



# A closest point method library for PDEs on surfaces with parallel domain decomposition solvers and preconditioners

Ian C. T. May<sup>1</sup> · Ronald D. Haynes<sup>2</sup> · Steven J. Ruuth<sup>3</sup>

Received: 17 April 2022 / Accepted: 27 September 2022 / Published online: 2 November 2022  
© The Author(s) 2022

## Abstract

The DD-CPM software library provides a set of tools for the discretization and solution of problems arising from the closest point method (CPM) for partial differential equations on surfaces. The solvers are built on top of the well-known PETSc framework, and are supplemented by custom domain decomposition (DD) preconditioners specific to the CPM. These solvers are fully compatible with distributed memory parallelism through MPI. This library is particularly well suited to the solution of elliptic and parabolic equations, including many reaction-diffusion equations. The software is detailed herein, and a number of sample problems and benchmarks are demonstrated. Finally, the parallel scalability is measured.

**Keywords** Surface PDEs · Reaction-diffusion equations · Domain decomposition · Parallel linear solvers

---

✉ Ian C. T. May  
imay1@ucsc.edu

Ronald D. Haynes  
rhaynes@mun.ca

Steven J. Ruuth  
sruuth@sfu.ca

- <sup>1</sup> Department of Applied Mathematics, University of California Santa Cruz, 95064 Santa Cruz, CA, USA
- <sup>2</sup> Department of Mathematics, Memorial University of Newfoundland, A1C 5S7 St. John's, Newfoundland, Canada
- <sup>3</sup> Department of Mathematics, Simon Fraser University, V5A 1S6 Burnaby, British Columbia, Canada

## 1 Introduction

The numerical solution of PDEs intrinsic to surfaces presents interesting challenges over their flat-space analogs. The closest point method (CPM) is an approach to discretize general surface-intrinsic PDEs over a wide class of surfaces. The software library presented here provides an MPI parallel implementation of the CPM built on top of PETSc [1]. This software is particularly well suited to solving surface intrinsic reaction-diffusion systems, generically of the form

$$\left(\frac{\partial}{\partial t} - \mu_i \Delta_S\right) u_i = f_i(t, \mathbf{x}, \mathbf{u}), \quad (1)$$

where  $\mathbf{u}$  are the different species in the system,  $\mu_i$  are the diffusion coefficients, and  $f_i$  are the nonlinear reaction terms coupling the systems together.

This software provides access to the extensive suite of tools within PETSc for use with the CPM, and additionally defines custom Schwarz-type domain decomposition solvers/preconditioners for equations of the form

$$(c - \Delta_S)u = f, \quad (2)$$

where  $c \in \mathbb{R}^+$  is a positive constant, and  $\Delta_S$  is the Laplace-Beltrami operator intrinsic to an embeddable surface  $S \subset \mathbb{R}^d$ . These custom domain decomposition (DD) methods were first presented in [2], and subsequently expanded upon in [3]. Solving (1) semi-implicitly requires the solution of many elliptic systems. The provided custom DD preconditioners are well suited to these linear systems. Additionally, the flexibility of PETSc is maintained, and the use of other preconditioners and solvers within PETSc is straightforward.

To date, most software for the CPM has been limited to the use of direct solvers on shared memory machines. The DD-CPM software detailed here provides the first known implementation of this method that can leverage larger distributed memory machines. This facilitates the application of the CPM to larger and more complicated problems. For elliptic PDEs, the dependence on direct solvers that has been present in most CPM implementations has also been lifted through the introduction of custom DD preconditioned Krylov methods compatible with the CPM. For reaction-diffusion equations, this software allows the use of fine grids to reliably capture either complex surface geometry or sharp features in the generated solutions. Indeed, the software has been kept general to allow many other equations to be posed and solved.

### 1.1 Existing software

There is little existing software for the CPM. An important resource is the `cp_matrices` repository [4], hosted by Prof. Colin Macdonald (coauthor of references [5–10]). This repository consists of MatLab and Python codes, and is mostly restricted to serial execution. For problems outside our framework, we suggest to start with the `cp_matrices` repository.

A large motivator for developing CPM-specific DD methods was the experience that algebraic preconditioners did not perform well in all cases, and required a great deal of problem-specific tuning. However, there may be cases where the built-in PETSc preconditioners perform well with little cost per iteration [11]. The back-end preconditioners and solvers used with the DD-CPM library can be substituted with PETSc options easily, and without re-compiling the library. An important goal of the development of this software is to maintain a high degree of flexibility and interoperability with other software packages.

## 2 Review of core methods

Before discussing the software, we will briefly review the closest point method (CPM), and the (optimized) restricted additive Schwarz (ORAS) domain decomposition method. To keep these reviews concrete and simple, they will be constrained to (eq:shiftPoisson); the consideration of more complicated equations will be deferred to Section 3.1.

### 2.1 Closest point method

Methods for the numerical solution of surface intrinsic PDEs generally take one of two approaches: discretize the surface itself, or find a solution in a higher dimensional embedding space. Surface parametrization methods [12] are very efficient and allow the use of familiar discretizations, but are limited to simple surfaces where the parametrization is known. With these methods the user may need to contend with coordinate singularities. Finite element methods acting on a triangulation of the surface [13] yield sparse symmetric systems for model equation (eq:shiftPoisson), but are sensitive to the quality of the triangulation. Level set methods [14] embed the surface in a higher dimensional flat space, treating the surface implicitly; however, the requisite artificial boundary conditions on the computational domain and the treatment of open surfaces are non-trivial.

The CPM [5, 15] is an embedding method that uses an implicit representation of the surface similar to the level set method. When implemented, the CPM is posed only over a small region of the embedding space,  $\mathbb{R}^d$ , near the surface. This dramatically reduces the number of unknowns in the discretization. The CPM has been applied to triangulated surfaces [9], surfaces of mixed codimension [15], moving surfaces [16], and even point cloud domains [7]. The DD-CPM library supports triangulated surfaces directly. Surfaces of mixed codimension and point cloud domains are both supported by allowing user-defined closest point functions to be used. This software does not currently support moving surfaces.

In this section, the CPM is described for (eq:shiftPoisson) over a simple closed surface to keep the presentation brief. The CPM utilizes familiar Cartesian discretizations of differential operators for surface PDEs by first extending the solution off the surface to a narrow tube in the embedding space. The extended solution is

