

Two-dimensional Fourier Continuation and applications

Oscar P. Bruno[†] and Jagabandhu Paul[†]

[†]Computing and Mathematical Sciences, Caltech, Pasadena, CA 91125

Abstract

This paper presents a “two-dimensional Fourier Continuation” method (2D-FC) for construction of bi-periodic extensions of smooth non-periodic functions defined over general two-dimensional smooth domains. The approach can be directly generalized to domains of any given dimensionality, and even to non-smooth domains, but such generalizations are not considered here. The 2D-FC extensions are produced in a two-step procedure. In the first step the *one-dimensional* Fourier Continuation method is applied along a discrete set of outward boundary-normal directions to produce, along such directions, continuations that vanish outside a narrow interval beyond the boundary. Thus, the first step of the algorithm produces “blending-to-zero along normals” for the given function values. In the second step, the extended function values are evaluated on an underlying Cartesian grid by means of an efficient, high-order boundary-normal interpolation scheme. A Fourier Continuation expansion of the given function can then be obtained by a direct application of the two-dimensional Fast Fourier Transform (FFT). Algorithms of arbitrarily high order of accuracy can be obtained by this method. The usefulness and performance of the proposed two-dimensional Fourier Continuation method are illustrated with applications to the Poisson equation and the time-domain wave equation within a bounded domain. As part of these examples the novel “Fourier Forwarding” solver is introduced which, *propagating plane waves as they would in free space* and relying on certain boundary corrections, can solve the time-domain wave equation and other hyperbolic partial differential equations *within general domains* at computing costs that grow *sublinearly* with the size of the spatial discretization.

Keywords: Two-dimensional Fourier Continuation, Poisson Equation, Wave Equation, FC Solver, Fourier Forwarding, FFT.

1 Introduction

This paper presents a “two-dimensional Fourier Continuation” method (2D-FC) for construction of bi-periodic extensions of smooth non-periodic functions defined over general two-dimensional smooth domains. The approach can be directly generalized to domains of any given dimensionality, and even to non-smooth domains, but such generalizations are not considered here. The usefulness and performance of the proposed two-dimensional Fourier Continuation method are illustrated with applications to the Poisson equation and the time-domain wave equation within a bounded domain. As part of these examples the novel “Fourier Forwarding” solver is introduced which, *propagating plane waves as they would in free space* and relying on certain boundary corrections, can solve the time-domain wave equation and other constant-coefficient hyperbolic partial differential equations *within general domains* at computing costs that grow *sublinearly* with the size of the spatial discretization.

The periodic-extension problem has actively been considered in the recent literature, in view, in particular, of its applicability to the solution of various types of Partial Differential Equations (PDE) [1, 3, 4, 6, 8, 13, 16, 20, 23]. The contributions [3, 4, 8, 20], in particular, utilize the Fourier Continuation (FC) method in one dimension in conjunction with dimensional splitting for the treatment of multidimensional

PDE problems. The dimensional splitting is also used in [11] to produce Fourier extensions to rectangular domains in two dimensions, where the Fourier Continuation is effected by separately applying the one-dimensional FC-Gram method [3, 4, 8] first to the columns and then to the rows of a given data matrix of function values. The method does assume that the given smooth function is known on a rectangular region containing the domain for which the continuation is sought.

The approach to periodic function extension presented in [6, 23] is based on the solution of a high-order PDE, where the extension shares the values and normal derivatives along the domain boundary. Reference [13], in turn, presents a function-extension method based on use of radial basis functions. In that approach, overlapping circular partitions, or patches, are placed along the physical boundary of the domain, and a local extension is defined on each patch by means of Radial Basis Functions (RBFs). A second layer of patches is placed outside the first, on which the local values are set to vanish. The zero patches are used in conjunction with a partition of unity function to blend the local extensions into a global counterpart. The choice of functions used to build-up the partition of unity determines the regularity of the extended function.

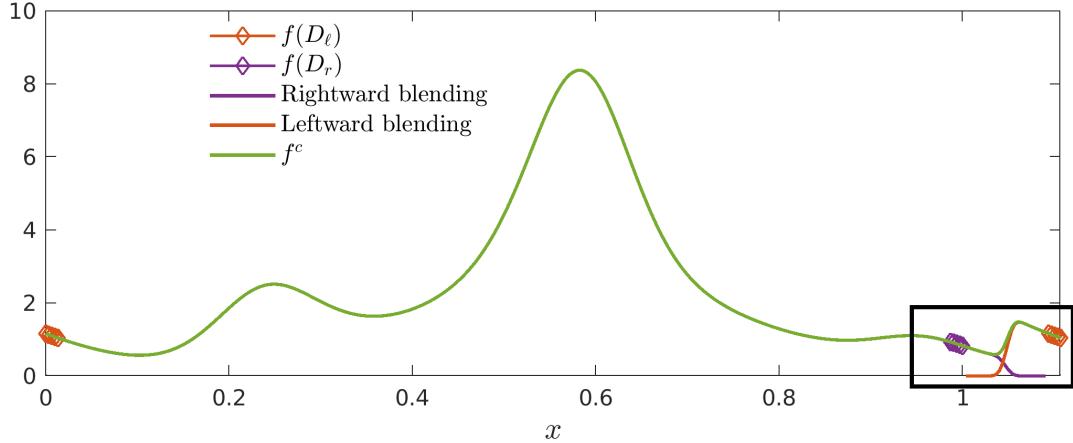
The 2D-FC extensions proposed in this paper are produced in a two-step procedure. In the first step the *one-dimensional* Fourier Continuation method [4] is applied along a discrete set of outward boundary-normal directions to produce, along such directions, continuations that vanish outside a narrow interval beyond the boundary. Thus, the first step of the algorithm produces “blending-to-zero along normals” for the given function values. In the second step, the extended function values are evaluated on an underlying Cartesian grid by means of an efficient, high-order boundary-normal interpolation scheme. A Fourier Continuation expansion of the given function can then be obtained by a direct application of the two-dimensional FFT algorithm. Algorithms of arbitrarily high order of accuracy can be obtained by this method. Since the continuation-along-normals procedure is a fixed cost operation, the cost of the method grows only linearly with the size of the boundary discretization.

As mentioned above, this paper demonstrates the usefulness of the proposed general-domain 2D-FC technique via applications to both, the Poisson problem for the Laplace equation and the time-domain wave equation. In the Poisson case the 2D-FC method is utilized to obtain a *particular solution* for a given right hand side; the boundary conditions are then made to match the prescribed boundary data by adding a solution of the Laplace equation which is produced by means of boundary-integral methods. The Fourier Forwarding approach, in turn, uses the 2D-FC method to solve the spatio-temporal PDE in the interior of the domain and it then corrects the solution values near the boundary by means of a classical time-stepping solver. The overall procedure, which utilizes large time-steps for the interior solver and small CFL-constrained time-steps for the near-boundary solver, runs in computing times that grow *sublinearly* with the size of the spatial discretization mesh.

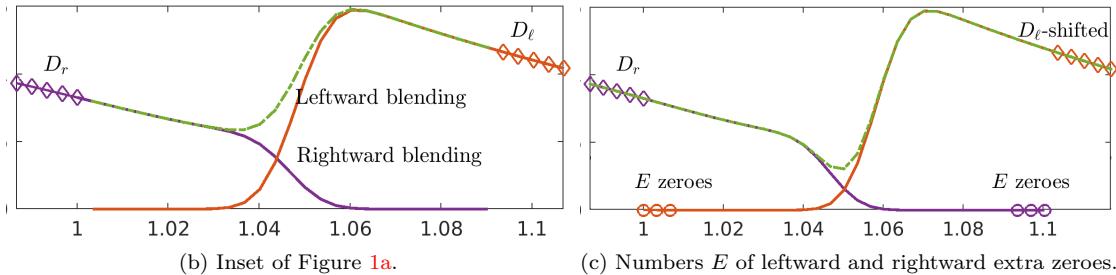
It is interesting to note that the primary continuation device in the 2D-FC method, namely, continuation along normals to the domain boundary, is a *one-dimensional procedure*. This one-dimensional continuation procedure can be utilized in a generalization of the method to n -dimensional domains with $n > 2$. This is in contrast with other extension methods mentioned above. For example, the RBF-based extension method [13] requires solution of boundary problems of increasing dimensionality as the spatial dimension grows, which, given the method’s reliance on dense-matrix linear algebra for the local-extension process, could have a significant impact on computing costs. Similar comments apply to PDE-based extension methods such as [23].

The proposed 2D-FC algorithm performs favorably in the context of existing related approaches. Specific comparisons with results presented in [13] are provided in Section 4.1.1 for a Poisson problem considered in that reference. The recent contribution [12], in turn, presents an FFT-based high-order solver for the Poisson problem for *rectangular* domains, namely, Cartesian products of one-dimensional intervals in either two- or three-dimensional space. The present 2D-FC based Poisson solver achieves, *for general domains*, a similar performance (similar accuracy and computing time) to that demonstrated in [12, Tables 3 and 4] under the Cartesian-domain assumption.

This paper is organized as follows. After a brief review of the 1D-FC method presented in Section 2, the proposed 2D-FC method is introduced in Section 3. The two main applications considered, namely, solution of the Poisson and Fourier-Forwarding for the wave equation, are presented in Sections 4.1 and 4.2. Finally our conclusions are presented in Section 5.



(a) Demonstration of the 1D-FC method. The continuation values are computed as the sum of the blended-to-zero rightward and leftward extensions.



(b) Inset of Figure 1a.

(c) Numbers E of leftward and rightward extra zeroes.

Figure 1: Illustration of the 1D-FC procedure. Figure 1a depicts the Fourier Continuation of the non-periodic function $\phi : [0, 1] \rightarrow \mathbb{R}$ given by $\phi(x) = \exp(\sin(5.4\pi x - 2.7\pi) - \cos(2\pi x)) - \sin(2.5\pi x) + 1$. Figure 1b presents a close up of the right continuation region $[1 - (d - 1)k, b]$. Subsequently Figure 1c illustrates the use of a number E of extra zeroes in the blending to zero process, to yield a continuation mesh containing FFT-friendly numbers (products of powers small prime numbers) of point values.

2 Background: 1D “blending-to-zero” FC algorithm

This section presents a brief review of the one-dimensional Fourier Continuation (1D-FC) method [3, 4], with an emphasis on one of its key components, the *blending-to-zero* procedure—which is employed in the normal-direction continuation portion of the proposed two-dimensional Fourier Continuation (2D-FC) approach presented in Section 3.

2.1 1D-FC algorithm: Outline

Given the vector $\phi_D = (\phi_0, \dots, \phi_{N-1})^t$ of values of a smooth function $\phi : [0, 1] \rightarrow \mathbb{C}$ on the equispaced grid $D = \{x_j = jk : 0 \leq j \leq N - 1\}$ of step-size $k = 1/(N - 1)$, the 1D-FC method [3, 4] of order d (with, e.g., $4 \leq d \leq 12$) produces, at first, an $(N + C)$ -dimensional vector ϕ^c of discrete continuation function values (including the N given function values) over an extended interval $[0, b]$, $b > 1$. To do this, the

algorithm utilizes the d -dimensional vectors $\phi_\ell = (\phi_0, \dots, \phi_{d-1})^t$ and $\phi_r = (\phi_{N-d}, \dots, \phi_{N-1})^t$ of values of the function ϕ on the left and right “matching-point” sets $D_\ell = \{x_0, \dots, x_{d-1}\}$ and $D_r = \{x_{N-d}, \dots, x_{N-1}\}$, respectively, each one of which is contained in a small subinterval of length $(d-1)k$ near the corresponding endpoint of the containing interval $[0, 1]$. In order to obtain the C necessary continuation values, the 1D-FC method blends ϕ_ℓ and ϕ_r to zero (see Section 2.2), towards the left and right, respectively, resulting in two zero-blending vectors of length C . The sum of these two vectors is then utilized as a rightward discrete continuation to the set $D^c = \{x_j = 1 + jk : 1 \leq j \leq C\}$ of points in the interval $(1, b]$ —as described in Section 2.3. As indicated in that section, the overall 1D-FC procedure is then completed via an application of the FFT algorithm to the $(N + C)$ -dimensional vector ϕ^c (cf. equation (2.7) below) of “smoothly-periodic” discrete continued function values. The following two subsections describe the blending-to-zero and 1D-FC approaches, respectively.

