approach is solve the 2D Padé approximation by row reduction, using Section 2, and solve this row reduction using an order basis computation as explained in [8]. This has the advantage of immediately carrying over the solutions from this paper for handling weights. In either case the order basis problem considered will have more rows than $\rho + \sigma$ and seemingly attempts to capture more information than we are interested in. The asymptotic complexity is completely unharmed, though, so one can therefore sensibly use order bases for solving 2D Padé approximations.

2. SOLUTION BY ROW REDUCTION

Let some unweighted 2D Padé approximation problem be given; i.e we are seeking the $\mathbf{s} = (\Lambda_1, \dots, \Lambda_\rho, \Omega_1, \dots, \Omega_\sigma)$ such that the congruences (1) are satisfied, and such that $\deg \mathbf{s}$ is minimal. If we are solving an asymmetric approximation, we also wish $\text{LP}(\mathbf{s}) \leq \rho$. Consider first the space \mathcal{M} of all vectors $(\lambda_1, \dots, \lambda_\rho, \omega_1, \dots, \omega_\sigma) \in \mathbb{F}[x]^{\rho+\sigma}$ which satisfy only the congruence equations:

$$\sum_{i=1}^{\rho} \lambda_i S_{i,j} \equiv \omega_j \pmod{G_j}, \quad \text{for } j = 1, \dots, \sigma \quad (2)$$

\mathcal{M} forms an $\mathbb{F}[x]$ -module. In fact, we have

PROPOSITION 1. \mathcal{M} is generated as an $\mathbb{F}[x]$ -module by the rows of $\mathbf{M} \in \mathbb{F}[x]^{(\rho+\sigma) \times (\rho+\sigma)}$ where

$$\mathbf{M} = \left(\begin{array}{c|c} \mathbf{I} & \mathbf{S} \\ \hline \mathbf{0} & \text{diag}(G_1, \dots, G_\sigma) \end{array} \right)$$

PROOF. Equation (2) is easily seen to be satisfied for each row of \mathbf{M} , so by $\mathbb{F}[x]$ -linearity, it is satisfied for every vector in \mathcal{M} . Contrarily, for any vector $\mathbf{v} = (\lambda_1, \dots, \lambda_\rho, \omega_1, \dots, \omega_\sigma)$ satisfying (2) there must exist $p_1, \dots, p_\sigma \in \mathbb{F}[x]$ such that $\sum_{i=1}^{\rho} \lambda_j S_{i,j} + p_i G_i = \omega_i$ wherefore $\mathbf{v} = \lambda_1 \mathbf{m}_1 + \dots + \lambda_\rho \mathbf{m}_\rho + p_1 \mathbf{m}_{\rho+1} + \dots + p_\sigma \mathbf{m}_{\rho+\sigma}$, where \mathbf{m}_i are the rows of \mathbf{M} . \square

To solve the 2D Padé approximation we can seek a vector of minimal degree in the row space of \mathbf{M} , and in the asymmetric case, also subject to having leading position at most ρ . We will see that in the weak Popov form for polynomial matrices, which was introduced in [10], such vectors will figure directly as rows:

DEFINITION 2. A full-rank matrix $\mathbf{V} \in \mathbb{F}[x]^{m \times m}$ is in weak Popov form if and only if the leading position of all rows are different. The orthogonality defect of \mathbf{V} is $\Delta(\mathbf{V}) \triangleq \text{rowdeg } \mathbf{V} - \deg \det \mathbf{V}$.

We will use “a weak Popov form of \mathbf{U} ” for some matrix \mathbf{U} to mean any unimodular equivalent matrix \mathbf{V} which is in weak Popov form. That such a \mathbf{V} always exists for any \mathbf{U} follows from e.g. the Mulders–Storjohann algorithm for computing it; note however that it is not canonical. That a matrix \mathbf{V} is “row reduced” means that $\text{rowdeg } \mathbf{V}$ is minimal over all matrices unimodular equivalent to \mathbf{V} . It follows from the lemma below that a matrix in weak Popov form is row reduced. Sometimes, “row reduced” is defined directly as the more useful definition of weak Popov form above.

The concept of orthogonality defect was introduced by Lenstra [9] for estimating the running time of his algorithm for row reduction; we will use it to a similar effect. The following lemma gives the foundations for such a use. We omit the easy proof which can be found in [12]:

LEMMA 3 ([12, PROPOSITION 2.5]). If a matrix \mathbf{V} over $\mathbb{F}[x]$ is in weak Popov form then $\Delta(\mathbf{V}) = 0$.

Clearly we always have $\Delta(\mathbf{V}) \geq 0$ so since the determinant is the same for any basis of the module for which \mathbf{V} is a basis, $\Delta(\mathbf{V})$ measures how much $\text{rowdeg } \mathbf{V}$ is greater than the minimal degree possible.

DEFINITION 4. The value of a vector $\psi : \mathbb{F}[x]^m \rightarrow \mathbb{N}_0$ is $\psi(\mathbf{v}) = m \deg \mathbf{v} + \text{LP}(\mathbf{v})$

The following well-known result is key to our method:

PROPOSITION 5. Let $\mathbf{V} \in \mathbb{F}[x]^{m \times m}$ be a basis in weak Popov form of a module \mathcal{V} . Any non-zero $\mathbf{b} \in \mathcal{V}$ satisfies $\deg \mathbf{v} \leq \deg \mathbf{b}$ where \mathbf{v} is the row of \mathbf{V} with $\text{LP}(\mathbf{v}) = \text{LP}(\mathbf{b})$.