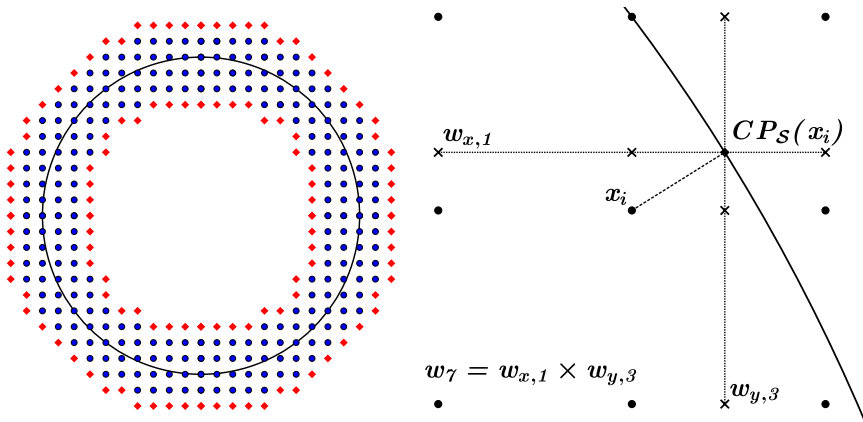
formed to be constant in the surface normal direction, and is obtained through the closest point function

$$\begin{aligned}
 cp_S : \mathbb{R}^d &\rightarrow \mathcal{S} \\
 \mathbf{x} &\mapsto \underset{y \in \mathcal{S}}{\operatorname{arg\,min}} \|\mathbf{x} - \mathbf{y}\|_2,
 \end{aligned}
 \tag{3}$$

which maps each point in the embedding space,  $\mathbf{x} \in \mathbb{R}^d$ , to the closest point on the surface. The extension operator  $E$  is defined as the composition of a function  $f : \mathcal{S} \rightarrow \mathbb{R}$  with  $cp_S$  over the embedding space, i.e.,  $(Ef)(\mathbf{x}) := f(cp_S(\mathbf{x}))$  for  $\mathbf{x} \in \mathbb{R}^d$ . Importantly, the restriction of the Laplacian of an extended function,  $\Delta Ef$ , to the surface  $\mathcal{S}$ , is equivalent to the Laplace-Beltrami operator acting on that surface bound function,  $\Delta_{\mathcal{S}} f$  [5, 8, 15]. The core idea of the CPM lies in discretizing  $\Delta E$  over an appropriate domain in the embedding space instead of discretizing  $\Delta_{\mathcal{S}}$  over the surface directly.

A structured grid with spacing  $h$  is placed over the embedding space, and points near the surface are collected into the set of active nodes  $\Sigma_A$ . The points neighboring the members of  $\Sigma_A$  are gathered into the set of ghost nodes  $\Sigma_G$ . The closest point function (eq:cpFunc) is continuous on a region of  $\mathbb{R}^d$  that lies within a distance  $\kappa_\infty^{-1}$  of the surface, where  $\kappa_\infty$  is an upper bound on the curvatures of  $\mathcal{S}$  [17]. The extent of the active nodes in  $\Sigma_A$  will be determined by the extension operator, defined next, and it is assumed that  $h$  is chosen small enough that all nodes in  $\Sigma_A$  and  $\Sigma_G$  lie within a distance  $\kappa_\infty^{-1}$  of the surface. The active and ghost nodes for a circle embedded in  $\mathbb{R}^2$  can be seen in the left panel of Fig. 1.

Given a function  $\tilde{f}$  sampled on  $\Sigma_A$ , the discrete extension operator  $\mathbf{E}$  produces a function on  $\Sigma_A$  that is constant in the surface normal direction and coincides with  $\tilde{f}$  on the surface. However, for  $\mathbf{x}_i \in \Sigma_A$ , the closest point  $cp_S(\mathbf{x}_i)$  will generally not be another grid point. The function value  $\tilde{f}(cp_S(\mathbf{x}_i))$  is therefore found by interpolating



**Fig. 1** The left panel shows the active nodes  $\Sigma_A$  (blue circles) and ghost nodes  $\Sigma_G$  (red diamonds) for a circle embedded in  $\mathbb{R}^2$ . The right panel shows a bi-quadratic extension stencil for the node marked  $x_i$ . The values at all nine active nodes contribute to the value interpolated to  $cp_S(x_i)$ , and thus the extension back to  $x_i$

$\tilde{f}$  over the points in  $\Sigma_A$  near  $cp_S(\mathbf{x}_i)$ . We use tensor product barycentric Lagrange interpolation [18] of degree  $p$ , and thus  $\Sigma_A$  contains the union of all  $(p + 1)^d$  cubes of grid points needed to perform this extension. The interpolation weights are independent of the data being interpolated, and  $\mathbf{E}$  can be written as a matrix where each row holds the interpolation weights for a given node in the computational domain (consisting of  $\Sigma_A$  and  $\Sigma_G$ ). An example extension stencil can be seen in the right panel of Fig. 1.

The ambient Laplacian,  $\Delta$ , over the embedding space is discretized by standard second-order centered finite differences, denoted  $\Delta^h$ . Along the edges of the tube of active nodes in  $\Sigma_A$ , there will be incomplete finite difference stencils. The ghost nodes in  $\Sigma_G$  complete these stencils.

Finally, the discrete form of  $\Delta_S$  will be given as the composition of  $\Delta^h$  and  $\mathbf{E}$ . To apply the approach to implicit time stepping of diffusive problems [5], and eigenvalue problems [6], we use a stabilized discretization given by

$$\Delta_S^h = -\frac{2d}{h^2}\mathbf{I} + \left(\Delta^h + \frac{2d}{h^2}\mathbf{I}\right)\mathbf{E}, \tag{4}$$

where the removal of the diagonal elements from  $\Delta^h$  avoids redundant self-interpolation.

### 2.2 Domain decomposition

Iterative methods for the solution of large linear systems are attractive for many reasons. Critical to this work, they may have greatly reduced time and memory requirements over direct solvers, and they are easy to parallelize on distributed memory machines. Domain decomposition (DD) methods seek to replace the solution of one large problem with the repeated solution of several smaller problems, and are particularly well suited to elliptic PDEs [19, 20]. These methods can be used as iterative solvers on their own, or embedded within a Krylov solver as a preconditioner [19]. The DD-CPM library implements restricted additive Schwarz (RAS) and optimized restricted additive Schwarz (ORAS) solvers and preconditioners that respect the unique needs of the CPM.

For the CPM, it will be beneficial to write the continuous formulation of the (O)RAS methods with respect to the surface intrinsic PDE (eq:shiftPoisson). Consider splitting the global domain  $\mathcal{S}$  into  $N_S$  disjoint subdomains  $\tilde{\mathcal{S}}_j$ . Each disjoint subdomain is then grown to form an overlapping set of subdomains  $\mathcal{S}_j$ . The boundary of each subdomain,  $\partial\mathcal{S}_j$  is split into parts lying in the nearby disjoint subdomains denoted as  $\Gamma_{jk} \equiv \partial\mathcal{S}_j \cap \mathcal{S}_k$ . Given an initial guess for the global solution  $u^{(0)}$ , defined at least on the artificial boundaries  $\Gamma_{jk}$ , we may solve the local problems

$$\begin{cases} (c - \Delta_S)u_j = f, & \text{in } \mathcal{S}_j, \\ \mathcal{I}_{jk}u_j = \mathcal{I}_{jk}u^{(n)}, & \text{on } \Gamma_{jk}, \forall k. \end{cases} \tag{5}$$

After the local solutions  $u_j$  are found on each overlapping subdomain, a new global solution may be formed with respect to the disjoint subdomains as

$$u^{(n+1)} = \sum_j u_j |_{\tilde{\mathcal{S}}_j}. \quad (6)$$

Then, with an updated solution on the artificial boundaries, the local problems may be solved again. The iteration then continues. The boundary operators  $\mathcal{T}_{jk}$  transmit data between the local problems and are thus referred to as *transmission operators*. The DD-CPM software considers transmission operators in one of two forms:

$$\mathcal{T}_{jk} = \text{identity}, \quad (7)$$

$$\mathcal{T}_{jk} = \left( \frac{\partial}{\partial \hat{\mathbf{q}}_{jk}} + \alpha \right), \quad (8)$$

where  $\hat{\mathbf{q}}_{jk}$  is the unit conormal vector along  $\Gamma_{jk}$ . The first option enforces Dirichlet conditions on the local problems, while the second enforces Robin conditions. The Robin conditions also provide a parameter  $\alpha$  which may be tuned to accelerate convergence.