2.2 Blending-to-zero Algorithm

In our description of the order- d blending-to-zero algorithm [4] we only present details for the *rightward* blending-to-zero technique, since the leftward blending-to-zero problem can easily be reduced to the rightward problem. Thus, given the column vector $\mathbf{F}_D = (F_0, \dots, F_{d-1})^t$ of values of a complex-valued function F on the set $\mathcal{D} = \{x_0, x_1, \dots, x_{d-1}\}$, the rightward blending-to-zero approach starts by producing a polynomial interpolant for F over the interval $[x_0, x_{d-1}]$ relying on the Gram polynomial basis

$$G_d = \{g_0(x), g_1(x), \dots, g_{d-1}(x)\} \quad (2.1)$$

for this interval. The functions $g_j(x)$ ($j = 0, \dots, d-1$) are the polynomials with real coefficients that are obtained as the Gram-Schmidt orthogonalization procedure is applied, in order of increasing degree, to the polynomials in the set $\{1, x, x^2, \dots, x^{d-1}\}$, with respect to the discrete scalar product

$$(g, h) = \sum_{j=0}^{d-1} g(x_j)h(x_j).$$

Discrete values of the Gram polynomials on the set \mathcal{D} can be computed on the basis of the QR factorization [14]

$$\mathbf{P} = \mathbf{Q}\mathbf{R} \quad \text{of the matrix} \quad \mathbf{P} = (x_i^{j-1})_{0 \leq i, j \leq d-1}. \quad (2.2)$$

(Note that the j^{th} column of \mathbf{Q} contains the values of the j^{th} Gram polynomial on the set \mathcal{D} .) Following [4] we obtain the necessary QR factorization by applying the stabilized Gram-Schmidt orthogonalization method to the matrix \mathbf{P} .

In order to closely approximate each one of the Gram polynomials in G_d , throughout the continuous interval $[0, (d-1)k]$ containing \mathcal{D} , by corresponding trigonometric polynomials, as described below, we use a certain “oversampled matching” method. According to this method the polynomials in G_d are oversampled to an equispaced set of discretization points with stepsize k/n_{os} (containing $n_{\text{os}}(d-1)+1$ points) where n_{os} denotes the oversampling factor, and where the oversampled values are used as part of a certain Singular Value Decomposition (SVD) matching procedure described in what follows. Note that the aforementioned oversampled values on the refined grid $\mathcal{D}_{\text{os}} := \{\tilde{x}_j = jk/n_{\text{os}} : 0 \leq j \leq n_{\text{os}}(d-1)\}$ coincide with the columns of the matrix

$$\mathbf{Q}_{\text{os}} = \mathbf{P}_{\text{os}} \mathbf{R}^{-1}, \quad (2.3)$$

where \mathbf{P}_{os} is the Vandermonde matrix of size $(n_{\text{os}}(d-1)+1) \times d$ corresponding to the oversampled discretization \mathcal{D}_{os} , and where \mathbf{R} is the upper triangular matrix in equation (2.2).

The aforementioned SVD matching procedure, which is one of the crucial steps in the FC approach [4], produces a band-limited Fourier series of the form

$$g_j^c(x) = \sum_{m=-M}^M a_m^j e^{\frac{2\pi i m x}{(d+2C+Z-1)k}} \quad (2.4)$$

for each polynomial $g_j \in G_d$ ($0 \leq j \leq d-1$), where C is the number of blending-to-zero values to be produced, and Z being the number of “zero-matching” points. The Fourier coefficients are selected so as to match, in the least square sense, both the oversampled polynomial values over the interval $[0, (d-1)k]$, and identically zero values on an equally fine discretization of the “zero matching” interval $[(d+C)k, (d+C+Z-1)k]$ of length $(Z-1)k$. The coefficients in (2.4) are taken to equal the solution \mathbf{a} of the minimization problem

$$\min_{\mathbf{a}=(a_{-M}, \dots, a_M)^T} \left\| \mathbf{B}_{\text{os}} \mathbf{a} - \begin{pmatrix} \mathbf{q}_{\text{os}}^j \\ \mathbf{0} \end{pmatrix} \right\|_2, \quad (2.5)$$

where \mathbf{B}_{os} is a matrix whose entries are values of (2.4) at all points in the set \mathcal{D}_{os} as well as all the set of k/n_{os} -spaced points in the “zero matching” interval mentioned above (which, in particular, contains the endpoints $(d+C)k$ and $(d+C+Z-1)k$). The minimizing Fourier coefficients \mathbf{a} are then found via an SVD-based [14] least-squares approach. Once the coefficients \mathbf{a} have been obtained, the resulting Fourier expansions (2.4) are used to produce a certain “continuation matrix” $\mathbf{A} \in \mathbb{C}^{C \times d}$, whose columns equal the values of the expression (2.4) at the C (unrefined) k -discretization points in the interval $[dk, (d+C-1)k]$ (cf. Remark 2.2 below). The desired vector \mathbf{F}^r of rightward blending-to-zero function values at the C continuation points in the interval $[dk, (d+C-1)k]$ is then given by the expression

$$\mathbf{F}^r = \mathbf{A} \mathbf{Q}^T \mathbf{F}. \quad (2.6)$$

2.3 1D-FC Algorithm

As outlined in section 2.1, the 1D-FC algorithm requires use of a certain rightward (resp. leftward) blending-to-zero vector ϕ_r^r (resp. ϕ_ℓ^ℓ) for a given matching-point vector ϕ_r (resp. ϕ_ℓ). In view of equation (2.6) we define $\phi_r^r = \mathbf{A} \mathbf{Q}^T \phi_r$. To obtain the leftward extension ϕ_ℓ^ℓ , in turn, we first introduce the “order reversion” matrix $R^e \in \mathbb{C}^{e \times e}$ ($e \in \mathbb{N}$) by

$$R^e(g_0, g_1, \dots, g_{e-2}, g_{e-1})^t = (g_{e-1}, g_{e-2}, \dots, g_1, g_0)^t,$$

and we then define $\phi_\ell^\ell = R^C \mathbf{A} \mathbf{Q}^T R^d \phi_\ell$. A vector ϕ^c containing both the N given values in the vector $\phi = (\phi_0, \phi_1, \dots, \phi_{N-1})^t$ as well as the C “continuation” function values is constructed by appending the sum $\phi_\ell^\ell + \phi_r^r$ at the end of the vector ϕ , so that we obtain

$$\phi_j^c = \begin{cases} \phi_j & \text{for } 0 \leq j \leq N-1 \\ (\phi_\ell^\ell + \phi_r^r)_{(j-N)} & \text{for } N \leq j \leq N+C-1. \end{cases} \quad (2.7)$$

Following the various stages of the construction of the vector ϕ^c it is easy to check that, up to numerical error, this vector contains point values of a smoothly periodic function defined over the interval $[0, b]$. An application of the FFT algorithm to this vector therefore provides the desired continuation function in the form of a trigonometric polynomial,

$$\phi^c(x) = \sum_{\ell=-(N+C)/2}^{(N+C)/2} \hat{\phi}_\ell^c e^{\frac{2\pi i \ell x}{b}}, \quad (2.8)$$

which closely approximates ϕ in the interval $[0, 1]$. In fact, as demonstrated in previous publications (including [3, 4]), for sufficiently smooth functions ϕ , the corresponding 1D-FC approximants converge to ϕ with order $\mathcal{O}(k^d)$ —so that, as expected, the number d of points used in the blending-to-zero procedures determines the convergence rate of the algorithm. (As discussed in these publications, further, in view of its spectral character the 1D-FC approach enjoys excellent dispersion characteristics as well.) The 2D-FC algorithm introduced in the following section also relies on the one-dimensional blending-to-zero procedure described in section 2.2, and its convergence in that case is once again of the order $\mathcal{O}(k^d)$.

It is important to note that, for a given order d , the matrices \mathbf{A} and \mathbf{Q} can be computed once and permanently stored on disc for use whenever an application of the blending-to-zero algorithm is required—as these matrices do not depend on the point spacing k . A graphical demonstration of various elements of the 1D-FC procedure is presented in Figure 1.

Remark 2.1 (Extra vanishing values). The 1D-FC implementations [3, 4] allow for an additional number $E \geq 0$ of identically zero “Extra” function values to be added on a (unrefined) k -discretization of the interval $[(d + C)k, (d + C + E - 1)k]$, as illustrated in Figure 1c, to obtain a desired overall number of discrete function values (including the given function values and the continuation values produced) such as e.g., a power of two or a product of powers of small prime numbers, for which the subsequent application of the fast Fourier transform is particularly well suited. The corresponding use of extra vanishing values for the 2D continuation problem is mentioned in Remark 3.2.

Remark 2.2 (Blending to zero on a refined grid). As indicated above in the present section, the two-dimensional Fourier continuation procedure introduced in Section 3 utilizes the 1D blending-to-zero strategy described above in this section to extend a function given on a two-dimensional domain Ω along the normal direction to Γ ; the continuation values obtained at all normals are then utilized to obtain the continuation function on the Cartesian grid by interpolation. As detailed in Section 3.4, in order to prevent accuracy loss in the 2D interpolation step we have found it necessary to use 1D normal-direction grids finer than the grids inherent in the blending-to-zero process itself. To easily provide the necessary fine-grid values, a modified fine-grid continuation matrix $\mathbf{A}_r \in \mathbb{C}^{C_r \times d}$ is constructed, where $C_r > C$ denotes the number of fine-grid points utilized. The modified continuation matrix \mathbf{A}_r can be built on the basis of the minimizing coefficients \mathbf{a} in (2.5): the corresponding columns of the fine-grid continuation matrix \mathbf{A}_r are obtained by evaluating (2.4) on the given fine-grid points in the interval $((d - 1)k, (d + C - 1)k]$. The necessary blending-to-zero function values at the C_r fine-grid points are given by $\mathbf{A}_r \mathbf{Q}^T \phi_D$.

3 Two-dimensional Fourier Continuation Method

This section presents the proposed volumetric Fourier continuation method on two-dimensional domains $\Omega \subset \mathbb{R}^2$ with a smooth boundary $\Gamma = \overline{\Omega} \setminus \Omega$, some elements of which are illustrated in Figure 2. Let a smooth function $f : \overline{\Omega} \rightarrow \mathbb{C}$ be given; we assume that values of f are known on a certain uniform Cartesian grid within $\overline{\Omega}$ as well as a grid of points on the boundary Γ . The 2D-FC algorithm first produces one-dimensional “blending-to-zero” values for the function f *along directions normal to the boundary* Γ , yielding continuation values on a certain two-dimensional tangential-normal curvilinear grid around Γ , as detailed in Sections 3.1 and 3.2 and illustrated in Figure 2. These continuation values, which are produced on the basis of the corresponding blending-to-zero procedure presented in section 2.2 in the context of the 1D-FC method, are then interpolated onto a Cartesian grid around the domain boundary, to produce a two-dimensional blending-to-zero continuation of the function f . The necessary interpolation from the curvilinear grid to the Cartesian grid is accomplished by first efficiently obtaining the foot of the normal that passes through a given Cartesian grid point $\mathbf{r} = (x, y)$ exterior to Ω and near the boundary Γ (Section 3.3), and then using a local two-dimensional interpolation procedure to produce the corresponding continuation value at the point \mathbf{r} (Section 3.4). Once the interpolated values have been obtained throughout the Cartesian mesh

around Γ , the desired two-dimensional Fourier continuation function

$$f^c(x, y) = \sum_{\ell=-N_x/2+1}^{N_x/2} \sum_{m=-N_y/2+1}^{N_y/2} \hat{f}_{\ell, m}^c e^{2\pi i \left(\frac{\ell x}{L_x} + \frac{my}{L_y} \right)} \quad (3.1)$$

(where L_x and L_y denotes the period in the x and y directions, respectively) is obtained by means of a 2D FFT. Following the algorithmic prescriptions presented in Sections 3.1 through 3.4, a summary of the overall 2D-FC approach is presented in Section 3.5.