PROOF. Since \mathbf{V} is a basis of \mathcal{V} , there exists $p_1, \dots, p_m \in \mathbb{F}[x]$ such that $\mathbf{b} = p_1 \mathbf{v}_1 + \dots + p_m \mathbf{v}_m$ where the \mathbf{v}_i are the rows of \mathbf{V} . But the \mathbf{v}_i all have different leading position, so the $p_i \mathbf{v}_i$ must as well for those $p_i \neq 0$, which in turn means that their values are all different. Notice that when for two $\mathbf{u}_1, \mathbf{u}_2$ with different value, then the value of $\mathbf{u}_1 + \mathbf{u}_2$ is that of either \mathbf{u}_1 or \mathbf{u}_2 . That implies that there is a j such that $\psi(\mathbf{b}) = \psi(p_j \mathbf{v}_j)$, and this gives $\text{LP}(\mathbf{b}) = \text{LP}(\mathbf{v})$ and $\deg \mathbf{b} = \deg \mathbf{v}_j + \deg p_j$. \square

For a 2D Padé approximation, we therefore see that a basis \mathbf{V} for \mathcal{M} in weak Popov form must contain a row \mathbf{s} which is a minimal solution: if the approximation is symmetric, it is simply the minimal degree row of \mathbf{V} , while if it is asymmetric, it is the minimal degree row of \mathbf{V} which has leading position at most ρ .

REMARK 6. For asymmetric 2D Padé approximations, it is easy to characterise the entire space of solutions, and not only those of minimal degree, from the weak Popov form, by extending Proposition 5: from the degrees of the coefficient polynomials p_i we can determine the leading position of any $\mathbf{b} \in \mathcal{M}$. This is a reflection of the deeper statement that if \mathbf{V} is in weak Popov form, then its rows form a Gröbner basis of its row space for a certain module monomial ordering. See [12, Section 2.1.2 and Theorem 2.35].

Any algorithm which brings $\mathbb{F}[x]$ -matrices to weak Popov form can thus be used to solve the approximation. There exist a number of such algorithms, and whichever is asymptotically fastest on generic, square matrices depends on the relative size of its dimension m and its degree d . The Mulders–Storjohann [10] has $O(m^3 d^2)$ but with a very low hidden constant. It admits a divide & conquer variant by Alekhovich [1] which has $O(\tilde{M}(m)P(md))$. Using the fast order basis algorithm by Giorgi, Jeannerod or Villard [8], or the one by Zhou and Labahn [20], we have a complexity of $O(\tilde{M}(m)P(d))$. Plugging the parameters of \mathbf{M} into the last expression, we directly get the corresponding entry in Table 2 for the unweighted case. However, for the two former algorithms, we will show in the following section how their complexity estimates can be refined to depend on the orthogonality defect of the input matrix. Due to its special form, \mathbf{M} has particularly low orthogonality defect:

LEMMA 7. $\Delta(\mathbf{M}) = \deg \mathbf{S} < \rho \max_j \{\deg G_j\}$

PROOF. Since \mathbf{M} is upper triangular $\det(\mathbf{M}) = \prod_{i=1}^{\sigma} G_i(x)$ and the lemma follows. \square

3. MULDETS–STORJOHANN

We will now review the Mulders–Storjohann algorithm [10] for bringing a matrix to weak Popov form. We consider only square matrices, though the algorithm also applies to non-square. This allows a new complexity analysis, tying the complexity of the algorithm to the orthogonality defect. This analysis was also presented in [13], where it was utilized in a fast multi-trial list-decoder for Reed–Solomon codes.

DEFINITION 8. *Applying a row reduction on a full-rank matrix over $\mathbb{F}[x]$ means to find two different rows $\mathbf{v}_i, \mathbf{v}_j$, $\deg \mathbf{v}_i \leq \deg \mathbf{v}_j$ such that $\text{LP}(\mathbf{v}_i) = \text{LP}(\mathbf{v}_j)$, and then replacing \mathbf{v}_j with $\mathbf{v}_j - \alpha x^\theta \mathbf{v}_i$ where $\alpha \in \mathbb{F}$ and $\theta \in \mathbb{N}_0$ are chosen such that the leading term of the polynomial $\text{LT}(\mathbf{v}_j)$ is cancelled.*

LEMMA 9. *If \mathbf{v}'_j is the vector replacing \mathbf{v}_j in a row reduction, then $\psi(\mathbf{v}'_j) < \psi(\mathbf{v}_j)$.*

PROOF. We can't have $\deg \mathbf{v}'_j > \deg \mathbf{v}_j$ since all terms of both \mathbf{v}_j and $\alpha x^\theta \mathbf{v}_i$ have degree at most $\deg \mathbf{v}_j$. If $\deg \mathbf{v}'_j < \deg \mathbf{v}_j$ we are done since $\text{LP}(\mathbf{v}'_j) < m$, so assume $\deg \mathbf{v}'_j = \deg \mathbf{v}_j$. Let $h = \text{LP}(\mathbf{v}_j) = \text{LP}(\mathbf{v}_i)$. By the definition of LP, all terms in both \mathbf{v}_j and $\alpha x^\theta \mathbf{v}_i$ to the right of h must have degree less than $\deg \mathbf{v}_j$, and so also all terms in \mathbf{v}'_j to the right of h satisfies this. The row reduction ensures that $\deg v'_{j,h} < \deg v_{j,h}$, so it must then be the case that $\text{LP}(\mathbf{v}'_j) < h$. \square

The elegant Mulders–Storjohann algorithm is now succinctly described as Algorithm 1.

Algorithm 1 Mulders–Storjohann

Input: $\mathbf{V} \in \mathbb{F}[x]^{m \times m}$, a basis of \mathcal{V} .

Output: A basis of \mathcal{V} in weak Popov form.