ORAS solvers and preconditioners are constructed by discretizing (eq:ddIter) via the CPM. Finding discretizations of the transmission operators  $\mathcal{T}_{jk}$ , such that the DD iteration converges to the discrete single domain solution, requires care. A full discussion of the requisite considerations, and guidance on setting the parameter  $\alpha$ , can be found in [2, 3].

### 3 Overview of the DD-CPM software

The DD-CPM software is hosted as a public git repository on BitBucket. Users may obtain the current stable version from the release branch by calling:

```
git clone -b release \
https://mayianm@bitbucket.org/mayianm/DD-CPM.git DD-CPM
```

Similarly, one can obtain the software as it existed at the time of publication by cloning as above, and then calling:

```
git checkout NumAlgoVersion
```

Detailed installation instructions can be found in the `README.md` file present in the top level of the repository.

#### 3.1 Specifying your own equations

Considering the `DDCPGrayScott` example, the essential ingredients in the driver code can be identified. A `ProblemDefinition` object, defined and setup

on lines 62 – 71, holds all of the settings relevant to the setup of all other objects. An object of the `CPPostProc` class (line 73) writes out all data files produced by the solver. The `GridFunc` and `DiffEq` objects created on lines 77 – 82 define the equation being solved. The usage of these classes is discussed in more detail below. The `CPMeshGlobal` class (used on line 84) defines the global computational domain and internally manages the partitioning of this domain into subdomains for the DD solvers to use. The `ProblemGlobal` object on line 86 builds the global system defined by the `DiffEq` object over the domain specified by the `CPMeshGlobal` object, and if DD solvers are desired, will also build the subproblems and manage communication internally. Finally, calling either the `solveTransient` method (line 91) or the `solveStationary` method (shown in the `DDCPPoisson` example) on the global problem will solve the defined equations.

### 3.1.1 Differential equations

The class `DiffEq` provides a container that combines one or more `GridFunc` objects, corresponding to the involved differential operators, with any forcing functions making up the differential equation to be solved. For equations without time dependence, the assumed structure is  $\mathcal{L}u = f(x)$  where  $\mathcal{L}$  is a differential operator expressed as a `GridFunc` object (discussed subsequently), and  $f(x)$  is a forcing function dependent on the spatial variable. For instance, line 67 in the `DDCPPoisson` example is:

```
DiffEq equation(pd,&shiftLap,&forcingFunc);
```

which encodes (eq:shiftPoisson) by applying the grid function `shiftLap` to the left side of the equation, and filling the forcing term from `forcingFunc` defined at the top of the driver file.

For time-dependent equations, the assumed structure is  $(\partial/\partial t + \mathcal{L})u = f(x, t, u)$ , where  $\mathcal{L}$  is a spatial differential operator represented by a `GridFunc` object, and  $f$  is the forcing function which can now depend on the solution  $u$  and the time  $t$  in addition to spatial variable  $x$ . Lines 68 and 69 in the `DDCPHeat` example are:

```
GridFunc laplacian("LaplacianSecondOrder",
                  pd.mesh.dim,pd.mesh.delta,-1.0);
DiffEq equation(pd,&laplacian,&icHeat,&rhsHeat);
```

where `laplacian` is a `GridFunc` object encoding the (negative) Laplace-Beltrami operator. Note now that in addition to the spatial operator and the grid function, the `DiffEq` object now requires a function specifying the initial condition.

Finally, a `std::vector` of grid functions can also be supplied to specify multicomponent equations. Consider lines 77 – 82 in the `DDCPGrayScott` example:

```

GridFunc lapU ("LaplacianSecondOrder",
               pd.mesh.dim, pd.mesh.delta, -gsPar_Du);
GridFunc lapV ("LaplacianSecondOrder",
               pd.mesh.dim, pd.mesh.delta, -gsPar_Dv);
std::vector<GridFunc*> grayScottOp {&lapU, &lapV};
DiffEq equation (pd, grayScottOp, icGrayScott, rhsGrayScott);

```

where the two separate Laplace-Beltrami operators are weighted by the different diffusivities of the  $u$  and  $v$  components in the Gray-Scott equations. These four lines, along with reaction terms and initial conditions, fully encode (eq:rdsys).

### 3.1.2 Grid functions

The class `GridFunc` provides the user with a way to define their own operators without needing the specific details of the computational domain or the distributed matrix that the operator will eventually be represented by. In a two-dimensional embedding space, the discrete Laplace-Beltrami operator shown in (eq:lapBel) can be obtained from the semidiscrete form

$$\Delta_S^h u(x, y) = -\frac{4}{h^2} u(x, y) + \frac{1}{h^2} [u(cp_S(x+h, y)) + u(cp_S(x-h, y)) + u(cp_S(x, y+h)) + u(cp_S(x, y-h))]. \quad (9)$$

Grid functions such as this are specified by three pieces of information: the stencil of nodes needed, the weight of each node in the stencil, and whether a node should be extended from the surface or not. The stencils for the extension operator are held internally, so this final item is simply a vector of booleans, and the full stencil including all nodes needed only for extension does not need to be specified.

Basic objects of the `GridFunc` class are formed from three `std::vector` objects holding these pieces of information. For ease, there are built-in grid functions for the identity, second-order accurate Laplace-Beltrami, and fourth-order accurate Laplace-Beltrami operators. Additionally, `GridFunc` objects can be added and composed to yield more complicated grid functions. For instance, the shifted Poisson (eq:shiftPoisson) is written in the `DDCPoisson` example as:

```

GridFunc laplacian ("LaplacianSecondOrder",
                  pd.mesh.dim, pd.mesh.delta, 1.0);
GridFunc shift ("Identity",
               pd.mesh.dim, pd.mesh.delta, 1.0);
GridFunc shiftLap = shift - laplacian;

```



where the final line encodes the total left-hand side operator.

The biharmonic operator can be written as:

```
GridFunc laplacian ("LaplacianSecondOrder",
                   pd.mesh.dim, pd.mesh.delta, 1.0);
GridFunc biharmonic = laplacian*laplacian;
```

which is much simpler than writing the entire biharmonic stencil. Crucially, building these grid functions avoids doing significantly more expensive operations on large distributed matrices after their assembly.

## 4 Running the software

The DD-CPM software acts primarily as a library. Each equation to be solved is specified in a brief driver code which calls the DD-CPM library. A number of sample driver programs are included to demonstrate how to interact with the library, and how to solve a few representative equations. A small plotting script is included in the `python` directory nested under the DD-CPM root directory. The relevant plotting command to visualize the solution from each example is given. The DD-CPM library and sample driver programs generate output both to the terminal and in the form of HDF5 data files. The interpretation and use of these outputs is discussed. With all of the sample driver programs as reference, the section is concluded with a guide towards specifying user-defined equations.

