ORIGINAL PAPER

Multidomain spectral method for the Gauss hypergeometric function

S. Crespo[1] · M. Fasondini[2] · C. Klein[1] ⓘ · N. Stoilov[1] · C. Vallée[1,3]

Abstract

We present a multidomain spectral approach for Fuchsian ordinary differential equations in the particular case of the hypergeometric equation. Our hybrid approach uses Frobenius' method and Moebius transformations in the vicinity of each of the singular points of the hypergeometric equation, which leads to a natural decomposition of the real axis into domains. In each domain, solutions to the hypergeometric equation are constructed via the well-conditioned ultraspherical spectral method. The solutions are matched at the domain boundaries to lead to a solution which is analytic on the whole compactified real line $\mathbb{R} \cup \infty$, except for the singular points and cuts of the Riemann surface on which the solution is defined. The solution is further extended to the whole Riemann sphere by using the same approach for ellipses enclosing the singularities. The hypergeometric equation is solved on the ellipses with the boundary data from the real axis. This solution is continued as a harmonic function to the interior of the disk by solving the Laplace equation in polar coordinates with an optimal complexity Fourier–ultraspherical spectral method. In cases where logarithms appear in the solution, a hybrid approach involving an analytical treatment of the logarithmic terms is applied. We show for several examples that machine precision can be reached for a wide class of parameters, but also discuss almost degenerate cases where this is not possible.

Keywords Hypergeometric function · Singular differential equations · Spectral methods

✉ C. Klein
   Christian.Klein@u-bourgogne.fr

Extended author information available on the last page of the article.

1 Introduction

Gauss' hypergeometric function $F(a, b, c, z)$ is arguably one of the most important classical transcendental functions in applications, see for instance [19] and references therein. Entire chapters are dedicated to it in various handbooks of mathematical functions such as the classical reference [1] and its modern reincarnation [16]. It contains large classes of simpler transcendental functions as degeneracies, for example the Bessel functions. Despite its omnipresence in applications, numerical computation for a wide range of the parameters a, b, c, z is challenging, see [18] for a comprehensive recent review with many references and a comparison of methods, and [3] for additional approaches to singular ODEs. This paper is concerned with the numerical evaluation of the hypergeometric function $F(a, b, c, z)$, treated as a solution to a Fuchsian equation. This class of equations further includes examples such as the Lamé equation, see [1], to which the method can be directly extended. The focus here is on the efficient computation of the hypergeometric function on the compactified real line $\mathbb{R} \cup \{\infty\}$ and on the Riemann sphere $\bar{\mathbb{C}}$, not just for individual values of z. The paper is intended as a proof of concept to study efficiently global (in the complex plane) solutions to singular ODEs (in these approaches, infinity is just a grid point and large values of the argument are treated as values in the vicinity of the origin) such as the Heun equation and Painlevé equations.

The hypergeometric function can be defined in many ways, see for instance [1]. In this paper, we construct it as the solution of the hypergeometric differential equation

$$x(1 - x)y'' + (c - (1 + a + b)x)y' - aby = 0 \tag{1}$$

with $F(a, b, c, 0) = 1$; here, $a, b, c \in \mathbb{C}$ are constant with respect to $x \in \mathbb{R}$. Equation (1) has regular (Fuchsian) singularities at 0, 1 and infinity. The Riemann symbol of the equation is given by

$$P \left\{ \begin{array}{ccc} 0 & 1 & \infty \\ 0 & 0 & a \quad z \\ 1 - c & c - a - b & b \end{array} \right\}, \tag{2}$$

where the three singularities are given in the first line. The second and third lines of the symbol (2) give the exponents in the generalized series solutions of (1): If ξ is a local parameter near any of these singularities, following Frobenius' method (see again [1]), the general solution can be written for sufficiently small $|\xi|$ in the form of generalized power series

$$y = \xi^{\kappa_1} \sum_{n=0}^{\infty} \alpha_n \xi^n + \xi^{\kappa_2} \sum_{m=0}^{\infty} \beta_m \xi^m \tag{3}$$

if the difference between the constants κ_1 and κ_2 (corresponding to the second and third lines of the Riemann symbol (2) respectively) is not integer; the constants (with respect to ξ) α_n and β_m in (3) are given for $n > 0$ and $m > 0$ in terms of α_0 and β_0, the last two being the only free constants in (3). It is known that generalized series of the form (3) have a radius of convergence equal to the minimal distance from the considered singularity to one of the other singularities of (1). Note that the form of the solution in (3) implies that at the singularities, only one condition can be given in

general though the equation is of second order where on generic points, two initial or boundary conditions need to be given in order to identify a unique solution.

Note also that logarithms may appear in the solution if the difference of the constants κ_1 and κ_2 is integer. In the first part of the paper, we do not consider such cases and concentrate on the generic case

$$c, c - a - b, a - b \notin \mathbb{Z}. \tag{4}$$

Solutions with logarithms will be discussed in Section 4 using a hybrid approach, i.e. the logarithmic will be addressed analytically, and just the analytic part of the solution is constructed numerically.

It is also well known that Moebius transformations (in other words elements of $PSL(2, \mathbb{C})$)

$$z \mapsto \frac{\alpha z + \beta}{\gamma z + \delta}, \quad \alpha, \beta, \gamma, \delta \in \mathbb{C}, \quad \alpha\delta - \beta\gamma \neq 0, \tag{5}$$

transform one Fuchsian equation into another Fuchsian equation, the hypergeometric one (1) into Riemann's differential equation, see [1]. Moebius transformations can thus be used to map any singularity of a Fuchsian equation to 0.

The goal of this paper is to construct numerically the hypergeometric function $F(a, b, c, x)$ for in principle arbitrary values of a, b, c, and for arbitrary $x \in \bar{\mathbb{C}}$, and this with an efficient numerical approach showing spectral accuracy, i.e. an exponential decrease of the numerical error with the number of degrees of freedom used. We employ a hybrid strategy, that is we use Moebius transformations (5) to map the considered singularity to 0, and we apply a change of the dependent variable so that the transformed solution is just the hypergeometric function in the vicinity of the origin, but with transformed values of the constants a, b, c. This is similar to Kummer's approach, see [1], to express the solutions to the hypergeometric equation in different domains of the complex plane via the hypergeometric function near the origin.[1] We thus obtain 3 domains covering the complex plane, each of them centred at one of the three singular points of (1). The dependent variable of the transformed (1) is then transformed as $y \mapsto \xi^{\kappa_i} y$, $i = 1, 2$, where the κ_i are the exponents in (3). This transformation implies that we get yet another form of the Fuchsian equation which has a solution in terms of a power series. This solution will then be constructed numerically.

This means that we solve one equation in the vicinity of $x = 0$, and two each in each of the vicinities of 1 and infinity. Instead of one form of the equation, we solve five $PSL(2, \mathbb{C})$ equivalent forms of (1). This will be first done for real values of x. Since power series are in general slowly converging because of *cancellation errors*, see for instance the discussions in [18], we solve instead each of the 5 equivalent formulations of (1) subject to the condition $y(\xi = 0) = 1$ (in an abuse of notation, we use the same symbol for the local variable ξ and the dependent variable y in all cases) with spectral methods. Spectral methods are numerical methods for the

---

[1]The Kummer group of symmetries of the hypergeometric equation is essentially generated by Moebius transformations which leave the set of singularities 0, 1, ∞ invariant; note that such a symmetry does not exist for Fuchsian equations with four singularities, but our approach can still be applied as discussed in Section 6.

global solution of differential equations that converge exponentially fast to analytic solutions. We shall use the efficient ultraspherical spectral method [17] which, as we shall see, can achieve higher accuracy than traditional spectral methods such as collocation methods because it is better conditioned. Solutions in each of the three domains are then matched (after multiplication with the corresponding factor ξ_i^κ, $i = 1, 2$) at the domain boundaries to the hypergeometric function constructed near $x = 0$ to obtain a function which is C^1 at these boundaries (being a solution of the hypergeometric equation then guarantees the function is analytical if it is C^1 at the domain boundaries). Thus. we obtain an analytic continuation of the hypergeometric function to the whole real line including infinity.

Solutions to Fuchsian equations can be analytically continued as meromorphic functions to the whole complex plane (more precisely, they are meromorphic functions on a Riemann surface as detailed below). Since the Frobenius approach (3) is also possible with complex x, techniques similar to the approach for the real axis can be retained: we consider again three domains, each of them containing exactly one of the three singularities and the union of which is covering the entire Riemann sphere $\bar{\mathbb{C}}$. On the boundary of each of these domains, we solve the 5 equivalent forms of (1) as on the real axis with boundary data obtained on $\mathbb{R}$. Then, the holomorphic function corresponding to the solution in the interior of the studied domain is obtained by solving the Laplace equation with the data obtained on the boundary. The advantage of the Laplace equation is that it is not singular in contrast to the hypergeometric equation. It is solved by introducing polar coordinates r, ϕ in each of the domains and then using the ultraspherical spectral method in r and a Fourier spectral approach in ϕ (the solution is periodic in ϕ). Since the matching has already been done on the real axis, one immediately obtains the hypergeometric function on $\bar{\mathbb{C}}$ in this way.

The paper is organized as follows: in Section 2, we construct the hypergeometric function on the real axis. In Section 3, the hypergeometric function is analytically continued to a meromorphic function on the Riemann sphere. Non-generic cases of hypergeometric equation are studied in Section 4. In Section 5, we consider examples for interesting values of a, b, c, x in (1). In Section 6, we add some concluding remarks.

2 Numerical construction of the hypergeometric function on the real line

In this section, we construct numerically the hypergeometric function on the whole compactified real line. To this end, we introduce the following three domains:

   domain I: local parameter x, $x \in [-1/2, 1/2]$,
  domain II: local parameter $t = 1 - x$, $t \in [-1/2, 1/2]$,
 domain III: local parameter $s = -1/(x - 1/2)$, $s \in [-1, 1]$.

In each of these domains, we apply transformations to the dependent variable such that the solutions are analytic functions on their domains. Since we shall approximate the solution using Chebychev polynomials that are defined on the unit interval

$[-1, 1]$, we map the above intervals $[x_l, x_r]$ to $[-1, 1]$ via $x_l(1-\ell)/2 + x_r(1+\ell)/2$, $\ell \in [-1, 1]$. For these equations, we look for the unique solutions with $y(0) = 1$ since the solutions are all hypergeometric functions (which is defined to be 1 at the origin) with transformed values of the parameters, as shown by Kummer [1]. This means we are always studying equations of the form

$$a_2(\ell)y'' + a_1(\ell)y' + a_0(\ell)y = 0, \qquad \ell \in [-1, 1], \qquad y(0) = 1, \qquad (6)$$

where $a_2(\ell)$, $a_1(\ell)$ and $a_0(\ell)$ are polynomials and $a_2(0) = 0$. At the domain boundaries, the solutions are matched by a C^1 condition on the hypergeometric function which is thus uniquely determined (the hypergeometric equation then implies that the solution is in fact analytic if it is C^1 at the domain boundaries). To illustrate this procedure, we use the fact that many hypergeometric functions can be given in terms of elementary functions, see for instance [1]. Here, we consider the example

$$F(a, b, c, x) = (1-x)^{-a}, \qquad a = -1/3, \ b = c = 1/2. \qquad (7)$$

The triple a, b, c is thus generic as per condition (4). More general examples are discussed quantitatively in Section 5, non-generic cases in Section 4.

2.1 The ultraspherical (US) spectral method

The recently introduced ultraspherical (US) spectral method [17] overcomes some of the weaknesses of traditional spectral methods such as dense and ill-conditioned matrices. The key idea underlying the US method is to change the basis of the solution expansion upon differentiation to the ultraspherical polynomials, which leads to sparse and well-conditioned matrices, as we briefly illustrate below.

Since the solutions we compute are analytic, they have convergent Chebychev series [22],

$$y = \sum_{j=0}^{\infty} y_j T_j(\ell), \qquad T_j(\ell) = \cos[j \arccos(\ell)], \qquad \ell \in [-1, 1]. \qquad (8)$$

The differentiation operators in the US method are based on the following relations involving the ultraspherical (or Gegenbauer) orthogonal polynomials, $C_j^{(\lambda)}(\ell)$:

$$\frac{dT_j}{d\ell} = \begin{cases} jC_{j-1}^{(1)} & j \geq 1 \\ 0 & j = 0 \end{cases}, \qquad \frac{dC_j^{(\lambda)}}{d\ell} = \begin{cases} 2\lambda C_{j-1}^{(\lambda+1)} & j \geq 1 \\ 0 & j = 0 \end{cases}, \qquad \lambda \geq 1.$$

Hence, differentiations of (8) give

$$y' = \sum_{j=0}^{\infty}(j+1)y_{j+1}C_j^{(1)}(\ell) \quad \text{and} \quad y'' = 2\sum_{j=0}^{\infty}(j+2)y_{j+2}C_j^{(2)}(\ell). \qquad (9)$$

Therefore, if we let y denote the (infinite) vector of Chebychev coefficients of y, then the coefficients of the $\{C_j^{(1)}\}$ and $\{C_j^{(2)}\}$ expansions of y' and y'' are given by $\mathcal{D}_1 y$ and $\mathcal{D}_2 y$, respectively, where

$$\mathcal{D}_1 = \begin{pmatrix} 0 & 1 & & & \\ & & 2 & & \\ & & & 3 & \\ & & & & \ddots \end{pmatrix} \quad \text{and} \quad \mathcal{D}_2 = 2 \begin{pmatrix} 0 & 0 & 2 & & \\ & & & 3 & \\ & & & & 4 \\ & & & & & \ddots \end{pmatrix}.$$

Notice that the expansions in (8) and (9) are expressed in different polynomial bases ($\{T_j\}$, $\{C_j^{(1)}\}$, $\{C_j^{(2)}\}$). The next steps in the US method are (i) substitute (8) and (9) into the differential (6) and perform the multiplications $a_2(\ell)y''$, $a_1(\ell)y'$ and $a_0(\ell)y$ in the $\{C_j^{(2)}\}$, $\{C_j^{(1)}\}$ and $\{T_j\}$ bases, respectively, and (ii) convert the expansions in the $\{T_j\}$ and $\{C_j^{(1)}\}$ bases to the $\{C_j^{(2)}\}$ basis.