3.1 Two-dimensional tangential-normal curvilinear grid

The necessary curvilinear grids around Γ can be produced on the basis of a (smooth) parametrization

$$\mathbf{r} = \mathbf{q}(\theta) = (x(\theta), y(\theta)), \quad 0 \leq \theta \leq 2\pi, \quad (3.2)$$

of the boundary Γ . In view of their intended application (blending to zero along the normal direction in accordance with section 2.2), the curvilinear grids are introduced within interior and exterior strips V^- and V^+ (illustrated in Figure 2) given by

$$\begin{cases} V^- = \{\mathbf{q}(\theta) - \mathbf{n}(\theta)\gamma : 0 \leq \theta \leq 2\pi \text{ and } 0 \leq \gamma \leq (d-1)k_1\}, \\ V^+ = \{\mathbf{q}(\theta) + \mathbf{n}(\theta)\gamma : 0 \leq \theta \leq 2\pi \text{ and } 0 \leq \gamma \leq Ck_1\}, \end{cases} \quad (3.3)$$

where $\mathbf{n}(\theta) = (n_x(\theta), n_y(\theta))$ denotes the unit normal to the boundary Γ at $\mathbf{q}(\theta)$, and where d , C and k_1 denote, respectively, the number of matching points, the number of blending-to-zero points, and the stepsize used, in the present application of the 1D blending-to-zero procedure described in Section 2.2. Using, in addition, the uniform discretization

$$I_B = \{\theta_p = pk_2 : 0 \leq p < B\}, \quad k_2 = \frac{2\pi}{B}, \quad (3.4)$$

of the interval $[0, 2\pi]$, we then construct a curvilinear two-dimensional discretization

$$V_{B,d}^- = \{\mathbf{r}_{p,q} : \mathbf{r}_{p,q} = \mathbf{q}(\theta_p) + \mathbf{n}(\theta_p)(q-d+1)k_1; 0 \leq p < B \text{ and } 0 \leq q \leq d-1\} \quad (3.5)$$

$$V_{B,C_r}^+ = \{\mathbf{s}_{p,q} : \mathbf{s}_{p,q} = \mathbf{q}(\theta_p) + \mathbf{n}(\theta_p)qk_1/n_r; 0 \leq p < B \text{ and } 0 \leq q \leq C_r\}, \quad (3.6)$$

within V^- and V^+ respectively, for the given stepsize k_1 where $C_r = Cn_r$ for certain integer (refinement factor) n_r ; note that the points in $V_{B,d}^-$ for $q = d-1$ and the points in V_{B,C_r}^+ for $q = 0$ coincide and that they lie on Γ . Here the constants d , C and n_r are independent of B . The continuation function is constructed so as to vanish at all points $\mathbf{s}_{p,q} \in V_{B,C_r}^+$ with $q = C_r$. Let now $\mathcal{R} = [a_0, a_1] \times [b_0, b_1]$ denote the smallest closed rectangle containing $\Omega \cup V^+$, and consider the equispaced Cartesian grid of stepsize h ,

$$H = \{\mathbf{z}_{i,j} = (x_i, y_j) : x_i = a_0 + ih; y_j = b_0 + jh : 0 \leq i < N_x, 0 \leq j < N_y\} \quad (3.7)$$

on \mathcal{R} , where the two-dimensional continuation function values are to be computed. We note that the size of the rectangle \mathcal{R} along with the strips V^- and V^+ decrease as the stepsize k_2 is decreased.

3.2 Computation of FC values on V_{B,C_r}^+

A continuation of the function f to the exterior of Ω is obtained via application of the blending-to-zero procedure presented in Section 2.2 (cf. Remark 2.2) along each one of the normal directions inherent in the definition of the set V_{B,C_r}^+ . For given p , the d equidistant points $\mathbf{s}_{p,q} \in V_{B,d}^-$ ($0 \leq q \leq d-1$), which

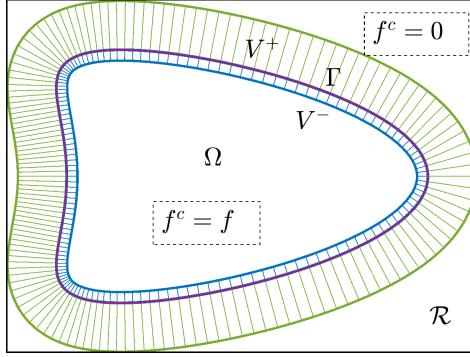


Figure 2: Geometrical constructions underlying the 2D-FC procedure, with reference to the various regions defined in Section 3.1.

are indicated by the solid circles in Figure 3, constitute a set \mathcal{D}_p of matching points that are used to effect the blending-to-zero procedure per the prescriptions presented in Section 2.2. To obtain the desired continuation function values it is necessary to first obtain the vector $\mathbf{f}_{\mathcal{D}_p}$ of the values of the function f (or suitable approximations thereof) on the set \mathcal{D}_p . In the proposed method, the needed function values $\mathbf{f}_{\mathcal{D}_p}$ are computed on the basis of a two-step polynomial interpolation scheme, using polynomials of a certain degree ($M - 1$), as briefly described in what follows.

With reference to the right image of Figure 3, and considering first the case $|n_x(\theta_p)| \geq |n_y(\theta_p)|$, the algorithm initially interpolates vertically the function values at M open-circle Cartesian points selected as indicated in the figure, onto the points of intersection, shown as red-stars, of the normal line and the vertical Cartesian grid lines. For intersection (red-star) points close enough to the boundary, boundary function values at boundary points shown as squares in the figure, are utilized in the one-dimensional interpolation process as well. Once the red-star function values are acquired, the function value at the matching solid-black point is effected by interpolation from the M red-star point values previously obtained, on the basis of a polynomial of degree ($M - 1$).

The case $|n_x(\theta_p)| < |n_y(\theta_p)|$ is treated similarly, substituting the initial interpolation along vertical Cartesian lines, by interpolation along horizontal Cartesian lines; the algorithm then proceeds in an entirely analogous fashion.

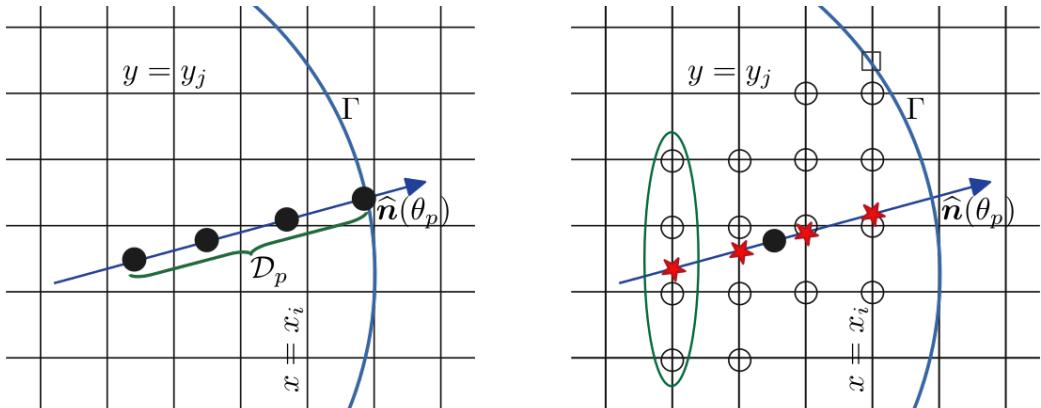


Figure 3: Interpolation scheme for evaluation of $\mathbf{f}_{\mathcal{D}_p}$ in the case $|n_x(\theta_p)| \geq |n_y(\theta_p)|$. Left: Black solid circles indicate the matching points that define the set \mathcal{D}_p . Right: Known function values at the Cartesian points (denoted by the empty circles) and, in some cases, at the point of intersection (represented by an empty square) of the vertical grid lines with Γ , are used to interpolate the matching function values at the red-star intersection points of the normal with the vertical grid lines. The function values at the red-points are then used to obtain, by interpolation, the function values at the matching points.

3.3 Proximity map

As described below in Section 3.4, the 2D FC algorithm interpolates the Fourier continuation values on V_{B,C_r}^+ onto the Cartesian mesh H . The interpolation algorithm used in that section relies on a certain “proximity map” $\mathcal{P} : H \cap V^+ \rightarrow V_{B,C_r}^+$ which associates a curvilinear grid point $s_{p,q} = \mathcal{P}(z_{i,j}) \in V_{B,C_r}^+$ in the “proximity” of each given Cartesian grid point $z_{i,j}$. The proximity function we use is obtained by first associating to each curvilinear discretization point $s_{p,q}$ the nearest Cartesian point, a procedure that results in a set $\mathcal{P}_0 \subseteq (H \cap V^+) \times V_{B,C_r}^+$ of pairs of points, one in the Cartesian grid and the other in the curvilinear grid. (The initial set \mathcal{P}_0 can easily be obtained by using the “integer part” *floor* ($\lfloor \cdot \rfloor$) and the *ceil* ($\lceil \cdot \rceil$) operators.) The set \mathcal{P}_0 is then modified by removing multiple associations for a given Cartesian point, and, if necessary, by adding a “next-nearest” curvilinear neighbor to Cartesian points that previously remained un-associated. The resulting set \mathcal{P} defines the desired function.

3.4 FC Values on the Cartesian Grid

Once Fourier continuation values on V_{B,C_r}^+ have been obtained, per the procedure presented in Section 3.2, the two-dimensional FC scheme can be completed by (a) Interpolation onto the set $H \cap V^+$ of outer Cartesian grid points; and, (b) Subsequent evaluation of the corresponding Fourier coefficients in equation Equation (3.1) by means of an FFT. (Note that since the continuation function f^c is a smooth function which vanishes outside $\bar{\Omega} \cup V^+$, this function can be viewed as the restriction of a smooth and bi-periodic function with periodicity rectangle \mathcal{R} —whose Fourier series approximates f^c , and therefore f , accurately.) The efficiency of the interpolation scheme is of the utmost importance in this context—since interpolation to a relatively large set $H \cap V^+$ of Cartesian points is necessary. An accurate and efficient interpolation strategy is obtained by combining two one-dimensional (local) interpolation procedures based on nearby normal directions. The first interpolation procedure produces the parameter value θ of the foot of the normal to Γ passing through a given Cartesian point; the second procedure then approximates the continuation function value utilizing the parameter value θ just mentioned and the continuation function values at the points in V_{B,C_r}^+ around the given Cartesian point. A detailed description of the combined interpolation methodology is presented in what follows. Specifically, we describe the strategy we use to interpolate the continuation function onto each point $Q = z_{i,j} \in H \cap V^+$ and, to do this, we first obtain the foot $\mathcal{F}(Q)$ of this point. Using the proximity map \mathcal{P} described in Section 3.3, the algorithm utilizes the curvilinear discretization point $s_{p,q} = \mathcal{P}(z_{i,j}) \in V_{B,C_r}^+$ as well as the corresponding boundary discretization parameter value θ_p ; according to equation (3.6), the point $q(\theta_p)$ equals the foot of the normal passing through $s_{p,q}$: $q(\theta_p) = \mathcal{F}(s_{p,q})$. The algorithm then seeks approximation of the foot $\mathcal{F}(Q)$ and the corresponding parameter value $\theta = \theta^Q$ via a preliminary interpolation step, as indicated in what follows. The foot $\mathcal{F}(Q)$ and the corresponding parameter value $\theta = \theta^Q$ are then used by the algorithm to produce the desired interpolated continuation value at Q .

In order to obtain $\mathcal{F}(Q)$ the algorithm uses the M boundary parameter values in the set

$$S_{\theta_p} = \{\theta_{p-K_\ell}, \theta_{p-K_\ell+1}, \dots, \theta_p, \theta_{p+1}, \dots, \theta_{p+K_r}\} \subset I_B \quad (3.8)$$

around θ_p (where $K_r + K_\ell + 1 = M$, and where $K_r = K_\ell$ if M is odd and $K_r = K_\ell + 1$ if M is even). Parameter values θ_k with negative values of k , which may arise in Equation (3.8), are interpreted by periodicity: $\theta_k = \theta_{B+k}$.

The algorithm then utilizes the line L_Q^\perp passing through Q that is orthogonal to the normal vector $\mathbf{n}(\theta_p)$ (see left image in Figure 4), together with the parametrization $\ell_Q^{\text{tr}}(\tau)$ of L_Q^\perp , where the parameter τ represents the signed distance of the points on L_Q^\perp from the point Q . Clearly, then, $\ell_Q^{\text{tr}}(0) = Q$. Each point of intersection of L_Q^\perp with the normals $\mathbf{n}(\theta_j)$ ($\theta_j \in S_{\theta_p}$), on the other hand, equals $\ell_Q^{\text{tr}}(\tau_j)$ where τ_j denotes the signed distance between Q and the corresponding intersection point. Thus, defining the

function $\theta = \mathcal{T}(\tau)$, where $\mathcal{T}(\tau)$ gives the parameter value of the foot of the normal through the point $\ell_Q^{\text{tr}}(\tau)$, we clearly have

$$\theta_j = \mathcal{T}(\tau_j); \quad p - K_\ell \leq j \leq p + K_r. \quad (3.9)$$

It follows that a 1D interpolation procedure on the function $\mathcal{T}(\tau)$ can be used to obtain the desired approximation of the value $\theta^Q = \mathcal{T}(0)$ of the parameter corresponding to the foot of the point $Q = \mathbf{z}_{i,j}$: $\mathcal{F}(Q) = \mathbf{q}(\theta^Q) = \mathbf{q}(\mathcal{T}(0))$.