### 4.1 Included examples

The DD-CPM library comes with several example programs to demonstrate usage of the library and the structure of typical driver codes. The example programs do not rely on any particular choice of surface or level of parallelism. The example programs are

- `DDCPoisson.ex`: Solves the shifted Poisson equation, demonstrating stationary solvers and Laplace-Beltrami operators.
- `DDCPBiharmonic.ex`: Solves the shifted biharmonic equation, demonstrating stationary solvers and composition of operators.
- `DDCPHeat.ex`: Solves the heat equation, demonstrating implicit time stepping.
- `DDCPFitzhughNagumo.ex`: Solves the Fitzhugh-Nagumo equation, demonstrating implicit/explicit time stepping and multi-component equations.
- `DDCPGrayScott.ex`: Solves the Gray-Scott equation.
- `DDCPSchnackenberg.ex`: Solves the Schnackenberg equation, and includes a user-defined surface.

After compilation, the binaries for these examples can be found in the `bin` directory, and the example source codes can be found in the `examples` directory.

Running each of the example programs follows the same format:

```
mpiexec -n <nprocs> bin/DDCPEXample.ex \
-infile <path/to/inputfile> \
<additional options>
```

where `DDCPEXample` is replaced by one of the above example programs, `<nprocs>` is replaced by the number of processes you want to use for the run, and `<path/to/inputfile>` is the path to a file defining a surface and various default values. Many options can also be set on the command line to override values supplied by the input file. A full list of the available options can be obtained by running the example with the flag `-help`. The `-help` flag also gives several sample calls, specific to each example, with interesting options set to familiarize the user.

For examples using the included preconditioners (via flag `-pc_type ddcpm`), it is required that the number of subdomains be divisible by the number of processes used. This may be changed depending on the resources available on your system. The subdomain count can easily be changed in all examples (via flag `-mesh_nparts <n>`) keeping in mind that the number of subdomains must be divisible by the number of processes used.

There is also a small plotting tool included in the `DD-CPM/python` directory. In the following examples, one call to the solver is shown alongside a command to generate a plot of the solution. All of the following commands are issued from inside the `DD-CPM/build` directory.

#### 4.1.1 DDCPPoisson

The shifted Poisson equation, (eq:shiftPoisson) above, makes an excellent benchmark problem to compare various linear solvers and preconditioners. Similarly, it is a useful problem to showcase the flexibility available in the DD-CPM software. This equation is solved on a torus with a grid resolution of  $h = 1/300$  using 64 processes and a variety of linear solvers.

In all cases, the software is invoked as:

```
mpiexec -n 64 bin/DDCPPoisson.ex \
-infile inputFiles/torus.icpm -mesh_res 300 \
-pp_out_all -mesh_npoll 800 \
-log_view -ksp_converged_reason \
< additional options >
```

where `<additional options>` is set according to the desired linear solver, and consists mostly of standard PETSc flags. The flag `-mesh_res 300` specifies the resolution, and for this surface results in a problem with 4,090,560 unknowns.

The flags `-pp_out_all` and `-mesh_npoll 800` control output to the polling surface, which is discussed in more detail below.

The PETSc geometric/algebraic multigrid (GAMG) [21] solver can be specified by setting the additional options as:

```
-mesh_nparts 1 -pc_type gamg -ksp_type richardson \
-pc_gamg_sym_graph true -pc_gamg_agg_nsmooths 2
```

Alternatively, GAMG can be used as a preconditioner for GMRES by changing the flag `-ksp_type richardson` to `-ksp_type gmres`, or omitting it entirely as GMRES is the default choice within the DD-CPM software.

The MUMPs parallel sparse direct solver [22, 23] can be used by instead setting the additional options to:

```
-mesh_nparts 1 -ksp_type preonly -pc_type lu \
-pc_factor_mat_solver_type mumps
```

The DD-CPM software includes custom DD preconditioners as discussed in Section 2.2. The solver can be set to GMRES using these preconditioners by setting the additional options as:

```
-mesh_nparts 64 -mesh_nover 4 \
-petscpartitioner_type parmetis \
-pc_type ddcpm -pc_ddcpm_robfo true
```

This partitions the mesh into 64 subdomains using the ParMETIS [24] mesh partitioner, though any partitioner included in the Petsc installation can be used. These subdomains are overlapped with a width of 4, and each process is responsible for the solution of one subproblem at each iteration. The flag `-pc_type ddcpm` enables the included DD preconditioner, and the subsequent flag enables the use of Robin transmission conditions. The parameters in the Robin transmission conditions default to  $\alpha = 4$  and  $\alpha^\times = 20$ , and can be overridden with the flags `-pc_ddcpm_osm_alpha <alpha>` and `-pc_ddcpm_osm_alpha_cross <alpha_cross>`. More information regarding these transmission conditions and the parameters therein can be found in [2, 3]. Finally, we note that the subproblems are solved using a one level incomplete LU (iLU) factorization by default. This, and all other options regarding the solution of the subproblems, can be overridden by using the `-sub_prefix` (e.g., `-sub_pc_type lu`). We emphasize that the parameter choices specified in the options above are almost certainly not optimal, and have instead been set somewhat naively to demonstrate the flexibility of the software.

The mesh resolution can be doubled from  $h = 1/300$  to  $h = 1/600$  with the flag `-mesh_res 600`, yielding a global system with 16,370,528 unknowns. The Richardson/GAMG and GMRES/GAMG solvers can be used as is at this higher resolution. Our novel DD preconditioners can also be used, but require minor modification.

For problems of this size, the use of incomplete LU (iLU) factorization can occasionally lead to stagnation of the solver. This stagnation effect arises from the interaction between iLU factorization of the local systems and the Robin transmission conditions and is not yet fully understood, although large system size has been found to be a decisive factor. Fortunately, the standard RAS preconditioner remains well behaved for large problems; Dirichlet transmission conditions are the default option, so this may be specified by omitting the flag `-pc_ddcpm_robfo true`. Finally, note that the provided DD preconditioners are defined only for a single level, and their efficacy starts to fall off on problems of the size considered here. The inclusion of a coarse grid correction to alleviate this effect will be investigated in future releases.

#### 4.1.2 DDCPBiharmonic

This example solves the shifted biharmonic equation

$$(c + \Delta_S^2)u = f(\mathbf{x}), \quad (10)$$

where, as before,  $c$  is a positive constant. The forcing function is given in spherical coordinates as

$$f(\theta, \phi) = 5000 \left( \theta^2 - \frac{\pi}{2} \theta \right) \cos(7\phi), \quad (11)$$

where here, and for any uses of spherical coordinates in the remainder,  $\phi \in [0, 2\pi]$  is the azimuthal coordinate and  $\theta \in [0, \pi]$  is the polar coordinate.