Concerning step (i), consider the term $a_0(\ell)y$ and suppose that $a_0(\ell) = \sum_{j=0}^{\infty} a_j T_j(\ell)$, then

$$a_0(\ell)y = \sum_{j=0}^{\infty} a_j T_j(\ell) \sum_{j=0}^{\infty} y_j T_j(\ell) = \sum_{j=0}^{\infty} c_j T_j(\ell),$$

where [17]

$$c_j = \begin{cases} a_0 y_0 + \frac{1}{2}\sum_{k=1}^{\infty} a_k y_k & j = 0 \\ \frac{1}{2}\sum_{k=0}^{j-1} a_{j-k} y_k + a_0 y_j + \frac{1}{2}\sum_{k=1}^{\infty} a_k y_{k+j} + \frac{1}{2}\sum_{k=0}^{\infty} a_{k+j} y_k & j \geq 1. \end{cases} \quad (10)$$

Expressed as a multiplication operator on the Chebychev coefficients y, (10) becomes $c = \mathcal{M}_0[a_0(\ell)]y$, where $\mathcal{M}_0[a_0(\ell)]$ is a Toeplitz plus an almost Hankel operator given by

$$\mathcal{M}_0[a_0(\ell)] = \frac{1}{2} \left[\begin{pmatrix} 2a_0 & a_1 & a_2 & a_3 & \cdots \\ a_1 & 2a_0 & a_1 & a_2 & \ddots \\ a_2 & a_1 & 2a_0 & a_1 & \ddots \\ a_3 & a_2 & a_1 & 2a_0 & \ddots \\ \vdots & \ddots & \ddots & \ddots & \ddots \end{pmatrix} + \begin{pmatrix} 0 & 0 & 0 & 0 & \cdots \\ a_1 & a_2 & a_3 & a_4 & \cdot^{\cdot^{\cdot}} \\ a_2 & a_3 & a_4 & a_5 & \cdot^{\cdot^{\cdot}} \\ a_3 & a_4 & a_5 & a_6 & \cdot^{\cdot^{\cdot}} \\ \vdots & \cdot^{\cdot^{\cdot}} & \cdot^{\cdot^{\cdot}} & \cdot^{\cdot^{\cdot}} & \cdot^{\cdot^{\cdot}} \end{pmatrix} \right]. \quad (11)$$

For all the equations considered in this section (i.e. for the computation of the hypergeometric function on the real axis), $a_0(\ell)$ is either a zeroth or first degree polynomial and hence $a_j = 0$ for $j > 0$ or $j > 1$. In the next section, when we consider the computation of $F(a, b, c, z)$ in the complex plane, $a_0(\ell)$ will be analytic (and entire), in which case $a_0(\ell)$ can be uniformly approximated to any desired accuracy by using only m coefficients a_j, $j = 0, \ldots, m - 1$ for sufficiently large m ($m = 36$ will be sufficient for machine precision accuracy in the next section). Therefore if $n > m$, then the multiplication operator (11) is banded with bandwidth $m - 1$ (i.e. $m - 1$ nonzero diagonals on either side of the main diagonal).

In a similar vein, the multiplication of the series $a_2(\ell)y''$ and $a_1(\ell)y'$ can be expressed in terms of the multiplication operators $\mathcal{M}_2[a_2(\ell)]$ and $\mathcal{M}_1[a_1(\ell)]$ operating on the coefficients of y'' (in the $\{C_j^{(2)}\}$ basis) and y' (in the $\{C_j^{(1)}\}$ basis), i.e. $\mathcal{M}_2[a_2(\ell)]\mathcal{D}_2 y$ and $\mathcal{M}_1[a_1(\ell)]\mathcal{D}_1 y$. If we represent or approximate $a_2(\ell)$ and $a_1(\ell)$ with m Chebychev coefficients, then, as with $\mathcal{M}_0[a_0(\ell)]$, the operators $\mathcal{M}_2[a_2(\ell)]$ and $\mathcal{M}_1[a_1(\ell)]$ are banded with bandwidth $m-1$ if $n > m$. The entries of these operators are given explicitly in [17].

For step (ii), converting all the series to the $\{C_j^{(2)}\}$ basis, we use

$$
T_j = \begin{cases} \frac{1}{2}\left(C_j^{(1)} - C_{j-2}^{(1)}\right) & j \geq 2 \\ \frac{1}{2}C_1^{(1)} & j = 1 \\ C_0^{(1)} & j = 0 \end{cases}, \qquad
C_j^{(1)} = \begin{cases} \frac{1}{1+j}\left(C_j^{(2)} - C_{j-2}^{(2)}\right) & j \geq 2 \\ \frac{1}{2}C_1^{(2)} & j = 1 \\ C_0^{(2)} & j = 0 \end{cases},
$$

from which the operators for converting the coefficients of a series from the basis $\{T_j\}$ to $\{C_j^{(1)}\}$ and from $\{C_j^{(1)}\}$ to $\{C_j^{(2)}\}$ follow, respectively:

$$
\mathcal{S}_0 = \begin{pmatrix} 1 & & -\frac{1}{2} & & \\ & \frac{1}{2} & & -\frac{1}{2} & \\ & & \frac{1}{2} & & -\frac{1}{2} \\ & & & \ddots & & \ddots \end{pmatrix}, \quad
\mathcal{S}_1 = \begin{pmatrix} 1 & & -\frac{1}{3} & & \\ & \frac{1}{2} & & -\frac{1}{4} & \\ & & \frac{1}{3} & & -\frac{1}{5} \\ & & & \ddots & & \ddots \end{pmatrix}.
$$

Thus, the linear operator

$$
\mathcal{L} := \mathcal{M}_2[a_2(\ell)]\mathcal{D}_2 + \mathcal{S}_1\mathcal{M}_1[a_1(\ell)]\mathcal{D}_1 + \mathcal{S}_1\mathcal{S}_0\mathcal{M}_0[a_0(\ell)],
$$

operating on the Chebychev coefficients of the solution, i.e. $\mathcal{L}y$, gives the coefficients of the differential (6) in the $\{C_j^{(2)}\}$ basis.

The solution (8) is approximated by the first n terms in its Chebychev expansion,

$$
y \approx \widetilde{y}_n := \sum_{j=0}^{n-1} y_j T_j(\ell),
$$

that satisfies the condition $\widetilde{y}_n(0) = 1$. To obtain an $n \times n$ linear system for these coefficients, the $\infty \times \infty$ operator $\mathcal{L}$ operator needs to be truncated using the $n \times \infty$ projection operator given by

$$
\mathcal{P}_n = (I_n, \mathbf{0}),
$$

where I_n is the $n \times n$ identity matrix. The $n-1 \times n$ truncation of $\mathcal{L}$ is $\mathcal{P}_{n-1}\mathcal{L}\mathcal{P}_n^\top$, which is complemented with the condition $\widetilde{y}_n(0) = 1$. Then, the $n \times n$ system to be solved is

$$
\begin{pmatrix} T_0(0) & T_1(0) & \cdots & T_{n-1}(0) \\ & \mathcal{P}_{n-1}\mathcal{L}\mathcal{P}_n^\top & & \end{pmatrix} \begin{pmatrix} y_0 \\ y_1 \\ \vdots \\ y_{n-1} \end{pmatrix} = \begin{pmatrix} 1 \\ 0 \\ \vdots \\ 0 \end{pmatrix}, \tag{12}
$$

where $T_j(0) = \cos(j\pi/2)$. We construct the matrix in (12) by using the functions provided in the Chebfun Matlab package [7]. Chebfun is also an ideal environment

for stably and accurately performing operations on the approximation $\widetilde{y}_n$ (e.g. evaluation (with barycentric interpolation [4]) and differentiation, which we require in the following section when matching the solutions at the domain boundaries).

2.2 Domain I

Figure 1 gives the results for solving (1) for the test example (7) on the interval $[-1/2, 1/2]$. For comparison purposes, a Chebychev collocation, or pseudospectral (PS) method [21] is also used. In contrast to the US spectral method in which the operators operate in coefficient space, in the PS method, the matrices operate on the solution values at collocation points (e.g. the Chebyshev points (of the second kind) $\cos(j\pi/n)$, $j = 0, \ldots, n$. These matrices can be constructed in Matlab with Chebfun or the Differentiation Matrix Suite [25].

The top-left frame shows the almost banded structure of the matrix in the system (12), which can be solved in only $\mathcal{O}\left(m^2 n\right)$ operations using the adaptive QR method in [17]. On the real axis, the system (12) can be solved in $\mathcal{O}\left(n\right)$ operations since then,
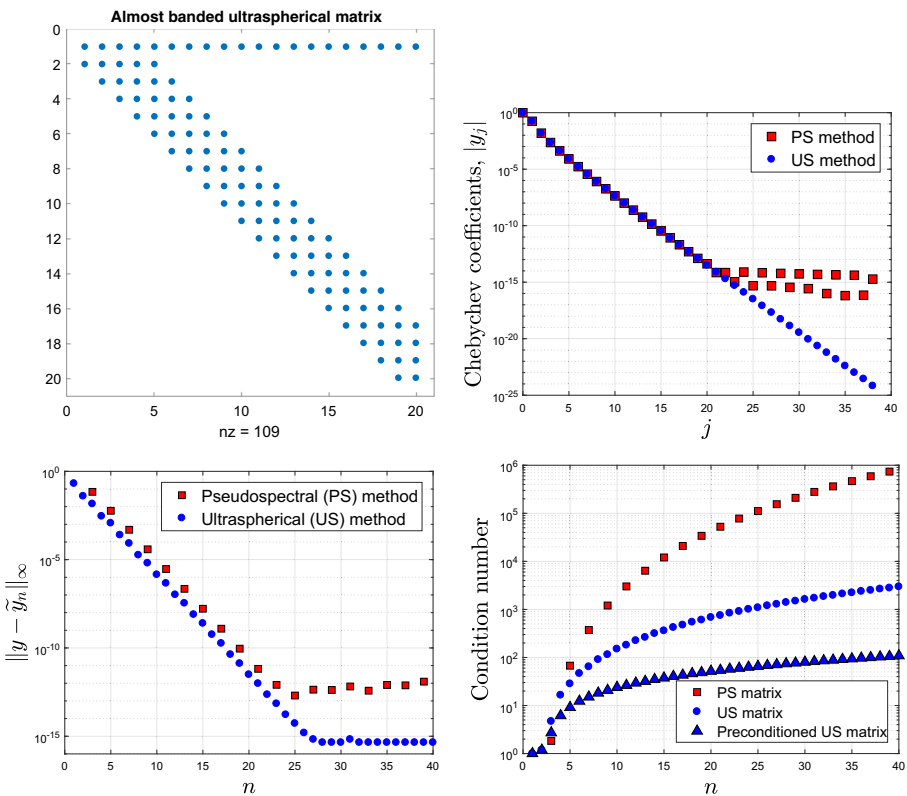

Fig. 1 The performance of the ultraspherical (US) spectral method and a Chebychev collocation, or pseudospectral (PS) method for the solution of the hypergeometric (1) on the interval $[-1/2, 1/2]$ with $y(0) = 1$, $a = -1/3$ and $b = c = 1/2$

as we shall see, $m \leq 4$; in the complex plane, the value of m depends on the required precision. Since the adaptive QR algorithm is not included in Matlab, we express the solution to the almost banded system (12) in terms of banded systems by using the Sherman–Morrison formula; the banded systems are then solved with Matlab's backslash command. The PS method, by comparison, yields dense linear systems.

The top-right frame shows the magnitude of the $n = 40$ Chebychev coefficients of the solution obtained with the US and PS methods. As expected, the magnitudes decrease exponentially with n since the solution is analytic on the interval.

The bottom-left frame shows the maximum error of the computed solution $\widetilde{y}_n$ on the interval $[-1/2, 1/2]$, which can be accurately approximated in Chebfun. The solution converges exponentially fast to the exact solution for both methods; however, the US method stably achieves almost machine precision accuracy (on the order of 10^{-16}) while the PS method reaches an accuracy of around 10^{-13} at $n = 25$ and then the error increases slightly as n is further increased.

The higher accuracy attainable by the US method and the numerical instability of the PS method are partly explained by the condition numbers of the matrices that arise in these methods (see the bottom-right frame). In [17], it is shown that, provided the equation has no singular points on the interval, the condition number of the US matrices grow linearly with n and with preconditioning the condition number can be bounded for all n. By contrast, the condition numbers of collocation methods increase as $\mathcal{O}\left(n^{2N}\right)$, where N is the order of the differential equation. Since (1) has a singular point on the interval, we find different asymptotic growth rates of the condition numbers (by doing a least squares fit on the computed condition numbers): $\mathcal{O}(n)$, $\mathcal{O}\left(n^2\right)$ and $\mathcal{O}\left(n^{4.17}\right)$ for, respectively, preconditioned US matrices (using the diagonal preconditioner in [17]), the US matrices in (12) (with no preconditioner) and PS matrices. We find that, as observed in [17], the accuracy achieved by the US method is much better than the most pessimistic bound based on the condition number of the matrix—hardly any accuracy is lost despite condition numbers on the order of 10^3. As pointed out in [17], while a diagonal preconditioner decreases the condition number of the US matrix, solving a preconditioned system does not improve the accuracy of the solution if the system is solved using QR. This agrees with our experience that the accuracy obtained with Matlab's backslash solver does not improve if some digits are lost by the US method. Hence, all the numerical results reported in this paper were computed without preconditioning.