Once we have the corresponding foot parameter value θ^Q for the given Cartesian point Q , the distance η_Q of the point Q to the boundary Γ is easily computed. Let $S_{\theta Q} \subset I_B$ be a set of M boundary parameter values, similar to the set S_{θ_p} defined in (3.8), but with $S_{\theta Q}$ “re-centered” around θ^Q . In order to obtain the continuation function value at the point Q using the continuation function values on V_{B,C_r}^+ , we employ a local two-dimensional polynomial interpolation scheme based on the set $S_{\theta Q}$ and the distance η_Q , as indicated in what follows and illustrated in the right image of Figure 4. First, the continuation values are obtained at each point (marked by blue asterisks in the figure) at a distance η_Q from the boundary Γ along the normal grid lines in V_{B,C_r}^+ that correspond to boundary parameter values in the set $S_{\theta Q}$. Each one of these values is obtained via one-dimensional interpolation of the continuation function values on V_{B,C_r}^+ along the corresponding normal grid line in V_{B,C_r}^+ . The desired continuation value at the point Q , then, is obtained via a final one-dimensional degree- $(M-1)$ interpolation step based on the parameter set $S_{\theta Q}$ and values at the “blue-asterisk” points just obtained. Finally, by applying the two-dimensional FFT to

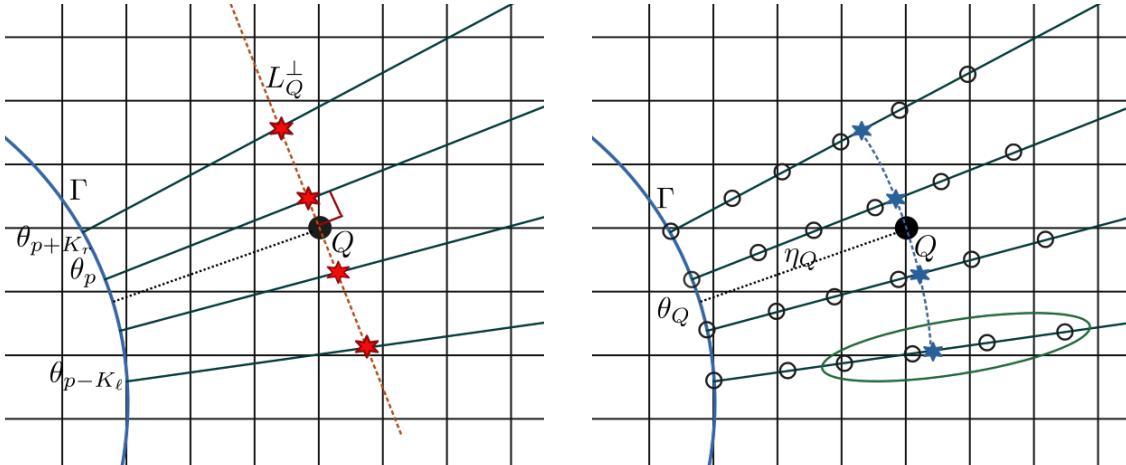


Figure 4: Interpolation schemes utilized to obtain the continuation function values on $H \cap V^+$ on the basis of the continuation values on V_{B,C_r}^+ . Left: Evaluation of the boundary parameter value for the foot of the normal line passing through Q , depicted as a finely dotted line passing through that point. The left image also displays the set of red-star interpolation points along the dashed-orange line L_Q^{\perp} . Right: Interpolation of continuation values from the curvilinear mesh to a point Q on the Cartesian mesh.

the continuation function values computed above we then obtain the desired Fourier series expression in (3.1) for the continuation function.

Remark 3.1 (Function values on Γ). I. For definiteness, in this paper we have assumed that the boundary data is provided in the form of values of the given function—which corresponds to the the Dirichlet boundary data in the PDE context. But the approach is also applicable in cases for which the boundary data is given as the normal derivative of the function, (Neumann boundary data), or even a combination of function and normal derivative values (Robin data) by relying on a slightly modified blending-to-zero procedure of the type presented in [4]. II. If no boundary data available, the two-dimensional Fourier

continuation method can still be utilized on the basis of interior data only, albeit with a certain reduction in accuracy near the boundary.

Remark 3.2 (Extra vanishing values in 2D). As in the 1D case, prior to the FFT procedure the grid H can be enlarged, with vanishing function values assigned to the added discretization points to obtain a discretization containing a number of discretization points equal to a power of two (or a product of powers of small prime numbers) along each Cartesian direction, which leads to specially fast evaluations by means of the fast Fourier transform.

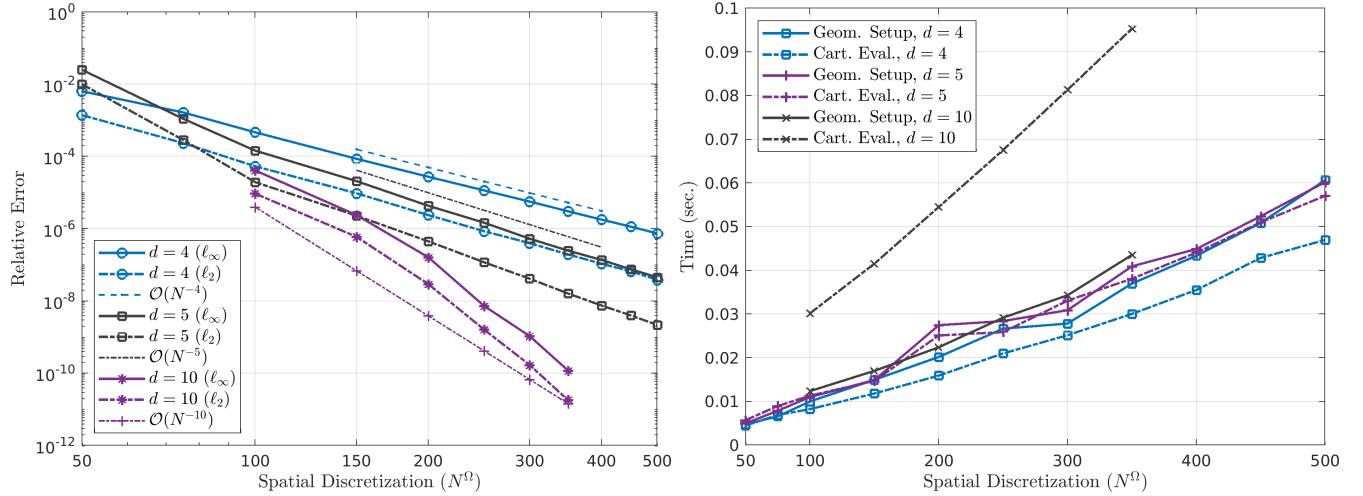


Figure 5: Numerical errors in log-log scale (left graph) and computing times required (right graph) in the 2D-FC approximation of the function f in (3.10) in the setting of Example 3.1. The interpolating polynomial degree $M = d + 3$ was used in all cases. The integer N^Ω equals the number of spatial grid points used over the diameter of the disc. Times reported correspond to averages over 10 runs.

3.5 Summary of the 2D-FC procedure

This section presents a summary of the 2D-FC procedure described in the Sections 3.1 through 3.4 for a function f given on a uniform Cartesian grid $H \cap \Omega$ within the domain of definition Ω , where H is a Cartesian mesh over the rectangle \mathcal{R} containing both Ω and the near-boundary outer region V^+ ; see Section 3.1 and, in particular, Figure 2. The construction of the continuation function f^c for the given function f relies on use of three main parameters associated with the 1D blending-to-zero approach presented in Section 2, namely d (number of points in the set \mathcal{D}), C (number of unrefined discrete continuation points), and n_{os} (oversampling factor for the 1D blending-to-zero FC procedure), together with the parameters n_r (refinement factor for the discrete continuation points, Remark 2.2 and Section 3.1), and, $M - 1$ (degree of the interpolating polynomials, Sections 3.2 and 3.4). Additionally, the 2D-FC procedure utilizes the precomputed matrices \mathbf{A}_r and \mathbf{Q} , which, with reference to Section 2.2, are obtained as per the description provided in Remark 2.2.

Using the aforementioned parameters and matrices, the algorithm proceeds in two main steps, namely step (a) A “*Geometrical Setup*” precomputation procedure (comprising points 1. through 4. below); and step (b) A “*Cartesian Evaluation*” procedure (comprising points 5. through 7. below). Part (a) only concerns geometry, and, for a given domain and configuration, could be produced once and used for evaluation of Fourier continuations for many different functions, as is often necessary in practice. The full 2D-FC algorithm thus proceeds according to the following seven-steps:

1. Discretize the boundary Γ using a smooth parametrization $\{\mathbf{q}(\theta) = (x(\theta), y(\theta)) : 0 \leq \theta \leq 2\pi\}$ of Γ and the uniform discretization $I_B = \{0 = \theta_0 < \theta_1, \dots, < \theta_{B-1} < 2\pi\}$ of the interval $[0, 2\pi]$ (Section 3.1).
2. Using the discretization I_B , construct two curvilinear meshes $V_{B,d}^-$ and V_{B,C_r}^+ in the near-boundary regions V^- (within $\bar{\Omega}$) and V^+ (outside Ω), respectively (equations (3.5) and (3.6)). Note that the discrete boundary points $\mathbf{q}(\theta_j)$ with $\theta_j \in I_B$ are common to both $V_{B,d}^-$ and V_{B,C_r}^+ .
3. Determine the set $H \cap V^+$ of Cartesian grid points and construct the proximity map $\mathcal{P} : H \cap V^+ \rightarrow V_{B,C_r}^+$ (Section 3.3).
4. For all $Q \in H \cap V^+$ obtain the parameter value $\theta^Q \in [0, 1]$ of the foot $\mathcal{F}(Q)$ of the normal through Q (Section 3.4 and the left image in Figure 4).
5. For each normal grid line (inherent in $V_{B,d}^-$ and V_{B,C_r}^+) given by the discretization I_B , compute the blending-to-zero function values along that normal (Section 3.2).
6. For all the points $Q \in H \cap V^+$, obtain the continuation function value at Q by local 2D interpolation (Section 3.4 and the right image in Figure 4).
7. Apply the two-dimensional FFT once to the continuation function values to obtain the desired Fourier series in (3.1).

3.6 2D-FC approximation: Numerical Results

This section demonstrates the accuracy and efficiency of the proposed 2D-FC method. Use of the 2D-FC method requires selection of specific values for each one of the following parameters (all of which are introduced in Sections 2 and 3):

- d : number of points in the boundary section (Section 2.2).
- C : number of continuation points (Section 2.2).
- Z : number of zero matching points (Section 2.2).
- n_{os} : oversampling factor used in the oversampled matching procedure (Section 2.2).
- n_r : refinement factor along the normal directions in V_{B,C_r}^+ (Remark 2.2).
- R : the smallest rectangle containing $\Omega \cup V^+$ (Section 3).
- $N = N_x \times N_y$: number of points in the uniform spatial grid H (Section 3).
- B : number of points in the boundary discretization (Section 3).
- $M - 1$: interpolating polynomial degree (Sections 3.2 and 3.4).

All the errors reported in this section were computed on a Cartesian grid of step size $h/2$ within Ω . In all of the numerical examples considered in this article the parameter selections were made in accordance with Remark 3.3. The computer system used, in turn, is described in Remark 3.4.

Remark 3.3 (Parameter selections). For a given step-size h in the two-dimensional Cartesian grid H , the normal and the boundary step-sizes k_1 and k_2 (Section 3) were taken to coincide with h : $k_1 = k_2 = h$. The parameter values $C = 27$, $n_{os} = 20$, $Z = 12$ and $n_r = 6$ were used in the evaluation of the matrices \mathbf{A}_r and \mathbf{Q} (see Section 2.2 and Remark 2.2). And, finally, with exception of the interpolation-degree experiments presented in Example 3.2 (Table 2) and Example 4.2 (Table 3), the interpolating-polynomial degree $(M - 1) = (d + 2)$ was used for the various matching-point numbers d considered.