- 1 Apply row reductions on the rows of \mathbf{V} until no longer possible.
 - 2 **return** \mathbf{V} .
-

LEMMA 10. *The Mulders–Storjohann algorithm is correct. It performs fewer than $K = m\Delta(\mathbf{V}) + \sum_{i=1}^m (\text{LP}\mathbf{v}_i - i) \leq m\Delta(\mathbf{V}) + \frac{1}{2}m^2$ row reductions on input \mathbf{V} with rows \mathbf{v}_i . It has complexity $O(mK \deg \mathbf{V}) \in O(m^2 \Delta(\mathbf{V}) \deg \mathbf{V})$.*

PROOF. Since we can apply a row reduction on a matrix if and only if it is not in weak Popov form, the algorithm must bring \mathbf{V} to weak Popov form in case it terminates. Since all we have then done is row operations, the resulting matrix must have the same row space as the input. Termination follows directly from Lemma 9 since the value of a row decreases each time a row reduction is performed. We will be more precise, though. Denote by \mathbf{V} the input to the algorithm and by \mathbf{U} the output in weak Popov form, and let $\mathbf{u}_i, \mathbf{v}_i$ denote the rows of these matrices. Then $\text{rowdeg}(\mathbf{U}) = \deg \det(\mathbf{U}) = \deg \det(\mathbf{V})$ and the $\text{LP}(\mathbf{u}_i)$ are all different. The total number of row reductions is then upper bounded by:

$$\begin{aligned} & \sum_{i=1}^m (\psi(\mathbf{v}_i) - \psi(\mathbf{u}_i)) \\ & \leq m(\text{rowdeg}(\mathbf{V}) - \text{rowdeg}(\mathbf{U})) + \sum_{i=1}^m (\text{LP}(\mathbf{u}_i) - \text{LP}(\mathbf{v}_i)) \\ & \leq m\Delta(\mathbf{V}) + \sum_{i=1}^m (\text{LP}\mathbf{v}_i - i) \end{aligned}$$

For the asymptotic complexity, note that during the algorithm, no polynomial in the intermediate matrix will have larger degree than $\deg \mathbf{V}$. One row reduction consists of m times scaling and adding two such polynomials. \square

For generic square matrices \mathbf{V} the orthogonality defect is $O(m \deg \mathbf{V})$; therefore the above complexity estimate relaxes to the same as in [10]. However, we have:

COROLLARY 11. *Solving an unweighted 2D Padé approximation by applying the Mulders–Storjohann algorithm to \mathbf{M} has complexity $O(\rho(\rho + \sigma)^2 \max_j \{\deg G_j^2\})$.*

Alekhovich showed in [1] how to accelerate the Mulders–Storjohann with fast multiplication. His algorithm calculates the same row reductions, but they are bundled and structured in a binary tree which allows matrix multiplication techniques to reduce the quadratic dependence on $\deg \mathbf{V}$. The reported complexity is $O(\tilde{M}(m)P(m \deg \mathbf{V}))$. However, by the observations on orthogonality defect as above, this can be immediately improved to $O(\tilde{M}(m)\Delta(\mathbf{V}))$. More details can be found in [11]. We therefore get the improved complexity as given in Table 2.

For $\rho \in \omega(1)$, this is still asymptotically worse than the fast algorithms GJV and Zhou–Labahn. However, $\rho \in O(1)$ is still an interesting special case, including e.g. multi-sequence linear-feedback register synthesis, and here one would need a deeper analysis and/or experimentation to determine which algorithm is fastest.

4. THE DEMAND–DRIVEN ALGORITHM

We will show how Mulders–Storjohann admits a wholly different variant, but only when applied to 2D Padé approximations. This variant, which we call the Demand–Driven algorithm, is an improvement when ρ is small and when all or most G_j are sparse, i.e. have few non-zero coefficients.

The intuitive idea is very simple: for any basis of \mathcal{M} , it is basically sufficient to keep track of only its ρ first columns, as the remaining columns can be computed from these. By keeping track of the leading position for each row, one then

computes the other entries only when the need arise. Formalising and proving the resulting algorithm is unfortunately somewhat technical. The result is Algorithm 2.

Algorithm 2 Demand-Driven algorithm

Input: $M = \left(\begin{array}{c|c} \mathbf{I} & \mathbf{S} \\ \hline 0 & \text{diag}(G_1, \dots, G_\sigma) \end{array} \right)$.

Let \mathbf{m}_i be the rows of this matrix.

Output: The first ρ columns of a basis of \mathcal{M} in weak Popov form.

```

1  $\mathbf{\Lambda} =$  first  $\rho$  columns of  $M$ . Denote its rows by  $\lambda_i$ .
2  $(\theta_i, \alpha_i) = (\text{deg}(\mathbf{m}_i), \text{LC}(\text{LT}(\mathbf{m}_i)))$  for  $i = 1, \dots, \rho + \sigma$ 
3  $W = \{(i, \text{LP}(\mathbf{m}_i)) \mid 1 \leq i \leq \rho \wedge \text{LP}(\mathbf{m}_i) \neq i\}$ 
4 while  $W \neq \emptyset$  do
5    $(i, h) = \text{pop}(W)$ 
6   if  $\theta_i < \theta_h$  then swap  $(\lambda_i, \alpha_i, \theta_i)$  and  $(\lambda_h, \alpha_h, \theta_h)$ 
7    $\lambda_i = \lambda_i - \frac{\alpha_i}{\alpha_h} x^{\theta_i - \theta_h} \lambda_h$ 
8   repeat
9      $(\theta_i, h) = \text{previous}(\theta_i, h)$ 
10     $\alpha_i =$  coefficient to  $x^{\theta_i}$  in
11       $\begin{cases} \lambda_{i,h} & \text{if } h \leq \rho \\ \sum_{j=1}^{\rho} \lambda_{i,j} S_{j,h-\rho} \bmod G_{h-\rho} & \text{otherwise} \end{cases}$ 
12  until  $\alpha_i \neq 0$ 
13  if  $i \neq h$  then
14    if  $(h, j) \in W$  for some  $j$  then
15      swap  $(\lambda_i, \alpha_i, \theta_i)$  and  $(\lambda_h, \alpha_h, \theta_h)$ 
16      replace  $(h, j)$  with  $(i, j)$  in  $W$ .
17    else
18      push( $W, (i, h)$ )
19 return  $\mathbf{\Lambda}$ 

