# Robust Padé Approximation via SVD* 

Pedro Gonnet ${ }^{\dagger}$ Stefan Güttel ${ }^{\ddagger}$ Lloyd N. Trefethen ${ }^{\S}$


#### Abstract

Padé approximation is considered from the point of view of robust methods of numerical linear algebra, in particular, the singular value decomposition. This leads to an algorithm for practical computation that bypasses most problems of solution of nearly-singular systems and spurious pole-zero pairs caused by rounding errors, for which a MATLAB code is provided. The success of this algorithm suggests that there might be variants of Padé approximation that are pointwise convergent as the degrees of the numerator and denominator increase to $\infty$, unlike traditional Padé approximants, which converge only in measure or capacity.


Key words. Padé approximation, SVD, regularization, Froissart doublet, Nuttall-Pommerenke theorem

AMS subject classifications. 41A21, 65F22
DOI. 10.1137/110853236
I. Introduction. Padé approximants, like rational approximants more generally, are well known to be fragile. One complication is that degeneracies may occur in which the numerator and denominator have less than the allowed degree, which leads to several entries in the Padé table being identical, some of them matching the Taylor series of the function being approximated to less than the expected order. Another complication is that even in theory, in the absence of rounding errors on a computer, Padé approximants are subject to the appearance of seemingly spurious pole-zero pairs or "Froissart doublets" in arbitrary locations that prevent pointwise convergence of the kind one might hope for. When rounding errors or other forms of noise are brought into the picture, such anomalies become almost ubiquitous.

In this paper we shall propose a theoretical variation on Padé approximation, and a corresponding numerical algorithm, that go a good way toward eliminating

[^0]these anomalies. Philosophically speaking, our approach is one of regularization of an ill-posed problem, and perhaps it is no surprise that it relies upon the singular value decomposition (SVD). The reason Padé approximation is ill-posed is that it is related to analytic continuation, since the aim is typically to gain information about a function in a region of the complex plane based on information at a single point.

This paper is an outgrowth of earlier work with Pachón, which proposed an algorithm for robust rational interpolation in roots of unity [15]. The present problem is a limiting case in which the points of interpolation degenerate to a single point. Here one can take advantage of the exceptionally clean theory of square blocks of degeneracies in the Padé table. Our algorithm runs in standard floating point computer arithmetic and treats the mathematically classical problem of computing the type $(m, n)$ Padé approximant from data given as a sequence of $m+n+1$ Taylor coefficients. This is in contrast to related algorithms in the signal processing and model reduction literatures, mentioned in the discussion at the end, which attempt to extract low-rank signals from potentially much longer sequences of noisy data.
2. Basic Algorithm: Padé Approximation in Exact Arithmetic. Let $f$ be a function analytic in a neighborhood of $z=0$ with Taylor series

$$
f(z)=c_{0}+c_{1} z+c_{2} z^{2}+\cdots
$$

alternatively, $f$ could be a formal power series. Given $n \geq 0$, let $P_{n}$ denote the set of polynomials of degree at most $n$, and, given $m \geq 0$ and $n \geq 0$, let $R_{m n}$ be the set of rational functions of type $(m, n)$, that is, functions that can be written as quotients $p / q$ with $p \in P_{m}$ and $q \in P_{n}$. If $\mu \leq m$ and $\nu \leq n$ denote the smallest integers such that $r \in R_{\mu \nu}$, then we say that $r$ is of exact type $(\mu, \nu)$ and has defect $\delta=\min \{m-\mu, n-\nu\} \geq 0$ with respect to $R_{m n}$. (In the special case $r=0$ we define $\mu=-\infty, \nu=0$.) The type $(m, n)$ Padé approximant to $f$ is the function $r_{m n} \in R_{m n}$ whose Taylor series at $z=0$ matches that of $f$ as far as possible:

$$
\begin{equation*}
r_{m n}(z)-f(z)=O\left(z^{\text {maximum }}\right) \tag{2.1}
\end{equation*}
$$

It is known that $r_{m n}$ exists and is unique and has the following characterization: if $r \in R_{m n}$ has defect $\delta$ with respect to $R_{m n}$, then $r=r_{m n}$ if and only if

$$
\begin{equation*}
r(z)-f(z)=O\left(z^{m+n+1-\delta}\right) \tag{2.2}
\end{equation*}
$$

Here the "big O" notation has its usual meaning: the first nonzero term in the Taylor series of $r(z)-f(z)$ is $C z^{k}$ for some $k \geq m+n+1-\delta$. If $\delta=0$, then $p$ and $q$ are uniquely determined up to a scale factor, whereas if $\delta>0$, then $p$ and $q$ are only determined up to a scale factor if we make the additional assumption that they have degrees $\mu$ and $\nu$ or, equivalently, that they have no common roots. For extensive information about Padé approximation, see the book by Baker and Graves-Morris [1]. However, that monograph uses an alternative definition according to which a Padé approximant only exists if $f$ can be matched to order $z^{m+n+1}$ or further and, in fact, the present paper is mathematically closer to the landmark review of Gragg [16], which uses the definition (2.1). For historical developments, see [5].

Equation (2.1) is nonlinear, but multiplying through by the denominator gives the linear condition

$$
\begin{equation*}
p(z)=f(z) q(z)+O\left(z^{\text {maximum }}\right) \tag{2.3}
\end{equation*}
$$

By itself, this condition is vacuous, since matching to all orders could be achieved by taking $p$ and $q$ identically zero. The condition becomes meaningful when $q$ is required to satisfy $q \not \equiv 0$. With this requirement, it is known that the matching condition can always be satisfied through degree $m+n$ or higher,

$$
\begin{equation*}
p(z)=f(z) q(z)+O\left(z^{m+n+1}\right) \tag{2.4}
\end{equation*}
$$

as we shall confirm after (2.8).
Finding $p$ and $q$ to satisfy (2.4) is a linear algebra problem. Suppose $\mathbf{a}$ and $\mathbf{b}$ are $(m+1)$ - and $(n+1)$-vectors of coefficients of polynomials $p \in P_{m}$ and $q \in P_{n}$, respectively:

$$
\begin{gathered}
\mathbf{a}=\left(\begin{array}{c}
a_{0} \\
a_{1} \\
\vdots \\
a_{m}
\end{array}\right), \quad \mathbf{b}=\left(\begin{array}{c}
b_{0} \\
b_{1} \\
\vdots \\
b_{n}
\end{array}\right), \\
p(z)=\sum_{k=0}^{m} a_{k} z^{k}, \quad q(z)=\sum_{k=0}^{n} b_{k} z^{k} .
\end{gathered}
$$

Then (2.4) can be written in matrix form, and it is here that our treatment of Padé approximation begins to depart from the usual. Usually one normalizes by a coefficient condition such as $b_{0}=1$, whereupon what remains in (2.4) is a system of linear equations that may be highly ill-conditioned or singular. Instead, following [15], we normalize by the condition

$$
\begin{equation*}
\|\mathbf{b}\|=1 \tag{2.5}
\end{equation*}
$$

where $\|\cdot\|$ is the vector 2 -norm. This normalization helps to eliminate problems of singularity and ill-conditioning.