Remark 3.4 (Computer system). All of the numerical results reported in this paper were run on a single core of a 3.40 GHz Intel Core i7-6700 processor with 15.4Gb of 2133 MHz memory.

Example 3.1 (Performance and efficiency of the 2D-FC method). In our first example we consider a problem of FC approximation of the function $f : \Omega \rightarrow \mathbb{R}$ given by

$$f(x, y) = -\sin(5\pi x) \sin(5\pi y) \quad (3.10)$$

h	N^Ω	B	T_P	T_F	T_V	T_H	Rel. Err. (ℓ_∞)
$2 \cdot 10^{-2}$	100	313	$2.1 \cdot 10^{-3}$	$2.6 \cdot 10^{-3}$	$1.7 \cdot 10^{-3}$	$6.5 \cdot 10^{-3}$	$4.7 \cdot 10^{-4}$
$1 \cdot 10^{-2}$	200	628	$4.2 \cdot 10^{-3}$	$5.2 \cdot 10^{-3}$	$3.8 \cdot 10^{-3}$	$1.2 \cdot 10^{-2}$	$2.7 \cdot 10^{-5}$
$5 \cdot 10^{-3}$	400	1250	$1.0 \cdot 10^{-3}$	$1.2 \cdot 10^{-2}$	$1.0 \cdot 10^{-4}$	$2.5 \cdot 10^{-2}$	$1.8 \cdot 10^{-6}$
$2 \cdot 10^{-2}$	100	313	$2.1 \cdot 10^{-3}$	$2.9 \cdot 10^{-3}$	$2.2 \cdot 10^{-3}$	$9.0 \cdot 10^{-3}$	$1.4 \cdot 10^{-4}$
$1 \cdot 10^{-2}$	200	628	$5.9 \cdot 10^{-3}$	$6.9 \cdot 10^{-3}$	$5.8 \cdot 10^{-3}$	$1.9 \cdot 10^{-2}$	$4.3 \cdot 10^{-6}$
$5 \cdot 10^{-3}$	400	1250	$1.1 \cdot 10^{-2}$	$1.3 \cdot 10^{-2}$	$1.2 \cdot 10^{-2}$	$3.1 \cdot 10^{-2}$	$1.4 \cdot 10^{-7}$
$2 \cdot 10^{-2}$	100	313	$2.2 \cdot 10^{-3}$	$3.9 \cdot 10^{-3}$	$4.2 \cdot 10^{-3}$	$2.6 \cdot 10^{-2}$	$4.1 \cdot 10^{-5}$
$1 \cdot 10^{-2}$	200	628	$4.3 \cdot 10^{-3}$	$6.6 \cdot 10^{-3}$	$8.8 \cdot 10^{-3}$	$4.7 \cdot 10^{-2}$	$1.6 \cdot 10^{-7}$

Table 1: Times (in sec.) required by the various tasks in the 2D-FC algorithm in the setting of Example 3.1. The times reported were calculated as time-averages over 10 runs. The integer N^Ω equals the number of spatial grid points used over the diameter of the disc.

on the unit disc $\Omega = \{(x, y) \in \mathbb{R}^2 : x^2 + y^2 \leq 1\}$. The left graph in Figure 5 displays the relative ℓ_∞ and ℓ_2 errors, in log-log scale, obtained from 2D-FC approximations of the function f , for three different values of polynomial degree d defined in Section 3.2, namely, $d = 4$, $d = 5$ and $d = 10$ —demonstrating the respective fourth, fifth and tenth orders of convergence expected. Higher rates of convergence, which are useful in some cases, can be achieved by using higher values of d , as demonstrated in the context of the Poisson solver in Section 4.1.1. The corresponding computing costs, including “Geometrical Setup” cost as well as the “Cartesian Evaluation” cost are presented in the right graph of Figure 5. The Geometrical setup cost combines the setup time for the grids (Section 3.1) $V_{B,d}^-$, V_{B,C_r}^+ and H ; the time T_P required for construction of the proximity map \mathcal{P} (Section 3.3); and the time T_F required for evaluation of the foot of the normal for all points in $H \cap V^+$ (Section 3.4). The Cartesian Evaluation time, in turn, equals the sum of the time T_V required for evaluation of the f^c values on the curvilinear grid V^+ and time T_H required for subsequent interpolation onto the Cartesian grid H . Table 1 reports additional details concerning computing times required by various tasks associated with the Geometrical Setup and Cartesian Evaluation for this example, including the times T_P , T_F , T_V and T_H . The Cartesian-interpolation time T_H dominates the overall Cartesian Evaluation step (cf. Table 1). In all of these cases we see that the computation time grows linearly with $1/h$. Also, the slope of the Cartesian evaluation cost depends on the degree d whereas the Geometrical setup cost, which remains similar in all the cases considered in this example, depends mainly on B and the refinement factor n_r .

Remark 3.5 (Use of higher degree Gram polynomials). Comparison of the various accuracy and timing values reported in Figure 5 suggests that use of lower 2D FC degrees such as $d = 4$ or $d = 5$ may provide the highest efficiency for approximation accuracies up to single precision.

Example 3.2 (Interpolation degree ($M-1$) for a given 2D-FC order d). Our next example concerns the approximation of the trigonometric function

$$f = -(x^6 + y^6) \sin(10\pi x) \sin(10\pi y), \quad (3.11)$$

defined over the kite shaped domain contained within the curve given by $x(\theta) = \cos(\theta) + 0.35 \cos(2\theta) - 0.35$; $y(\theta) = 0.7 \sin(\theta)$ for $0 \leq \theta \leq 2\pi$. A convergence study for this test case is presented in Table 2 for the values $d = 4$ and $d = 5$, and with $M = d + 1$, $M = d + 2$ and $M = d + 3$. Clearly, the selections $M = d + 2$ and $M = d + 3$ provide similar accuracy in most of the 2D-FC approximation cases considered. The value $M = d + 3$, which yields somewhat better interpolation accuracy for larger step-sizes (e.g. $h = 10^{-2}$), and which, as illustrated in Table 3, gives rise to some improvements for all step-sizes in the Poisson-problem applications considered in Section 4.1, was used in all of the numerical experiments presented in

this article (except for the cases specifically designed to test the dependence of the accuracy on variations of the parameter M).

h	N_x	N_y	$M = d + 1$		$M = d + 2$		$M = d + 3$	
			Abs. Err.	Order	Abs. Err.	Order	Abs. Err.	Order
$1 \cdot 10^{-2}$	297	197	$1.8 \cdot 10^{-3}$	—	$1.4 \cdot 10^{-3}$	—	$9.2 \cdot 10^{-4}$	—
$5 \cdot 10^{-3}$	537	337	$2.0 \cdot 10^{-4}$	3.2	$9.0 \cdot 10^{-5}$	3.9	$3.1 \cdot 10^{-5}$	4.9
$2.5 \cdot 10^{-3}$	1017	617	$1.0 \cdot 10^{-5}$	4.3	$4.6 \cdot 10^{-6}$	4.3	$2.3 \cdot 10^{-6}$	3.8
$1.25 \cdot 10^{-3}$	1977	1177	$6.2 \cdot 10^{-7}$	4.0	$1.4 \cdot 10^{-7}$	5.0	$1.4 \cdot 10^{-7}$	4.0
$6.25 \cdot 10^{-4}$	3897	2297	$3.4 \cdot 10^{-8}$	4.2	$9.0 \cdot 10^{-9}$	4.0	$9.0 \cdot 10^{-9}$	4.0
$d = 4$								
$1 \cdot 10^{-2}$	297	197	$4.4 \cdot 10^{-3}$	—	$2.3 \cdot 10^{-3}$	—	$2.5 \cdot 10^{-4}$	—
$5 \cdot 10^{-3}$	537	337	$3.3 \cdot 10^{-4}$	3.7	$4.2 \cdot 10^{-5}$	5.8	$1.5 \cdot 10^{-5}$	4.1
$2.5 \cdot 10^{-3}$	1017	617	$1.8 \cdot 10^{-5}$	4.2	$4.3 \cdot 10^{-7}$	6.6	$2.6 \cdot 10^{-7}$	5.8
$1.25 \cdot 10^{-3}$	1977	1177	$4.1 \cdot 10^{-7}$	5.4	$6.4 \cdot 10^{-9}$	6.1	$4.1 \cdot 10^{-9}$	6.0
$6.25 \cdot 10^{-4}$	3897	2297	$1.2 \cdot 10^{-8}$	5.0	$1.3 \cdot 10^{-10}$	5.6	$1.3 \cdot 10^{-10}$	5.0
$d = 5$								

Table 2: Convergence table for the 2D-FC method in the setting of Example 3.2. For both $d = 4$ and $d = 5$, we observe the expected fourth and fifth orders of convergence, respectively, for all the three choices of M , namely, $M = d + 1$, $M = d + 2$ and $M = d + 3$. The value $M = d + 3$ leads to somewhat improved accuracy.

Example 3.3 (Graphical illustration of the 2D-FC method). Figure 6 demonstrates the 2D-FC extension method for the function defined by

$$f(x, y) = 4 + (1 + x^2 + y^2)(\sin(2.5\pi x - 0.5) + \cos(2\pi y - 0.5)),$$

over the kite shaped domain considered in Example 3.2. Both the original function and its extension are presented in Figure 6.

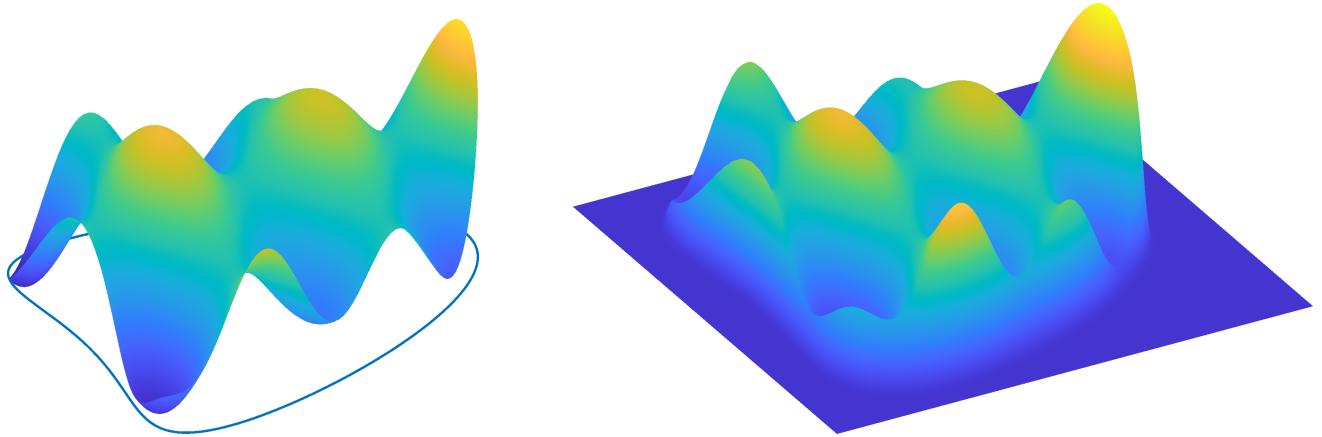


Figure 6: Demonstration of the 2D-FC procedure. Continuation of an oscillatory function defined on a kite shaped domain, as detailed in Example 3.3. The left and right images display the original and 2D-FC function values, respectively. The blue curve in the left image indicates the boundary of the kite-shaped domain Ω . Note the narrowness of the region wherein the transition to zero takes place.

4 Applications of the 2D-FC method

The 2D-FC method introduced in Section 3 can be used to facilitate spectral treatment in cases for which iterated use of the 1D Fourier expansions does not suffice, but for which use of full 2D Fourier expansions is beneficial. Sections 4.1 and 4.2 briefly describe two such cases, one concerning solution of the Poisson problem via Fourier expansions, and, the other, the solution of the wave equation via a novel Fourier Forwarding (FF) technique. A few numerical examples are presented illustrating the character of the resulting numerical solvers.