```

To begin, we will overload ψ to $\mathbb{N}_0 \times \{1, \dots, \rho + \sigma\} \rightarrow \mathbb{N}_0$ by $\psi(\theta, h) = (\rho + \sigma)\theta + h$, i.e. for any non-zero $\mathbf{v} \in \mathbb{F}[x]^{\rho + \sigma}$, $\psi(\mathbf{v}) = \psi(\text{deg } \mathbf{v}, \text{LP}(\mathbf{v}))$. We will also use the helper function $\text{previous}(\theta, h)$ which gives the degree and leading position for the “previous” possible value, that is:

$$\text{previous}(\theta, h) = \begin{cases} (\theta, h - 1) & \text{if } h > 1 \\ (\theta - 1, \rho + \sigma) & \text{if } h = 1 \end{cases}$$

We will prove the correctness of Algorithm 2 by showing that the computations correspond to a possible run of a slight variant of Mulders–Storjohann; first we need a technical lemma on this variant:

LEMMA 12. *Consider the Mulders–Storjohann with input \mathbf{M} . Consider now a variant of the algorithm where we, when replacing some \mathbf{v}_j with \mathbf{v}'_j in a row reduction, instead replace it with*

$$\mathbf{v}''_j = (v'_{j,1}, \dots, v'_{j,\rho}, v'_{j,\rho+1} \bmod G_1, \dots, v'_{j,\rho+\sigma} \bmod G_\sigma)$$

This does not change correctness of the algorithm or the upper bound on the number of row reductions performed.

PROOF. Correctness follows if we can show that each of the σ modulo reductions could have been achieved by a series of $\mathbb{F}[x]$ row operations on the current matrix \mathbf{V} after the row reduction producing \mathbf{v}' , since then \mathbf{V} would remain a basis of \mathcal{M} .

Consider the modulo reduction on the $(\rho + h)$ th position for some h . This could be achieved by adding a multiple of the vector $\mathbf{g}_h = (0, \dots, 0, G_h, 0, \dots, 0)$, with position $\rho + h$

non-zero, to \mathbf{v}' . Since \mathbf{g}_h is a row in \mathbf{M} , then as long as this has not yet been row reduced, the $(\rho + h)$ th position reduction is allowed. Notice that if $\psi(\mathbf{v}') < \psi(\mathbf{g}_h)$ then the modulo reduction on the $(\rho + h)$ th position is void.

Introduce now a loop invariant involving $J_h = \{\mathbf{g}_h\}$, a subset of the current rows in \mathbf{V} having two properties: that \mathbf{g}_h can be constructed as an $\mathbb{F}[x]$ -linear combination of the rows in J_h ; and that each $\mathbf{v} \in J_h$ has $\psi(\mathbf{v}) \leq \psi(\mathbf{g}_h)$. After row reductions on rows not in J_h , the $(\rho + h)$ th modulo reduction is therefore allowed, since \mathbf{g}_h can be constructed by the rows in J_h . On the other hand, after a row reduction on a row $\mathbf{v} \in J_h$ by some \mathbf{v}_k resulting in \mathbf{v}' , the h th modulo reduction has no effect since $\psi(\mathbf{v}') < \psi(\mathbf{v}) \leq \psi(\mathbf{g}_h)$. Afterwards, J_h is updated as $J_h = J_h \setminus \{\mathbf{v}\} \cup \{\mathbf{v}', \mathbf{v}_k\}$ and the loop invariant is kept since $\psi(\mathbf{v}_k) \leq \psi(\mathbf{v})$.

Since $\psi(\mathbf{v}'_j) \leq \psi(\mathbf{v}_j)$ the proof of Lemma 10 shows that the number of row reductions performed is not greater than in the original Mulders–Storjohann. \square

We will say for a matrix \mathbf{U} that there is a “collision on (i, j) ” if one could perform a row reduction involving \mathbf{u}_j and \mathbf{u}_i , i.e. $i \neq j$, $\text{LP}(\mathbf{u}_i) = \text{LP}(\mathbf{u}_j)$; we will also say that these rows are “involved in a collision”. Now the proof of correctness; it essentially just establishes that W is a “work-list” containing indices of rows involved in a collision, which means we should at some point row reduce it. All indices i not in W have been “parked” such that $\text{LP}(\mathbf{v}_i) = i$, and they will only be used further in the algorithm if there is a row named in W with leading position i .

THEOREM 13. *Algorithm 2 is correct.*

PROOF. Let \mathbf{V} be the matrix continually changing in the variant of Mulders–Storjohann described in Lemma 12. For notational convenience, and simplicity in the description of Algorithm 2, we will consider a further variant of Mulders–Storjohann where we sometimes swap two rows, to be described momentarily.

Let us write $(i, \star) \notin W$ to mean $\neg \exists j. (i, j) \in W$, and similarly for $(\star, i) \notin W$. We will then demonstrate that each iteration of Algorithm 2 corresponds to exactly one row reduction on \mathbf{V} , and that the following loop invariants hold:

1. $\mathbf{\Lambda}$ is the first ρ columns of \mathbf{V} ;
2. $\alpha_i x^{\theta_i}$ is the leading monomial of $\text{LT}(\mathbf{v}_i)$;
3. If there is a collision (i, j) then $\exists k. (i, k) \in W \vee (j, k) \in W$;
4. If $(i, h) \in W$ then $i \neq h$ and $\text{LP}(\mathbf{v}_i) = \text{LP}(\mathbf{v}_h) = h$;
5. For any i , there is at most one pair in W with first position i .
6. If $(i, \star) \notin W$ then $\text{LP}(\mathbf{v}_i) = i$;

The invariants are confirmed to be true after initialisation; assume now they are true on entry of an iteration, and we will show they are true on exit. Once Algorithm 2 terminates, by Invariant 3 then there are no collisions, so \mathbf{V} is in weak Popov form, so by Invariant 1 the result is correct.

