

THE IMPLICIT CLOSEST POINT METHOD FOR THE NUMERICAL SOLUTION OF PARTIAL DIFFERENTIAL EQUATIONS ON SURFACES

COLIN B. MACDONALD* AND STEVEN J. RUUTH†

Abstract. Many applications in the natural and applied sciences require the solutions of partial differential equations (PDEs) on surfaces or more general manifolds. The Closest Point Method is a simple and accurate embedding method for numerically approximating PDEs on rather general smooth surfaces. However, the original formulation is designed to use explicit time stepping. This may lead to a strict time-step restriction for some important PDEs such as those involving the Laplace-Beltrami operator or higher-order derivative operators. To achieve improved stability and efficiency, we introduce a new implicit Closest Point Method for surface PDEs. The method allows for large, stable time steps while retaining the principal benefits of the original method. In particular, it maintains the order of accuracy of the discretization of the underlying embedding PDE, it works on sharply defined bands without degrading the accuracy of the method, and it applies to general smooth surfaces. It also is very simple and may be applied to a rather general class of surface PDEs. Convergence studies for the in-surface heat equation and a fourth-order biharmonic problem are given to illustrate the accuracy of the method. We demonstrate the flexibility and generality of the method by treating flows involving diffusion, reaction-diffusion and fourth-order spatial derivatives on a variety of interesting surfaces including surfaces of mixed codimension.

Key words. Closest Point Method, surface computation, implicit surfaces, partial differential equations, implicit time stepping, Laplace-Beltrami operator, biharmonic operator, surface diffusion

AMS subject classifications. 65M06, 58J35, 65M20

1. Introduction. Partial differential equations are ubiquitous throughout the sciences and applied sciences. Because analytical solutions are rarely possible, the practical importance of accurate and efficient numerical methods for PDEs cannot be overemphasized. There has been a great effort made to develop numerical methods for many important classes of PDEs. These efforts often focus on finding solutions in \mathbb{R}^1 , \mathbb{R}^2 and \mathbb{R}^3 . But problems involving general differential equations also arise on manifolds (such as a two-dimensional curved surface or a one-dimensional curved filament) in some embedding space (such as \mathbb{R}^3). For example, in material science one might wish to examine phase change of a material on a surface [36]. In biological modeling, one might study wound healing on skin [27], or the evolution of a pattern on an animal coat [24]. Computer graphics and image processing also frequently use PDEs on surfaces; for example, they might use such methods to place a texture on a surface [40], segment out objects defined in surface textures [37] or restore a damaged pattern [5] on a surface such as a vase.

PDEs on surfaces are traditionally handled by one of several techniques. These include finite element discretization on a triangulation of the surface [18, 9], discretizations based on some parameterization of the surface [11, 35, 16], and embedding techniques [6, 8, 7, 14, 15, 10, 25] which solve some embedding PDE in a small region near the surface. Naturally, each class of method has its own advantages and disadvantages. The finite element approach can treat elliptic equations, but is often

*Department of Mathematics, University of California Los Angeles, CA 90095, USA (cbm@math.ucla.edu). The work of this author was supported by a grant from NSERC Canada, an NSERC postdoctoral fellowship and NSF grant number CCF-0321917.

†Department of Mathematics, Simon Fraser University, Burnaby, British Columbia, V5A 1S6 Canada (sruuth@sfu.ca). The work of this author was partially supported by a grant from NSERC Canada.

less effective when convective or advective terms arise. Parameterization methods can be effective in certain specialized situations, but have the deficiency of introducing distortions and singularities into the method through the parameterization [11]. Methods within the embedding class have the attractive feature of being able to work in standard Cartesian coordinates, using Cartesian grid methods. Frequently, however, such methods are built around level set representations of the surface, a property that complicates the treatment of general open surfaces or surfaces without orientation, because a standard level set representation needs a well-defined inside and outside. Furthermore, level-set-based embedding methods typically impose some artificial boundary conditions when computations are carried out in a small region surrounding the surface. The imposition of such artificial boundary conditions leads to low-order accurate results for some important PDEs [14].