If $m \geq n$, (2.4) takes the Toeplitz form

$$
\left(\begin{array}{c}
a_{0}  \tag{2.6}\\
a_{1} \\
\vdots \\
a_{n} \\
\vdots \\
a_{m} \\
\hline a_{m+1} \\
\vdots \\
a_{m+n}
\end{array}\right)=\left(\begin{array}{cccc}
c_{0} & & & \\
c_{1} & c_{0} & & \\
\vdots & \vdots & \ddots & \\
c_{n} & c_{n-1} & \ldots & c_{0} \\
\vdots & \vdots & & \vdots \\
c_{m} & c_{m-1} & \ldots & c_{m-n} \\
\hline c_{m+1} & c_{m} & \ldots & c_{m+1-n} \\
\vdots & \vdots & \ddots & \vdots \\
c_{m+n} & c_{m+n-1} & \ldots & c_{m}
\end{array}\right)\left(\begin{array}{c}
b_{0} \\
b_{1} \\
\vdots \\
b_{n}
\end{array}\right)
$$

coupled with the condition

$$
\begin{equation*}
a_{m+1}=\cdots=a_{m+n}=0 \tag{2.7}
\end{equation*}
$$

In other words, $\mathbf{b}$ must be a (right) null vector of the $n \times(n+1)$ matrix displayed below the horizontal line. The coefficients $a_{0}, \ldots, a_{m}$ of $p$ are then obtained by multiplying out the matrix-vector product above the line.

If $m \leq n$, the essence of the matter remains the same, though it is worth displaying the matrix anew to make its form clear:

$$
\left(\begin{array}{c}
a_{0} \\
a_{1} \\
\vdots \\
a_{m} \\
a_{m+1} \\
\vdots \\
a_{n} \\
\vdots \\
a_{m+n}
\end{array}\right)=\left(\begin{array}{ccccccc}
c_{0} & & & & & & \\
c_{1} & c_{0} & & & & & \\
\vdots & \vdots & \ddots & & & & \\
c_{m} & c_{m-1} & \cdots & c_{0} & & & \\
\hline c_{m+1} & c_{m} & \cdots & c_{1} & c_{0} & & \\
\vdots & \vdots & & & \ddots & \ddots & \\
c_{n} & c_{n-1} & & \ddots & & c_{1} & c_{0} \\
\vdots & \vdots & & & & & \vdots \\
c_{m+n} & c_{m+n-1} & & \cdots & & & c_{m}
\end{array}\right)\left(\begin{array}{c}
b_{0} \\
b_{1} \\
\vdots \\
b_{n}
\end{array}\right)
$$

Again we require (2.7) to hold, meaning that $\mathbf{b}$ should again be a null vector of the matrix below the line, with a again obtained by multiplying above the line. Since an $n \times(n+1)$ matrix always has a nontrivial null vector, we have confirmed (2.4), as promised. There is no assurance that $\mathbf{a}$ and $\mathbf{b}$ are unique, and we shall spell out the details of nonunique solutions in the next section.

Let us now focus on the portion of the linear algebra below the line, involving the $n \times(n+1)$ matrix. This equation takes the form

$$
\begin{equation*}
\mathbf{0}=\widetilde{C} \mathbf{b} \tag{2.9}
\end{equation*}
$$

where $\widetilde{C}$ is the $n \times(n+1)$ Toeplitz matrix

$$
\widetilde{C}=\left(\begin{array}{cccc}
c_{m+1} & c_{m} & \ldots & c_{m+1-n}  \tag{2.10}\\
\vdots & \vdots & \ddots & \vdots \\
c_{m+n} & c_{m+n-1} & \ldots & c_{m}
\end{array}\right)
$$

For simplicity in cases with $n>m+1$, we have adopted the convention that $c_{k}=0$ for $\underset{\sim}{k}<0$. Let $C$ denote the square $n \times n$ matrix obtained by deleting the first column of $\widetilde{C}$ :

$$
C=\left(\begin{array}{ccc}
c_{m} & \ldots & c_{m+1-n}  \tag{2.11}\\
\vdots & \ddots & \vdots \\
c_{m+n-1} & \cdots & c_{m}
\end{array}\right)
$$

Many treatments of Padé approximation work with this matrix, solving a linear system of equations if its determinant is nonzero and bypassing it if its determinant is zero. Often $C$ is flipped so that its structure is Hankel rather than Toeplitz.

We shall make use of the SVD of $\widetilde{C}$, a factorization

$$
\begin{equation*}
\widetilde{C}=U \Sigma V^{*}, \tag{2.12}
\end{equation*}
$$

where $U$ is $n \times n$ and unitary, $V$ is $(n+1) \times(n+1)$ and unitary, and $\Sigma$ is an $n \times(n+1)$ real diagonal matrix with diagonal entries $\sigma_{1} \geq \sigma_{2} \geq \cdots \geq \sigma_{n} \geq 0$.

Suppose $\sigma_{n}>0$. Then $\widetilde{C}$ has rank $n$ and the final column of $V$ provides a unique nonzero null vector $\mathbf{b}$ of $\widetilde{C}$ up to a scale factor. This null vector defines the coefficients of $q$, with two subcases of special interest. If the square submatrix $C$ is singular, then
necessarily $b_{0}=0$. From (2.6) or (2.8) we see that this also implies $a_{0}=0$. Thus $p$ and $q$ share a common factor $z$, or possibly $z^{\lambda}$ for some $\lambda>1$, and this factor can be divided out at the end. If $C$ is nonsingular, then $b_{0}$ must be nonzero, and the defect is $\delta=0$.