If $W \neq \emptyset$, then by Invariant 4 there is a collision (i, h) , and the considered variant of Mulders–Storjohann could have

chosen to perform a row reduction on this. Invariant 4 implies that $(\star, i) \notin W$ and $(h, \star) \notin W$, and that Invariant 5 then implies that no other element in W contains i .

Now comes the promised swap in our Mulders–Storjohann variant: we will possibly perform a swap of rows of \mathbf{V} such that the row to be reduced is \mathbf{v}_i ; i.e. if $\deg \mathbf{v}_i < \deg \mathbf{v}_j$ we swap, which is exactly when $\theta_i < \theta_j$ in Line 6. We note that the Invariants 1 and 2 are maintained since everything involved is swapped. The remaining invariants never distinguish the two swapped rows, which means all non-violations and violations are kept intact. Thus for the remainder of the proof, we can assume that no swap was performed since this will be indistinguishable from the other case.

In Line 7 we perform the actual row reduction, and Invariant 1 is seen to be maintained, due to Invariant 2. The ensuing loop is for maintaining Invariant 2: it will go through possible degrees and leading positions of the updated row \mathbf{v}_i in descending order of ψ until it finds one with non-zero coefficient. Note that $\psi(\mathbf{v}'_i) < \psi(\mathbf{v}_i)$, so the loop initiates correctly. The calculation in Line 11 is also exactly correct for the considered variant of Mulders–Storjohann, cf. Lemma 12.

Refer to \mathbf{V}' with rows \mathbf{v}'_i as \mathbf{V} after this row operation, and refer to h' as the updated value of h . Thus $\text{LP}(\mathbf{v}'_i) = h'$. The last thing is then to maintain Invariants 3–6. Note that all possible violations must involve i since only for i we have $\mathbf{v}'_i \neq \mathbf{v}_i$. We distinguish the three cases of the if-statements:

- $i = h'$: Then \mathbf{v}'_i cannot be involved in any collision: for if there was a $j \neq i$ such that (j, i) is a collision in \mathbf{V}' , then $\text{LP}(\mathbf{v}_j) = \text{LP}(\mathbf{v}'_j) = \text{LP}(\mathbf{v}'_i) = i$ which means that $(j, i) \notin W$ due to Invariant 4 holding on entry of the iteration (at which time $\text{LP}(\mathbf{v}_i) \neq i$); but then $(j, \star) \notin W$ which means $\text{LP}(\mathbf{v}_j) = j$ by Invariant 6. Thus, we must have $j = h' = i$ but this was false by assumption. Since then \mathbf{v}'_i is involved in no collisions, the remaining invariants are seen to hold since: all possible violations to them must involve i ; since i is no longer in any pair of W ; and since $\text{LP}(\mathbf{v}'_i) = i$.
- $i \neq h'$ and $\exists j. (h', j) \in W$: That means $(\star, h') \notin W$ by Invariant 4 at iteration entry, and since $(\star, i) \notin W$, we do not create violations to the invariants by swapping \mathbf{v}'_i and $\mathbf{v}'_{h'}$ when also replacing (h', j) with (i, j) in W . After this swap, we have $\text{LP}(\mathbf{v}'_{h'}) = h'$, and invariants involving i all hold. There now can't be any collision involving $\mathbf{v}'_{h'}$: for if there was a $k \neq h'$ with $\text{LP}(\mathbf{v}'_k) = \text{LP}(\mathbf{v}'_{h'}) = h'$, then also $\text{LP}(\mathbf{v}_k) = h'$. But at iteration entry, by Invariant 6, since $k \neq h'$ then $(k, \star) \in W$, which means by Invariant 4 that $(k, h') \in W$ and $\text{LP}(\mathbf{v}_{h'}) = h'$. This contradicts $\text{LP}(\mathbf{v}_{h'}) = j$ from $(h', j) \in W$. Summarily, since we then have removed the only element in W involving h' and since $\text{LP}(\mathbf{v}'_{h'}) = h'$, the invariants are again maintained.
- $i \neq h'$ and $(h', \star) \notin W$: By Invariant 6, we have $\text{LP}(\mathbf{v}_{h'}) = h'$ and so (i, h') is a collision, meaning the invariants are maintained by adding (i, h') to W .

□

The algorithm is very straight-forward to analyse, and we can even include the leading constant:

PROPOSITION 14. *Algorithm 2 performs at most*

$$2(\rho + 1)^2(\rho + \sigma)\gamma P_G(\gamma) + O((\rho + \sigma)\gamma)$$

field multiplications, where $\gamma = \max_j \{\deg G_j\}$, and $P_G(\gamma)$ is the cheaper of $P(\gamma)$ and $\gamma \max_j \{\#\text{supp}(G_j)\}$. It uses $2\rho(\rho + \sigma)\gamma + O((\rho + \sigma)\gamma)$ field elements of memory.

PROOF. Notice that the complexity of executing Line 11 is at most $(\rho + 1)P_G(\gamma)$: for we can compute this one coefficient either by ρ multiplications and one division of polynomials of degree at most γ ; or we can in each product $\lambda_{i,j}S_{j,h-\rho}$ compute only the $\text{supp}(G_{h-\rho})$ single coefficients influencing the x^{θ_i} -coefficient of the remainder.

Each iteration through the repeat loop will decrease the upper bound on $\psi(\mathbf{v}_i)$, so by arguments exactly like those of the proof of Lemma 10, the total number of times through the repeat loop is at most the number of row reductions by Mulders–Storjohann, i.e. $(\rho + \sigma)(\Delta(M) + \rho)$. Apart from the Line 11, only Line 7 is not $O(1)$, and this has complexity at most $O(\rho \deg(M)) \subset \rho P_G(\gamma)$. Multiplying the number of iterations with the cost per iteration yields the result. The main memory requirement is for $\mathbf{\Lambda}$ and \mathbf{S} , and otherwise only asymptotically lesser amounts. □

REMARK 15. Algorithm 2 bears a striking resemblance to the Berlekamp–Massey algorithm [4] and various generalisation, e.g. for synthesising multi-sequence LFSRs [16], or for solving Roth–Ruckenstein “key equations” [15]. The Mulders–Storjohann algorithm can be considered a generalisation of the Extended Euclidean algorithm, so this resemblance generalise the long-known correspondence between the Euclidean algorithm and the Berlekamp–Massey [5, 7] to the much more general case of 2D Padé approximations.