Within the embedding class, a method (the Closest Point Method) was developed and analyzed in [30] which uses a closest point representation of the underlying surface \mathcal{S} . A closest point representation assumes that for any point \mathbf{x} in the embedding space \mathbb{R}^d containing \mathcal{S} , a (possibly non-unique) point $\text{cp}(\mathbf{x}) \in \mathcal{S}$ is known which is closest to \mathbf{x} in Euclidean distance. For smooth surfaces, the closest point representation is unique in a narrow band around the surface, the width of which depends on, for example, the principal curvatures. The closest point representation allows for general surfaces with boundaries and does not require the surface to have an inside/outside. The surface can be of any codimension, or even of mixed codimension.

Given a closest point representation, a *closest point extension* which propagates data off the surface \mathcal{S} into the surrounding space \mathbb{R}^d is given by $u(\text{cp}(\mathbf{x}))$. Closest point extensions lead to simplified derivative calculations in the embedding space because they result in functions which are constant in the direction normal to the surface, at least within a neighborhood of a smooth surface. The relationship between surface derivatives and derivatives in the embedding space may be made precise using two fundamental principles [30]. Let ∇ denote the “standard” gradient in \mathbb{R}^d and let $\nabla_{\mathcal{S}}$ denote the gradient intrinsic to the smooth surface \mathcal{S} , then:

PRINCIPLE 1 (Gradients). *For points \mathbf{x} on a smooth surface, $\nabla_{\mathcal{S}}u(\mathbf{x}) = \nabla(u(\text{cp}(\mathbf{x})))$ because the function $u(\text{cp}(\mathbf{x}))$ is constant in the normal direction and therefore only varies along the surface. In other words, at points \mathbf{x} on the surface, intrinsic surface gradients $\nabla_{\mathcal{S}}u(\mathbf{x})$ are the same as gradients of $u(\text{cp}(\mathbf{x}))$.*

PRINCIPLE 2 (Divergence). *Let $\nabla_{\mathcal{S}}$ denote the divergence operator intrinsic to a smooth surface \mathcal{S} and let \mathbf{v} be any vector field on \mathbb{R}^d that is tangent at \mathcal{S} and also tangent at all surfaces displaced by a fixed distance from \mathcal{S} (i.e., all surfaces defined as level sets of the distance function to \mathcal{S}). Then at points \mathbf{x} on the surface $\nabla \cdot \mathbf{v}(\mathbf{x}) = \nabla_{\mathcal{S}} \cdot \mathbf{v}(\mathbf{x})$.*

Combinations of these two principles may be made, to encompass very general second-order differential operators, including the Laplace–Beltrami operator $\Delta_{\mathcal{S}} = \nabla_{\mathcal{S}} \cdot \nabla_{\mathcal{S}}$, in-surface curvature, and nonlinear diffusion operators [30]. Indeed, even higher-order and more general derivative replacements may be considered by carrying out multiple closest point extensions as first described in [30] and demonstrated in practice in Sections 5 and 6. Notice that the introduction of a closest point extension step leads us to consider standard PDEs in the embedding space for important classes of surface PDEs. This in turn enables the use of standard discretizations for the underlying embedding PDEs.

The Closest Point Method has been applied to a variety of time-dependent problems including in-surface advection, diffusion, reaction-diffusion, and Hamilton-Jacobi

equations [30, 23, 22]. In these previous works, time stepping was performed explicitly and the numerical solution was propagated by alternating between two steps:

1. Evolution of an embedding PDE on a regular Cartesian grid enveloping the surface using an explicit time-stepping scheme.
2. Extension of the surface data to the surrounding grid points with a closest point extension, i.e., for grid point \mathbf{x} , assign $u(\mathbf{x}) = u(\text{cp}(\mathbf{x}))$. As $\text{cp}(\mathbf{x})$ is typically not a grid point, this extension is an interpolation step.