Note that the included domain decomposition preconditioners do not support biharmonic equations. Instead, one can either use preconditioners included in PETSc or use a parallel direct solver. We take the latter approach here. (eq:biharm) can be solved on a sphere with 64 processes and the MUMPS [22, 23] parallel sparse direct solver by calling:

```
mpiexec -n 64 bin/DDCPBiharmonic.ex \
-infile inputFiles/sphere.icpm \
-mesh_res 200 -mesh_nparts 1 -ksp_type preonly \
-pc_type lu -pc_factor_mat_solver_type mumps
```

where the flags `-mesh_res 200` and `-mesh_nparts 1` set a grid spacing of  $1/200$  and inform the library that no mesh partitioning is necessary. The remaining flags are all standard PETSc options to enable the MUMPS solver.

The grid function defining this operator can be obtained directly from the operator for the Laplacian by composition. Crucially, this is done before the global operator is constructed to avoid expensive matrix-matrix multiplication. This composition is quite simple, and is discussed in more detail in Section 3.1.2.

The solution can be visualized on the CPM point cloud by calling the included plotting tool as:

```
python ../python/cpmpplot.py \
data/Sphere200_2_1_3/Sphere000000.h5 cloud sol
```

or over the surface itself by:

```
python ../python/cpmpplot.py \
data/Sphere200_2_1_3/Sphere000000.h5 poll sol
```

where `poll` indicates that the solution should be displayed on the *polling surface* discussed in more detail below. The solution can be seen in Fig. 2.

### 4.1.3 DDCPHeat

This example solves the heat equation

$$\frac{\partial u}{\partial t} - \Delta_S u = f(\mathbf{x}, t), \tag{12}$$

where the time-dependent forcing function is given in spherical coordinates as

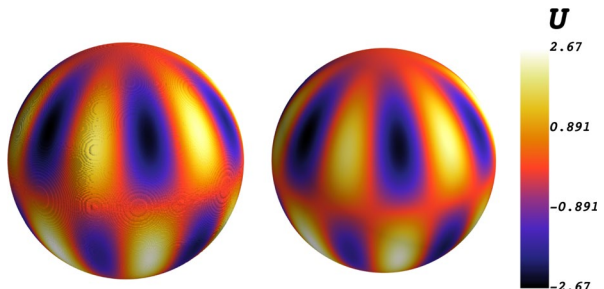
$$f(\theta, \phi, t) = \cos(t) \sin(3\phi). \tag{13}$$

The initial condition is  $u(\mathbf{x}, 0) = 0$ .

This can be solved on a sphere using 4 processes and our ORAS preconditioner over 12 subdomains by calling:

```
mpiexec -n 4 bin/DDCPHeat.ex \
-infile inputFiles/sphere.icpm \
-mesh_res 30 -mesh_nparts 12 -mesh_nover 4 \
-petscpartitioner_type parmetis \
-pc_type ddcpm -pc_ddcpm_robfo true \
-time_final 0.2 -ts_monitor
```

**Fig. 2** The left and right panels show the solution of the shifted biharmonic (eq:biharm) on the point cloud and polling surface respectively. The grid spacing is  $h = 1/200$ , and the forcing function is given in (eq:bhForce)



The default time integrator is the PETSc TSARKIMEX method, and the linear systems that arise are solved with GMRES supplemented by a 12 subdomain DD-CPM ORAS preconditioner. As above, the flags `-petscpartitioner_type parmetis -pc_type ddcpm -pc_ddcpm_robfo true` indicate that the mesh should be partitioned using ParMETIS, and that Robin transmission conditions should be used with their default parameters of  $\alpha = 4$  and  $\alpha^\times = 20$ . Again, we note that iLU factorization is used in the solution of all subproblems, which is generally useful for transient problems.

The solution can be visualized on the CPM point cloud by calling the included plotting tool as:

```
python ../python/cpmplot.py \
data/Sphere30_2_12_4/Sphere000006.h5 cloud sol
```

or over the polling surface by:

```
python ../python/cpmplot.py \
data/Sphere30_2_12_4/Sphere000006.h5 poll sol
```

both of which can be seen in Fig. 3.

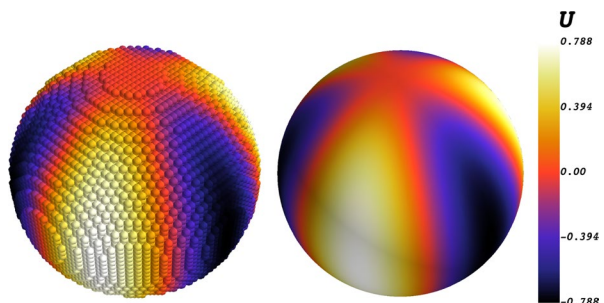
#### 4.1.4 DDCPFitzhughNagumo

This example solves the Fitzhugh-Nagumo equations

$$\begin{cases} \frac{\partial u}{\partial t} - D_u \Delta_S u = (0.1 - u)(u - 1)u - v, \\ \frac{\partial v}{\partial t} - D_v \Delta_S v = 0.01(0.5u - v) \end{cases} \quad (14)$$

where the diffusivities of each species are set as  $D_u = 10^{-4}$  and  $D_v = 10^{-7}$ . The initial condition is chosen to be

**Fig. 3** The left and right panels show the solution of the heat (eq:heat) on the point cloud and polling surface respectively. The grid spacing is  $h = 1/30$ , the final time is  $t = 0.2$ , and the forcing function is given in (eq:htForce)



$$u(\mathbf{x}, 0) = \begin{cases} 1, & x > 0, y > 0, z > 0 \\ 0, & \text{otherwise} \end{cases} \quad (15)$$

$$v(\mathbf{x}, 0) = \begin{cases} 1, & x > 0, y < 0, z > 0 \\ 0, & \text{otherwise} \end{cases}. \quad (16)$$

Time-dependent problems require more setup, though most of the flags are self-explanatory. This can be solved on a triangulated surface using 12 processes and the PETSc built-in block Jacobi preconditioner:

```
mpirun -n 12 bin/DDCPFitzhughNagumo.ex \
-infile inputFiles/triang.icpm \
-mesh_res 40 -pc_type bjacobi \
-time_final 250 -ts_monitor
```

The triangulated surface is supplied in the form of a PLY file which is specified in the `triang.icpm` input file. The default option is the Stanford Bunny [25], though several other PLY files are included in the `DD-CPM/plyFiles` directory.

The  $u$  component of the solution can be visualized by calling:

```
python ../python/cpmplot.py \
data/Triangulated40_2_12_3/Triangulated000176.h5 \
poll sol 0
```

and similarly the  $v$  component of the solution can be visualized by calling:

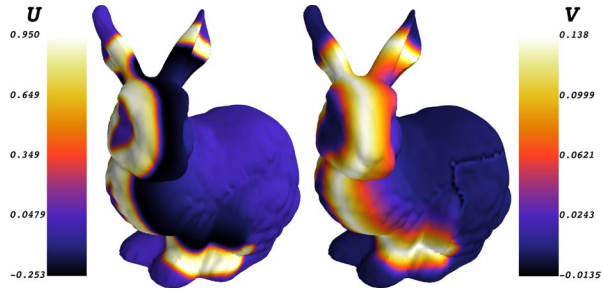
```
python ../python/cpmplot.py \
data/Triangulated40_2_12_3/Triangulated000176.h5 \
poll sol 1
```

each of which can be seen in Fig. 4. Note that the solution here is not shown on the point cloud, but rather on the surface itself. This is specified by the `poll` option. When the DD-CPM library produces output files, it will try to project the solution back onto the surface for later visualization. This requires that some triangulation of the surface is known. Of course, this triangulation is *not* used to solve the equations, but is needed to produce these visualizations. The solution on the point cloud can be seen by replacing `poll` with `cloud` in the above plotting calls.