5. ADDING WEIGHTS

We now study the weighted variant of the problem. We are therefore again seeking $\mathbf{s} = (\Lambda_1, \dots, \Lambda_\rho, \Omega_1, \dots, \Omega_\sigma) \in \mathbb{F}[x]^{\rho+\sigma}$ such that the congruence (2) holds, but now such that $\max_{i,j} \{\nu \deg \Lambda_i + \eta_i, \nu \deg \Omega_j + \mu_j\}$ is minimised. If we are solving an asymmetric approximation, we furthermore wish that $\max_i \{\nu \deg \Lambda_i + \eta_i\} > \max_j \{\nu \deg \Omega_j + \mu_j\}$.

We will introduce two injective mappings for matrices such that finding a weak Popov form of the image of \mathbf{M} under either will solve the weighted approximation problem. The first is a straightforward embedding of the weights but carries a computational penalty if one can not handle it specifically within the row reduction algorithm. The second we derive from the first to mitigate this problem.

First the straightforward map $\Phi : \mathbb{F}[x]^{\rho+\sigma} \mapsto \mathbb{F}[x]^{\rho+\sigma}$:

$$\Phi((v_1, \dots, v_{\rho+\sigma})) = (x^{\eta_1} v_1(x^\nu), \dots, x^{\eta_\rho} v_\rho(x^\nu), \\ x^{\mu_1} v_{\rho+1}(x^\nu), \dots, x^{\mu_\sigma} v_{\rho+\sigma}(x^\nu))$$

Extend Φ element-wise to sets of vectors, and extend Φ row-wise to $(\rho + \sigma) \times (\rho + \sigma)$ matrices such that the i th row of $\Phi(\mathbf{V})$ is $\Phi(\mathbf{v}_i)$, where \mathbf{v}_i are the rows of \mathbf{V} . Note that $\Phi(\mathcal{M})$ is a free $\mathbb{F}[x^\nu]$ -module of dimension $\rho + \sigma$, and that any basis of it is by Φ^{-1} sent back to a basis of \mathcal{M} .

PROPOSITION 16. *A minimal solution to the weighted 2D Padé approximation problem is the Φ^{-1} -image of a vector in $\Phi(\mathcal{M})$ with minimal degree, and if the approximation is asymmetric, of only those vectors with LP at most ρ .*

PROOF. This follows immediately since for any vector $\mathbf{v} = (v_1, \dots, v_{\rho+\sigma}) \in \mathbb{F}[x]^{\rho+\sigma}$, then $\deg \Phi(\mathbf{v}) = \max_{i,j} \{\nu \deg v_i + \eta_i, \nu \deg v_{\rho+j} + \mu_j\}$, and $\text{LP}(\Phi(\mathbf{v})) \leq \rho \iff \max_i \{\nu \deg v_i + \eta_i\} > \max_j \{\nu \deg v_{\rho+j} + \mu_j\}$. □

The above proposition together with Proposition 5 therefore immediately gives a method for solving the weighted approximation problems: simply find a weak Popov form³ of $\Phi(\mathbf{M})$, where \mathbf{M} is from Proposition 1. However, since $\deg \Phi(\mathbf{M}) \geq \nu \deg \mathbf{M}$ we should expect a penalty in performance of at least a factor ν from doing this. To mitigate this problem, we now derive a second mapping. It can already be mentioned, though, since we will need it for the Demand–Driven algorithm later, that the Mulders–Storjohann can be implemented to not suffer this ν penalty, even under the direct Φ -mapping.

For notational convenience, let $\mathbf{w} = (w_1, \dots, w_{\rho+\sigma}) = (\eta_1, \dots, \eta_\rho, \mu_1, \dots, \mu_\sigma)$. For the improved mapping, consider first the permutation π of $[1, \dots, m]$ indirectly defined by

$$\pi(i) > \pi(j) \iff (w_i \bmod \nu) > (w_j \bmod \nu) \vee ((w_i \bmod \nu) = (w_j \bmod \nu) \wedge i > j)$$

This is seen to be well-defined. Extend now π to elements of $\mathbb{F}[x]^{\rho+\sigma}$ such that π permutes the positions of such vectors. Our desired mapping is now Ψ :

$$\Psi((v_1, \dots, v_{\rho+\sigma})) = \pi((x^{\lfloor w_1/\nu \rfloor} v_1, \dots, x^{\lfloor w_{\rho+\sigma}/\nu \rfloor} v_{\rho+\sigma}))$$

PROPOSITION 17. *For any $\mathbf{v} \in \mathbb{F}[x]^{\rho+\sigma}$ then*

$$(\pi^{-1} \circ \text{LP} \circ \Psi)(\mathbf{v}) = (\text{LP} \circ \Phi)(\mathbf{v})$$

PROOF. Let v_i be the elements of \mathbf{v} , and $h = (\text{LP} \circ \Phi)(\mathbf{v})$. We will prove that no index but $\pi(h)$ can be $(\text{LP} \circ \Psi)(\mathbf{v})$. Consider first some $j > h$. By the definition of h then

$$\nu \deg v_h + w_h > \nu \deg v_j + w_j, \quad \text{i.e.} \quad (3)$$

$$\deg v_h + \lfloor w_h/\nu \rfloor + \frac{w_h \bmod \nu}{\nu} > \deg v_j + \lfloor w_j/\nu \rfloor + \frac{w_j \bmod \nu}{\nu}$$

So either $\deg v_h + \lfloor w_h/\nu \rfloor > \deg v_j + \lfloor w_j/\nu \rfloor$, or they are equal and $w_h \bmod \nu > w_j \bmod \nu$. In the first case, then clearly $\pi(j)$ can't be $(\text{LP} \circ \Psi)(\mathbf{v})$ due to degrees. In the second case the degrees of $\Psi(\mathbf{v})$ at positions $\pi(h)$ and $\pi(j)$ are tied, but we have $\pi(h) > \pi(j)$, which means that $\pi(j)$ can't be the LP.