This explicit approach to the Closest Point Method works very well for nonstiff problems such as arise from Hamilton–Jacobi equations. The method achieves high-order accuracy [23] using standard spatial discretizations and explicit Runge–Kutta time-stepping methods with time-step sizes $\Delta t = O(\Delta x)$, i.e., on the order of the spatial grid. It also avoids introducing artificial boundary conditions when computations are limited to a neighborhood of the surface [30, 23], thereby enabling the method to localize computations while maintaining the accuracy of the underlying discretization.

Unfortunately, explicit time-stepping methods are a poor choice for many important PDEs. For example, a severe stability time-step restriction may arise when spatially discretized diffusion, biharmonic or other stiff terms appear. Such time-step restrictions can lead to inefficient or impractical algorithms (for example, when biharmonic terms arise we might anticipate a step-size restriction like $\Delta t \sim \Delta x^4$). Severe time-step restrictions are also possible when local grid refinement is carried out because the time step will scale like some power of the minimal grid spacing. On the other hand, the use of implicit time-stepping methods typically leads to mild stability time-step restrictions, making them natural candidates for treating the stiff systems arising from certain PDEs. This motivates the development and analysis of an implicit Closest Point Method.

The remainder of this paper unfolds as follows. In Section 2, the implicit Closest Point Method is presented. It uses implicit linear multistep or Runge–Kutta time-stepping schemes and allows for large time steps. When the implicit Closest Point Method is applied to linear problems such as the in-surface heat equation, each time step requires a linear system solve: Section 3 discusses the properties of that system. Numerical results in Sections 4 and 5 demonstrate the high-order accuracy of the method on the in-surface heat equation and a biharmonic problem. Several applications are presented, including blurring an image on a triangulated surface, modelling heat flow on a complicated surface with components of varying codimension, in-surface pattern formation, and fourth-order “surface diffusion” for an interface on a surface. Section 6 extends the method to odd-order differential terms.

2. The implicit Closest Point Method. In this section we derive the implicit Closest Point Method. Before giving a more general construction, we illustrate the main ideas for the simple, but important case of in-surface heat flow.

2.1. The Closest Point Method for in-surface heat flow. To motivate the method, let \mathcal{S} be a surface embedded in \mathbb{R}^d and consider the in-surface heat equation

$$u_t = \Delta_{\mathcal{S}} u, \tag{2.1a}$$

$$u(0, \mathbf{x}) = u_0(\mathbf{x}). \tag{2.1b}$$

Based on the Principles 1 and 2, we may replace the Laplace–Beltrami operator $\Delta_{\mathcal{S}}$ by the standard Laplacian Δ in the embedding space \mathbb{R}^d , provided the argument of the right hand side is replaced by the closest point operator, $\text{cp}(\mathbf{x})$. This leads to the

embedding PDE

$$u_t(t, \mathbf{x}) = \Delta(u(t, \text{cp}(\mathbf{x}))), \quad (2.2a)$$

$$u(0, \mathbf{x}) = u_0(\text{cp}(\mathbf{x})), \quad (2.2b)$$

which has the property that the solutions of (2.1) and (2.2) agree on the surface [30].

The implicit Closest Point Method is based on discretizing (2.2) directly. For ease of exposition, in this example we assume the surface \mathcal{S} is a curve embedded in 2D and consider a second-order centered finite difference scheme applied to the Laplacian Δ in (2.2). This yields

$$\begin{aligned} \frac{\partial}{\partial t} u(t, x, y) = \frac{1}{\Delta x^2} & \left(-4u(t, \text{cp}(x, y)) + u(t, \text{cp}(x + \Delta x, y)) + u(t, \text{cp}(x - \Delta x, y)) \right. \\ & \left. + u(t, \text{cp}(x, y + \Delta x)) + u(t, \text{cp}(x, y - \Delta x)) \right) + O(\Delta x^2), \end{aligned} \quad (2.3a)$$