4.1 Application Example I: Poisson Problem

In this section we present a 2D-FC based method for the solution of the 2D Dirichlet Poisson problem

$$\begin{cases} \Delta u(x, y) = f(x, y), & (x, y) \in \Omega \\ u(x, y) = g(x, y), & (x, y) \in \Gamma; \end{cases} \quad (4.1)$$

the corresponding problem under Neumann or Robin boundary conditions can be treated similarly. Only a brief sketch is presented here, as an illustration of the 2D-FC approach. A complete description of the method, including a detailed geometrical treatment needed for reliable evaluation of the solution near boundaries, is presented in [10].

The proposed Poisson solver obtains the solution u within the prescribed tolerance as a sum

$$u = u_p + v \quad (4.2)$$

of a “particular solution” u_p , produced by means of the 2D-FC method, which satisfies the Poisson equation $\Delta u = f$ (but which generically does not satisfy the boundary conditions), and a solution v of the “homogeneous problem”, produced by means of a boundary integral equation, which satisfies the Dirichlet boundary value problem for Laplace’s equation

$$\begin{cases} \Delta v(x, y) = 0, & (x, y) \in \Omega \\ v(x, y) = g_{\text{hom}}(x, y), & (x, y) \in \Gamma, \end{cases} \quad (4.3)$$

where

$$g_{\text{hom}}(x, y) = g(x, y) - u_p(x, y)|_{\Gamma}. \quad (4.4)$$

A particular solution u_p for the problem (4.1) can easily be obtained from the 2D-FC expansion $f^c(x, y)$ of the right-hand function $f(x, y)$ (equation (3.1))—in view of the diagonal character of the Laplace operator in Fourier space. We thus obtain

$$u_p(x, y) = -\hat{f}_{0,0}^c(x^2 + y^2)/4 + \sum_{\ell=-N_x/2+1}^{N_x/2} \sum_{m=-N_y/2+1}^{N_y/2} b_{\ell,m} e^{2\pi i \left(\frac{\ell x}{L_x} + \frac{m y}{L_y} \right)} \quad (4.5)$$

where

$$b_{\ell,m} = \begin{cases} 0, & \text{if } (\ell, m) = (0, 0) \\ \frac{-\hat{f}_{\ell,m}^c}{(2\pi\ell/L_x)^2 + (2\pi m/L_y)^2}, & \text{if } (\ell, m) \neq (0, 0), \end{cases} \quad (4.6)$$

where one of a variety of possible selections was made for the constant Laplacian term. In view the asymptotically small factors that relate the Fourier coefficients $b_{\ell,m}$ to the original FC coefficients $\hat{f}_{\ell,m}^c$ it can be shown [10] that, as illustrated in Section 4.1.1, the rate of convergence in the overall numerical 2D-FC based solution u is of $\mathcal{O}(h^{d+2})$ if a 2D-FC algorithm of $\mathcal{O}(h^d)$ is utilized to compute the particular

solution u_p (provided a sufficiently accurate method is subsequently used for evaluation of the homogeneous solution).

Values $u_p(\mathbf{r})$ of the particular solution at points \mathbf{r} on the boundary Γ are required as an input (via (4.4)) in the boundary-value problem (4.3) for the Laplace solution v . It is therefore necessary to utilize an efficient method for evaluation of u_p at points \mathbf{r} that are not part of the Cartesian mesh H . The straightforward procedure based on direct addition of all terms in (4.5) for each discretization point on Γ does not match the optimal $\mathcal{O}(N \log(N))$ cost asymptotics enjoyed by all the other elements of the algorithm and is therefore avoided. Instead, the proposed algorithm first obtains the values of the right hand side of (4.5) for all $\mathbf{r} \in H$ via a direct application of the FFT algorithm, and, then, using these values, it produces the values for $\mathbf{r} \in \Gamma$ via iterated one-dimensional interpolation, as described in [21, Sec. 3.6.1]. In order to match the overall order $(d + 2)$ accuracy of the overall Poisson solution, one-dimensional polynomial interpolants of degree $(M_P - 1) \geq (d + 2)$ (cf. Example 4.2) are used in this context for both the x and y interpolation directions.

The numerical solution of the Laplace equation in (4.3), in turn, can be obtained rapidly and efficiently on the basis of the boundary integral method (see e.g. [18]). Relying on the boundary parametrization (3.2), the proposed algorithm incorporates an integral equation with a smooth kernel together with the simple and effective Nyström algorithm presented in [18, Sec. 12.2]. Based on trapezoidal-rule quadrature, this algorithm results in highly accurate solutions: in view of the periodicity and smoothness of the solution and the kernel, the approach yields super-algebraically small errors provided the boundary and right-hand side g_{hom} are both smooth. The associated linear system is solved by means of the iterative linear algebra solver *GMRES* [22]. Note that the integrand exhibits a near singular behavior for evaluation points that are near the boundary Γ but that are not on Γ . In order to address this difficulty, the proposed method uses a scheme (some elements of which were introduced in [2, 7]) which, based on local mesh refinement and subsequent interpolation using polynomial of degree $(M_P - 1)$, successfully resolves this difficulty. A detailed description of this and other aspects concerning the 2D-FC based Poisson solver is presented in the forthcoming contribution [10].

Once the particular and homogeneous solutions u_p and v have been obtained, the solution u of the Poisson problem is given by (4.2). The numerical convergence rate of the solution produced by the algorithm is mainly determined by the order d of the 2D-FC algorithm used. In all, the method is fast and highly accurate; a few illustrations, including accuracy and timing comparisons with leading solvers, are presented in the following section.

4.1.1 Numerical Illustrations for the Poisson Problem

The numerical illustrations presented in this section utilize the 2D-FC parameter selections presented in Remark 3.3, with various choices of the order parameter d . In addition, the size of the uniform boundary discretization used by the trapezoidal-rule based Nyström method is taken, for simplicity, to equal N_x —but, of course, in view of the super-algebraic convergence of the trapezoidal-rule quadrature, a smaller discretization size could have been used without sacrificing accuracy. The ℓ_2 and ℓ_∞ errors reported in this section were computed over the Cartesian grid $H \cap \Omega$ unless indicated otherwise. The first Poisson-solver example concerns a simple problem previously considered in [13].

Example 4.1 (High-order 2D-FC based Poisson solution). We consider the Poisson problem 4.1 in the domain $\Omega = \{(x, y) \in \mathbb{R}^2 : x^2 + y^2 \leq 1\}$ with $f = -\sin(2\pi x) \sin(2\pi y)$. The left portion of Figure 7 presents the numerical errors in the solutions produced by the 2D-FC based Poisson solvers for $d = 4$, $d = 6$ and $d = 8$ for $f = -\sin(2\pi x) \sin(2\pi y)$. The observed rates of convergence for all the three cases match the expected increased rates of convergence, as discussed in Section 4.1, that is, rates convergence of orders 6, 8 and 10, respectively. This problem was also considered in [13]. Comparison of the results presented in [13, Fig. 8] and those on the left graph in Figure 7 suggests the 2D-FC based Poisson solver performs favorably for high accuracies. For instance, a number $N^\Omega = 100$ of spatial grid points over the diameter of

Ω , that is to say, $N_x = 154$ grid points over one length of the rectangular computational domain, provides, as shown in the Figure 7, an ℓ_2 error $1.4 \cdot 10^{-12}$ whereas, in [13, Fig. 8], a similar discretization provides ℓ_2 errors close to 10^{-9} . The 10^{-12} error is achieved in that reference at a number of approximately 275 points in spatial discretization points in each spatial direction. For a different test case in this problem setting we now take $f = -\sin(5\pi x)\sin(5\pi y)$ (a function that was also used for the convergence study of the 2D-FC algorithm as presented in Example 3.1), and we report, on the right graph in Figure 7, the numerical errors in the solution produced by the solvers for higher values of d , namely, $d = 10$ and $d = 12$. Once again the expected convergence rates (in this case, of orders 12 and 14, respectively) are observed in practice.

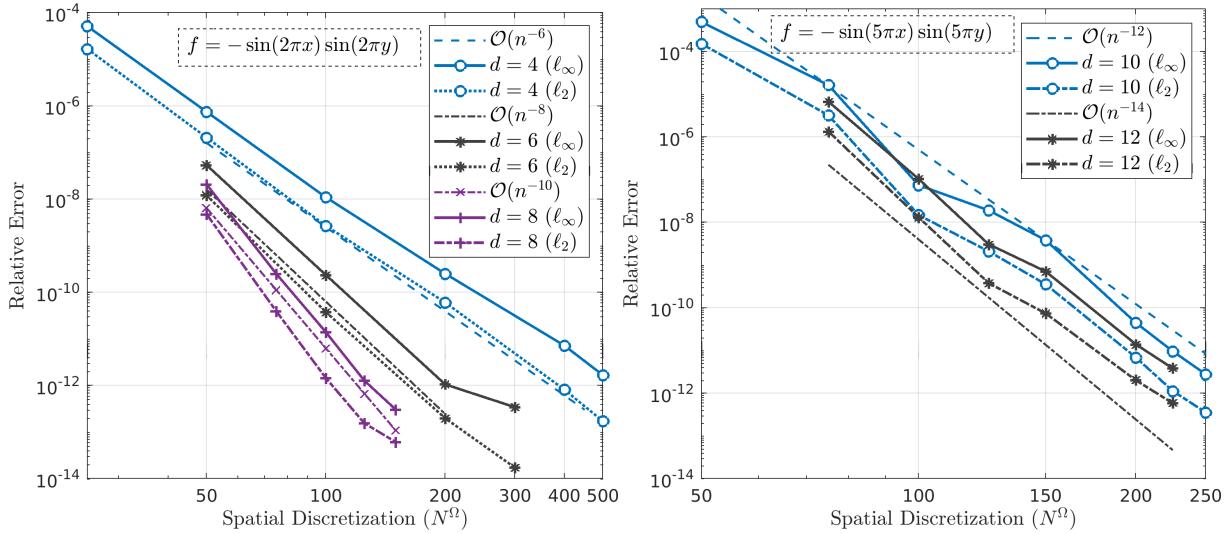


Figure 7: Numerical solution errors, in log-log scale, resulting from use of the 2D-FC based Poisson solver over the unit disc (Example 4.1). Left graph: solution errors for 2D-FC methods of orders $d = 4$, $d = 6$ and $d = 8$ for the function $f = -\sin(2\pi x)\sin(2\pi y)$. Right graph: solution errors with $d = 10$ and $d = 12$ for the function $f = -\sin(5\pi x)\sin(5\pi y)$. The parameter N^Ω denotes the number of spatial grid points used over the diameter of the disc. A total of $N_x = N_y = N^\Omega + 2C$ points ($C = 27$) were used over each dimension of the periodic square R . As discussed in the text, a convergence rate of order $(d+2)$ results from use of a 2D-FC approximation of order d .

h	N_x	N_y	$M = M_P = d + 1$		$M = M_P = d + 2$		$M = M_P = d + 3$	
			Rel. Err.(ℓ_2)	Order	Rel. Err.(ℓ_2)	Order	Rel. Err.(ℓ_2)	Order
$4 \cdot 10^{-2}$	117	93	$6.0 \cdot 10^{-6}$	—	$7.5 \cdot 10^{-7}$	—	$4.1 \cdot 10^{-7}$	—
$2 \cdot 10^{-2}$	177	127	$1.7 \cdot 10^{-7}$	5.1	$1.0 \cdot 10^{-8}$	6.2	$6.3 \cdot 10^{-9}$	6.0
$1 \cdot 10^{-2}$	297	197	$5.2 \cdot 10^{-9}$	5.0	$1.5 \cdot 10^{-10}$	6.1	$8.7 \cdot 10^{-11}$	6.2
$5 \cdot 10^{-3}$	537	337	$1.6 \cdot 10^{-10}$	5.1	$2.6 \cdot 10^{-12}$	5.9	$1.2 \cdot 10^{-12}$	6.2
$4 \cdot 10^{-3}$	657	407	$5.2 \cdot 10^{-11}$	4.9	$6.8 \cdot 10^{-13}$	6.0	$3.0 \cdot 10^{-13}$	6.2

Table 3: Convergence of the 2D-FC based solution of the Poisson problem described in Example 4.2.