2.3 Domain II

Next we address domain II with $x \in [0.5, 1.5]$, where we use the local parameter $t = 1 - x$, in which (1) reads

$$t(1 - t)u'' + (a + b + 1 - c - (1 + a + b)t)u' - abu = 0. \qquad (13)$$

The solution corresponding to the exponent 0 in the symbol (2) is constructed as in domain I with the US method. There does not appear to be an elementary closed form of this solution for the studied example, which is plotted in Fig. 2. The results obtained for (13), and also for the three equations remaining to be solved, (14), (16)
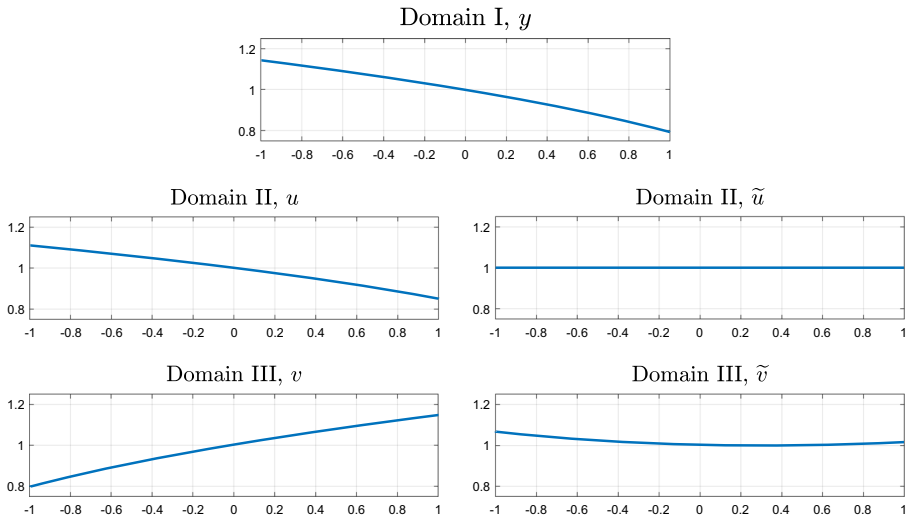
Fig. 2 The computed solutions to (1), (13), (14), (16) and (17), all with the condition $y(0) = 1$ and parameter values $a = -1/3$, $b = c = 1/2$, the building blocks for the numerical construction of the hypergeometric function on the real axis. Accuracy on the order of machine precision is achieved for each of these solutions with, respectively, $n = 28$, $n = 28$, $n = 1$, $n = 30$ and $n = 27$ Chebychev coefficients

and (17), are qualitatively the same as those in Fig. 1 and therefore we do not plot the results again.

The solution proportional to t^{c-a-b} is constructed by writing $u = t^{c-a-b}\widetilde{u}(t)$. Equation (13) implies for $\widetilde{u}$ the equation

$$t(1-t)\widetilde{u}'' + (c - a - b + 1 - t(2c - a - b + 1))\widetilde{u}' - (b - c)(a - c)\widetilde{u} = 0. \quad (14)$$

The solution to (14) with $\widetilde{u}(0) = 1$ is $\widetilde{u} = 1$ since $b = c$ for the test problem and it is recovered exactly by the US method.

Remark 2.1 The appearance of the root t^{c-a-b} in $u = t^{c-a-b}\widetilde{u}$ indicates that the solution as well as the hypergeometric function will in general not be single valued on $\mathbb{C}$, but on a Riemann surface. If the genericness condition (4) is satisfied, this surface will be compact. If this is not the case, logarithms can appear which are only single valued on a non-compact surface. To obtain a single valued function on a compact Riemann surface, monodromies have to be traced which can be numerically done as in [11]. This is beyond the scope of the present paper. Here, we only construct the solutions to various equations which are analytic and thus single valued on their domains. The roots and logarithms appearing in the representation of the hypergeometric function built from these single valued solutions are taken to be branched along the negative real axis. Therefore, cuts may appear in the plots of the hypergeometric function.

2.4 Domain III

For $x \sim \infty$, we use the local coordinate $s = -1/(x - 1/2)$ with $s \in [-1, 1]$. In this case, we get for the hypergeometric equation

$$\frac{s^2}{4}(s-2)(s+2)y'' + sy' \left[\frac{s^2}{2} + \left(c - \frac{a+b+1}{2} \right) s + a + b - 1 \right] - aby = 0. \quad (15)$$

Writing $y = s^a v$, we get for (15)

$$\begin{aligned} &\frac{s}{4}(s-2)(s+2)v'' + v' \left[(a+1)\frac{s^2}{2} + \left(c - \frac{a+b+1}{2} \right) s + b - a - 1 \right] \\ &+ a \left[c - \frac{a+b+1}{2} + \frac{s}{4}(a+1) \right] v = 0. \end{aligned} \quad (16)$$

The hypergeometric (1) is obviously invariant with respect to an interchange of a and b. Thus, we can always consider the case $\Re b > \Re a$. The solution of (15) proportional to s^b can be found by writing $y = s^b \widetilde{v}$ and exchanging a and b in (16) ($b - a$ is not an integer here because of (4)),

$$\begin{aligned} &\frac{s}{4}(s-2)(s+2)\widetilde{v}'' + \widetilde{v}' \left[(b+1)\frac{s^2}{2} + \left(c - \frac{a+b+1}{2} \right) s + a - b - 1 \right] \\ &+ b \left[c - \frac{a+b+1}{2} + \frac{s}{4}(b+1) \right] \widetilde{v} = 0. \end{aligned} \quad (17)$$

The solution to (16) with $v(0) = 1$ and $a = -1/3$, $b = c = 1/2$ is $v = (1 + s/2)^{1/3}$, see Fig. 2. The solution $\widetilde{v}$ to (17) with $\widetilde{v}(0) = 1$, also plotted in Fig. 2, does not appear to have a simple closed form.

2.5 Matching at the domain boundaries

In this section, we have so far shown (for the studied example) that we can compute solutions for each respective domain to essentially machine precision with about $n = 30$ Chebychev coefficients per domain. These solutions are analytic functions and are also the building blocks for the general solution to a Fuchsian equation, as per Frobenius (3). Note that in this approach, infinity is just a normal point on the compactified real axis; thus, large values of $|x|$ are not qualitatively different from points near $x = 0$.

The construction of these analytic solutions allows us to continue the solution into domain I, which is just the hypergeometric function, to the whole real line, even to points where it is singular. This is done as follows: the general solution to the hypergeometric (1) in domain II has the form

$$y_{II}(t) = \alpha u(t) + \beta t^{c-a-b} \widetilde{u}(t), \quad (18)$$

where α and β are constants. These constants are determined by the condition that the hypergeometric function is differentiable at the boundary $x = 1/2$ (which corresponds to $t = 1/2$ since $t = 1 - x$) between domains I and II:

$$y(1/2) = y_{II}(1/2), \qquad y'(1/2) = y'_{II}(1/2)\frac{dt}{dx} = -y'_{II}(1/2). \quad (19)$$

The derivative of the numerical solutions at the endpoints of the interval can be computed using the formulæ $T'_j(1) = j^2$ and $T'_j(-1) = (-1)^{j+1}j^2$. Alternatively, Chebfun can be used, which implements the recurrence relations in [15] for computing the derivative of a truncated Chebychev expansion. In the example studied, we find as expected $\alpha = 0$ and $\beta = 1$ up to a numerical error of 10^{-16}.

Remark 2.2 If $c - a - b \in \mathbb{Z}$ which is here excluded by (4), it is possible that the second solution is not linearly independent. This would lead to non-unique values of α and β in (19).

In the same way, the hypergeometric function can be analytically continued along the negative real axis. The general solution in domain III can be written as

$$y_{III} = \gamma s^a v(s) + \delta s^b \widetilde{v}(s), \tag{20}$$

with γ and δ constants. Again linear independence of these solutions is assured by condition (4). The matching conditions at $x = -1/2$ (which corresponds to $s = 1$ since $s = -1/(x - 1/2)$) are

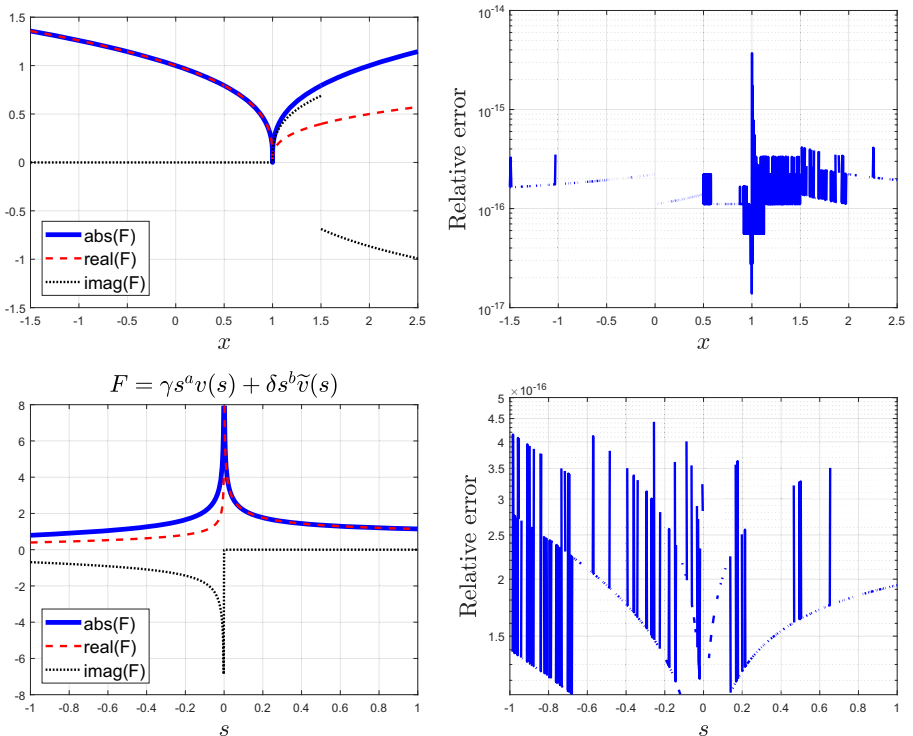$$y(-1/2) = y_{III}(1), \qquad y'(-1/2) = y'_{III}(1)\left[\frac{ds}{dx}\right]_{x=-1/2} = y'_{III}(1). \tag{21}$$

Fig. 3 The hypergeometric function $F(-1/3, 1/2, 1/2, x)$ in the x and s planes (column 1), numerically constructed from the solutions in Fig. 2, and the relative error (column 2)

For the studied example, we find as expected $\gamma = 1$ and $\delta = 0$ with an accuracy better than 10^{-16}. Note that the hypergeometric function is in this way analytically continued also to positive values of $x \geq 3/2$, but this does not imply that the function is continuous at $x = 3/2$ as can be seen in Fig. 3. The reason is the appearance of roots in the solutions, see Remark 2.1, which leads to different branches of the hypergeometric function (Matlab chooses different branches of the functions $t^{c-a-b} = t^{1/3}$ and $s^a = s^{-1/3}$ in domains II and III, respectively, causing the discontinuity in the imaginary part of the solution in Fig. 3).

It can be seen in the top-right frame that full precision is attained on the interval except in the vicinity of $x = 1$ since the singularity causes function evaluation to be ill-conditioned in this neighborhood. The second row of Fig. 3 illustrates the computed hypergeometric function and the error in the s-plane (recall that $s \in [-1, 0]$ is mapped to $x \in [3/2, +\infty)$ and $s \in (0, 1]$ is mapped to $x \in (-\infty, -1/2])$). We have thus computed the hypergeometric function for the test example on the whole compactified real line to essentially machine precision. To recapitulate, this required the solution of five almost banded linear systems of the form shown in Fig. 1, followed by the imposition of continuous differentiability at the domain boundaries as in (19) and (21).

3 Numerical construction of the hypergeometric function in the whole complex plane

In this section, the approach of the previous section is extended to the whole complex plane. Again three domains are introduced each of which exactly contains one of the three singular points 0, 1 and infinity, and which cover now the whole complex z plane. To keep the number of domains to three and in order to have simply connected boundaries, we choose ellipses as shown in Fig. 4:

- domain I: interior of the ellipse given by $\Re z = A \cos\phi$, $\Im z = B \sin\phi$, $\phi \in [-\pi, \pi]$.
- domain II: interior of the ellipse given by $\Re z = 1 + A \cos\phi$, $\Im z = B \sin\phi$, $\phi \in [-\pi, \pi]$.
- domain III: exterior of the circle $\Re z = 1/2 + R \cos\phi$, $\Im z = R \sin\phi$, $\phi \in [-\pi, \pi]$, where $R = B\sqrt{1 - 1/(4A^2)}$.

Thus, each of the ellipses is centred at one of the singular points. The goal is to stay away from the other singular points since the equation to be solved is singular there which might lead to numerical problems if one gets too close. As discussed in [12, 13], a distance of the order of 10^{-3} is not problematic with the used methods, but slightly larger distances can be handled with less resolution. We choose A and B such that the shortest distances between the boundaries of domains I, II and III and the singular points are equal. In the s-plane, the singular points $z = 0$ and $z = 1$ are mapped to, respectively, $s = 2$ and $s = -2$ and the domain boundary is a circle of radius $1/R$ centred at the origin, and thus we require, for a given $1/2 < A < 1$, that B is chosen such that

$$1 - A = 2 - \frac{1}{R}, \quad R = B\sqrt{1 - 1/(4A^2)} \quad \Rightarrow \quad B = \frac{1}{(A + 1)\sqrt{1 - 1/(4A^2)}}. \tag{22}$$
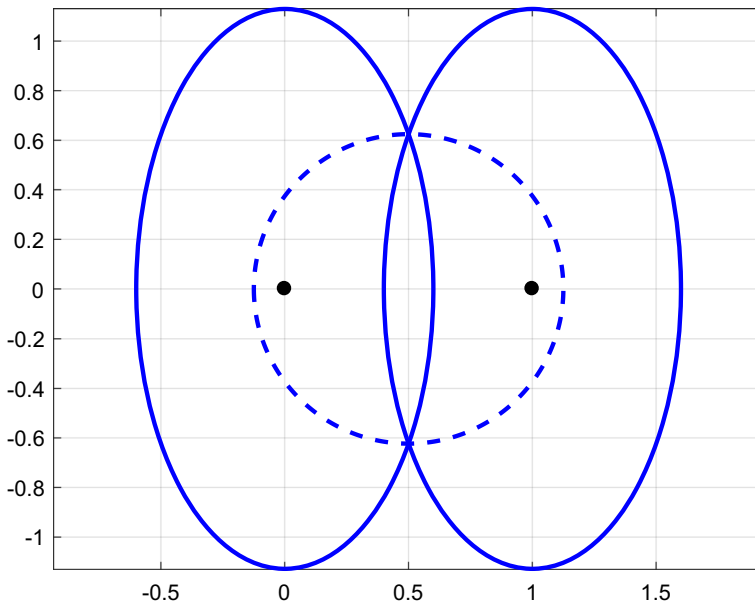
Fig. 4 Domains for the computation of the hypergeometric function. Domain I is the interior of the ellipse centred at $z = 0$, domain II is the interior of the ellipse centred at $z = 1$, and domain III is the exterior of the dashed circle centred at $z = 1/2$

For example, in Fig. 4, $A = 0.6$ (the parameter value we use throughout) and thus the shortest distance from any domain boundary to the nearest singularity is 0.4. This allows us to cover the whole complex plane whilst staying clear of the singularities. There are parts of the complex plane covered by more than one domain; the important point is, however, that the whole plane is covered.