$$u(0, x, y) = u_0(\text{cp}(x, y)), \quad (2.3b)$$

where x and y are still continuous variables (this is not yet a spatial discretization). The spatial discretization can be completed by choosing a grid $\{(x_i, y_j)\}$, approximating the solution with nodal values u_{ij} , and using interpolation to replace the value of $u(\text{cp}(x_i, y_j))$ by a linear combination of nearby nodal values. As the value of $\text{cp}(x_i, y_j)$ is known beforehand, the interpolation weights can be precomputed using barycentric Lagrange interpolation [4] (in a dimension-by-dimension fashion). These interpolation weights and the differentiation weights $\frac{1}{\Delta x^2} \{-4, 1, 1, 1, 1\}$ can then be combined into a matrix (see Section 2.2 for details) which approximates $\Delta(u(t, \text{cp}(x, y)))$. Thus, this method-of-lines approach yields a linear system of coupled ODEs in time for $u_{ij}(t) \approx u(t, x_i, y_j)$. The discretization may then be completed by applying standard implicit Runge–Kutta or linear multistep methods.

As we shall see in Section 2.2.3, the system resulting from (2.3) is unstable. An alternate consistent formulation which benefits from improved diagonal dominance and improved stability arises when the redundant closest point operators are removed from the discretization (where diagonal terms appear in the discretization consistency does not require an interpolation to the closest point on the surface). This leads to the following stabilized form:

$$\begin{aligned} \frac{\partial}{\partial t} u(t, x, y) = \frac{1}{\Delta x^2} & \left(-4u(t, x, y) + u(t, \text{cp}(x + \Delta x, y)) + u(t, \text{cp}(x - \Delta x, y)) \right. \\ & \left. + u(t, \text{cp}(x, y + \Delta x)) + u(t, \text{cp}(x, y - \Delta x)) \right) + O(\Delta x^2), \end{aligned} \quad (2.4a)$$

$$u(0, x, y) = u_0(\text{cp}(x, y)), \quad (2.4b)$$

where the only change is that the diagonal entries no longer involve the closest point operator cp . Note that this equation (2.4) and the previous (2.3) agree at the surface because for points (x, y) on the surface we have $\text{cp}(x, y) = (x, y)$. This ensures that the solution of (2.4) is consistent with that of (2.1).

2.2. Matrix formulation of the Closest Point Method. The general matrix formulation of the Closest Point Method introduces a matrix operator (the “extension matrix”) to discretize all instances of the closest point operator. Our approach explicitly constructs the extension matrix. This enables the use of either direct or iterative matrix solvers, and makes the analysis of the ODE system straightforward. The alternative approach of computing an iterative solution of the implicit equations




Fig. 2.1: Example of the lists of grid points L (indicated by \bullet) and G (indicated by \circ) where the curve \mathcal{S} is a circle. The interpolation stencil is a 4×4 grid (arising, for example, from using degree $p = 3$ barycentric Lagrange interpolation) and the shaded regions illustrate the use of this stencil at the two points on the circle indicated by \diamond . Five-point differentiation stencils are shown for two example points in L , in one case illustrating the use of ghost points.

without explicitly constructing the extension matrix has the advantage of being closer with respect to implementation to the explicit Closest Point Method. This alternative approach is discussed briefly in Section 3.2, and will be treated in detail as part of future work.

We now give the details on the general matrix formulation. Suppose the surface \mathcal{S} is embedded in \mathbb{R}^d . Assume we have an interpolation scheme where the interpolated value is a linear combination of the values at neighbouring grid points in, for example, a hypercube in \mathbb{R}^d . These neighbouring grid points form the *interpolation stencil* of the scheme. Various interpolation techniques are possible. Similar to the explicit Closest Point Method in [30], we will use barycentric Lagrange interpolation [4] of degree p , in a dimension-by-dimension fashion, so the hypercube has $p + 1$ points in each coordinate direction.

Now consider two discrete ordered lists of points in the embedding space. The first is $L = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_m\}$. This list contains every grid point which can appear in the interpolation stencil for some point on the surface \mathcal{S} .¹ The grid points in L form the computational band on which the numerical solution is defined. Specifically, to approximate the solution function u at all points in L we introduce the vector $\mathbf{u} \in \mathbb{R}^m$ with entries $u_i \approx u(\mathbf{x}_i)$ for each grid point $\mathbf{x}_i \in L$. The second list of points $G = \{\mathbf{x}_{m+1}, \mathbf{x}_{m+2}, \dots, \mathbf{x}_{m+m_g}\}$ is disjoint from L and contains m_g additional “ghost points” along the edges of the computational band. These points will be used in the derivation below but their values are not propagated in time. Figure 2.1 shows an example of the two lists L and G where the curve \mathcal{S} consists of a circle embedded in \mathbb{R}^2 . Algorithms for the construction of these lists are given in Appendix A.