On the other hand, suppose $\sigma_{n}=0$. In this case $\widetilde{C}$ has rank $\rho<n$, with zero singular values $\sigma_{\rho+1}=\cdots=\sigma_{n}=0$. Then $C$ must have rank $\rho$ or $\rho-1$, so it is singular. In particular, the submatrix of $\widetilde{C}$ consisting of its last $\rho+1$ columns must be rank-deficient, implying that $\widetilde{C}$ has a nonzero null vector that is zero in its first $n-\rho$ positions. The defect of the corresponding rational function is at least $n-\rho$, and we can reduce $n$ to $\rho$ and $m$ to $m-(n-\rho)$ (or to 0 in the special case $r=0$ ) and start again.

These observations suggest the following SVD-based algorithm for the calculation of the unique type $(m, n)$ Padé approximant to a function $f$ defined by its Taylor series.

Algorithm 1. Pure Padé approximation in exact arithmetic.
Input: $m \geq 0, n \geq 0$, and Taylor coefficients $c_{0}, \ldots, c_{m+n}$ of a function $f$.
Output: Polynomials $p(z)=a_{0}+\cdots+a_{\mu} z^{\mu}$ and $q(z)=b_{0}+\cdots+b_{\nu} z^{\nu}, b_{0}=1$, of the minimal degree type $(m, n)$ Padé approximation of $f$.

1. If $c_{0}=\cdots=c_{m}=0$, set $p=0$ and $q=1$ and stop.
2. If $n=0$, set $p(z)=c_{0}+\cdots+c_{m} z^{m}$ and $q=1$ and go to step 8 .
3. Compute the $S V D(2.12)$ of the $n \times(n+1)$ matrix $\widetilde{C}$. Let $\rho \leq n$ be the number of nonzero singular values.
4. If $\rho<n$, reduce $n$ to $\rho$ and $m$ to $m-(n-\rho)$ and return to step 2.
5. Get $q$ from the null right singular vector $\mathbf{b}$ of $\widetilde{C}$ and then $p$ from the upper part of (2.6) or (2.8).
6. If $b_{0}=\cdots=b_{\lambda-1}=0$ for some $\lambda \geq 1$, which implies also $a_{0}=\cdots=a_{\lambda-1}=0$, cancel the common factor of $z^{\lambda}$ in $p$ and $q$.
7. Divide $p$ and $q$ by $b_{0}$ to obtain a representation with $b_{0}=1$.
8. Remove trailing zero coefficients, if any, from $p(z)$ or $q(z)$.

This algorithm produces the unique Padé approximant $r_{m n}$ in a minimal-degree representation of type $(\mu, \nu)$ with $b_{0}=1$. We state this result as a theorem, whose proof is part of a fuller discussion in the next section.

Theorem 2.1. Algorithm 1 (in exact arithmetic) converges in a finite number of steps to the unique normalized minimal-degree representation of the type ( $m, n$ ) Padé approximant $r_{m n} \in R_{m n}$ to $f$ : polynomials $p$ and $q \neq 0$ of exact degrees $\mu$ and $\nu$ with no common factors and $b_{0}=1$. The number of times that step 3 is executed is no greater than $2+\log _{2}(\delta+1)$.
3. Square Blocks and Proof of Theorem 2.I. An understanding of how Algorithm 1 works requires a discussion of block structure in the Padé table, by which we mean the array of Padé approximants $r_{m n}$ for various $m, n \geq 0$ associated with a given function $f[1,11,16,23]$ :

$$
\left(\begin{array}{cccc}
r_{00} & r_{10} & r_{20} & \ldots \\
r_{01} & r_{11} & r_{21} & \ldots \\
r_{02} & r_{12} & r_{22} & \ldots \\
\vdots & \vdots & \vdots & \ddots
\end{array}\right)
$$

Suppose $r$ is a nonzero rational function of exact type $(\mu, \nu)$ that is the type $(m, n)$ Padé approximant to $f$ for at least one pair $(m, n)$. Then it is known that there is an integer $k \geq 0$ such that $r$ is the type $(m, n)$ Padé approximant to $f$ if and only if
$\mu \leq m \leq \mu+k$ and $\nu \leq n \leq \nu+k$ [16, sec. 3], [30]. In other words, $r$ is the Padé approximant to $f$ precisely in the following $(k+1) \times(k+1)$ square block of the Padé table:

$$
\left(\begin{array}{ccc}
r_{\mu \nu} & \ldots & r_{\mu+k, \nu}  \tag{3.1}\\
\vdots & & \vdots \\
r_{\mu, \nu+k} & \ldots & r_{\mu+k, \nu+k}
\end{array}\right)
$$

We have already defined the defect $\delta$ of $r$ as a function of type $(m, n)$ with $m \geq \mu$ and $n \geq \nu$. Showing the case $k=5$ for illustration, $\delta$ takes the following values within a block:

$$
\operatorname{defect} \delta: \quad\left(\begin{array}{cccccc}
0 & 0 & 0 & 0 & 0 & 0  \tag{3.2}\\
0 & 1 & 1 & 1 & 1 & 1 \\
0 & 1 & 2 & 2 & 2 & 2 \\
0 & 1 & 2 & 3 & 3 & 3 \\
0 & 1 & 2 & 3 & 4 & 4 \\
0 & 1 & 2 & 3 & 4 & 5
\end{array}\right)
$$

Similarly we define the deficiency $\lambda$ of $r$ as the distance below the cross-diagonal in the square block, $\lambda=\max \{0,(m-\mu)+(n-\nu)-k\}$, with the following pattern:

$$
\text { deficiency } \lambda: \quad\left(\begin{array}{cccccc}
0 & 0 & 0 & 0 & 0 & 0  \tag{3.3}\\
0 & 0 & 0 & 0 & 0 & 1 \\
0 & 0 & 0 & 0 & 1 & 2 \\
0 & 0 & 0 & 1 & 2 & 3 \\
0 & 0 & 1 & 2 & 3 & 4 \\
0 & 1 & 2 & 3 & 4 & 5
\end{array}\right)
$$

Finally, the rank deficiency $\chi$ of $r$ is defined by the formula

$$
\begin{equation*}
\chi=\delta-\lambda \tag{3.4}
\end{equation*}
$$

and corresponds to the distance from the border of the square block:

$$
\operatorname{rank} \text { deficiency } \chi: \quad\left(\begin{array}{cccccc}
0 & 0 & 0 & 0 & 0 & 0  \tag{3.5}\\
0 & 1 & 1 & 1 & 1 & 0 \\
0 & 1 & 2 & 2 & 1 & 0 \\
0 & 1 & 2 & 2 & 1 & 0 \\
0 & 1 & 1 & 1 & 1 & 0 \\
0 & 0 & 0 & 0 & 0 & 0
\end{array}\right) .
$$

The following characterization of Pade approximants, which explains why $\chi$ is called the rank deficiency, is equivalent to results derived in section 3 of [16].