The solution in each of the ellipses is then constructed in 3 steps:

(i)   The code for the real axis described in the previous section is run on larger domains than needed for a computation on the real axis only in order to obtain the boundary values of the five considered forms of (1) at the intersections of the ellipses with the real axis.

(ii) On the ellipse, the equivalent forms of (1) of the previous section are solved in the considered domain with boundary values given on the real axis, again with the US spectral method.

(iii) The obtained solutions on the ellipses serve as boundary values for the solution to the Laplace equation in the interior of the respective domains. In this way, the solutions on the real axis are analytically continued to the complex domains. As described below, the Laplace equation is solved by representing the solution in the Chebychev–Fourier basis, which reduces the problem to a coupled (on an ellipse) or uncoupled (on a disk) system of second-order boundary value problems (BVPs) which we again solve with the US method.

In the last step, the matching described in Section 2.5 provides the hypergeometric function on the whole Riemann sphere built from the constructed holomorphic function in the three domains. As detailed below, this can be achieved as before with spectral accuracy as will be again discussed for the example $F(-1/3, 1/2, 1/2, z)$.

3.1 Domain I

In domain I, the task is to give the solution to (1) with $y(0) = 1$. In step (i) the solution is first constructed on the interval $x \in [-A, A]$ which yields $F(a, b, c, A)$ with the US method detailed in the previous section. We find for the example problem that $n = 40$ is sufficient to compute the solution to machine precision.

In step (ii), the ODE (1) with x replaced by z is solved on the ellipse

$$z(\phi) = \frac{A+B}{2} \exp(i\phi) + \frac{A-B}{2} \exp(-i\phi), \quad \phi \in [-\pi, \pi] \tag{23}$$

as an ODE in ϕ,

$$z_\phi z(1-z)y_{\phi\phi} + \left[z^2(1-z) + (c-(1+a+b)z)z_\phi^2\right]y_\phi - abz_\phi^3 y = 0, \tag{24}$$

where an index ϕ denotes the derivative with respect to ϕ, and where z is given by (23). We seek the solution to this ODE with the boundary conditions $y(\phi = -\pi) = y(\phi = \pi) = F(a, b, c, -A)$. Since the solution is periodic in ϕ, it is natural to apply Fourier methods to solve (24). The Fourier spectral method, which we briefly present, is entirely analogous to the US spectral method—indeed it served as the inspiration for the US spectral method—but simpler since it does not require a change of basis. Note that all the variable coefficients in (24) are band-limited functions of the form $\sum_{j=-m}^{m} a_j e^{ij\phi}$ with $m = 3$. As in the US method, we require multiplication operators to represent the differential equation in coefficient space. Hence, suppose $a(\phi) = \sum_{j=-m}^{m} a_j e^{ij\phi}$ and $y(\phi) = \sum_{j=-\infty}^{\infty} y_j e^{ij\phi}$, then

$$a(\phi)y(\phi) = \sum_{j=-m}^{m} a_j e^{ij\phi} \sum_{j=-\infty}^{\infty} y_j e^{ij\phi} = \sum_{j=-\infty}^{\infty} c_j e^{ij\phi},$$

where

$$c_j = \sum_{k=j-m}^{j+m} a_{j-k} y_j, \quad j \in \mathbb{Z},$$

or

$$c = \mathcal{M}[a(\phi)]y := \begin{pmatrix} \ddots & & \ddots & & \\ & a_m & \cdots & a_{-m} & \\ & & a_m & \cdots & a_{-m} & \\ & & & a_m & \cdots & a_{-m} \\ & & & & \ddots & & \ddots \end{pmatrix} \begin{pmatrix} \vdots \\ y_{-1} \\ y_0 \\ y_1 \\ \vdots \end{pmatrix}.$$

In the Fourier basis, the differential operators are diagonal:

$$\mathcal{D}_1 = i \begin{pmatrix} \ddots & & & \\ & -1 & & \\ & & 0 & \\ & & & 1 \\ & & & & \ddots \end{pmatrix}, \qquad \mathcal{D}_2 = \mathcal{D}_1^2,$$

and thus in coefficient space (24) without the boundary conditions becomes $\mathcal{L}y = 0$, where

$$\mathcal{L} := \mathcal{M}[a_2(\phi)]\mathcal{D}_2 + \mathcal{M}[a_1(\phi)]\mathcal{D}_1 + \mathcal{M}[a_0(\phi)],$$

and $a_2(\phi)$, $a_1(\phi)$ and $a_0(\phi)$ denote the variable coefficients in (24). To find the $2n+1$ coefficients of the solution, y_j, $j = -n, \ldots, n$, we need to truncate $\mathcal{L}$, for which we define the $(2n + 1) \times \infty$ operator

$$\mathcal{P}_{-n,n} = \left(\mathbf{0}, I_{-n,n}, \mathbf{0}\right).$$

The subscripts of the $(2n + 1) \times (2n + 1)$ identity matrix $I_{-n,n}$ indicate the indices of the vector on which it operates, e.g. $\mathcal{P}_{-n,n}y = \left[y_{-n}, \ldots, y_n\right]^\top$. Then, the system to be solved to approximate the solution of (24) is

$$\begin{pmatrix} (-1)^{-n} & (-1)^{1-n} & \cdots & (-1)^n \\ & \mathcal{P}_{-n,n-1}\mathcal{L}\mathcal{P}_{-n,n}^\top & \end{pmatrix} \begin{pmatrix} y_{-n} \\ \vdots \\ y_n \end{pmatrix} = \begin{pmatrix} F(a, b, c, -A) \\ 0 \\ \vdots \\ 0 \end{pmatrix}. \qquad (25)$$

In Fig. 5, (24) is also solved with a Fourier pseudospectral (PS) method [25], which operates on solution values at equally spaced points on $\phi \in [-\pi, \pi)$. The Fourier and Chebyshev PS methods used in Fig. 5 lead to dense matrices whereas the Fourier and Chebyshev spectral methods give rise to almost banded linear systems
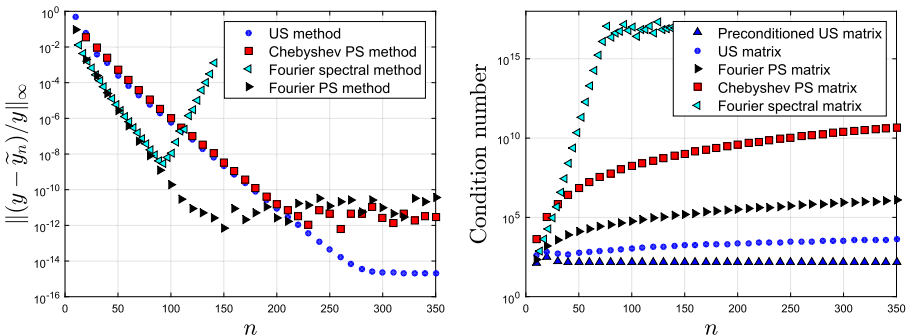

Fig. 5 Numerical solution of the hypergeometric function $F(-1/3, 1/2, 1/2, x)$ on the ellipse (23); on the left, the maximum relative error on the ellipse decreases exponentially with n, and on the right, the growth of condition numbers of the matrices of the preconditioned US method, the US method, the Fourier PS method, the Chebychev PS method and the Fourier spectral method are, respectively, bounded for all n, linear, $\mathcal{O}\left(n^{2.4}\right)$, $\mathcal{O}\left(n^{4.5}\right)$ and $\mathcal{O}\left(e^{0.56N}\right)$, where $N = 2n + 1$

with bandwidths 3 and 35, respectively. The almost banded Fourier spectral matrices (25) have a single dense top row whereas the US matrices have two dense top rows (one row for each of the conditions $y(\phi = -\pi) = y(\phi = \pi) = F(a, b, c, A)$).

Note that the Fourier methods converge at a faster rate than the Chebychev methods. This is to be expected since generally for periodic functions, Fourier series converge faster than Chebychev series by a factor of $\pi/2$ [20]. Again the US method achieves the best accuracy and, as before, this is due to the difference in the conditioning of the methods, as shown in the right frame of Fig. 5.

Unlike the equations solved in Section 2, (24) has no singular points on its domain, and thus the condition numbers of the Chebyshev PS and US matrices grow at different rates than in Fig. 1. We find that, as shown in [17], the condition numbers of the US matrices grow linearly with n and the preconditioned US matrices have condition numbers that are bounded for all n. The condition numbers of the Fourier PS, Chebychev PS and Fourier spectral matrices grow as, respectively, $\mathcal{O}\left(n^{2.4}\right)$, $\mathcal{O}\left(n^{4.5}\right)$ and $\mathcal{O}\left(e^{0.56N}\right)$ (where $N = 2n + 1$, the number of Fourier coefficients of the solution in $(25)^2$), according to a least squares fit of the data. Figure 5 shows that the exponential ill-conditioning of the Fourier spectral matrices results in a rapid loss of accuracy for large enough N. The ill-conditioning could be due to the boundary condition imposed in the first row of (25), which is not standard in periodic problems; however, we do not have a satisfactory explanation.

In step (iii), in order to analytically continue the hypergeometric function to the interior of the ellipse, we use the fact that the function is holomorphic there and thus harmonic. Therefore, we can simply solve the Laplace equation in elliptic coordinates

$$z(r, \phi) = \frac{A + B}{2} r \exp(i\phi) + \frac{A - B}{2} r \exp(-i\phi), \quad r \in [0, 1], \quad \phi \in [0, 2\pi]. \quad (26)$$

In these coordinates, the Laplace operator reads

$$\begin{aligned} r^2 \Delta &= \left(\tfrac{1}{A^2} - \tfrac{1}{B^2}\right)\left[\tfrac{1}{2}\cos 2\phi\left(r^2\partial_{rr} - r\partial_r - \partial_{\phi\phi}\right) + \sin 2\phi\left(\partial_\phi - r\partial_{r\phi}\right)\right] \\ &\quad + \tfrac{1}{2}\left(\tfrac{1}{A^2} + \tfrac{1}{B^2}\right)\left(r^2\partial_{rr} + r\partial_r + \partial_{\phi\phi}\right). \end{aligned} \quad (27)$$

Notice that (27) simplifies considerably on the disk (if $A = B$), which results in a more efficient numerical method. However, using ellipses allows us to increase the distance between the domain boundaries and the nearest singularities and we have found that, if the closest singularity is sufficiently strong, this yields more accurate solutions compared with using disks. For the test problem (7), where the exponent of the singularity at $z = 1$ is $c - a - b = 1/3$, we have found that using an ellipse as opposed to a disk improves the accuracy only by a factor slightly more than two. However, for an example to be considered in Section 5 (the first three rows of Table 2) where the exponent at $z = 1$ is $c - a - b = -0.6$, using an ellipse (with $A = 0.6$ in (22)) yields a solution that is more accurate than the solution obtained on a disk (with parameters $A = B = 0.7574\ldots$, obtained by solving (22) with $A = B$) by two orders of magnitude.

---

[2]In Fig. 5, the results for the Fourier spectral method are plotted against N and not n, as the axis label indicates.

Another possibility, which combines the advantages of ellipses (better accuracy) and disks (more efficient solution of the Laplace equation), is to conformally map disks to ellipses as in [2]. However, we found that computing this map (which involves elliptic integrals) to machine precision for $A = 0.6$ in (22) requires more than 1200 Chebyshev coefficients. This is about four times the number of Chebyshev coefficients required to resolve the solution in Fig. 5. In addition, the first and second derivatives of the conformal map, which are needed to solve the hypergeometric (1) and also (13)–(14) on ellipses, involve square roots and this requires that the right branches be chosen. Hence, due to the expense and complication of this approach, we did not pursue it further.

Yet another alternative is to use rectangular domains, where the boundary data have to be generated by solving ODEs on the 4 sides of each rectangle. Then, the solution can be expressed as a bivariate Chebychev expansion. A disadvantage of this approach, noted in [5], is that the grid clusters at the four corners of the domain which decreases the efficiency of the method.