Finally, we require two other vectors over the list L and G for use in the derivation. Let $u(\text{cp}(L)) \in \mathbb{R}^m$ denote the vector with components $u(\text{cp}(\mathbf{x}_i))$ for $\mathbf{x}_i \in L$ and let $u(\text{cp}(G)) \in \mathbb{R}^{m_g}$ denote the vector with components $u(\text{cp}(\mathbf{x}_{m+i}))$ for $\mathbf{x}_{m+i} \in G$.

2.2.1. The extension matrix. The closest point extension is an operator which assigns a value of $u(\text{cp}(\mathbf{x}))$ to $u(\mathbf{x})$. In a discrete setting, $\text{cp}(\mathbf{x}_i)$ is generally not a grid point and interpolation is used to approximate $u(\text{cp}(\mathbf{x}_i))$. Because each interpolation is a linear combination of values in \mathbf{u} , this operation can be expressed as a matrix multiplication using the *extension matrix* \mathbf{E} defined as follows

¹This intuitive condition is sufficient. A slightly smaller computational band is often possible by taking L to be the list of all interpolation nodes over just the interpolated points $\text{cp}(\mathbf{x}_i)$, $\mathbf{x}_i \in L \cup G$ rather than over all of the surface \mathcal{S} . This is the approach used in Appendix A to construct L .

DEFINITION 1. [**extension matrix \mathbf{E}**] Given the vectors \mathbf{u} , $u(\text{cp}(L))$, $u(\text{cp}(G))$ and an interpolation scheme as described above, the extension matrix is a $(m+m_g) \times m$ matrix operator \mathbf{E} such that

$$\begin{pmatrix} u(\text{cp}(L)) \\ u(\text{cp}(G)) \end{pmatrix} \approx \mathbf{E}\mathbf{u} = \begin{array}{c} \boxed{\phantom{\mathbf{E}}} \\ \mathbf{E} \end{array} \begin{array}{c} \boxed{\phantom{\mathbf{u}}} \\ \mathbf{u} \end{array}. \quad (2.5)$$

The nonzero entries in the i^{th} row of \mathbf{E} consist of the weights in the interpolation scheme for $u(\text{cp}(\mathbf{x}_i))$. That is, the components of the matrix $\mathbf{E} = [\gamma_{ij}]$ are

$$\gamma_{ij} = \begin{cases} w_j & \text{if } \mathbf{x}_j \text{ is in the interpolation stencil for } \text{cp}(\mathbf{x}_i), \\ 0 & \text{otherwise,} \end{cases}$$

with w_j denoting the weight associated with the grid point x_j in the interpolation scheme for point $\text{cp}(\mathbf{x}_i)$.

2.2.2. The discrete differential operator. We are now in a position to introduce the discrete differential operator. Our exposition will treat the Laplace–Beltrami operator Δ_S and the associated in-surface heat equation. More general operators are certainly possible, however, the Laplace–Beltrami operator is particularly instructive because its spatial discretization may be stiff and because it is the key building block for the nonlinear and high-order flows we shall consider in subsequent sections.