Theorem 3.1. Let $f$ and $m, n \geq 0$ be given, let $\mu, \nu, k, \delta, \lambda$, $\chi$ be the parameters defined above for the type $(m, n)$ Padé approximant $r_{m n}$ to $f$, and let $\hat{p}$ and $\hat{q} \neq 0$ be polynomials of exact degrees $\mu$ and $\nu$ with $r_{m n}=p / q$. Then the matrix $\widetilde{C}$ of (2.10) has rank $n-\chi$, and two polynomials $p \in P_{m}$ and $q \in P_{n}, q \neq 0$, satisfy (2.4) if and only if

$$
\begin{equation*}
p(z)=z^{\lambda} w(z) \hat{p}(z), \quad q(z)=z^{\lambda} w(z) \hat{q}(z) \tag{3.6}
\end{equation*}
$$

for some $w \in P_{\delta-\lambda}$.

It is worth summarizing in words some of the implications of (3.6). The minimaldegree polynomials $\hat{p}$ and $\hat{q} \neq 0$ of a Padé approximant $r$ are unique apart from a scalar multiple. The defect $\delta$ determines how many additional degrees of freedom there are in the polynomials $p$ and $q$ of degrees $m \geq \mu$ and $n \geq \nu$ that represent $r$ if the condition (2.3) is used. In the upper-left half of a square block, this is the same as the number of degrees of freedom in the polynomials $p$ and $q$ representing $r$ by (2.4). In the lower-right half of the block, however, representations (2.4) are more constrained than those of (2.3), with the number of degrees of freedom being equal to $\chi=\delta-\lambda$.

We are now equipped to prove Theorem 2.1. The reader may find it helpful to review the example summarized in Figure 3 in parallel with reading this proof.

Proof of Theorem 2.1. Each time step 3 of Algorithm 1 is executed, either $\rho=n$ or $\rho<n$. In the first instance the algorithm executes steps $5-8$ and then stops, whereas in the second it returns to step 2 with smaller values of $m$ and $n$. Thus it terminates in at most $\delta$ steps. We must show that it terminates with $p$ and $q \neq 0$ of exact degrees $\mu$ and $\nu$, that is, at the upper-left corner of the square block.

One way the algorithm can terminate is with $n=0$ at step 2 , followed by step 8 . In this case $q=1$, with exact degree $\nu=0$, and $p$ must have exact degree $\mu$ since trailing zeros are removed at step 8.

The other way the algorithm can terminate is with $\rho=n$ at step 5 followed by steps $6-8$. In this case $\widetilde{C}$ is of full rank $n$, so $\chi=0$ and $(m, n)$ lies along the edge of the square block. If it lies in the left column or top row, then step 8 brings it to the upper-left corner as required. If it lies in the right column or bottom row, then the cancellation of the common factor $z^{\lambda}$ in step 6 moves it to the top row or left column, respectively, and step 8 brings it to the upper-left corner again.

This completes the proof except for the final assertion that the number of executions of step 3 is at most $2+\log _{2}(\delta+1)$. If $(m, n)$ lies in the top row or left column, then the number of such executions is 0 or 1 , and, since $\delta=0$, this is less than $2+\log _{2}(\delta+1)$. If $(m, n)$ lies elsewhere in the upper-left half of the square block, then the number of steps is at most 2 , which again is less than $2+\log _{2}(\delta+1)$. Higher numbers become possible if $(m, n)$ begins in the lower-right half of the square block, below the cross-diagonal and away from the boundary (hence always with $\delta>0$ ). Here successive steps in the worst case might make $(m, n)$ hop from one position to the next with $\chi=1,2,4,8, \ldots$ until eventually the upper-left half of the block is reached; since $\chi$ can attain values no greater than $\delta / 2$, the maximum number of such steps is $\log _{2} \delta$. At this point two further steps must complete the process, giving at most $2+\log _{2} \delta$ steps in total.
4. Modified Algorithm for Noisy Data. It is well known that the singular values of a matrix are well-conditioned functions of its entries. Specifically, if $\widetilde{C}$ is perturbed by a matrix of 2 -norm at most $\varepsilon$, then each singular value is perturbed by at most $\varepsilon$. Thus Algorithm 1 immediately suggests a variant for the case in which the approximation problem is subject to noise, whether intrinsic to the data or introduced by rounding errors. The modified algorithm is this: carry out the operations as before, but treat a singular value as zero if it is less than tol $\cdot\|\mathbf{c}\|_{2}$, where $\mathbf{c}=\left(c_{0}, \ldots, c_{m+n}\right)$ and $t o l$ is a relative tolerance. The tolerance is also applied in the detection of zero coefficients in steps 1,6 , and 8 . For most purposes involving problems perturbed just by rounding errors, we take tol $=10^{-14}$.

In addition, there is another matter to consider for practical computation. Algorithm 1 is implicitly tied to a scaling associated with the unit disk through the use
of the 2-norm to measure the coefficient vector $\mathbf{b}$. In exact arithmetic this does not matter, but in practical applications it means the algorithm will be most effective when the coefficients $\left\{c_{j}\right\}$ of $f$ are of roughly comparable sizes, neither decreasing nor increasing at a rapid geometric rate. Thus it may be beneficial first to pick a parameter $\gamma>0$ and compute the Padé approximant $\hat{r}_{m n}(z)=r_{m n}(z / \gamma)$ to $\hat{f}(z)=f(z / \gamma)$. An algorithm for automatic determination of such parameters can be found in [9].

Algorithm 2. Robust Padé approximation for noisy data or floating POINT ARITHMETIC.

Input: $m \geq 0, n \geq 0$, Taylor coefficients $c_{0}, \ldots, c_{m+n}$ of a function $f$, and relative tolerance tol $\geq 0$.

Output: Polynomials $p(z)=a_{0}+\cdots+a_{\mu} z^{\mu}$ and $q(z)=b_{0}+\cdots+b_{\nu} z^{\nu}, b_{0}=1$, of the minimal-degree type $(m, n)$ Padé approximation of a function close to $f$.