To obtain the numerical solution of the Laplace equation on the ellipses, we use the ideas behind the optimal complexity Fourier–ultraspherical spectral method in [23, 24] for the disk. Since the solution is periodic in the angular variable ϕ, it is approximated by a radially dependent truncated Fourier expansion:

$$y \approx \widehat{y}_m(r, \phi) := \sum_{k=-m/2}^{m/2-1} u_k(r) e^{ik\phi}, \qquad r \in [-1, 1], \qquad \phi \in [-\pi, \pi). \quad (28)$$

As suggested in [21, 23], we let $r \in [-1, 1]$ instead of $r \in [0, 1]$ to avoid excessive clustering of points on the Chebychev–Fourier grid near $r = 0$. With this approach, the origin $r = 0$ is not treated as a boundary. Since (r, ϕ) and $(-r, \phi+\pi)$ are mapped to the same points on the ellipse, we require that

$$\widehat{y}_m(-r, \phi + \pi) = \widehat{y}_m(r, \phi). \quad (29)$$

On the boundary of the ellipse, we specify $\widehat{y}_m(1, \phi) = \widetilde{y}_n(\phi)$, where $\widetilde{y}_n(\phi)$ is the approximate solution of (24) obtained with the US method. Suppose that $\widetilde{y}_n(\phi)$ has the Fourier expansion $\widetilde{y}_n(\phi) = \sum_{k=-\infty}^{\infty} \gamma_k e^{ik\phi}$. Using the property (29), the boundary condition $\widehat{y}_m(1, \phi) = \widetilde{y}_n(\phi)$ becomes

$$u_k(1) = \gamma_k, \qquad u_k(-1) = (-1)^k \gamma_k, \qquad k = -m/2, \ldots, m/2 - 1. \quad (30)$$

Substituting (28) into (27), we find that the Laplace equation $r^2 \Delta \widehat{y}_m = 0$ reduces to the following coupled system of BVPs, with boundary conditions given by (30):

$$\frac{1}{4}\left(\frac{1}{A^2} - \frac{1}{B^2}\right)\left\{r^2 u_{k-2}'' - [1 + 2(k-2)]r u_{k-2}' + (k-2)[(k-2)+2]u_{k-2}\right\}$$
$$+\frac{1}{2}\left(\frac{1}{A^2} + \frac{1}{B^2}\right)\left\{r^2 u_k'' + r u_k' - k^2 u_k\right\} +$$
$$\frac{1}{4}\left(\frac{1}{A^2} - \frac{1}{B^2}\right)\left\{r^2 u_{k+2}'' - [1 + 2(k+2)]r u_{k+2}' + (k+2)[(k+2)-2]u_{k+2}\right\} = 0, \quad (31)$$

for $k = -m/2, \ldots, m/2 - 1$, where $u_k = 0$ if $k < -m/2$ or $k > m/2 - 1$. Note that on a disk $(A = B)$, the system (31) reduces to a decoupled system of m BVPs. The BVPs are solved using the US method, as in Section 2. Let

$$\mathcal{T}_0 = \mathcal{S}_1 \mathcal{S}_0, \qquad \mathcal{T}_1 = \mathcal{S}_1 \mathcal{M}_1[r]\mathcal{D}_1, \qquad \mathcal{T}_2 = \mathcal{M}_2[r^2]\mathcal{D}_2, \qquad (32)$$

where the operators in (32) are defined in Section 2. Let $\boldsymbol{u}^{(k)}$ denote the infinite vector of Chebychev coefficients of u_k, then in coefficient space (31) becomes

$$\underbrace{\frac{1}{4}\left(\frac{1}{A^2} - \frac{1}{B^2}\right)\{\mathcal{T}_2 - [1 + 2(k-2)]\mathcal{T}_1 + (k-2)[(k-2) + 2]\mathcal{T}_0\}\boldsymbol{u}^{(k-2)}}_{=\mathcal{L}^{(k-2)}}$$

$$+ \underbrace{\frac{1}{2}\left(\frac{1}{A^2} + \frac{1}{B^2}\right)\{\mathcal{T}_2 + \mathcal{T}_1 - k^2\mathcal{T}_0\}\boldsymbol{u}^{(k)}}_{=\mathcal{M}^{(k)}} +$$

$$\underbrace{\frac{1}{4}\left(\frac{1}{A^2} - \frac{1}{B^2}\right)\{\mathcal{T}_2 - [1 + 2(k+2)]\mathcal{T}_1 + (k+2)[(k+2) - 2]\mathcal{T}_0\}\boldsymbol{u}^{(k+2)} = \boldsymbol{0}.}_{=\mathcal{R}^{(k+2)}}$$

(33)

The operators $\mathcal{L}^{(k-2)}$, $\mathcal{M}^{(k)}$ and $\mathcal{R}^{(k+2)}$ defined (33) are truncated and the boundary conditions (30) are imposed as follows to obtain a linear system for the first n Chebychev coefficients of u_k, i.e. $\mathcal{P}_n\boldsymbol{u}^{(k)}$, for $k = -m/2, \ldots, m/2 - 1$:

$$\left(L_n^{(k-2)} \middle| M_n^{(k)} \middle| R_n^{(k+2)}\right)\begin{pmatrix}\mathcal{P}_n\boldsymbol{u}^{(k-2)} \\ \mathcal{P}_n\boldsymbol{u}^{(k)} \\ \mathcal{P}_n\boldsymbol{u}^{(k+2)}\end{pmatrix} = \begin{pmatrix}\gamma_k \\ (-1)^k\gamma_k \\ 0 \\ \vdots \\ 0\end{pmatrix}, \qquad (34)$$

where

$$L_n^{(k-2)} = \begin{pmatrix} 0 & \cdots & \cdots & 0 \\ 0 & \cdots & \cdots & 0 \\ & \mathcal{P}_{n-2}\mathcal{L}^{(k-2)}\mathcal{P}_n^\top & \end{pmatrix}, \qquad R_n^{(k+2)} = \begin{pmatrix} 0 & \cdots & \cdots & 0 \\ 0 & \cdots & \cdots & 0 \\ & \mathcal{P}_{n-2}\mathcal{R}^{(k+2)}\mathcal{P}_n^\top & \end{pmatrix},$$

and

$$M_n^{(k)} = \begin{pmatrix} T_0(1) & T_1(1) & \cdots & T_{n-1}(1) \\ T_0(-1) & T_1(-1) & \cdots & T_{n-1}(-1) \\ & \mathcal{P}_{n-2}\mathcal{M}^{(k)}\mathcal{P}_n^\top & \end{pmatrix}.$$

The (34) can be assembled into two $nm/2 \times nm/2$ block tridiagonal linear systems: one for even k and another for odd k (recall that $\mathcal{P}_n\boldsymbol{u}^{(k)} = \boldsymbol{0}$ for $k < -m/2$ or $k > m/2 - 1$). The systems can be further reduced by a factor of 2 by using the fact that the function $u_k(r)$ has the same parity as k [23] because of the property (29). That is, if k is even/odd, then u_k is an even/odd function and hence only the even/odd-indexed Chebychev coefficients of u_k are nonzero (and thus one of the two top rows imposing the boundary conditions in (34) may also be omitted). Then, the (34) are reduced to two $nm/4 \times nm/4$ block tridiagonal linear systems in which each off-diagonal block is tridiagonal and the diagonal block is almost banded with

bandwidth one and a single dense top row. On a disk, e.g. on domain III, only the diagonal blocks of the system remain and the equations reduce to m times $n/2 \times n/2$ tridiagonal plus rank one systems, which can be solved in $\mathcal{O}(n)$ operations with the Sherman–Morrison formula [23] resulting in a total computational complexity of $\mathcal{O}(mn)$.

Solving the above system, the first n Chebychev coefficients of u_k, where $k = -m/2, \ldots, m/2 - 1$, are obtained which are stored in column k of an $n \times m$ matrix X of Chebychev–Fourier coefficients. Then, the solution expansion (28) is approximated by

$$\widehat{y}_m(r, \phi) \approx \sum_{k=-m/2}^{m/2-1} \sum_{j=0}^{n-1} X_{j,k} T_j(r) e^{ik\phi}, \qquad r \in [-1, 1], \qquad \phi \in [-\pi, \pi). \quad (35)$$

Figure 6 shows the exponential decrease in the magnitude of $X_{j,k}$ for the solution on domain I of the test problem (7) with $m = 188$. Notice that k ranges over only $k = -20, \ldots 93 = m/2 - 1$ instead of $k = -m/2, \ldots, m/2 - 1$ since we only set up the systems (34) for k such that $|\gamma_k|$ is above machine precision. We use Chebfun (which uses the Fast Fourier Transform (FFT)) to compute the Fourier coefficients γ_k of the function $\widetilde{y}_n$ on the domain boundary obtained in step (ii).
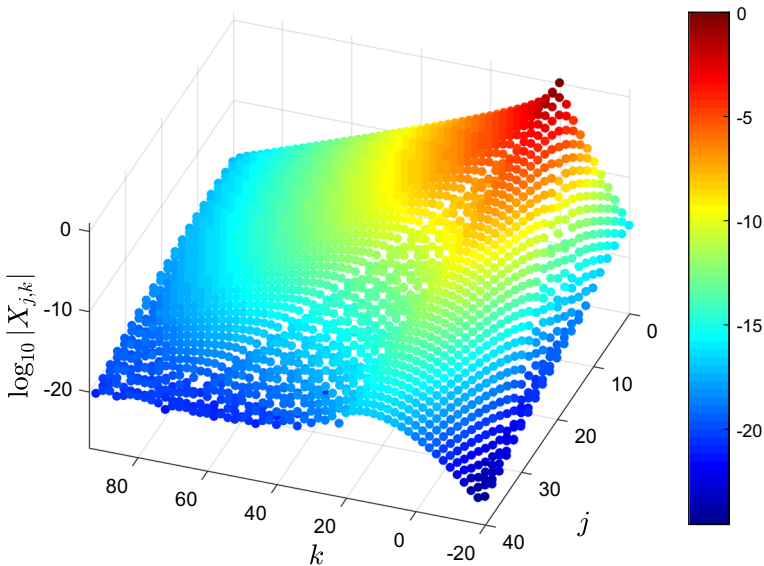

Fig. 6 The magnitude of the Chebychev–Fourier coefficients of the hypergeometric function $F(-1/3, 1/2, 1/2, z)$, obtained by solving the Laplace equation on domain I

To evaluate the Chebychev–Fourier expansion (35) at the set of $n_r n_\phi$ points (r_i, ϕ_j), $i = 1, \ldots, n_r$, $j = 1, \ldots, n_\phi$, where $0 \leq r_i \leq 1$, $-\pi \leq \phi_j < \pi$, we form the $n_r \times 1$ and $n_\phi \times 1$ vectors $\boldsymbol{r}$ and $\boldsymbol{\phi}$ and compute the $n_r \times n_\phi$ matrix

$$
\left(
\begin{array}{c|c|c|c}
\, & \, & \, & \, \\
T_0(\boldsymbol{r}) & T_1(\boldsymbol{r}) & \cdots & T_{n-1}(\boldsymbol{r}) \\
\, & \, & \, & \,
\end{array}
\right)
X
\left(
\begin{array}{c}
\exp(-m/2\boldsymbol{\phi}^\top) \\
\hline
\exp((-m/2+1)\boldsymbol{\phi}^\top) \\
\hline
\vdots \\
\hline
\exp((m/2-1)\boldsymbol{\phi}^\top)
\end{array}
\right).
$$

The columns $T_j(\boldsymbol{r})$ are computed using the three-term recurrence relation $T_{j+1} = 2rT_j - T_{j-1}$, with $T_0 = 1$ and $T_1(r) = r$. Alternatively, the expansion can be evaluated using barycentric interpolation [4] in both the Chebychev and Fourier bases (which also requires the Discrete Cosine Transform and FFT to convert the Chebychev–Fourier coefficients to values on the Chebychev–Fourier grid) or by using Clenshaw's algorithm in the Chebychev basis [15] and Horner's method in the Fourier basis.

Figure 7 shows the maximum relative error on domain I, as measured on a 500×500 equispaced grid on $(r, \phi) \in [0, 1] \times [-\pi, \pi)$, as a function of n, the number of Chebychev coefficients of $u_k(r)$, $k = -m/2, \ldots, m/2 - 1$, for $m = 188$.
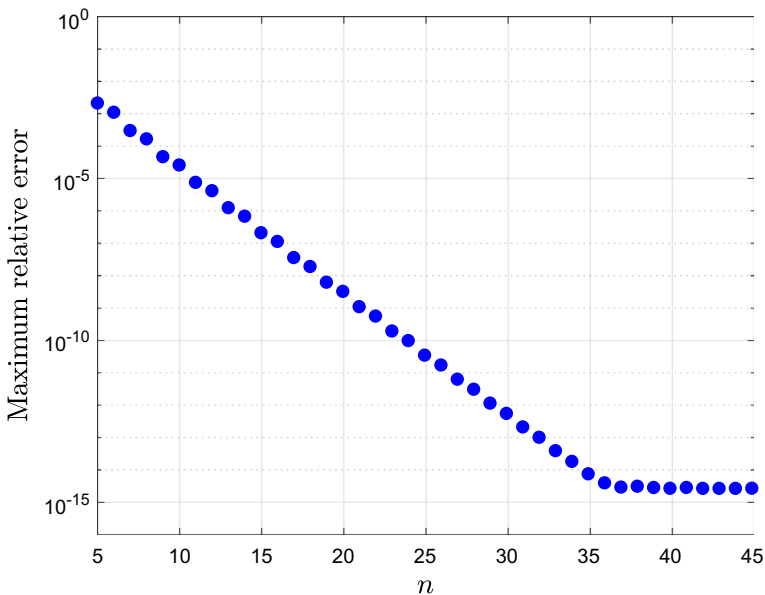

Fig. 7 Spectral convergence of the ultraspherical–Fourier spectral method to $F(-1/3, 1/2, 1/2, z)$ on domain I

3.2 Domains II and III

For the remaining domains and equations, the approach is the same: (13) and (14) are first solved on $[-A, A]$ (with $u(0) = 1 = \widetilde{u}(0)$), then on the ellipse centred at $z = 1$ shown in Fig. 4 and finally the Laplace equation is solved twice on the same ellipse but with different boundary data. Equations (16) and (17) are first solved on $[-1/R, 1/R]$ (with $v(0) = 1 = \widetilde{v}(0)$), then on the disk centred at $z = 1/2$ shown in Fig. 4, which is mapped to a disk of radius $1/R$ in the s-plane, and finally the Laplace equation is solved twice on a disk in the s-plane with different boundary data. The results are very similar to those obtained in Figs. 5 and 7.

Since the solutions constructed in the present section are just analytic continuations of the ones on the real axis of the previous section, the hypergeometric function is built from it as in Section 2.5. Even the values of α, β in (18)–(19) and γ, δ in (20)–(21) are the same and can be taken from the computation on the real axis. Thus, we have obtained the hypergeometric function in the three domains of Fig. 4 which cover the whole Riemann sphere. The computational cost in constructing it is essentially given by inverting five times the matrix approximating the Laplace operator (27), which can be performed in parallel (the one-dimensional computations are in comparison for free).

The relative error is plotted in Fig. 8 for the test problem in the z and s planes. Note that in the left frame that the error is largest close to the singular point $z = 1$ (due to the ill-conditioning of function evaluation in the vicinity of the singularity, as mentioned in the previous section).