Example 4.2 (Poisson solution interpolation degree ($M_P - 1$)). Once again we consider the problem (4.1) over a kite shaped domain as in Example 3.2 with $f = -\sin(2\pi x)\sin(2\pi y)$. The errors in the solutions produced by the 2D-FC based Poisson solver and the corresponding convergence rates for $d = 4$ and three different values of $M_P (= M)$, namely, $M_P = d + 1$, $M_P = d + 2$ and $M_P = d + 3$, are presented

in Table 3. The observed rates of convergence for $M_P = d+2$ and $M_P = d+3$ show the increased $d+2 = 6$ -th order convergence rate whereas the selection $M_P = d+1 = 5$ shows a fifth order convergence as the overall error in the Poisson solution is dominated by the error associated with the order-five interpolation process. The value $M_P = d+3$ is utilized for all Poisson-problem numerical results presented in this paper.

Example 4.3 (Highly oscillatory Poisson problem). Here we consider the problem (4.1) over the kite shaped domain considered in Example 3.2 with the highly oscillatory right hand side $f = -\sin(40\pi x)\sin(40\pi y)$. In this example, where we have used $d = 10$ for the 2D-FC particular-solution algorithm, the overall convergence rate, as reported in Table 4, is close to the expected convergence rate of order $(d+2)$. In order to avoid near-singular integration problems which arise, under the fine discretizations considered in this example, as the numerical solution is evaluated at points very near the boundary Γ , here we report the error at all the Cartesian points within the computational domain that are at a greater distance from Γ than 0.2. Near boundary evaluation algorithms capable of evaluating the solution at points arbitrarily close to Γ are presented in [10].

h	N_x	N_y	Rel. Err. (ℓ_∞)	Order	Rel. Err. (ℓ_2)	Order
$5 \cdot 10^{-3}$	537	337	$9.3 \cdot 10^{-4}$	—	$3.4 \cdot 10^{-4}$	—
$2.5 \cdot 10^{-3}$	1017	617	$6.0 \cdot 10^{-8}$	13.9	$1.7 \cdot 10^{-8}$	14.3
$1.25 \cdot 10^{-3}$	1977	1177	$5.8 \cdot 10^{-12}$	13.3	$1.6 \cdot 10^{-12}$	13.4

Table 4: Interior errors in the numerical solutions produced by the 2D-FC based Poisson solver of order $d = 10$, in the setting of Example 4.3. (Errors are computed over points lying at a distance 0.2 from the domain boundary; see Example 4.3.) A scaling-error even better than the expected order $(d+2) = 12$ was observed in this case.

4.2 Application Example II: Fourier Forwarding (FF) method for Wave propagation problems

This section presents the 2D-FC based Fourier-Forwarding method (FF) for the solution of the wave equation and other constant coefficient hyperbolic problems. Only a brief sketch of the FF approach, for demonstration purposes, is presented here; a detailed account of this methodology, including a variety of techniques designed for treatment of boundary conditions, and with application to linear hyperbolic systems, including, e.g., treatment of propagation in anisotropic media, is presented in the forthcoming contribution [9]. In what follows attention is restricted to the initial boundary value problem for the wave equation in 2D, namely

$$u_{tt} = c^2(u_{xx} + u_{yy}), \quad \text{for } (x, y, t) \in \Omega \times \mathbb{R}_+ \quad (4.7)$$

with initial conditions $u(x, y, t)|_{t=0} = f(x, y)$ and $u_t(x, y, t)|_{t=0} = g(x, y)$ for $(x, y) \in \Omega$, and with appropriate boundary conditions on Γ .

In order to obtain a solution of this problem, the FF method 1) Exploits the fact (also used in [15, 24] in the context of bi-periodic problems) that the solution of the wave equation in all of \mathbb{R}^2 with the initial data $u(\mathbf{r}, 0) = ae^{i\kappa \cdot \mathbf{r}}$ and $u_t(\mathbf{r}, 0) = be^{i\kappa \cdot \mathbf{r}}$ is itself given in closed form as a combination of two time-domain plane waves; and, 2) Constructs auxiliary solutions $U_{FC}(x, y, t; T_j)$ of the form

$$U_{FC}(x, y, t; T_j) = \sum_{\ell=-N_x/2+1}^{N_x/2} \sum_{m=-N_y/2+1}^{N_y/2} a_{\ell m}(t; T_j) e^{2\pi i \left(\frac{\ell x}{L_x} + \frac{m y}{L_y} \right)}, \quad (4.8)$$

of equation (4.7) utilizing the 2D-FC algorithm on certain initial solution values at various times: $T_j = j\Delta T$ ($j = 0, \dots, n$) with a “large” time step value ΔT (that is selected so as to optimize the overall computational

cost of the FF algorithm), and for an arbitrary user-prescribed positive integer n . In view of the limited domain of dependence of solutions of the wave equation [17], the auxiliary solution $U_{FC}(x, y, t; T_j)$, for $T_j \leq t \leq T_{j+1}$, provides a valid numerical approximation of the solution $u(x, y, t)$ over a certain subset $\Omega_{\Delta T} = \{\mathbf{r} = (x, y) \in \Omega : \text{dist}(\mathbf{r}, \Gamma) \geq c\Delta T\}$, away from the boundary Γ , of the domain Ω . To compute the solution $U_B(x, y, t; T_j)$ on the region $\Omega \setminus \Omega_{\Delta T}$ adjacent to the physical boundary, the FF method uses a classical time-stepping scheme, with spatial derivatives obtained by means of the 1D-FC method, and using a (typically much smaller) time step Δt , which should be small enough so to ensure stability (as dictated by the CFL condition) and to yield an accuracy level consistent with that inherent in the 2D-FC approximation used. In what follows we discuss the evaluation of the solution

$$u(x, y, t) = \begin{cases} U_{FC}(x, y, t; T_0) & \text{for } (x, y, t) \in \Omega_{\Delta T} \times [T_0, T_1], \\ U_B(x, y, t; T_0) & \text{for } (x, y, t) \in \Omega \setminus \Omega_{\Delta T} \times [T_0, T_1] \end{cases} \quad (4.9)$$

for the time interval $T_0 = 0 \leq t \leq T_1 = \Delta T$; a similar procedure can be used to evaluate, inductively, the solution at all other time intervals $T_j \leq t \leq T_{j+1}$. For notational simplicity the argument T_0 is suppressed in what follows.

In order to obtain the auxiliary solution $U_{FC}(x, y, t)$ the method utilizes the Cartesian grid H (equation (3.7)) together with the 2D-FC expansions

$$\begin{cases} F(x, y) = \sum_{\ell=-N_x/2+1}^{N_x/2} \sum_{m=-N_y/2+1}^{N_y/2} \hat{f}_{\ell m}^c e^{2\pi i \left(\frac{\ell x}{L_x} + \frac{my}{L_y} \right)}, \\ G(x, y) = \sum_{\ell=-N_x/2+1}^{N_x/2} \sum_{m=-N_y/2+1}^{N_y/2} \hat{g}_{\ell m}^c e^{2\pi i \left(\frac{\ell x}{L_x} + \frac{my}{L_y} \right)}, \end{cases} \quad (4.10)$$

(cf. equation (3.1)) of the initial data $f(x, y)$ and $g(x, y)$. (Note that while F and G are obtained from the given initial conditions in the present case $j = 0$, they are produced from the numerical values of the solution $u(x, y, t)|_{t=T_j}$ and its time derivative $u_t(x, y, t)|_{t=T_j}$ in the case $j > 0$.) Clearly, provided the functions $a_{\ell m}(t)$ in (4.8) satisfy the equations

$$\begin{cases} a''_{\ell m}(t) + \alpha_{\ell m} a_{\ell m}(t) = 0, \\ a_{\ell m}(t)|_{t=0} = \hat{f}_{\ell m}^c, \\ \frac{\partial a_{\ell m}(t)}{\partial t}|_{t=0} = \hat{g}_{\ell m}^c \end{cases} \quad (4.11)$$

for $-N_x/2 + 1 \leq \ell \leq N_x/2$ and $-N_y/2 + 1 \leq m \leq N_y/2$ (where $\alpha_{\ell m} = (2\pi c)^2 [(\ell/L_x)^2 + (m/L_y)^2]$), the function $U_{FC}(x, y, t)$ satisfies (4.7) as well as the initial conditions $U_{FC}(x, y, t)|_{t=0} = F(x, y)$ and $\frac{\partial U_{FC}(x, y, t)}{\partial t}|_{t=0} = G(x, y)$ for $(x, y) \in \mathbb{R}^2$. Substituting the explicit solutions

$$a_{\ell m}(t) = \begin{cases} \hat{f}_{\ell m}^c + t \hat{g}_{\ell m}^c, & \text{for } (\ell, m) = (0, 0) \\ \hat{f}_{\ell m}^c \cos(\alpha_{\ell m} t) + \frac{\hat{g}_{\ell m}^c}{\alpha_{\ell m}} \sin(\alpha_{\ell m} t), & \text{for } (\ell, m) \neq (0, 0). \end{cases} \quad (4.12)$$

of the ODE (4.11) into (4.8) the solution $U_{FC}(x, y, t)$ for $(x, y) \in \mathbb{R}^2$ is obtained. An application of the two-dimensional spatial inverse FFT over the Cartesian grid H to the coefficients $a_{\ell m}(T_1)$ then yields, per the discussion above concerning domains of dependence, a numerical approximation of the solution $u(x, y, t)$ for $T_0 \leq t \leq T_1$ and for all $(x, y) \in H \cap \Omega_{\Delta T}$. Note that the accuracy of the auxiliary solution in its domain of validity $\Omega_{\Delta T}$ at time t ($T_0 \leq t \leq T_1$), is only limited by the accuracy of the underlying 2D-FC approximation of the functions f and g by the FC expansions F and G throughout Ω , respectively.

To obtain the near boundary solution $U_B(x, y, t)$, on the other hand, the method uses a classical explicit time stepping scheme in a certain open set $\Omega_B \supset \Omega \setminus \Omega_{\Delta T}$ adjacent to the boundary. In the proposed near-boundary evaluation algorithm (which is based on use of the small time step Δt on a Cartesian grid on

Ω_B) various time stepping schemes, including the Adams-Bashforth [19] and the Taylor series [5] methods, can be utilized; the required spatial derivatives, in turn, are computed, with spectral accuracy and without dispersion, by means of the 1D-FC method (Section 2.2; cf. also [4, 8]). In the present $j = 0$ case the initial values for U_B are obtained from the initial conditions $f(x, y)$ and $g(x, y)$; for subsequent time intervals the solution process for U_B is simply continued forward in time: no additional initial values are needed for U_B at the start of the time intervals $[T_j, T_{j+1}]$ for $j > 0$. Boundary conditions for U_B must be enforced at all boundary points in $H \setminus \Omega_{\Delta T}$ —including those near Γ and those near $\partial\Omega_{\Delta T}$. The boundary condition at Cartesian points near Γ is enforced, with high-order accuracy, as proposed in [5], on the basis of a certain interpolation procedure which utilizes the given boundary values on Γ as well as previously obtained solution values on interior points of $H \cap \Omega$. The corresponding boundary condition at Cartesian points near $\partial\Omega_{\Delta T}$, on the other hand, are produced, for efficiency, by means of a special procedure [9] which avoids a full evaluation of the FC expansion for U_{FC} at each small- Δt time interval and each boundary point, and which constructs and uses, instead, solutions similar to (4.8) (without imposition of boundary conditions) but over some small square regions contained in Ω and centered at points on the interior boundary of $\Omega \setminus \Omega_{\Delta T}$.

Combining the 2D-FC forwarded solution U_{FC} and the near boundary solution U_B according to (4.9) the desired FF numerical approximation of the solution u throughout the Cartesian set $H \cap \Omega$ is thus obtained up to time $t = \Delta T$. Repeating this procedure as many times as necessary, the solution can be advanced up to $t = n\Delta T$ for arbitrarily large values of n , and, thus, up to an arbitrary final time T .