1. Rescale $f(z)$ to $f(z / \gamma)$ for some $\gamma>0$ if desired to get a function whose Taylor coefficients $c_{0}, \ldots, c_{m+n}$ do not vary too widely.
2. Define $\tau=$ tol $\cdot\|\mathbf{c}\|_{2}$. If $\left|c_{0}\right|=\cdots=\left|c_{m}\right| \leq \tau$, set $p=0$ and $q=1$ and stop.
3. If $n=0$, set $p(z)=c_{0}+\cdots+c_{m} z^{m}$ and $q=1$ and go to step 7 .
4. Compute the $S V D(2.12)$ of the $n \times(n+1)$ matrix $\widetilde{C}$. Let $\rho \leq n$ be the number of singular values of $\widetilde{C}$ that are greater than $\tau$.
5. If $\rho<n$, reduce $n$ to $\rho$ and $m$ to $m-(n-\rho)$ and return to step 3 .
6. Obtain $q$ from the null right singular vector $\mathbf{b}$ of $\widetilde{C}$ and then $p$ from the upper part of (2.6) or (2.8).
7. If $\left|b_{0}\right|, \ldots,\left|b_{\lambda-1}\right| \leq$ tol for some $\lambda \geq 1$, zero the first $\lambda$ coefficients of $p$ and $q$ and cancel the common factor $z^{\lambda}$.
8. If $\left|b_{n+1-\lambda}\right|, \ldots,\left|b_{n}\right| \leq$ tol for some $\lambda \geq 1$, remove the last $\lambda$ coefficients of $q$. If $\left|a_{m+1-\lambda}\right|, \ldots,\left|a_{m}\right| \leq \tau$ for some $\lambda \geq 1$, remove the last $\lambda$ coefficients of $p$.
9. Divide $p$ and $q$ by $b_{0}$ to obtain a representation with $b_{0}=1$.
10. Undo the scaling of step 1 by redefining $\gamma^{j} a_{j}$ as $a_{j}$ and $\gamma^{j} b_{j}$ as $b_{j}$ for each $j$.

A MATLAB code padeapprox implementing Algorithm 2 is shown in Figure 1 and is freely available as part of Chebfun [34]. Three lines of this code, marked by comments beginning and ending with the word "reweighting," go beyond Algorithm 2. These lines compute the final null vector by QR factorization of a column-reweighted matrix rather than by the SVD, an alternative that does a better job of taking advantage of sparsity in Toeplitz matrices. The effect is that the blocks produced in regions of a table corresponding to approximation accuracies close to tol more often come out exactly square.
5. Examples of Computed Padé Tables and Noise Removal. Figure 2 shows Padé tables computed by padeapprox with tol $=10^{-14}$ for the functions $\exp (z)$, $\cos (z),\left(z^{5}-1\right) /\left(z^{5}+1\right)$, and $\log \left(5+z^{5}\right)$. Each figure is based on the computation of 441 distinct Padé approximants and took about 1 second to produce in MATLAB on an AMD Phenom II X3 720 processor running at 2.8 GHz . As described in the caption, the images clearly show the block structures for the various functions. Since the blocks arise from 441 independent computations for various $(m, n)$, this is a visual confirmation of the reliability of Algorithm 2.

A departure from the theoretically expected block structure is apparent in the lower-right part of the tables for the first two functions. Here, $m$ and $n$ are larger than needed for resolution to machine precision, and the algorithm automatically reduces them by equal amounts, moving up and left along a diagonal to smaller values of $m$ and $n$ (step 5 of Algorithm 2). Thus these diagonal stripes are indications of the robustness of Algorithm 2 in the presence of rounding errors.

end

Fig. I MATLAB code padeapprox for robust Padé approximation, based on Algorithm 2. The input function can be either a vector of Taylor coefficients or a function handle. The code is freely available as part of Chebfun [34].
$\exp (z)$


$$
\left(z^{5}-1\right) /\left(z^{5}+1\right)
$$


$\cos (z)$

$\log \left(5+z^{5}\right)$


Fig. 2 Padé tables computed numerically by Algorithm 2, with $m$ and $n$ on the horizontal and vertical axes, respectively. Each square $(m, n)$ is marked by a color determined by the exact type $(\mu, \nu)$ of the corresponding approximant, so that each square block appears in a single color. For $\exp (z)$, all the entries lie in $1 \times 1$ blocks until the function is resolved to machine precision, after which the numerator and denominator degrees are systematically reduced as far as possible, causing diagonal stripes (step 5 of Algorithm 2). For the even function $\cos (z), 2 \times 2$ square blocks appear. For $\left(z^{5}-1\right) /\left(z^{5}+1\right)$, we get an infinite square block since the function is resolved exactly for $m, n \geq 5$. For $\log \left(5+z^{5}\right)$, there are $5 \times 5$ blocks all the way down.

Figure 3 shows another Padé table as in Figure 2, but now with some additional numbers displayed. The purpose of this figure is to illustrate how the algorithm works, as explained in the caption.

Figure 4 shows three more Padé tables, all corresponding to approximations of the same noisy Taylor series but with different levels of the parameter tol. This example is mentioned at the beginning of [13]. When tol is above the noise level, padeapprox detects the underlying rational function reliably.


Fig. 3 The numerically computed Padé table for $f(z)=1+z+z^{8}+z^{20}+z^{30}$, with numbers showing the path taken by Algorithm 2 for the particular case $(m, n)=(14,9)$. Starting at this position of the table (label 1), the algorithm finds that the $9 \times 10$ Toeplitz matrix (2.10) has rank deficiency $\chi=2$, so it moves to position (12,7) (label 2, step 5 of Algorithm 2). This $7 \times 8$ matrix has rank deficiency 4, causing a move to (8,3) (label 3, step 5). At this point three trailing zeros of $\mathbf{b}$ are discarded (step 8 ), bringing us to the final position $(8,0)$ (label 4).


Fig. 4 An explicit example of noise removal by Algorithm 2. Here the Taylor series is defined by coefficients $c_{j}=1+10^{-6} s_{j}$, where $\left\{s_{j}\right\}$ are independent samples from the standard normal distribution. This corresponds to the function $1 /(1-z)$ plus noise on a scale of $10^{-6}$. If padeapprox is run with tol $=10^{-8}$ or a lower value such as the default $10^{-14}$, the noise has the effect of making the Padé approximants distinct, and we see $1 \times 1$ blocks. With a tolerance above the noise level, tol $\geq 10^{-5}$, the code detects that this is essentially a rational function of type $(0,1)$.
6. Examples of the Elimination of Froissart Doublets. As we have mentioned, rounding errors or other perturbations commonly introduce Froissart doublets in computed Padé approximations that neither reflect genuine information about $f$ nor contribute to the quality of the approximation. We now give a few examples to illustrate
how Algorithm 2 removes such effects by reducing the degrees $m$ and $n$. Following a pattern employed in [15], each example is presented in the form of a two-part figure showing results from Algorithm 2 in its nonrobust form with tol $=0$ on the left, and in its robust form with the default value tol $=10^{-14}$ on the right. The unit circle is marked by a dotted line. The pair $(m, n)$ is listed on the upper right. The lower left of each plot lists the exact type $(\mu, \nu)$ returned by padeapprox and the elapsed time for computing this approximation on the AMD Phenom II X3 processor specified earlier.