Consider now some $j < h$, so we have the same inequality (3) but with $>$ replaced by \geq . If sharp inequality really holds, then we can continue as before, so assume instead that equality holds. That implies both $\deg v_h + \lfloor w_h/\nu \rfloor = \deg v_j + \lfloor w_j/\nu \rfloor$ and $w_h \equiv w_j \pmod{\nu}$. So the degrees of $\Psi(\mathbf{v})$ at positions $\pi(h)$ and $\pi(j)$ are tied, but then since $h > j$, we have $\pi(h) > \pi(j)$. Again, $\pi(j)$ is not the LP. \square

COROLLARY 18. *For any $\mathbf{V} \in \mathbb{F}[x]^{\rho+\sigma}$, then $\Phi(\mathbf{V})$ is in weak Popov form if and only if $\Psi(\mathbf{V})$ is in weak Popov form.*

The algorithm is then clear: to solve the weighted 2D Padé approximation problem, simply compute a weak Popov form of $\Psi(\mathbf{M})$. The row with minimal degree—and in the asymmetric case, among only those with leading position among $\pi(1), \dots, \pi(\rho)$ —is the solution. This works immediately for the general row reduction matrices. The following lemma can be used for calculating the resulting complexities:

³Since $\Phi(\mathcal{M})$ is an $\mathbb{F}[x^\nu]$ -module, and not $\mathbb{F}[x]$, then it is not immediately obvious that computing a weak Popov form of $\Phi(\mathbf{M})$ will not result in a matrix with rows outside $\Phi(\mathcal{M})$. However this turns out to not be the case, which can easily be proved for Mulders–Storjohann and then generalised, see [12, Proposition 2.26]

LEMMA 19.

$$\deg \Psi(\mathbf{M}) = \gamma \quad \Delta(\Psi(\mathbf{M})) \leq \deg \Psi([\mathbf{I} \mid \mathbf{S}]) \leq \rho\gamma$$

where $\gamma = \max_{i,j} \{\lfloor \eta_i/\nu \rfloor, G_j + \lfloor \mu_j/\nu \rfloor\}$.

PROOF. The degree is obvious. For the orthogonality defect, the determinant follows from $\Psi(\mathbf{M})$ having been an upper triangular matrix before the column permutations. \square

5.1 Handling weights Demand–Drivenly

For Algorithm 2, however, the permutations performed by Ψ would make it messy to describe. Instead, we are going to use the more direct mapping Φ , and prove that for this algorithm, there is no performance penalty for ν .

Only certain ψ -values are possible for the vectors of $\Phi(\mathcal{M})$, so first the definition of previous needs to be generalised to still provide the degree and leading position of the previous possible value:

$$\text{previous}(\theta, h) = \arg \max_{\theta', h'} \{\psi(\theta', h') \mid \psi(\theta', h') < \psi(\theta, h)$$

$$\wedge \theta' \equiv w_{h'} \pmod{\nu}\}$$

Just as with π , this function is easily calculable though it is cumbersome to write down a direct description.

The Mulders–Storjohann with input $\Phi(\mathbf{M})$ exhibits structurally the same behaviour as with input \mathbf{M} ; it is therefore not surprising that we can make a demand–driven recasting for this case also. In fact, apart from the more general definition of previous, this recast is *exactly* as in Algorithm 2 except for the following details:

- Input is now $\Phi(\mathbf{M})$ and the \mathbf{m}_i are its rows.
- The output of the algorithm is now the first ρ columns of a basis of $\Phi(\mathbf{M})$ in weak Popov form.
- In Line 11, the polynomial from which to choose the coefficient to x^{θ_i} for the case $h > \rho$ is now:

$$\sum_{j=1}^{\rho} x^{-\eta_j} \lambda_{i,j} \tilde{S}_{j,h-\rho} \bmod \tilde{G}_{h-\rho},$$

$$\tilde{S}_{j,h-\rho} = x^{\mu_j} S_{j,h-\rho}(x^\nu) \quad \text{and} \quad \tilde{G}_{h-\rho} = x^{\mu_h} G_{h-\rho}(x^\nu).$$

The proof that this amended algorithm works correctly is exactly like in Theorem 13: first we get a statement equivalent to Lemma 12 but for $\Phi(\mathbf{M})$, and then all the invariants used in Theorem 13 are again shown to be fulfilled. The change to Line 11 is correct since the \mathbf{V} will at position (i, h) have exactly this element.

For complexity, we first need to bound the number of iterations through the loop; this is done by bounding how many row reductions Mulders–Storjohann with input $\Phi(\mathbf{M})$ would do. The fact that only certain ψ -values are possible implies that this number is roughly $1/\nu$ 'th of what the degree of $\Phi(\mathbf{M})$ would have us think:

LEMMA 20. *The number of row reductions performed by Mulders–Storjohann on $\Phi(\mathbf{M})$ is less than $\rho(\rho + \sigma)(\gamma + 1)$ where γ is as in Lemma 19.*

PROOF. Observe that for any $\mathbf{v} \in \mathcal{M}$, if we let $h = \text{LP}(\Phi(\mathbf{v}))$ we have

$$\begin{aligned} \psi(\Phi(\mathbf{v})) &= (\rho + \sigma)(\nu \deg v_h + w_h) + h \\ &\equiv (\rho + \sigma)w_h + h \pmod{(\rho + \sigma)\nu} \end{aligned}$$