We begin by recalling that for points \mathbf{x} on the surface, $\Delta_S u(\mathbf{x})$ is consistent with $\Delta(u(\text{cp}(\mathbf{x})))$. We approximate the function $\Delta(u(\text{cp}(\mathbf{x})))$ with a vector of m values, with the i^{th} entry corresponding to an approximation of $\Delta(u(\text{cp}(\mathbf{x}_i)))$ for $\mathbf{x}_i \in L$. The vector is given by the right-hand side of

$$\Delta u(\text{cp}(\mathbf{x})) \approx \Delta_h \begin{pmatrix} u(\text{cp}(L)) \\ u(\text{cp}(G)) \end{pmatrix}, \quad (2.6)$$

where Δ_h is an $m \times (m+m_g)$ matrix which approximates the Laplacian operator in \mathbb{R}^d using a linear finite difference scheme (e.g., second-order or fourth-order centered differences). The finite difference scheme is applied at each point in L by taking a combination of neighbouring grid points appearing in a *differentiation stencil*. Some of these neighbouring grid points will be in L with corresponding values of $u(\text{cp}(\mathbf{x}))$ in the vector $u(\text{cp}(L))$; the remaining grid points will be ghost points in G , and hence $u(\text{cp}(G))$ also appears in (2.6). For example, Figure 2.1 shows two differentiation stencils, one of which includes a ghost point.

From (2.5), approximations of the vectors $u(\text{cp}(L))$ and $u(\text{cp}(G))$ are obtained by left multiplying the vector \mathbf{u} by the extension matrix \mathbf{E} . Applying this to (2.6) yields

$$\Delta u(\text{cp}(\mathbf{x})) \approx \Delta_h \begin{pmatrix} u(\text{cp}(L)) \\ u(\text{cp}(G)) \end{pmatrix} \approx \Delta_h \mathbf{E}\mathbf{u} = \begin{array}{c} \boxed{} \\ \Delta_h \end{array} \begin{array}{c} \boxed{\phantom{\mathbf{E}}} \\ \mathbf{E} \end{array} \begin{array}{c} \boxed{\phantom{\mathbf{u}}} \\ \mathbf{u} \end{array} = \begin{array}{c} \boxed{\phantom{\tilde{\mathbf{M}}}} \\ \tilde{\mathbf{M}} \end{array} \begin{array}{c} \boxed{\phantom{\mathbf{u}}} \\ \mathbf{u} \end{array} = \tilde{\mathbf{M}}\mathbf{u},$$

where $\tilde{\mathbf{M}} = \Delta_h \mathbf{E}$ is an $m \times m$ matrix. Thus a spatial discretization of the in-surface heat equation $u_t = \Delta_S u$ is

$$\frac{\partial}{\partial t} \mathbf{u} = \tilde{\mathbf{M}}\mathbf{u}. \quad (2.7)$$

Note that the ghost points and $u(\text{cp}(G))$ were used only in the derivation; the system (2.7) is defined on the list of grid points L .

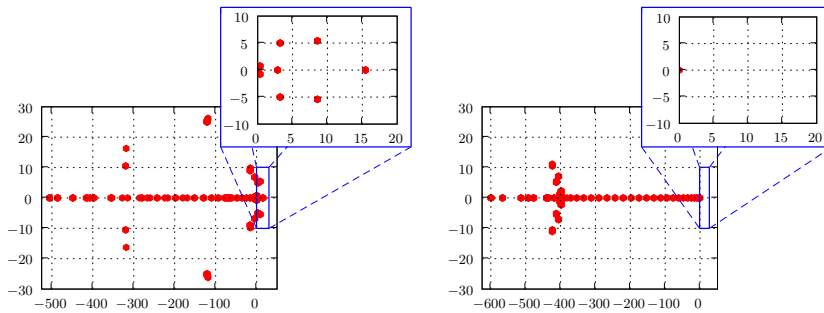


Fig. 2.2: Spectra of the \tilde{M} (left) and M (right) matrices. The geometry is a unit circle in 2D, using interpolation matrix E with biquartic interpolation ($p = 4$), and a mesh spacing of $\Delta x = 0.1$. Observe that \tilde{M} has eigenvalues in the right half plane.

2.2.3. Stabilizing the Closest Point Method. By examining the spectra of \tilde{M} given in Figure 2.2 (left), we find that the matrix has some eigenvalues with positive real parts. In the solution of the semi-discrete system (2.7), these positive eigenvalues correspond to exponential growth in the associated eigenvectors. But the Laplace–Beltrami operator is diffusive so these exponentially growing components correspond to instability, and in fact, this instability is observed in practice. For example, Figure 2.3(a) displays the oscillatory results obtained for an in-surface heat flow calculation using the implicit Closest Point Method with \tilde{M} .