Each plot also lists a number Err, equal to the maximum of $|f(z)-r(z)|$ over the discrete grid of 1976 points in the disk $|z| \leq 0.5$ whose real and imaginary coordinates are odd multiples of 0.01 . This gives a rough indication of the success of the algorithm in computing an effective approximation to $f$. The radius is chosen less than 1 to stay away from Froissart doublets clustered near the unit circle [13], but there is no special significance to the particular choice 0.5.

Finally, each plot also shows the poles of $r$, marked by dots, following a scheme suggested by Grady Wright of Boise State University. The absolute value of the residue at each pole of $r$, evaluated by a finite difference, is indicated by the following color code:

$$
\mid \text { residue } \left\lvert\, \in \begin{cases}{\left[10^{-3}, \infty\right),} & \text { blue } \\ {\left[10^{-12}, 10^{-3}\right),} & \text { green } \\ {\left[0,10^{-12}\right),} & \text { red }\end{cases}\right.
$$

Thus a blue pole has a good chance of being genuine and useful for approximation, whereas red poles are likely to be artifacts introduced by rounding errors. As it happens, no green poles appear in any of the figures of this paper.

Figures $5-7$ show Padé approximations to the function $f(z)=\tan \left(z^{4}\right)$, which has poles outside the unit circle lying along eight rays emanating from the origin. For larger $m$ and $n$, the removal of Froissart doublets by Algorithm 2 is striking.

Figure 8 shows results for a function with a branch point, $f(z)=\log (1.2-z)$. According to a theory of Stahl, most of the poles of Padé approximants to such functions line up along certain branch curves determined by a capacity-minimization condition in the $z^{-1}$-plane [27]. (There is no assurance that all the poles must lie near the branch curves, merely a fraction of them approaching 1.) In this case the curve in question is the interval $[1.2, \infty)$, and the figure shows blue poles lining up as expected.

Figure 9 shows results for a function with an essential singularity, $f(z)=\exp ((z+$ $1.5)^{-2}$ ).
7. III-Posedness and Stability. It is easy to show that Padé approximation problems are sometimes ill-posed. For example, here is a function and its type $(1,1)$ Padé approximant:

$$
f(z)=1+z^{2}, \quad r_{11}(z)=1
$$

The approximant matches $f$ only through power $z^{1}$, but that is enough, according to (2.2), since the defect is 1 . An arbitrarily small perturbation of $f$, however, changes $r_{11}$ completely:

$$
f(z)=1+\varepsilon z+z^{2}, \quad r_{11}(z)=\frac{1-\left(1-\varepsilon^{2}\right) z / \varepsilon}{1-z / \varepsilon} \quad(\varepsilon \neq 0)
$$

This perturbed function, which now matches $f$ through order $z^{2}$, has a Froissart doublet with a pole at $z=\varepsilon$ of residue $-\varepsilon^{3}$ and a zero at $\varepsilon /\left(1-\varepsilon^{2}\right)$.


Fig. 5 Type (20,20) approximation of $\tan \left(z^{4}\right)$, a function with poles along 8 rays emanating from the origin. Both the nonrobust and robust algorithms place poles near the innermost 16 poles of $f$, the inner 8 matching the poles of $f$ to six digits and the next 8 to two digits. There is little difference between the two algorithms except that the robust one reduces the type from $(20,20)$ to $(20,16)$, removing four poles with absolute value about $2 \times 10^{5}$.


Fig. 6 Type $(100,100)$ approximation of $\tan \left(z^{4}\right)$. Now each algorithm places four poles along each of the 8 rays, which match the poles of $f$ to approximately 14, 6, 3, and 2 digits (from inside out). These very accurate agreements of computed poles show how powerful Padé approximation can be for extracting information about a Taylor series beyond its circle of convergence. The nonrobust algorithm also produces 64 Froissart doublets near the unit circle, as well as four additional doublets of absolute value about $3 \times 10^{3}$, all of which are removed by the robust algorithm.


Fig. 7 The type $(20,100)$ approximation of $\tan \left(z^{4}\right)$ shows the kind of effects that may arise with $n>m$. Here both the robust and nonrobust approximations place some poles along a circle outside the unit disk.



Fig. 8 Type $(20,20)$ approximation of $\log (1.2-z)$. This function has a branch cut $[1.2, \infty)$ that attracts poles in keeping with the theory of Stahl [27]: 13 in the nonrobust approximation, 10 in the robust one. The nonrobust algorithm scatters six additional spurious poles near the left half of the unit circle as well as two other negative real poles off the scale of the plot at about -6 and -8 .


Fig. 9 Type $(60,60)$ approximation of $f(z)=\exp \left((z+1.5)^{-2}\right)$, a function with an essential singularity at $z=-1.5$.

The example points the way to the general case: a Padé approximation problem is ill-posed if and only if $r_{m n}$ has defect $\delta>0$. The reason is that an arbitrarily small perturbation could fracture the block, forcing $r_{m n}$ to match $f$ to a higher order than before. For details see [30, 33].

Our point of view as numerical analysts is that it is not reasonable to ask an algorithm to find nearly the exact solution of an ill-posed problem. A more reasonable expectation is that an algorithm should be stable, finding nearly the exact solution of a slightly perturbed problem [32]. We hope to discuss stability of Algorithm 2 in a future publication. Note that a stable algorithm need not always produce square blocks, since different choices of $m$ and $n$ will correspond to different perturbations of the data.

The idea of regularization, which applies across a wide range of ill-posed problems, involves a balance between the accuracy of a solution and its other properties such as smoothness [18]. In eliminating Froissart doublets, we are slightly reducing the accuracy of our solution to a certain matrix problem, but the benefit is that the resulting function $r_{m n}$ will be pole-free in most regions where $f$ itself is pole-free. Consequently, in an application, it may be a better solution to the scientific problem that lies behind the ill-posed linear algebra problem.
8. A Modified Padé Approximant with Pointwise Convergence. The main purpose of this paper has been to propose a regularized algorithm for computing Padé approximations in floating point arithmetic or for problems with noise. However, one of the interesting features of the SVD-based approach is that it may also be applicable in the theoretical setting of exact arithmetic for problems with no noise, because this situation too is afflicted by seemingly spurious pole-zero pairs.