#### 4.1.5 DDCPGrayScott

This example solves the Gray-Scott equations

**Fig. 4** The left and right panels show the  $u$  and  $v$  components of the solution to the Fitzhugh-Nagumo (eq:fitznag) at time  $t = 250$ . The initial conditions are set according to (eq:fnICs), and the grid spacing used is  $h = 1/40$



$$\begin{cases} \frac{\partial u}{\partial t} - D_u \Delta_S u = (1 - u)F - uv^2 \\ \frac{\partial v}{\partial t} - D_v \Delta_S v = uv^2 - (F + K)v \end{cases} \quad (17)$$

where the diffusivities of each species are chosen as  $D_u = 6 \times 10^{-5}$  and  $D_v = 3 \times 10^{-5}$ , the feed rate is set at  $F = 0.03$ , and the kill rate is given in spherical coordinates by the function

$$k(r, \theta, \phi) = 0.06 + 0.001 \cos(\phi) \quad (18)$$

where  $\phi$  is the azimuthal angle. The initial condition is chosen to be

$$u(r, \theta, \phi, 0) = 1 - e^{-5 \sin^2(5(\phi+0.5))} \quad (19)$$

$$v(r, \theta, \phi, 0) = e^{-5 \sin^2(5(\phi-0.5))}. \quad (20)$$

As in the previous example, the nonlinear coupling terms are advanced through time explicitly. To solve the Gray-Scott equation on a Dupin cyclide with 128 processes, one could call:

```
mpiexec -n 128 bin/DDCPGrayScott.exe \
-infile inputFiles/triang.icpm -mesh_res 300 \
-mesh_nparts 128 -mesh_nover 4 \
-petscpartitioner_type parmetis \
-pc_type ddcpm -pc_ddcpm_robfo true \
-time_final 15000 -pp_plotfreq 100 -ts_monitor
```

This generates a mesh with 3,739,898 active nodes globally. The diffusion terms are treated implicitly, and the linear systems that arise are solved with GMRES preconditioned by a 128 subdomain DD-CPM ORAS preconditioner. The time stepper defaults to using the PETSc TSARKIMEX3 method, and the present flags set a final time of  $T = 15,000$ . Finally, the flag `-pp_plotfreq 100` informs the post-processor to write a data file every 100 time steps.

These solutions can be visualized by calling the included plotting tool. The  $u$  component of the solution can be visualized by calling



```
python ../python/cpmplot.py \
data/Triangulated300_2_128_4/Triangulated000510.h5 \
poll sol 0
```

while the  $v$  component can be visualized via

```
python ../python/cpmplot.py \
data/Triangulated300_2_128_4/Triangulated000510.h5 \
poll sol 1
```

In each call, `poll` specifies that the polling surface should be used for visualization. The solutions can be seen in Fig. 5. To see the solution on the CPM point cloud, one would need to replace `poll` by `cloud` in the visualization command.

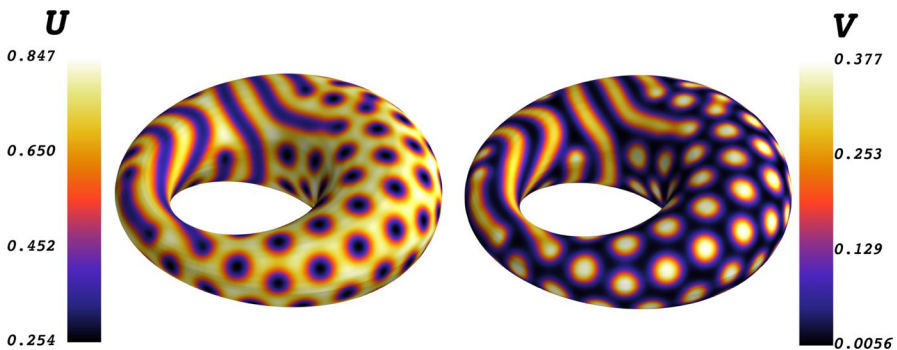
### 4.1.6 DDCPSchnackenberg

The final example solves the Schnackenberg equations

$$\begin{cases} \frac{\partial u}{\partial t} - D_u \Delta_S u = \delta_1 - u + u^2 v + \lambda v, \\ \frac{\partial v}{\partial t} - D_v \Delta_S v = \delta_2 - u^2 v - \lambda v \end{cases} \quad (21)$$

where the parameters are chosen to be  $D_u = 0.001$ ,  $D_v = 0.1$ ,  $\delta_1 = 0.003$ ,  $\delta_2 = 0.7$ , and  $\lambda = 0.06$  inspired in part by [26]. The initial conditions are set as

$$u = 0.02 + \mathcal{U} \sin(5x) \cos(3y), \quad (22)$$



**Fig. 5** The left and right panels show the  $u$  and  $v$  components of the solution to the Gray-Scott (eq:grayscott) at time  $t = 15,000$ . The initial conditions and parameters are set according to (eq:gsICs) and (eq:gsPars), and the grid spacing used is  $h = 1/400$

$$v = 4 + \mathcal{V} \sin(5xy) \sin(5x^3y) \quad (23)$$

where  $\mathcal{U} \sim \mathcal{N}(0, 0.01)$  and  $\mathcal{V} \sim \mathcal{N}(0, 1)$  are normally distributed random values.

This example also demonstrates how to implement user-defined surfaces. Inside the driver code `DDCPSchnackenberg.cpp`, there are functions defining the closest point function, a distance function, and a surface normal function, which are all handed to the library. In this case, these functions correspond to a unit hemisphere.

The boundary conditions along the circular boundary are mixed between homogeneous Neumann and homogeneous Dirichlet. Taking  $\phi$  as the azimuthal angle, any point that maps to the boundary where  $\sin(3\phi) < -0.5$  is mirrored to enforce second order accurate homogeneous Dirichlet boundary conditions. The remainder of the circular boundary utilizes the natural homogeneous Neumann boundary conditions. See [6] for details of the Dirichlet boundary formulation.

This equation can be solved on the user-defined surface using 12 processes and our ORAS preconditioner over 12 subdomains by calling:

```
mpirun -n 12 bin/DDCPSchnackenberg.ex \
-infile inputFiles/user.icpm -cp_pc_ras -mesh_res 30 \
-mesh_nparts 12 -mesh_nover 4 \
-petscpartitioner_type parmetis \
-pc_type ddcpm -pc_ddcpm_robfo true \
-time final 100 -ts monitor
```

The  $u$  component of the solution can be visualized by calling:

```
python ../python/cpmplot.py \
data/User30_2_12_4/User000832.h5 cloud sol 0
```

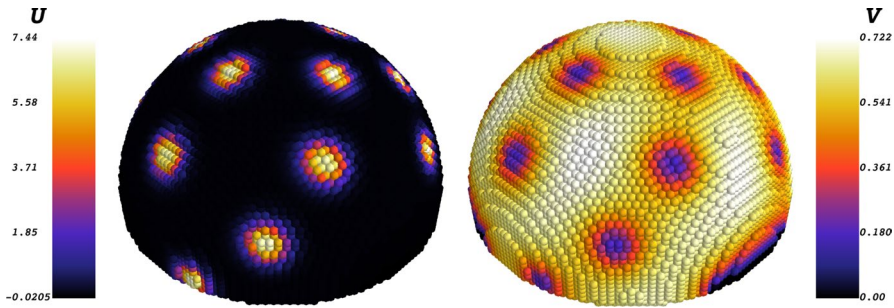
and similarly the  $v$  component of the solution can be visualized by calling:

```
python ../python/cpmplot.py \
data/User30_2_12_4/User000832.h5 cloud sol 1
```

where each can be seen in Fig. 6. Note that each of these calls visualizes the solution over the point cloud. Currently there is no support for polling surfaces corresponding to user-defined surfaces.