That is, on any given interval of size $(\rho + \sigma)\nu$, then $\psi(\Phi(\mathbf{v}))$ can attain at most $\rho + \sigma$ of the values, depending on the

value of $\text{LP}(\Phi(\mathbf{v}))$. Continuing as in the proof of Lemma 10, the number of row reductions is at most $\sum_{i=1}^{\rho+\sigma} \nu^{-1}(\psi(\mathbf{v}_i) - \psi(\mathbf{u}_i))$, which simplifies to the statement. \square

PROPOSITION 21. *The amended Algorithm 2 on input $\Phi(\mathbf{M})$ performs at most*

$$2(\rho + 1)^2(\rho + \sigma)\gamma P_G(\gamma) + O(\rho(\rho + \sigma)\gamma)$$

field multiplications, where γ is as in Lemma 19, and $P_G(\gamma)$ is the cheaper of $P(\gamma)$ and $\gamma \max_j \{\#\text{supp}(G_j)\}$. It uses $2\rho(\rho + \sigma)\gamma + O((\rho + \sigma)\gamma)$ field elements of memory.

PROOF. As in the proof of Proposition 14, then Lemma 20 gives an upper bound on the number of iterations of the main loop as well as executions of Line 11. This critical line costs $P_G(\gamma)$ by exactly the same arguments as in Proposition 14 combined with the observation that all the involved polynomials are sparse with only every ν coefficient possibly non-zero; therefore computations with them cost as if their degree was only $1/\nu$ 'th. \square

6. CONCLUSION

We introduced the 2D Padé approximations, a general form of approximations for polynomials over fields which includes several earlier types as special cases. Inspired by various applications we introduce two *types* of these: the traditional “symmetric” Padé approximation, and the “asymmetric” form. We furthermore consider both types in a weighted form, where the priority of the various terms of an approximation can essentially be “shifted” by rational numbers.

Using off-the-shelf row reduction algorithms, we showed how these approximations can be solved efficiently. It speaks for the naturalness of the approximation’s definition that the matrices to be row reduced have low dimension; lower than had we solved them using one of the already studied, general approximation types matrix-Padé or order bases.

The Mulders–Storjohann is such a row reduction algorithm, and we showed that one can in general describe its complexity based on the orthogonality defect of the input matrix. As a corollary, both it and its divide & conquer variant by Alekhovich performs faster than initially expected on 2D Padé approximations.

In any basis of the lattice induced by a 2D Padé approximation problem, if it is seen as a matrix, there is a large redundant part. With this observation, we recast the Mulders–Storjohann to perform the computations on demand. This yields a “Demand–Driven” algorithm which often has lower time and memory complexity. This algorithm has been implemented in Sage [17] and is available at <http://jsrn.dk/codinglib>

For future work, it seems possible that when applying the Mulders–Storjohann for solving asymmetric 2D Padé approximation, one can carefully order the row reductions in such a way that once a solution appears, this must be minimal, and one can stop the computation. For the Demand–Driven algorithm, this would be particularly interesting, since one could then bound the complexity by the degree of the smallest solution, instead of the degree of the input matrix.

7. REFERENCES

- [1] M. Alekhovich. Linear diophantine equations over polynomials and soft decoding of Reed–Solomon codes. *IEEE Trans. Inf. Theory*, 51(7):2257–2265, July 2005.
- [2] G. Baker and P. Graves-Morris. *Padé approximants*, volume 59. Cambridge Univ. Press, 1996.
- [3] 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, 1994.
- [4] E. R. Berlekamp. *Algebraic Coding Theory*. Aegean Park Press, 1968.
- [5] J. Dornstetter. On the equivalence between berlekamp’s and euclid’s algorithms. *IEEE Trans. Inf. Theory*, 33(3):428–431, 1987.
- [6] 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.
- [7] P. Fitzpatrick. On the key equation. *IEEE Trans. Inf. Theory*, 41(5):1290–1302, 1995.
- [8] P. Giorgi, C. Jeannerod, and G. Villard. On the complexity of polynomial matrix computations. In *Proc. of ISSAC*, page 135–142, 2003.
- [9] A. Lenstra. Factoring multivariate polynomials over finite fields. *J. Comp. Syst. Sc.*, 30(2):235–248, 1985.
- [10] T. Mulders and A. Storjohann. On lattice reduction for polynomial matrices. *J. Symb. Comp.*, 35(4):377–401, 2003.
- [11] J. S. R. Nielsen. Generalised multi-sequence shift-register synthesis using module minimisation. In *Proc. of IEEE ISIT*, 2013.
- [12] J. S. R. Nielsen. *List Decoding of Algebraic Codes*. PhD thesis, Technical University of Denmark, 2013. Available at jsrn.dk.
- [13] J. S. R. Nielsen and A. Zeh. Multi-trial guruswami–sudan decoding for generalised reed–solomon codes. In *Proc. of WCC*, 2013.
- [14] R. Roth. *Introduction to Coding Theory*. Cambridge Univ. Press, 2006.
- [15] R. Roth and G. Ruckenstein. Efficient decoding of reed–solomon codes beyond half the minimum distance. *IEEE Trans. Inf. Theory*, 46(1):246–257, 2000.
- [16] V. Sidorenko and G. Schmidt. A linear algebraic approach to multisequence shift-register synthesis. *Problems of Information Transmission*, 47(2):149–165, 2011.
- [17] W. A. Stein et al. *Sage Mathematics Software*. <http://www.sagemath.org>.
- [18] Y. Sugiyama, M. Kasahara, S. Hirasawa, and T. Namekawa. A method for solving key equation for decoding goppa codes. *Information and Control*, 27(1):87–99, 1975.
- [19] A. Zeh, C. Gentner, and D. Augot. An interpolation procedure for list decoding reed-solomon codes based on generalized key equations. *IEEE Trans. Inf. Theory*, 57(9):5946–5959, 2011.
- [20] W. Zhou and G. Labahn. Efficient algorithms for order basis computation. *J. Symb. Comp.*, 47(7):793–819, July 2012.