How can an exactly correct pole-zero pair be spurious? The phenomenon in question has been recognized for a century, going back at least to Perron in 1913 [25], and a formal definition of spurious poles is given in [28]. In the simplest case, suppose $f$ is a meromorphic function in the complex plane $\mathbb{C}$, and consider the behavior of $f-r_{n n}$ as $n \rightarrow \infty$. A natural hope might be that on any compact set $E$ disjoint from the poles of $f$, the supremum norm of $f-r_{m n}$ over $E,\left\|f-r_{m n}\right\|_{E}$, should converge to zero. However, this does not happen in general. The Padé approximants $r_{m n}$ can have poles in arbitrary locations in $\mathbb{C}$, and as $n \rightarrow \infty$, although the residues of these poles will decrease, they need never disappear entirely. In fact it may even happen that the type $(n, n)$ Padé approximants to a fixed entire function $f$ have so many spurious poles that the sequence of approximants is unbounded at every nonzero point in the complex plane [36].

In the face of this fact about Padé approximation, theorems have been developed that assert a weaker kind of convergence: convergence in measure or convergence in capacity. For meromorphic functions there is the Nuttall-Pommerenke theorem [1, 21, 26], and for functions with branch points there is a beautiful generalization by Stahl mentioned in the last section of [27]. For the simpler problem of approximations of type ( $m, n$ ) with $m \rightarrow \infty$ while $n$ remains fixed, there is the de Montessus de Ballore theorem [1, 20], and even in this case one must be careful to exclude spurious poles.

Our SVD-based algorithm suggests another possibility. Can one define a modified Padé approximant $\tilde{r}_{m n}$ that is guaranteed to converge, with no exceptional sets? For each $n$, one could define $\tilde{r}_{n n}$ to be the rational function computed in exact arithmetic by an algorithm like Algorithm 2 with tolerance $\mathrm{tol}=\mathrm{tol}_{n}$, where $\mathrm{tol} l_{n}$ is a function of $n$ that decreases to 0 as $n \rightarrow \infty$. Analogous definitions could be developed for approximations $r_{m n}$ with $m \neq n$. Might such an approximation scheme be pointwise convergent? This is a challenge for theorists.
9. Discussion. Padé approximations arise in many applications, and Algorithm 2 will not be appropriate for all of them. Our framework is numerical approximation theory, in the spirit of [31] and [34]. Applications we have in mind include analytic continuation and extrapolation of sequences and series, both based on data at a single point in the complex plane. For data on a circle, we recommend the leastsquares approach of [15], and for data on an interval, see the ratinterp command in Chebfun [34].

Many methods that have been proposed for extrapolation, such as the Aitken, Shanks, eta, and epsilon methods, are mathematically equivalent to the evaluation of certain Padé approximants but have traditionally been formulated algorithmically by fast recurrence relations related to continued fractions, in the interests of speed. Indeed, the Toeplitz structure of the Padé problem makes possible some remarkably fast algorithms, in theory [4]. This focus on fast recurrences may be the right strategy for some applications, but we suspect that for many problems nowadays it is better to use more robust algorithms, whose $O\left(n^{3}\right)$ complexity will rarely be a problem.

A central feature of our algorithm is that it removes spurious Froissart doublets as a by-product of the use of numerical ranks computed with the SVD. Other
methods for removing doublets have also been proposed, typically based on explicit computation of zeros and poles, sometimes with the use of extended precision arithmetic $[2,3,6,8,12,13,14,29]$. On the other hand, Perotti, Vrinceanu, and Bessis argue that Froissart doublets can actually be beneficial in the sense that perturbations in their distribution from a regular pattern along the unit circle may reveal locations of underlying genuine poles [24].

An example of a different set of applications related to Padé approximation are problems of signal processing or model reduction, where the aim is to extract a loworder model from a sequence of hundreds of possibly noisy data values [2, 7, 17, 22]. In such problems the sampled data are typically assumed to come from a stable system, which would correspond for our problem to poles only outside the unit disk, and lowrank fits to sufficiently long data sequences will inherit this stability property. That is a different setting from the one we have focused on here, which starts from arbitrary Taylor coefficients without an a priori assumption of scaling to the unit disk.

Padé approximants also appear in the study of Krylov subspace iterations for the solution of large matrix problems. Here the issue of degeneracies and square blocks has been of considerable interest, and "look-ahead" variants of some algorithms have been devised to jump around square blocks [10]. A feature of such problems is that since the scales are large, one is necessarily working with partial information obtained from recurrences, so the systematic use of the SVD as in Algorithm 2 is not an option.

The algorithm we have proposed works well, but we do not regard it as the last word on this subject. Other choices could have been made, and, in particular, our method of hopping across a square block to find the upper-left corner is not the only reasonable one. For example, one could reduce just $n$ rather than both $m$ and $n$ in cases of rank deficiency, and this would lead to vertical rather than diagonal stripes in the upper panels of Figure 2. (See [35] for an algorithm to find the upper-left corner of a square block in the case of Carathéodory-Fejér approximation.) The choices we have made seem to be effective, and one must remember that for most functions, few reductions of $(m, n)$ will take place at all, so getting to the upper-left corner of a block by absolutely the shortest route possible is not a high priority. The true priority is robustness, and we believe Algorithm 2 meets this need effectively.

Acknowledgments. We have had fruitful discussions with a number of colleagues that have led to improvements in this article. We are grateful in particular for contributions from Daniel Bessis, Adhemar Bultheel, Jaček Gilewicz, Ricardo Pachón, Ed Saff, Marc Van Barel, Joris Van Deun, and Grady Wright. We are also grateful to Olga Ibryaeva for bringing her new related work [19] to our attention. The Ibryaeva-Adukov approach to Padé approximation is quite close to our own, and we recommend readers with a serious interest in this subject to study that paper too.