## 5 Profiling

The DD-CPM library yields great parallel scalability for moderate and large problems using hundreds of cores spread across a few nodes of a compute cluster. To demonstrate this scalability, we focus on the shifted Poisson equation due to its ubiquity and relative simplicity. We remark before presenting the results that mesh



**Fig. 6** The left and right panels show the  $u$  and  $v$  components of the solution to the Schnackenberg (eq:schnack) at time  $t = 50$ . The initial conditions are given by (eq:schnackIcV), the grid spacing is  $h = 1/30$ , and the parameters used are  $D_u = 0.001$ ,  $D_v = 0.1$ ,  $\delta_1 = 0.003$ ,  $\delta_2 = 0.7$ , and  $\lambda = 0.06$

partitioning can be done by any of the PETSc supported graph partitioners, but does have a minor serial component afterward. This does present an early bottleneck when running large problems. However, this only occurs once in any run and does not pose any issue for problems with millions of active grid points. This bottleneck will be removed in a future release of the library. All tests were performed on whole nodes of the Graham compute cluster managed by Compute Canada. Each node has two 16-core Intel E5-2683 v4 Broadwell 2.1 GHz processors and 125G of memory.

Each run is broken into three major phases: mesh construction and partitioning, global and local operator construction, and the GMRES solve. As discussed above, the meshing phase is not expected to scale due to the serial bottleneck there. The time required for each phase is recorded using the logging features of PETSc. Each test of scalability is performed using  $N_p = 64, 128, 256$  processes, corresponding to 2, 4, and 8 full compute nodes on Graham respectively.

To assess the strong scalability, where a problem of fixed size is solved with progressively greater amounts of parallelism, we consider the torus described in Section 4.1.1 with a grid spacing of  $h = \frac{1}{300}$  and tri-quadratic interpolation. This combination yields a problem with 4,090,560 unknowns. The linear system is solved with GMRES using the DD-CPM RAS preconditioner with an overlap width of  $N_O = 4$ . All local problems inside the RAS preconditioner are solved inexactly using a one level incomplete LU factorization.

The number of subdomains used must be divisible by the number of processes, giving two interesting options. First, Table 1 presents results from setting  $N_S = N_p$ . For  $N_S = N_p = 64$ , each disjoint subdomain contains approximately 64,000 active nodes, and each overlapping subdomain contains approximately 73,000 nodes. For  $N_S = N_p = 128$ , each disjoint subdomain contains approximately 32,000 active nodes, and each overlapping subdomain contains approximately 38,000 nodes. Finally, for  $N_S = N_p = 256$ , the disjoint and overlapping subdomains contain approximately 16,000 and 20,000 active nodes respectively. In this case, the preconditioner is stronger for smaller process counts, but more expensive to apply due to the larger local problems. The operator construction

**Table 1** The time in seconds is shown for each phase in the solution of the shifted Poisson equation on a torus with  $h = 300^{-1}$  and 4,090,560 unknowns

Phase	$N_p = 64, N_S = 64$	$N_p = 128, N_S = 128$	$N_p = 256, N_S = 256$
Meshing	33	33	34
Operators	20	11	8
GMRES	33	22	12
Total	85	65	55
Iterations	391	496	481
Time per iteration	0.084	0.044	0.025

The required number of iterations and time per iteration is listed in the lower portion of the table. Here  $N_S = N_p$  is fixed and good strong scalability is observed apart from the meshing phase

and GMRES solve scale quite well, though the disparity in the size of the local problems is visible when comparing  $N_p = 64$  to  $N_p = 128$ .

Alternatively, the number of subdomains can be fixed at  $N_S = 256$  irrespective of the number of processes. This ensures that all of the local problems are roughly the same size, and that the preconditioner being applied is the same in all cases. In this case, each disjoint subdomain contains approximately 16,000 active nodes, and each overlapping subdomain contains approximately 20,000 nodes. Note that for  $N_p = 64$  and  $N_p = 128$  each process is responsible for constructing and factoring multiple local operators. As shown in Table 2, the now weakened preconditioner (for the smaller process counts) leads to longer GMRES solution times. Excellent strong scalability is observed apart from the meshing phase.

Weak scalability is assessed by fixing  $N_S = N_p$ , and varying the resolution of the problem such that the sizes of the disjoint partitions are roughly constant. For  $N_S = N_p = 64$ , the resolution is  $h = \frac{1}{300}$  for a total of 4,090,560 unknowns with disjoint and overlapping subdomains of approximate size 64,000 and 73,000 respectively. For  $N_S = N_p = 128$ , the resolution is increased to  $h = \frac{1}{420}$  for a total of 8,060,600 unknowns, and for  $N_S = N_p = 256$  the resolution is further increased to  $h = \frac{1}{595}$  for a total of 1,608,2068 unknowns. As can be seen in Table 3, the operator construction phase scales almost perfectly. The time required for the GMRES solution slowly climbs, which is to be expected since the included preconditioner incorporates only a single level. This is demonstrated concisely by the growing number of required iterations to convergence. Excellent weak scalability is observed in the time per iteration.

## 6 Conclusion

The DD-CPM library provides a general purpose implementation of the closest point method with an emphasis on distributed memory parallelism. The DD-CPM library is particularly useful for surface intrinsic reaction-diffusion equations, where the included DD preconditioners can efficiently handle the stiff diffusion terms implicitly. The DD-CPM library also handles domains in flat space with

**Table 2** The time in seconds is shown for each phase in the solution of the shifted Poisson equation on a torus with  $h = 300^{-1}$  and 4,090,560 unknowns

Phase	$N_p = 64, N_s = 256$	$N_p = 128, N_s = 256$	$N_p = 256, N_s = 256$
Meshing	33	34	34
Operators	21	12	8
GMRES	46	22	12
Total	99	67	55
Iterations	554	480	481
Time per iteration	0.083	0.046	0.025

Here  $N_s = 256$  is fixed independent of the number of processes (compare to Table 1). Excellent strong scalability is observed apart from the meshing phase

**Table 3** The time in seconds is shown for each phase in the solution of the shifted Poisson equation on a torus on a sequence of finer meshes

Phase	$N_p = 64, N_s = 64$	$N_p = 128, N_s = 128$	$N_p = 256, N_s = 256$
Meshing	33	65	139
Operators	20	21	25
GMRES	33	77	121
Total	85	164	288
Iterations	391	877	1438
Time per iteration	0.084	0.088	0.084

Resolutions of  $h = 300^{-1}$ ,  $h = 420^{-1}$ , and  $h = 595^{-1}$  were used to keep the number of unknowns per process roughly constant. Overall solution times increase somewhat due to increasing times in the meshing phase (which contains a serial component) and an increase in the number of GMRES iterations (which is as expected due to increasing subdomain size). Excellent weak scalability is observed in the time per iteration

otherwise difficult boundaries, and with some development could be applied to coupled bulk-surface equations. The software could be similarly extended to handle fully implicit discretizations in time, with the given preconditioners aiding in the solution of the linear systems that arise inside the Newton iterations. Additionally, the DD-CPM library could make an effective back end for eigenvalue problems due to the interoperability of PETSc and SLEPc.