In view of the fact that auxiliary solutions $U_{\text{FC}}(x, y, t)$ need to be computed only once per large time step ΔT , a significant improvement in the asymptotic global computational cost per small time step Δt results over the cost required by classical finite-difference and finite-element spatial discretizations. Indeed, calling δ the thickness of the boundary region $\Omega \setminus \Omega_{\Delta T}$, letting $n_c = [\delta/h] > 0$ and assuming the Cartesian mesh $H \cap \Omega$ contains a total of $\mathcal{O}(N)$ discretization points, it follows that $\Omega \setminus \Omega_{\Delta T}$ contains a total of $\mathcal{O}(\sqrt{N}n_c)$ grid points. As shown in [9], the optimum value of n_c is $\mathcal{O}(N^{1/4})$, so that the computational cost per time step of the overall FF algorithm is $\mathcal{O}(N^{3/4} \log N)$ operations. As shown in that contribution, further, owing to a certain large multiplicative constant in front of the asymptotic cost estimate for the time-stepping portion in the boundary region $\Omega \setminus \Omega_{\Delta T}$, large increases in N are necessary for the optimal n_c value to increase by one or a few units. Thus, in the numerical examples considered in the present paper, for all of which we have $N \leq 4 \cdot 10^6$, the value of n_c is set to a constant. This selection leads to an overall cost estimate of approximately $\mathcal{O}(\sqrt{N})$ operations for the cases considered in this paper as the asymptotically large $\mathcal{O}(N \log N)$ FFT cost incurred by the algorithm has in fact a limited impact in such cases. A detailed discussion in this regard is presented in [9]. The performance of the resulting Fourier Forwarding method for a number of test cases is demonstrated in Section 4.2.1 below.

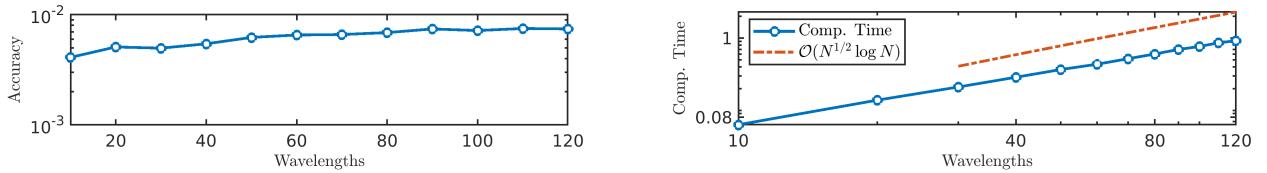


Figure 8: Accuracies and computing times (in seconds) for the problem considered in Example 4.4. Left graph: maximum absolute errors at time $t = 2$ for various wavelengths. Right graph: computing times required for each small- Δt time step. Fifteen points per wavelength and $n_c = 6$ points across the boundary region were used for this numerical experiment.

4.2.1 Numerical Examples for the Fourier Forwarding method

Two examples presented in this section demonstrate the character of the FF method. The numerical examples presented in this section were produced using Fourier continuations of order $d = 4$ for both the

1D-FC (for time stepping in the boundary region) and the 2D-FC (for interior FC-forwarding) algorithms. In both the cases, the 2D-FC parameter selections were made in accordance with Remark 3.3 and the simulations were run in the computer described in Remark 3.4. For time-stepping in the boundary region we have utilized second order Taylor series method with the CFL number 0.125.

Example 4.4 (Accuracy and efficiency of the FF method). In this example, we demonstrate the accuracy in the wave equation solution obtained via the FF method as well as other properties, namely, dispersionlessness and sublinear computing costs, enjoyed by the FF algorithm (Section 4.2) via the method of manufactured solutions. Here we consider the wave equation (4.7) with the initial and (Dirichlet) boundary conditions taken such a way that the exact solution is given by

$$u(x, y, t) = \cos(2\kappa(x + t)/3) + \cos(\kappa(y + t)), \quad (4.13)$$

on the unit disc $\Omega = \{(x, y) \in \mathbb{R}^2 : x^2 + y^2 \leq 1\}$. A fixed fifteen spatial points per wavelength, and a fixed number $n_c = 6$ of points across the boundary region ($\delta = n_c h$) have been utilized for this numerical experiment. In left graph of Figure 8, we report the maximum absolute error (computed over all the Cartesian grid points within Ω) of the solution, produced by the FF method, at time $t = 2$. The required computational times for each small- Δt time step for various spatial frequencies κ (cf. (4.13)) are reported in right graph of Figure 8. As discussed at the end of Section 4.2, we have observed an $\mathcal{O}(N^{1/2} \log N)$ growth in the computational time as the size N of the spatial discretization grows.

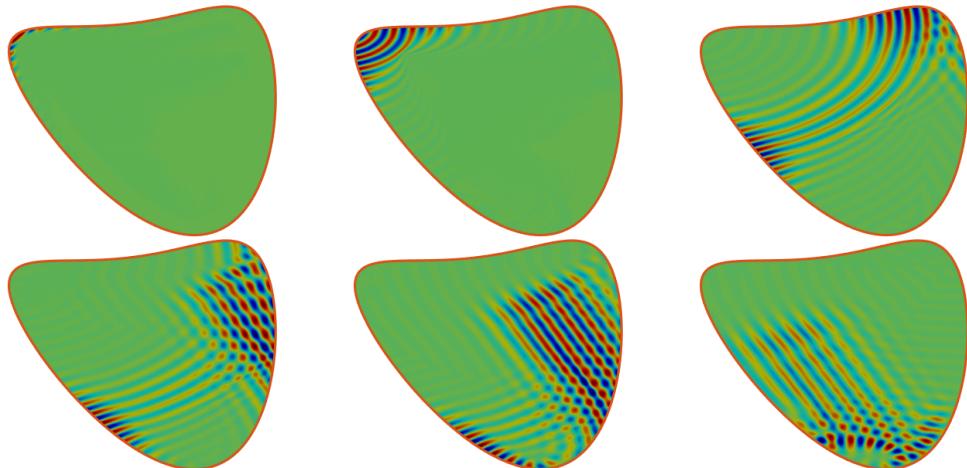


Figure 9: Fourier Forwarding method applied to an interior-domain wave propagation problem described in Example 4.5. The solution is shown, from top left to bottom right, after 100, 500, 2050, 2750, 3200 and 4300 timesteps.

Example 4.5 (Interior-domain wave scattering). For an interior-domain graphical demonstration of the FF method, we consider the wave equation problem with boundary condition

$$u = \cos(\kappa s) \exp(-s^2/\sigma^2) g(t), \quad (4.14)$$

where

$$s = (t - |\mathbf{r} - \mathbf{r}_0|)/|\mathbf{r} - \mathbf{r}_0|,$$

and

$$g(t) = \begin{cases} 0, & \text{for } t = 0, \\ 1 - \exp\left(2.0 \exp\left(\frac{-t_0/t}{t/t_0 - 1.0}\right)\right), & \text{for } t < t_0, \\ 1, & \text{for } t \geq t_0, \end{cases} \quad (4.15)$$

with $\mathbf{r}_0 = (-1.1, -0.72)$, $t_0 = 0.05$ and $\kappa = 20\pi$. The PDE domain is the interior of the kite-shaped curve considered in Example 3.2. A fixed number $n_c = 8$ of points across the boundary region and slightly over 20 points per wavelength have been utilized in this numerical experiment. Vanishing initial values of u and u_t at $t = 0$ were used; note that both the imposed boundary values (4.14) vanish for $t \leq t_0 > 0$. For clear visibility, a version of the computed solution values at certain selected times scaled by the maximum value at that specific point in time are displayed in the left portion of Figure 9. (The scaling values range approximately between 1 and 3, and they are nearly equal to 1 for the first four images, all three in the upper row, and the leftmost image in the lower row.)

5 Conclusions

This paper introduced a novel two-dimensional Fourier continuation (2D-FC) method, for bi-periodic extension of functions defined on arbitrary smooth two-dimensional domains. Applications to the Poisson and wave-equation problem, including the development of the Fourier Forwarding method, have resulted in numerical PDE solvers of high orders of accuracy, and, most importantly, of extremely low numerical dispersion. Extensions of these methodologies to problems in higher dimensions, and to problems on non-smooth domains, are left for future work.

Acknowledgments

This work was supported by NSF and DARPA through contracts DMS-1714169 and HR00111720035, and by the NSSEFF Vannevar Bush Fellowship under contract number N00014-16-1-2808.

References

- [1] B. ADCOCK, D. HUYBRECHS, AND J. MARTÍN-VAQUERO, *On the numerical stability of Fourier extensions*, Foundations of Computational Mathematics, 14 (2014), p. 635–687.
- [2] E. AKHMETGALIYEV, O. P. BRUNO, AND N. NIGAM, *A boundary integral algorithm for the Laplace Dirichlet–Neumann mixed eigenvalue problem*, Journal of Computational Physics, 298 (2015), pp. 1 – 28.
- [3] N. ALBIN AND O. P. BRUNO, *A spectral FC solver for the compressible Navier–Stokes equations in general domains I: Explicit time-stepping*, Journal of Computational Physics, 230 (2011), pp. 6248 – 6270.
- [4] F. AMLANI AND O. P. BRUNO, *An FC-based spectral solver for elastodynamic problems in general three-dimensional domains*, Journal of Computational Physics, 307 (2016), pp. 333–354.
- [5] D. APPELÖ AND N. A. PETERSON, *A fourth-order accurate embedded boundary method for the wave equation*, SIAM Journal of Scientific Computing, 34 (2012), pp. 2982–3008.
- [6] T. ASKHAM AND A. J. CERFON, *An adaptive fast multipole accelerated Poisson solver for complex geometries*, Journal of Computational Physics, 344 (2017), pp. 1 – 22.
- [7] O. P. BRUNO AND B. DELOURME, *Rapidly convergent two-dimensional quasi-periodic Green function throughout the spectrum—including Wood anomalies*, Journal of Computational Physics, 262 (2014), p. 262–290.
- [8] O. P. BRUNO AND M. LYON, *High-order unconditionally stable FC-AD solvers for general smooth domains I. Basic elements*, Journal of Computational Physics, 229 (2010), pp. 2009 – 2033.

- [9] O. P. BRUNO AND J. PAUL, *2D-FC based Fourier-Forwarding wave equation solver*, In Preparation.
- [10] ———, *2D-FC based Poisson solver*, In Preparation.
- [11] K. T. ELGINDY, *A high-order embedded domain method combining a predictor–corrector-Fourier-Continuation-Gram method with an integral Fourier pseudospectral collocation method for solving linear partial differential equations in complex domains*, Journal of Computational and Applied Mathematics, 361 (2019), pp. 372 – 395.
- [12] H. FENG AND S. ZHAO, *FFT-based high order central difference schemes for three-dimensional Poisson's equation with various types of boundary conditions*, Journal of Computational Physics, 410 (2020), p. 109391.
- [13] F. FRYKLUND, E. LEHTO, AND A.-K. TORNBERG, *Partition of unity extension of functions on complex domains*, Journal of Computational Physics, 375 (2018), pp. 57 – 79.
- [14] G. H. GOLUB AND C. F. V. LOAN, *Matrix Computations*, Matrix Computations, Johns Hopkins University Press, 2012.
- [15] I. HABER, R. LEE, H. KLEIN, AND J. BORIS, *Advances in electromagnetic simulation techniques*, in Proc. Sixth Conf. on Num. Sim. Plasmas, Berkeley, CA, 1973, pp. 46–48.
- [16] D. HUYBRECHS, *On the Fourier extension of nonperiodic functions*, SIAM Journal on Numerical Analysis, 47 (2010), p. 4326–4355.
- [17] F. JOHN, *Partial differential equations*, Applied Mathematical Sciences, (1982).
- [18] R. KRESS, *Linear Integral Equations*, Applied Mathematical Sciences, Springer Science and Business Media, third ed., 2014.
- [19] R. J. LEVEQUE, *Finite Difference Methods for Ordinary and Partial Differential Equations*, Society for Industrial and Applied Mathematics, 2007.
- [20] M. LYON AND O. P. BRUNO, *High-order unconditionally stable FC-AD solvers for general smooth domains II. Elliptic, parabolic and hyperbolic PDEs; theoretical considerations*, Journal of Computational Physics, 229 (2010), pp. 3358 – 3381.
- [21] W. H. PRESS, S. A. TEUKOLSKY, W. T. VETTERLING, AND B. P. FLANNERY, *Numerical Recipes 3rd Edition: The Art of Scientific Computing*, Cambridge University Press, USA, 3 ed., 2007.
- [22] Y. SAAD, *Iterative Methods for Sparse Linear Systems*, Society for Industrial and Applied Mathematics, second ed., 2003.
- [23] D. B. STEIN, R. D. GUY, AND B. THOMASES, *Immersed boundary smooth extension: A high-order method for solving PDE on arbitrary smooth domains using Fourier spectral methods*, Journal of Computational Physics, 304 (2016), pp. 252 – 274.
- [24] J.-L. VAY, I. HABER, AND B. B. GODFREY, *A domain decomposition method for pseudo-spectral electromagnetic simulations of plasmas*, Journal of Computational Physics, 243 (2013), pp. 260 – 268.