## References

[1] G. A. Baker, Jr. and P. R. Graves-Morris, Padé Approximants, 2nd ed., Cambridge University Press, Cambridge, UK, 1996.
[2] D. Belkić and K. Belkić, Padé-Froissart exact signal-noise separation in nuclear magnetic resonance spectroscopy, J. Phys. B, 44 (2011), 125003.
[3] D. Bessis, Padé approximations in noise filtering, J. Comput. Appl. Math., 66 (1996), pp. 8588.
[4] R. P. Brent, F. G. Gustavson, and D. Y. Y. Yun, Fast solution of Toeplitz systems of equations and computation of Padé approximants, J. Algorithms, 1 (1980), pp. 259-295.
[5] C. Brezinski, History of Continued Fractions and Padé Approximants, Springer, Berlin, 1991.
[6] C. Brezinski and M. Redivo-Zaglia, Padé-type rational and barycentric interpolation, Numer. Math., to appear.
[7] P. Feldmann and R. W. Freund, Efficient linear circuit analysis by Padé approximation via the Lanczos process, IEEE Trans. Computer-Aided Design of Integrated Circuits and Systems, 14 (1995), pp. 639-649.
[8] C. FitzGerald, Confirming the accuracy of Padé table approximants, in Padé and Rational Approximation: Theory and Applications, E. B. Saff and R. S. Varga, eds., Academic Press, New York, 1976, pp. 51-60.
[9] B. Fornberg, Numerical differentiation of analytic functions, ACM Trans. Math. Software, 7 (1981), pp. 512-526.
[10] R. W. Freund, M. H. Gutknecht, and N. M. Nachtigal, An implementation of the lookahead Lanczos algorithm for non-Hermitian matrices, SIAM J. Sci. Comput., 14 (1993), pp. 137-158.
[11] G. Frobenius, Über Relationen zwischen den Näherungsbrüchen von Potenzreihen, J. Reine Angew. Math., 90 (1881), pp. 1-17.
[12] M. Froissart, Approximation de Padé: Application à la physique des particules élémentaires, in RCP, Programme 25, v. 9, CNRS, Strasbourg, 1969, pp. 1-13.
[13] J. Gilewicz and M. Pindor, Padé approximants and noise: A case of geometric series, J. Comput. Appl. Math., 87 (1997), pp. 199-214.
[14] J. Gilewicz and M. Pindor, Padé approximants and noise: Rational functions, J. Comput. Appl. Math., 105 (1999), pp. 285-297.
[15] P. Gonnet, R. Pachón, and L. N. Trefethen, Robust rational interpolation and leastsquares, Electron. Trans. Numer. Anal., 38 (2011), pp. 146-167.
[16] W. B. Gragg, The Padé table and its relation to certain algorithms of numerical analysis, SIAM Rev., 14 (1972), pp. 1-62.
[17] S. Gugercin and A. C. Antoulas, Model reduction of large-scale systems by least-squares, Linear Algebra Appl., 415 (2006), pp. 290-321.
[18] P. C. Hansen, Rank-Deficient and Discrete Ill-Posed Problems: Numerical Aspects of Linear Inversion, SIAM, Philadelphia, 1998.
[19] O. L. Ibryaeva and V. M. Adukov, An algorithm for computing a Padé approximant with minimal degree denominator, J. Comput. Appl. Math., 237 (2013), pp. 529-541.
[20] R. de Montessus de Ballore, Sur les fractions continues algébriques, Bull. Soc. Math. France, 30 (1902), pp. 28-36.
[21] J. Nuttall, The convergence of Padé approximants of meromorphic functions, J. Math. Anal. Appl., 31 (1970), pp. 147-153.
[22] E. A. O'Sullivan and C. F. N. Cowan, Modelling room transfer functions using the decimated Padé approximant, IET Signal Process., 2 (2008), pp. 49-58.
[23] H. Padé, Sur la representation approchée d'une fonction par des fractions rationelles, Thesis, Ann. École Nor. (3), 9 (suppl.) (1892), pp. 1-93.
[24] L. Perotti, D. Vrinceanu, and D. Bessis, Beyond the Fourier transform: Signal symmetry breaking in the complex plane, IEEE Signal Process. Lett., 19 (2012), pp. 865-867.
[25] O. Perron, Die Lehre von den Kettenbrüchen, BG Teubner, Stuttgart, 1913.
[26] Ch. Pommerenke, Padé approximants and convergence in capacity, J. Math. Anal. Appl., 41 (1973), pp. 775-780.
[27] H. Stahl, The convergence of Padé approximants to functions with branch points, J. Approx. Theory, 91 (1997), pp. 139-204.
[28] H. Stahl, Spurious poles in Padé approximation, J. Comput. Appl. Math., 99 (1998), pp. 511527.
[29] P. Stoica and T. Söderström, Common factor detection and estimation, Automatica, 33 (1997), pp. 985-989.
[30] L. N. Trefethen, Square blocks and equioscillation in the Padé, Walsh, and CF tables, in Rational Approximation and Interpolation, P. R. Graves-Morris, E. B. Saff, and R. S. Varga, eds., Lecture Notes in Math. 1105, Springer, Berlin, 1984, pp. 170-181
[31] L. N. Trefethen, Approximation Theory and Approximation Practice, SIAM, Philadelphia, 2012.
[32] L. N. Trefethen and D. Bau, III, Numerical Linear Algebra, SIAM, Philadelphia, 1997.
[33] L. N. Trefethen and M. H. Gutknecht, On convergence and degeneracy in rational Padé and Chebyshev approximation, SIAM J. Math. Anal., 16 (1985), pp. 198-210.
[34] L. N. Trefethen et al., Chebfun software package, www.maths.ox.ac.uk/chebfun/.
[35] J. Van Deun and L. N. Trefethen, A robust implementation of the Carathéodory-Fejér method, BIT, 51 (2011), pp. 1039-1050.
[36] H. Wallin, On the convergence theory of Padé approximants, in Linear Operators and Approximation, Internat. Ser. Numer. Math. 20, Birkhäuser, Basel, 1972, pp. 461-469.


[^0]:    *Received by the editors October 28, 2011; accepted for publication (in revised form) April 2, 2012; published electronically February 7, 2013.
    http://www.siam.org/journals/sirev/55-1/85323.html
    $\dagger$ School of Engineering and Computing Sciences, Durham University, South Road, Durham, DH1 3LE, UK (pedro.gonnet@durham.ac.uk). The work of this author was supported by Swiss National Science Foundation Individual Support Fellowship PBEZP2-127959.
    ${ }^{\ddagger}$ School of Mathematics, University of Manchester, Oxford Rd., Manchester M13 9PL, UK (stefan. guettel@manchester.ac.uk). The work of this author was supported by Deutsche Forschungsgemeinschaft Fellowship GU 1244/1-1.
    §Oxford University Mathematical Institute, 24-29 St Giles', Oxford OX1 3LB, UK (trefethen@ maths.ox.ac.uk). The work of this author was supported by the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013)/ERC grant agreement 291068. The views expressed in this article are not those of the ERC or the European Commission, and the European Union is not liable for any use that may be made of the information contained here.