**Funding** The authors gratefully acknowledge the financial support of NSERC Canada (RGPIN 2016-04361 and RGPIN 2018-04881). This research was enabled in part by support provided by ACEnet (acenet.ca) and Compute Canada ([www.computecanada.ca](http://www.computecanada.ca)).

**Data availability** All data generated as part of this study is available upon request or can be generated using the described library and the calling sequences listed in the paper.

**Declarations**

**Conflict of interest** The authors declare no competing interests.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

1. Balay, S., Gropp, W.D., McInnes, L.C., Smith, B.F.: Efficient management of parallelism in object-oriented numerical software libraries, pp 163–202. Birkhäuser Boston (1997). [https://doi.org/10.1007/978-1-4612-1986-6\\_8](https://doi.org/10.1007/978-1-4612-1986-6_8)
2. May, I.C.T., Haynes, R.D., Ruuth, S.J.: Domain decomposition for the closest point method. In: Domain Decomposition Methods in Science and Engineering XXV. Lecture Notes in Computational Science and Engineering ; 138, pp 458–465. Springer, (2020). International Conference on Domain Decomposition Methods in Science and Engineering (25th : 2018 : St. John's, N.L.)
3. May, I.C.T., Haynes, R.D., Ruuth, S.J.: Schwarz solvers and preconditioners for the closest point method. *SIAM J. Sci. Comput.* **42**(6), 3584–3609 (2020). <https://doi.org/10.1137/19M1288279>
4. Macdonald, C.B.: CP\_Matrices. GitHub (2018). [https://github.com/cbm755/cp\\_matrices](https://github.com/cbm755/cp_matrices)
5. Macdonald, C.B., Ruuth, S.J.: The implicit closest point method for the numerical solution of partial differential equations on surfaces. *SIAM J. Sci. Comput.* **31**(6), 4330–4350 (2010)
6. Macdonald, C.B., Brandman, J., Ruuth, S.J.: Solving eigenvalue problems on curved surfaces using the closest point method. *J. Comput. Phys.* **230**(22), 7944–7956 (2011)
7. Macdonald, C.B., Merriman, B., Ruuth, S.J.: Simple computation of reaction-diffusion processes on point clouds. *Proc. Natl. Acad. Sci.* **110**(23), 9209–9214 (2013). <https://doi.org/10.1073/pnas.1221408110>
8. März, T., Macdonald, C.B.: Calculus on surfaces with general closest point functions. *SIAM J. Numer. Anal.* **50**(6), 3303–3328 (2012)
9. Macdonald, C.B., Ruuth, S.J.: Level set equations on surfaces via the closest point method. *J. Sci. Comput.* **35**(2), 219–240 (2008)
10. Chen, Y., Macdonald, C.B.: The closest point method and multigrid solvers for elliptic equations on surfaces. *SIAM J. Sci. Comput.* **37**(1), 134–155 (2015)
11. May, D.A., Sanan, P., Rupp, K., Knepley, M.G., Smith, B.F.: Extreme-scale multigrid components within PETSc. In: Proceedings of the platform for advanced scientific computing conference. PASC '16. Association for computing machinery, (2016). <https://doi.org/10.1145/2929908.2929913>
12. Floater, M.S., Hormann, K.: Surface parameterization: a tutorial and survey, pp 157–186. Springer, (2005)
13. Dziuk, G., Elliott, C.M.: Surface finite elements for parabolic equations. *J. Comput. Math.* **25**(4), 385–407 (2007)
14. Bertalmio, M., Cheng, L.-T., Osher, S., Sapiro, G.: Variational problems and partial differential equations on implicit surfaces. *J. Comput. Phys.* **174**(2), 759–780 (2001)
15. Ruuth, S.J., Merriman, B.: A simple embedding method for solving partial differential equations on surfaces. *J. Comput. Phys.* **227**(3), 1943–1961 (2008)
16. Petras, A., Ruuth, S.J.: PDEs on moving surfaces via the closest point method and a modified grid based particle method. *J. Comput. Phys.* **312**, 139–156 (2016)
17. Chu, J., Tsai, R.: Volumetric variational principles for a class of partial differential equations defined on surfaces and curves. *Res. Math. Sci.* **5**(2), 19 (2018)
18. Berrut, J.-P., Trefethen, L.N.: Barycentric lagrange interpolation. *SIAM Rev.* **46**(3), 501–517 (2004)
19. Dolean, V., Jolivet, P., Nataf, F.: An introduction to domain decomposition methods: algorithms, theory, and parallel implementation. SIAM, (2015)
20. Toselli, A., Widlund, O.: Domain Decomposition Methods—Algorithms and Theory. Springer Series in Computational Mathematics, 34. Springer, (2005)

21. Balay, S., Abhyankar, S., Adams, M.F., Benson, S., Brown, J., Brune, P., Buschelman, K., Constantinescu, E.M., Dalcin, L., Dener, A., Eijkhout, V., Gropp, W.D., Hapla, V., Isaac, T., Jolivet, P., Karpeev, D., Kaushik, D., Knepley, M.G., Kong, F., Kruger, S., May, D.A., McInnes, L.C., Mills, R.T., Mitchell, L., Munson, T., Roman, J.E., Rupp, K., Sanan, P., Sarich, J., Smith, B.F., Zampini, S., Zhang, H., Zhang, H., Zhang, J.: PETSc Web page. <https://petsc.org/> (2022)
22. Amestoy, P.R., Duff, I.S., L'Excellent, J.-Y., Koster, J.: A fully asynchronous multifrontal solver using distributed dynamic scheduling. *SIAM J. Matrix Anal. Appl.* **23**(1), 15–41 (2001)
23. Amestoy, P.R., Guermouche, A., L'Excellent, J.-Y., Pralet, S.: Hybrid scheduling for the parallel solution of linear systems. *Parallel Computing* **32**(2), 136–156 (2006)
24. Karypis, G., Kumar, V.: A parallel algorithm for multilevel graph partitioning and sparse matrix ordering. *J Parallel Distrib Comput.* **48**(1), 71–95 (1998)
25. Turk, G., Levoy, M.: Zippered polygon meshes from range images. In: Proceedings of the 21st annual conference on computer graphics and interactive techniques. SIGGRAPH '94, pp 311–318. ACM, (1994)
26. Noufaey, K.S.A.: Stability analysis for Selkov-Schnakenberg reaction-diffusion system. *Open Math.* **19**(1), 46–62 (2021). <https://doi.org/10.1515/math-2021-0008>