4 Degenerate cases

In this section, we consider cases where the genericness condition (4) is not satisfied, i.e. where at least one of $c, c - a - b, b - a$ is an integer. We shall also discuss an example in which these parameters are a distance of $0 < \epsilon \ll 1$ from integer
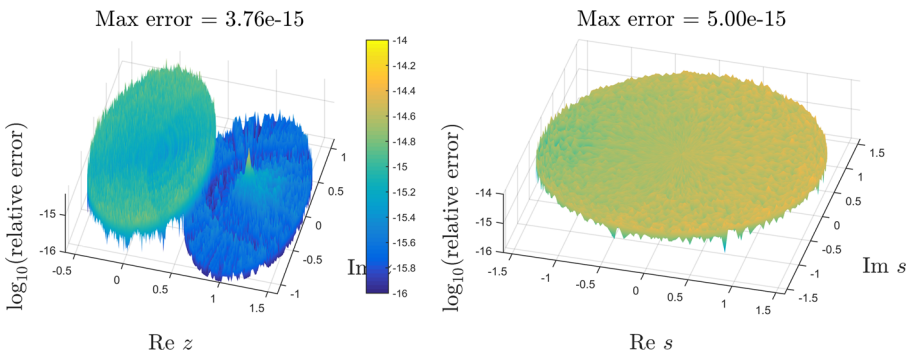

Fig. 8 The relative error on domain I and domain II in the z-plane (left) and on domain III in the s-plane (right) for the approximation of $F(-1/3, 1/2, 1/2, z)$ computed with the multidomain spectral method. Recall that the s-plane corresponds to large z, the 'far field', since $s = -1/(z - 1/2)$

values. In the degenerate cases, logarithms can appear in the solution which are not well approximated by the polynomials applied in the previous sections. Therefore, we propose in this section a hybrid approach based on an analytical treatment of the logarithmic terms and an approximation as before of the analytic parts of the solution. The approach is illustrated with an example.

4.1 General approach for degenerate cases

The hypergeometric (1) can be written in standard form

$$y''(x) + p(x)y'(x) + q(x)y(x) = 0, \tag{36}$$

where

$$p(x) = \frac{c}{x} + \frac{c - (1 + a + b)}{1 - x}, \quad q(x) = \frac{-ab}{x(1 - x)}. \tag{37}$$

Let $x = 0$ be a regular singularity and κ_1, κ_2 with $\Re\kappa_1 \geq \Re\kappa_2$ in (3) be the roots of the characteristic equation. Then, it is well known (see for instance [16, section 2.7] and references therein) that for $M := \kappa_1 - \kappa_2 \in \mathbb{Z}$, the solutions do not necessarily have the form (3). There is always a solution of the form

$$y_1(x) = x^{\kappa_1} \sum_{n=0}^{\infty} \alpha_n x^n, \tag{38}$$

where one can always choose $\alpha_0 = 1$. This solution can be obtained numerically as discussed in the previous sections. A linearly independent solution y_2 to (37) can be obtained for a given solution y_1 via

$$y_2(x) = y_1(x) \int_{x_0}^{x} \frac{ds}{y_1^2(s)} \exp\left(-\int_{x_0}^{s} p(t)dt\right), \tag{39}$$

where c_1, c_2 and x_0 are constants. The integral in (39) takes the form

$$\int^{x} \frac{ds}{y_1^2(s)} \exp\left(-\int^{s} p(t)dt\right) = a_0 \frac{x^{-M}}{-M} + \ldots + a_M \ln x + \sum_{n=0}^{\infty} a_{M+1+n} \frac{x^{n+1}}{n + 1}. \tag{40}$$

Remark 4.1 Formula (39) is not convenient for a spectral method in numerical applications if c is not an integer since the integrand is not an analytic function of x. As can be seen in (40), the integrand is singular for $x = 0$ if $c \in \mathbb{Z}$ and might have further singularities at the zeros of y_1. Thus, formula (39) is not directly suited for a numerical approach, but of theoretical importance (see, however, the remarks in Section 6).

The solution y_2 has in this case the well known form

$$y_2(x) = r(x) + Cy_1(x) \ln x, \tag{41}$$

where $r(x)/x^{\kappa_2}$ is an analytic function, and where C is a constant.

This allows us to use the following approach: we put $y = x^{\kappa_2}\tilde{y}$ where $\tilde{y}$ satisfies with (36) the equation

$$\tilde{y}'' + \tilde{p}(x)\tilde{y}' + \tilde{q}(x)\tilde{y} = 0, \tag{42}$$

where $\tilde{p}(x) = p(x) + 2\kappa_2/x$ and $\tilde{q}(x) = q(x) + p(x)\kappa_2/x + \kappa_2(\kappa_2 - 1)/x^2$. Similarly, we write $y_1(x) = x^{\kappa_1}\tilde{y}_1(x)$ where we choose $\tilde{y}_1(0) = 1$. Then, we can put $\tilde{y} = \tilde{r} + Cx^M\tilde{y}_1(x)\ln x$. Equation (42) implies that $\tilde{r}$ satisfies

$$\tilde{r}'' + \tilde{p}(x)\tilde{r}' + \tilde{q}(x)\tilde{r} = CS(x), \tag{43}$$

where

$$S(x) = (1 - 2\kappa_1 - p(x)x)\, x^{M-2}\tilde{y}_1 - 2x^{M-1}\tilde{y}_1'. \tag{44}$$

For the two cases, we have to consider this implies:

- $M = 0$: the only boundary condition to be imposed is $\tilde{r}(0) = 0$, the constant C in (43) is $C = 1$.
- $M > 0$: we impose the boundary conditions

$$\tilde{r}(0) = 1, \quad \tilde{r}^{(M)}(0) = 0. \tag{45}$$

The constant C in (44) follows from a Taylor expansion of the integrand in (40),

$$\frac{x^{M+1}}{y_1^2(x)} \exp\left(-\int^s p(t)dt\right) = \alpha_0 + \alpha_1 x + \ldots + \alpha_M x^M + \ldots, \quad C = -\frac{M\alpha_M}{\alpha_0}. \tag{46}$$

The Taylor coefficients can be computed in standard way via differentiation. Since y_1 can have zeros, this is best done locally in the vicinity of $x = 0$ by using again barycentric interpolation.

Alternatively, C can be computed by substituting the following expansions into (43) and matching coefficients:

$$p(x) = \sum_{k=-1}^{\infty} p_{k+1}x^k, \quad q(x) = \sum_{k=-2}^{\infty} q_{k+2}x^k, \quad \tilde{r} = \sum_{k=0}^{\infty} \tilde{r}_k x^k, \quad \tilde{y}_1 = \sum_{k=0}^{\infty} \hat{y}_k x^k.$$

Here, we have used the fact that since (36) is a Fuchsian ODE, $p(x)$ and $q(x)$ have a pole of order at most one and two, respectively, at the origin. By also using that the critical exponents κ_1 and κ_2 are roots of the indicial equation

$$Q(\kappa) = \kappa\,(\kappa - 1) + p_0\kappa + q_0,$$

we find that C can be computed as follows for $M \geq 1$:

$$C = -\frac{1}{M}\left(\sum_{j=0}^{M-2}(M - 1 - j)p_{j+1}\tilde{r}_{M-1-j} + \sum_{j=-1}^{M-2}(q_{j+2} + \kappa_2 p_{j+2})\tilde{r}_{M-2-j}\right), \tag{47}$$

where we have set $\tilde{r}_0 = 1$, $\tilde{r}_M = 0$ and $\tilde{r}_k$, $k = 1, \ldots, M - 1$ are computed from the recursion

$$\tilde{r}_{k+2} = \frac{-1}{(k+2)(k-M+2)}\left(\sum_{j=0}^{k}(k+1-j)p_{j+1}\tilde{r}_{k+1-j} + \sum_{j=-1}^{k}(q_{j+2} + \kappa_2 p_{j+2})\tilde{r}_{k-j}\right),$$

for $k = -1, \ldots, M - 3$.

Remark 4.2 The important point of this approach is that $\tilde{r}$ is an analytic function of x, and that $S(x)$ does not contain logarithms. Thus, one has to solve as before a singular ODE, this time with an analytic inhomogeneity $S(x)$ in order to obtain analytical solutions. The logarithms which can appear in the solutions to the original ODE (36) have been taken care of analytically. A numerical approximation is as in the previous sections just needed for analytic functions.

4.2 Hypergeometric solution in the various domains

Below we apply this approach to cases for the hypergeometric equation where logarithms may appear in the various domains introduced before.

4.2.1 Domain I (vicinity of 0)

In this case, logarithms can appear if $c \in \mathbb{Z}$. Since we are interested in the hypergeometric function $F(a, b, c, z)$ defined as a solution to (1) with $F(a, b, c, 0) = 1$, the only not yet addressed case is c a non-positive integer. Note that the hypergeometric series in terms of which the hypergeometric function can be given is not defined in this case. But we can still identify a solution to the hypergeometric equation with the condition $F(a, b, c, 0) = 1$ in this case. We put $M_I = 1 - c$ and $y = x^{M_I} \tilde{y}$. The hypergeometric (1) implies for $\tilde{y}$

$$x(1-x))\tilde{y}'' + (2-c-(a+b+3-2c)x)\tilde{y}' + ((c-1)(1+a+b-c)-ab)\tilde{y} = 0. \quad (48)$$

In the considered case, $\tilde{y}$ has an analytic solution with $\tilde{y}(0) = 1$ which we construct as before. Then, we put $y = r_I + C_I x^{M_I} \tilde{y} \ln x$ and solve

$$x(1-x)r_I'' + (c - (1 + a + b)x)r_I' - ab r_I = C_I S_I, \quad (49)$$

with the conditions $r_I(0) = 1$, $r_I^{(M_I)}(0) = 0$, where

$$S_I = -2(1-x)x^{M_I}\tilde{y}' - x^{M_I-1}\tilde{y}(M_I - (a + b + 2M_I)x). \quad (50)$$

The constant C_I follows as in (46) or (47) from a Taylor expansion.

4.2.2 Domain II (vicinity of $x = 1$)

We again introduce the variable $t = 1 - x$. According to (2), logarithms can appear in the solutions of the hypergeometric equation near $x = 1$ if $c - a - b = M_{II}$ is an integer. There are two cases to be distinguished:

If $M_{II} \leq 0$, (13) has an analytic solution $u_1(t)$ with $u_1(0) = 1$. We put $\tilde{u} = \tilde{r}_{II} + \tilde{C}_{II} t^{-M_{II}} u_1(t) \ln t$ and get from (14)

$$t(1-t)\tilde{r}_{II}'' + (c-a-b+1-t(2c-a-b+1))\tilde{r}_{II}' - (b-c)(a-c)\tilde{r}_{II} = \tilde{C}_{II} \tilde{S}_{II}, \quad (51)$$

where

$$\tilde{S}_{II} = -2t^{-M_{II}}(1-t)u_1' + t^{-M_{II}-1}(M_{II} + (c - M_{II})t)u_1 \quad (52)$$

and solve the equation with the condition $u(0) = 0$ for $M_{II} = 0$ and $u(0) = 1$, $u^{(-M_{II})}(0) = 0$ for $M_{II} < 0$. The constant $\tilde{C}_{II}$ is computed as in (46) or (47) via a Taylor expansion.

If $M_{II} > 0$, (14) has an analytic solution $\tilde{u}_1(t)$ with $\tilde{u}_1(0) = 1$. We put $u = r_{II} + C_{II}t^{M_{II}}\tilde{u}_1 \ln t$ and get from (49)

$$t(1-t)r''_{II} + (a+b+1-c-(1+a+b)t)r'_{II} - abr_{II} = C_{II}S_{II}, \qquad (53)$$

with the conditions $r_{II}(0) = 1$, $r_{II}^{(M_{II})}(0) = 0$, where

$$S_{II} = -2t^{M_{II}}(1-t)u'_1 + t^{M_{II}-1}((c+M_{II})t - M_{II})u_1. \qquad (54)$$

The constant C_{II} is determined as in (46) or (47) via a Taylor expansion.

4.2.3 Domain III (vicinity of infinity)

In domain III, we use again the local variable $s = -1/(x - 1/2)$. Logarithms can appear in the solutions to the hypergeometric equation if $M_{III} = b - a$ is an integer (we recall that we choose $\Re b = \Re a > 0$ without loss of generality). In this case, (17) has an analytic solution $\tilde{v}_1(s)$ with $\tilde{v}_1(0) = 1$. We put $v = r_{III} + C_{III}s^{M_{III}}\tilde{v}_1 \ln s$ and get from (16)

$$\begin{aligned}
&\tfrac{s}{4}(s-2)(s+2)r''_{III} + r'_{III}\left[(a+1)\tfrac{s^2}{2} + \left(c - \tfrac{a+b+1}{2}\right)s + b - a - 1\right]\\
&+a\left[c - \tfrac{a+b+1}{2} + \tfrac{s}{4}(a+1)\right]r_{III} = C_{III}S_{III},
\end{aligned} \qquad (55)$$

with the conditions $r_{III}(0) = 0$ if $M_{III} = 0$ and $r_{III}(0) = 1$, $r_{III}^{(M_{III})}(0) = 0$ if $M_{III} > 0$, where

$$S_{III} = -1/2s^{M_{III}}(s^2-4)\tilde{v}'_1 - s^{M_{III}-1}\tilde{v}_1(s^2/4(2M_{III}+2a+1)+(c-(a+b+1)/2)s - M_{III}). \qquad (56)$$

The constant C_{III} is determined as in (46) or (47) via a Taylor expansion.

4.3 Example $F(2, 3, 5, x)$

The above approach can be illustrated with the explicit solution for the hypergeometric (1) for $a = 2$, $b = 3$ and $c = 5$. It can be shown by direct calculation that the equation has in this case the general solution

$$y(x) = C_1\frac{2x-3}{x^4} + C_2\left(\frac{8(2x-3)}{x^4}\ln(1-x) + \frac{4x^2-6x-27}{x^4}\right), \qquad (57)$$

where C_1 and C_2 are arbitrary constants. As can be seen, the general solution has logarithmic singularities for $x = 1$ and $x = \infty$. We discuss the solutions in the various domains of the previous subsection for this example.

Domain I:

Since $c > 1$, there are no contributions to the hypergeometric functions proportional to $\ln x$. Choosing C_1 and C_2 in an appropriate way, one finds

$$F(2, 3, 5, x) = \frac{12(2x-3)}{x^4}\ln(1-x) + \frac{6x-36}{x^3}. \qquad (58)$$

Note that this is an analytic function near $x = 0$ with $F(2, 3, 5, 0) = 1$.

Domain II:

We have $c - a - b = M_{II} = 0$. The analytic solution $u_1(t)$ reads in this case

$$u_1(t) = \frac{2t + 1}{(t - 1)^4}. \tag{59}$$

Writing $\tilde{u} = \tilde{r}_{II} + u_1(t) \ln t$ (one has $\tilde{C}_{II} = 1$), we find

$$\tilde{r}_{II}(t) = -\frac{t(t + 14)}{2(t - 1)^4}. \tag{60}$$

Domain III:

We have $b - a = M_{III} = 1$. The analytic solution $\tilde{v}_1$ reads

$$\tilde{v}_1(s) = \frac{16(s + 1)}{(s - 2)^4}. \tag{61}$$

Writing $v = r_{III} + C_{III}\tilde{v}_1(t) \ln s$ with $C_{III} = 4$, we find

$$r_{III}(s) = -4\frac{16s(s + 1)\ln(1 + s/2) + 39s^2 + 8s - 4}{(s - 2)^4}. \tag{62}$$

4.4 Numerical implementation and tests

In Section 4.2, we presented a hybrid approach to degenerate cases for the hypergeometric equations which resulted in a set of equations with regular singularities as studied in the previous sections. The same approaches as there can and will be applied since we only construct analytic solutions for which the used spectral methods are very efficient. The only difference to the not degenerate cases is that the studied ODEs now have source terms which are, however, also analytic.

Thus, we apply the ultraspherical method also here which means that equations of the form (43) are approximated by a linear system (cf. (12))

$$\begin{pmatrix} T_0(0) & T_1(0) & \cdots & T_{n-1}(0) \\ T_0^{(M)}(0) & T_1^{(M)}(0) & \cdots & T_{n-1}^{(M)}(0) \\ & & \mathcal{P}_{n-2}\mathcal{L}\mathcal{P}_n^\top & \end{pmatrix} \begin{pmatrix} y_0 \\ y_1 \\ \vdots \\ y_{n-1} \end{pmatrix} = \begin{pmatrix} 1 \\ 0 \\ \mathcal{P}_{n-2}\mathcal{S}_1\mathcal{S}_0 C \mathbf{S} \end{pmatrix}, \tag{63}$$

where $\mathbf{S}$ denotes the Chebychev coefficients of the source term of the equation. If $M = 0$, the first equation of this linear system is removed and $\mathcal{P}_{n-2}$ is replaced by $\mathcal{P}_{n-1}$ so that one more row of the truncated operators is included.

For the example $F(2, 3, 5, x)$ detailed in the previous subsection, we achieve accuracy of the order of machine precision for the computation of (58), (59) and (61), which are solutions of homogeneous equations. Figures 9 and 10 show the source terms of the inhomogeneous (51) and (55) and the accuracy of the computed solutions.

We found (by doing an empirical fit of the condition numbers) that the condition numbers of the matrices in (63) grow as $\mathcal{O}\left(n^3\right)$ and $\mathcal{O}\left(n^{2.7}\right)$ for the solutions in Figs. 9 (domain II) and 10 (domain III), respectively. For the homogeneous equations on domain II (13) and domain III (17), the condition numbers grew as $\mathcal{O}\left(n^{1.5}\right)$ and $\mathcal{O}\left(n^{1.8}\right)$, respectively.
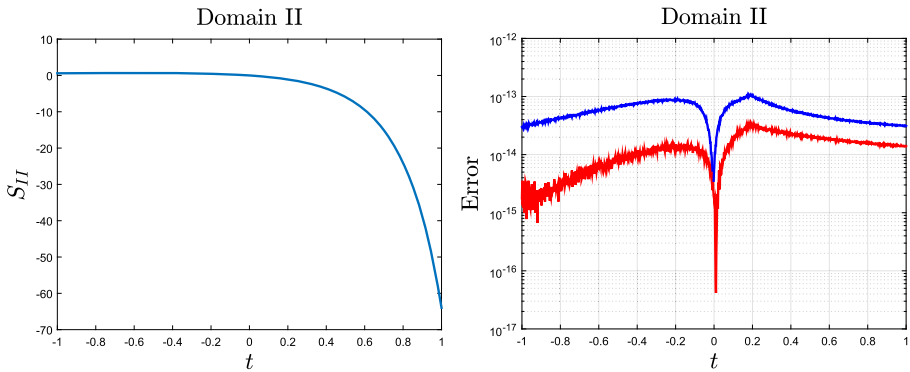
Fig. 9 Left: the source terms of (51), S_{II}, with $a = 2$, $b = 3$ and $c = 5$. Right: the error in the computed solution; the red curve is the exact error while the blue curve is an estimate of the error obtained by subtracting the computed solution with $n-1$ Chebychev coefficients from the one with n coefficients. The error at a point on the interval is taken to be the minimum of the absolute and relative error at that point

A notable feature of the error curves in Fig. 9 and especially Fig. 10 is that there is a noticeable difference between the exact error (in blue) and the estimated error obtained from the computed solutions. In addition, the exact error is smooth while the estimated error contains more random noise (indicative of rounding error). This suggests that the source terms magnify the error by a factor that is determined by the magnitude of the source terms.

The matching of the solutions can be done as in Section 2.5, the only change being that the logarithms appearing have to be included in the matching conditions. Since the logarithms are all bounded at the domain boundaries, there is no problem in determining them and their derivatives there. The continuation of the analytic functions r to the complex plane is done as in Section 3: the equations of Section 4.2 are solved on the ellipses introduced in Fig. 4 to get boundary data there. With these data, the
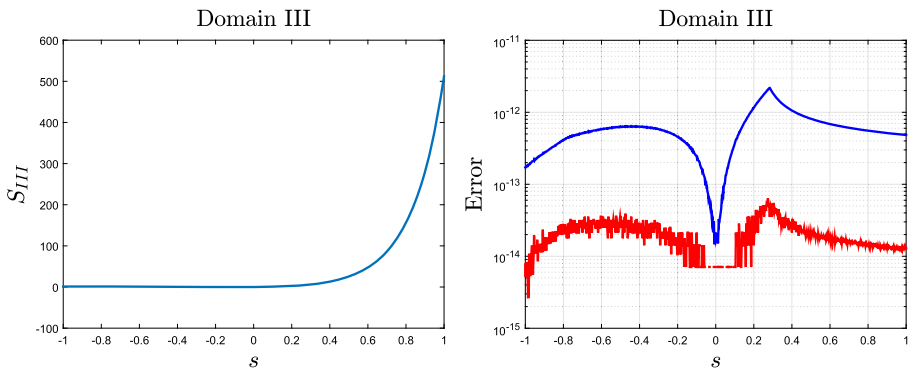

Fig. 10 The same as Fig. 9 but for the solution of (55) on domain III
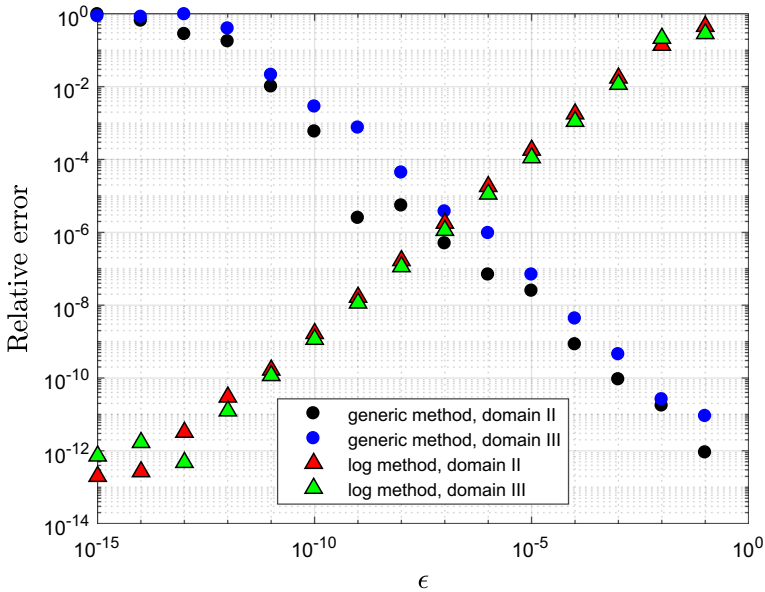
Fig. 11 A comparison of the accuracy of the generic or non-degenerate approach (described in the previous sections) and the method for degenerate cases (outlined in this section) for nearly degenerate cases. The nearly degenerate parameters are $a = 2 + \epsilon$, $b = 3$ and $c = 5$. The error is measured at the points $x = 3/4$ (domain II) and $x = -1$ (domain III). Domain I is not considered because for these parameter values, the methods are identical

Laplace equation is solved for the interior of the ellipses which gives an analytic continuation of the functions r to the complex plane. The complete solution follows then from (41) by including the logarithmic terms.

Although we do not consider nearly degenerate cases in general (i.e. when at least one of the critical exponents in (4) is at most a distance of $0 < \epsilon \ll 1$ from an integer), we consider in Fig. 11 what happens if we perturb the parameter $a \mapsto a + \epsilon$ in the example above ($a = 2$, $b = 3$, $c = 5$). If we use the generic approach described in the previous sections, then as ϵ becomes smaller, accuracy is lost because the condition numbers of the linear systems for the Chebyshev coefficients (12) and for the connection constants α, β, γ, δ, (18)–(21) can become very large.[3] If we use the method described in this section for nearly degenerate cases, we in effect ignore ϵ and round the parameters to the nearest integer. The accuracy of this approach depends on the sensitivity of the solution to perturbations of the parameters. For the example in Fig. 11, the solution changes linearly with ϵ and the crossover point where the method for degenerate cases is more accurate than the method for generic, non-degenerate

---

[3]The next section gives an example where the condition numbers of the system (12) become large. For the example considered in Fig. 11, for $\epsilon = 10^{-1}$, 10^{-5}, 10^{-10}, the condition numbers of the system (18)–(19) are on the order of 10^4, 10^8 and 10^{13}, respectively, and those of (20)–(21) are on the order of 10^6, 10^{14} and 10^{24}, respectively.

cases is close to $\epsilon = 10^{-7}$. At this crossover point, our approaches only attain 6-digit accuracy; thus, the methods in [18] should be considered for higher accuracy.

5 Examples

In this section, we consider further examples. The interesting paper [18] discussed challenging tasks for different numerical approaches and gave a table of 30 test cases for 5 different methods with recommendations when to use which. Note that our approach is complementary to [18]: we want to present an efficient approach to compute a solution to a Fuchsian equation, here the hypergeometric one, not for a single value, but on the whole compactified real line or on the whole Riemann sphere, and this for a wide range of the parameters a, b, c. To treat particularly problematic values of a, b, c, z (e.g. nearly degenerate and large parameter values), it is better to use the codes discussed in [18]. For generic values of the parameters, the present approach and the codes discussed in [18] (http://datashare.is.ed.ac.uk/handle/10283/607) produce similar results.

Below we present cases of [18] with parameters a, b, c which are generic within the given numerical precision and that can thus be treated with the present code together with additional examples along the lines of [18]. We define $\Delta F :=$ $|F_{num}(a, b, c, z) - F_{ex}(a, b, c, z)|/|F_{ex}(a, b, c, z)|$, where we use Maple with 30 digits as the reference solution.

We first address examples with real z and give in Table 1 the first 3 digits of the exact solutions, the quantity ΔF and the number of Chebychev coefficients n. It can

Table 1 Examples for the hypergeometric function compared with a multiprecision computation in Maple for real z

a, b, c, z	$F(a, b, c, z)$	ΔF	n
$-0.1, 0.2, 0.3, 0.5$	0.956	$1.2 * 10^{-16}$	30
$-0.1, 0.2, 0.3, 1.5$	$0.904 + 0.179i$	$6.1 * 10^{-16}$	30
$-0.1, 0.2, 0.3, 100$	$1.365 + 0.400i$	$4.7 * 10^{-16}$	30
$2 + 8i, 3 - 5i, \sqrt{2} - i\pi, 0.25$	$-3.670 - 4.764i$	$7.9 * 10^{-15}$	50
$2 + 8i, 3 - 5i, \sqrt{2} - i\pi, 0.75$	$6882.463 - 6596.555i$	$8.3 * 10^{-15}$	50
$2 + 8i, 3 - 5i, \sqrt{2} - i\pi, -10$	$-0.0166 - 0.0067i$	$7.5 * 10^{-15}$	50
$2 + 200i, 5 - 100i, 10 + 500i, 0.8$	$-4.103 + 6.013i$	$5.9 * 10^{-15}$	70
$2.25, 3.75, -0.5, -1$	-0.631	$4.3 * 10^{-12}$	50
$2 + 200i, 5, 10, 0.6$	$(1.4997 + 5.771i) * 10^{-7}$	$2.4 * 10^{-10}$	160
$1, 3, 7, 0.25$	1.122	$5.0 * 10^{-16}$	30
$1, 3, 7, 0.75$	1.537	$1.3 * 10^{-12}$	50
$1, 3, 7, -3$	0.462	$9.5 * 10^{-12}$	50

Table 2 Examples for the hypergeometric function compared with a multiprecision computation in Maple for complex z

a, b, c, z	$F(a, b, c, z)$	ΔF
$0.1, 0.2, -0.3, -0.5 + 0.5i$	$1.027 - 0.013i$	$2.3 * 10^{-16}$
$0.1, 0.2, -0.3, 1 + 0.5i$	$1.037 - 0.153i$	$6.4 * 10^{-16}$
$0.1, 0.2, -0.3, 5 + 5i$	$1.102 + 0.0288i$	$1.6 * 10^{-15}$
$4, 1.1, 2, \exp(i\pi/3)$	$-0.461 + 0.487i$	$4.0 * 10^{-14}$
$4, 1.1, 2, 1 + 5i$	$-0.0183 + 0.0436i$	$9.1 * 10^{-14}$
$4, 1.1, 2, -5 + 5i$	$0.0216 + 0.0255i$	$9.1 * 10^{-14}$
$2/3, 1, 4/3, \exp(i\pi/3)$	$0.883 + 0.50998i$	$4.0 * 10^{-15}$
$2/3, 1, 4/3, 2i$	$0.562 + 0.373i$	$7.1 * 10^{-15}$
$2/3, 1, 4/3, 1 + i$	$0.740 + 0.740i$	$4.5 * 10^{-15}$
$2/3, 1, 4/3, 100i$	$0.041 + 0.0609i$	$8.7 * 10^{-15}$

be seen that a relative accuracy of the order of 10^{-10} can be reached even when the modulus of the hypergeometric function is of the order of 10^{-7}.

For the results in Table 2 in which the argument z is complex, the number of (i) Chebychev and (ii) Fourier coefficients of the solutions on the ellipses and (iii) the number Chebychev coefficients of the radial Fourier coefficients $u_k(r)$ are in the same ballpark as those required for the test problem (roughly 300, 110 and 40 for (i), (ii) and (iii), respectively (see Figs. 5, 6 and 7)).

To illustrate some of the limitations of our approach, we consider the parameter values $a = 1/10 - 10i$, $b = 1/5 - 10i$, and $c = 3/10 + 10i$. For these parameters, machine precision is achieved on domains I and II; however, on domain III, only 5-digit accuracy is obtained for the solution of (16), as shown in Fig. 12. We estimate the relative error of the computed solution with n Chebyshev coefficients, $v_n(s)$, by comparing it against $v_{80}(s)$ for $n = 10, \ldots, 60$. The accuracy of the estimated errors was confirmed by comparing them with errors obtained from a few high-precision values of $v(s)$ computed with Maple.

For these parameter values, the (16) is ill-conditioned in the sense that small perturbations in the initial conditions and parameter values lead to much larger changes in the solution and this shows up in the large condition numbers in Fig. 12. One could reduce the severe ill-conditioning by using more sub-domains.

Not only is (16) highly ill-conditioned but also the connection relations (20) and (21) for γ and δ. With high-precision computing, it was found that the 2×2 system (21) has a condition number on the order of 10^{24} and thus all accuracy is lost for the connection constants (and consequently also for the solution on domain III) in IEEE arithmetic.

For larger parameter values such as these, computing the hypergeometric function as the solution to a differential equation can lead to highly ill-conditioned systems
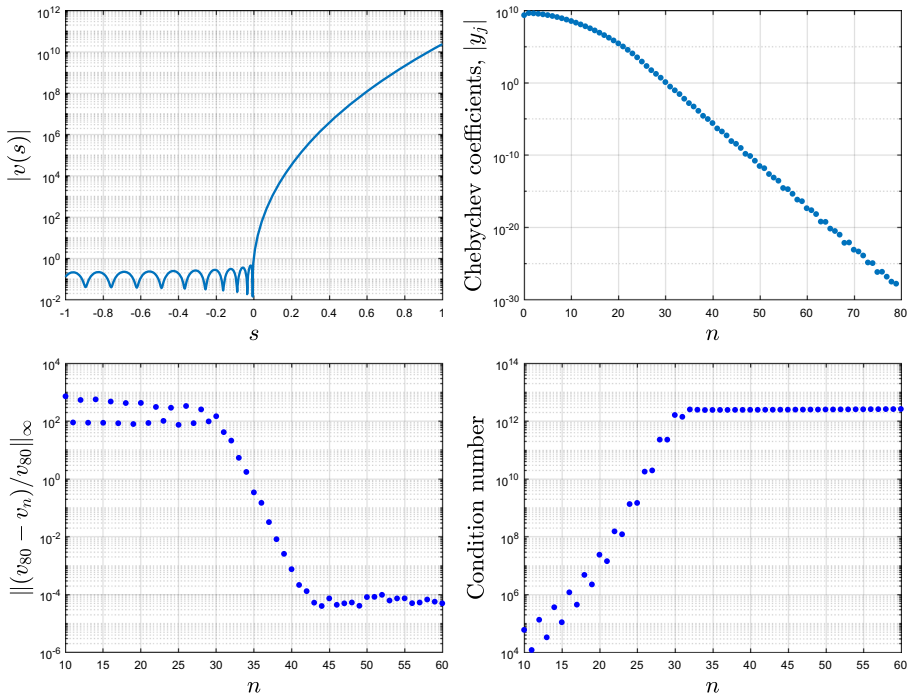
Fig. 12 The limited accuracy of the ultraspherical method for the solution of (16) for the parameter values $a = 1/10 - 10i$, $b = 1/5 - 10i$ and $c = 3/10 + 10i$; the results for (17) are similar

in which case it is better to use the alternative approaches (e.g. recurrence relations) in [18].

6 Outlook

In this paper, we presented a spectral approach for the construction of the Gauss hypergeometric function on the whole Riemann sphere. We also addressed degenerate cases not satisfying the conditions (4) where logarithms can appear in the solution. The approach we give is based on a Taylor expansion of a solution of the hypergeometric equation, see Section 4. In a finite precision setting, even with the well conditioned ultraspherical differentiation matrices, this approach is essentially limited to roughly 10th derivatives. If higher values of an integer difference between the characteristic roots are to be treated, different techniques are necessary. One possibility is given by the Puiseux expansion technique of [13] based on the Cauchy formula which can directly be applied to the integrand of (40). This can be applied in the vicinity of the singularity to essentially machine precision. On the remaining part of the studied domain, which does by definition not contain any singularities, a standard ODE can be solved, and the matching of the solutions is done as discussed

before. This approach will be explored elsewhere in order to keep the presentation here coherent.

One ingredient of our approach was essentially Kummer's relations, see chapter 15.10 of [16], which allows for the representation of the solution to the hypergeometric function in the vicinity of each of the singularities 0, 1, ∞ via the hypergeometric function near 0. Kummer's group of 24 transformations is generated by

$$z \mapsto \frac{z}{z-1}, \quad z \mapsto 1 - z \tag{64}$$

and appropriate rescalings of the hypergeometric function. Though these mappings in the z-plane are not equivalent to the ones used in this paper, the latter can be of course generated by the Kummer transformations. The important point is that Kummer's group uses Moebius transforms (5) which leave the points 0, 1 and infinity invariant as a set. We did not use this property which leaves the approach open to generalizations to Fuchsian equations with 4 singularities.

One motivation of this work was to present an approach for general Fuchsian equations such as the Heun equation, see [16],

$$w''(z) + \left(\frac{\gamma}{z} + \frac{\delta}{z-1} + \frac{\epsilon}{z-a}\right) w'(z) + \frac{\alpha\beta z - q}{z(z-1)(z-a)} w(z) = 0, \tag{65}$$

where $\alpha + \beta + 1 = \gamma + \delta + \epsilon$ and the Lamé equation which is a special case of the Heun equation (see [16]). The latter equation represents a significant challenge with rich potential benefits—see for example [9] for problems related to computation of the Heun function and its application to general relativity. In the case of a Fuchsian equation with 4 singularities, the Moebius degree of freedom can only be used to fix 3 of the singularities to be 0, 1 and infinity, and the fourth denoted by a in (65) can only be restricted to for instance a disk of radius $1/2$ around the origin. Thus, there is no analogue to Kummer's symmetry group for such equations, but our approach can still be applied with the only change that different Heun equations (i.e. with a singularity $\tilde{a}$ as a result of the used Moebius transform) need to be solved. The main change here is that a fourth singularity makes the introduction of a fourth domain necessary which in addition has to be treated for all possible values of a. Thus, it will be necessary to address also almost degenerate situations where $a \sim 0$ which can be done in principle as in the context of Riemann surfaces in [11].

The techniques used to study the hypergeometric function as a meromorphic function on the Riemann sphere are also applicable to Painlevé transcendents as discussed in [8, 14]. These nonlinear ordinary differential equations (ODEs) also have a wide range of applications, see [6] and references therein. The similarity is due to the fact that Painlevé transcendents are meromorphic functions on the complex plane as is the case for the solutions of Fuchsian equations. Note that nonlinearities only affect the solution process on the real line and on the ellipses in the complex plane, i.e. one-dimensional problems. The only truly two-dimensional method, the solution of the Laplace equation for the interior of the ellipses, is unchanged for the Painlevé transcendents since the latter will be in general meromorphic as well. This replaces the task of solving a nonlinear ODE in the complex plane (which ultimately requires the solution of a system of nonlinear algebraic equations) with a linear PDE

(which requires the solution of a linear system). The study of such transcendents, also on domains containing poles in the complex plane as in [10], with the techniques outlined in this paper will be also subject to further research. Combining the compactification techniques of the present paper and the Padé approach of [10], it should be possible to study domains with a finite number of poles.

References

1. Abramowitz, M., Stegun, I. (eds.): Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables. National Bureau of Standards (1970)
2. Atkinson, K., Han, W.: On the numerical solution of some semilinear elliptic problems. Electron Trans. Numer. Anal. **17**, 206–217 (2004)
3. Auzinger, W., Karner, E., Koch, O., Weinmüller, E.B.: Collocation methods for the solution of eigenvalue problems for singular ordinary differential equations. Opuscula Math. **26**, 29–41 (2006)
4. Berrut, J.-P., Trefethen, L.N.: Barycentric Lagrange interpolation. SIAM Rev. **46**(3), 501–517 (2004)
5. Boyd, J.P., Yu, F.: Comparing seven spectral methods for interpolation and for solving the Poisson equation in a disk: Zernike polynomials, Logan–Shepp ridge polynomials, Chebyshev–Fourier series, cylindrical Robert functions, Bessel–Fourier expansions, square-to-disk conformal mapping and radial basis functions. J. Comput. Phys. **230**(4), 1408–1438 (2011)
6. Clarkson, P.A.: Painlevé Equations – Nonlinear Special Functions, Lecture Notes in Mathematics, p. 1883. Springer, Berlin (2006)
7. Driscoll, T.A., Hale, N., Trefethen, L.N. (eds.): Chebfun Guide. Pafnuty Publications, Oxford (2014)
8. Dubrovin, B., Grava, T., Klein, C.: On universality of critical behaviour in the focusing nonlinear Schrödinger equation, elliptic umbilic catastrophe and the *tritronquée* solution to the Painlevé-I equation. J. Nonl. Sci. **19**(1), 57–94 (2009)
9. Fiziev, P.P., Staicova, D.R.: Solving systems of transcendental equations involving the Heun functions. Am. J. Comput. Math. **2**, 2 (2012)
10. Fornberg, B., Weideman, J.A.C.: A numerical methodology for the Painlevé equations. J. Comp. Phys. **230**, 5957–5973 (2011)
11. Frauendiener, J., Klein, C.: Computational approach to hyperelliptic Riemann surfaces. Lett. Math. Phys. **105**(3), 379–400 (2015). https://doi.org/10.1007/s11005-015-0743-4
12. Frauendiener, J., Klein, C. In: A. Bobenko, C. Klein (eds.) Computational Approach to Riemann Surfaces, Lecture Notes in Mathematics, vol. 2013. Springer (2011)
13. Frauendiener, J., Klein, C.: Computational approach to compact Riemann surfaces. Nonlinearity **30**(1), 138–172 (2017)
14. Klein, C., Stoilov, N.: Numerical approach to Painlevé transcendents on unbounded domains. SIGMA **14**, 68–78 (2018)
15. Mason, J.C., Hanscomb, D.C.: Chebyshev Polynomials. Chapman and Hall/CRC (2002)
16. Olver, F.W.J., Olde Daalhuis, A.B., Lozier, D.W., Schneider, B.I., Boisvert, R.F., Clark, C.W., Miller, B.R., Saunders, B.V.: NIST digital library of mathematical functions, https://dlmf.nist.gov, Release 1.0.22 of 2019-03-15
17. Olver, S., Townsend, A.: A fast and well-conditioned spectral method. SIAM Rev. **55**(3), 462–489 (2013)
18. Pearson, J.W., Olver, S., Porter, M.A.: Numerical methods for the computation of the confluent and Gauss hypergeometric functions. Numer Algor. **74**, 821–866 (2017)
19. Seaborn, J.B.: Hypergeometric Functions and their Applications. Springer (1991)

20. Trefethen, L.N., Weideman, J.A.C.: Two results on polynomial interpolation in equally spaced points. J. Approx. Theory **65**(3), 247–260 (1991)
21. Trefethen, L.N.: Spectral Methods in Matlab. SIAM, Philadelphia (2000)
22. Trefethen, L.N.: Approximation Theory and Approximation Practice, vol. 128. SIAM (2013)
23. Wilber, H.D.: Numerical Computing with Functions on the Sphere and Disk. Boise State University, Master's thesis (2016)
24. Wilber, H.D., Townsend, A., Wright, G.B.: Computing with functions in spherical and polar geometries II. The disk. SIAM J. Sci. Comput. **39**(3), C238–C262 (2017)
25. Weideman, J.A.C., Reddy, S.C.: A Matlab differentiation matrix suite. ACM TOMS **26**, 465–519 (2000)

Affiliations

S. Crespo[1] · M. Fasondini[2] · C. Klein[1] [ID] · N. Stoilov[1] · C. Vallée[1,3]

M. Fasondini
M.Fasondini@kent.ac.uk

N. Stoilov
Nikola.Stoilov@u-bourgogne.fr

C. Vallée
corentin.vallee@irisa.fr

[1] Institut de Mathématiques de Bourgogne, Université de Bourgogne-Franche-Comté, 9 avenue Alain Savary, BP 47870, 21078 Dijon Cedex, France

[2] School of Mathematics, Statistics and Actuarial Science, University of Kent, Sibson Building, Parkwood Road, Canterbury, Kent, CT2 7FS, UK

[3] Present address: Univ. Rennes, Inria, CNRS, IRISA, Inserm, Empenn U1228, Campus de Beaulieu, F-35042 Rennes, France