Fortunately, a stable modification of the implicit Closest Point Method procedure is straightforward to construct: when approximating $\Delta_S u$ at the point \mathbf{x}_i with a finite difference scheme, we map only the neighbouring points \mathbf{x}_j of the stencil back to their closest points $\text{cp}(\mathbf{x}_j)$ and use u_i itself instead of $u(\text{cp}(\mathbf{x}_i))$ (see also (2.4)). This special treatment of the diagonal elements yields a new $m \times m$ matrix

$$M = \text{stab}(\Delta_h, E) := \text{diag } \Delta_h + (\Delta_h - \text{diag } \Delta_h)E, \quad (2.8a)$$

or diagrammatically

$$\begin{array}{c} \boxed{\text{M}} \\ \text{M} \end{array} = \begin{array}{c} \boxed{\text{diag } \Delta_h} \\ \text{diag } \Delta_h \end{array} + \begin{array}{c} \boxed{\Delta_h - \text{diag } \Delta_h} \\ \Delta_h - \text{diag } \Delta_h \end{array} \begin{array}{c} \boxed{\text{E}} \\ \text{E} \end{array}, \quad (2.8b)$$

and thus the spatial discretization of the in-surface heat equation is

$$\frac{\partial}{\partial t} \mathbf{u} = M\mathbf{u}. \quad (2.9)$$

As explained in Section 2.1, this splitting eliminates redundant interpolations and does not impact the consistency of u at the surface. Intuitively, we expect improved stability from this splitting because the magnitude of the diagonal of the operator has generally increased relative to the off-diagonal elements (see [22, Section 3.3.3]). Indeed, a study of the spectra of M for a variety of surfaces [22, Section 3.3.1] indicates that all eigenvalues of M have negative real parts. See Figure 2.2(b) for the spectra of the heat equation on a unit circle in 2D and Figure 2.3(b) for the (stable) results from a related numerical experiment.




Fig. 2.3: Stable and unstable solutions of the in-surface heat equation on a unit circle embedded in 2D, with $\Delta x = 0.1$, degree $p = 4$ interpolation, and backward Euler time stepping until $t = 0.6$ with $\Delta t = \frac{1}{8}\Delta x$. Note oscillations indicating instability in (a) whereas the numerical solution essentially overlaps the exact solution in (b).

Remark. The two steps of the explicit Closest Point Method for solving $u_t = \Delta_S u$ can also be formulated using the matrices $\mathbf{\Delta}_h$ and \mathbf{E} as follows: (a) perform a forward Euler time step $\tilde{\mathbf{u}}^{n+1} = \mathbf{u}^n + \Delta t \mathbf{\Delta}_h \mathbf{u}^n$; and (b) perform a closest point extension $\mathbf{u}^{n+1} = \mathbf{E} \tilde{\mathbf{u}}^{n+1}$. Combining these two steps yields

$$\mathbf{u}^{n+1} = \mathbf{E} \mathbf{u}^n + \Delta t \mathbf{\Delta}_h \mathbf{E} \mathbf{u}^n,$$

which is notably *not* the forward Euler time-stepping method applied to either (2.7) or (2.9). However, it suggests another possible implicit discretization based on backward Euler:

$$\mathbf{u}^{n+1} = \mathbf{E} \mathbf{u}^n + \Delta t \mathbf{\Delta}_h \mathbf{E} \mathbf{u}^{n+1}. \quad (2.10)$$

Unfortunately, numerical experiments indicate that (2.10) is often unstable when time steps larger than those possible with forward Euler are used.

2.3. Truncation error analysis. The error from approximating the Laplace–Beltrami operator using the implicit Closest Point Method has two components, one arising from the polynomial interpolation and a second arising from the discretization of the Laplacian operator $\mathbf{\Delta}_h$.

Assume that p is the degree of interpolation to be used in the closest point extension procedure. Also assume an order- q discretization of the underlying Laplacian operator. In [30] it is shown that if an interpolation of degree $p \geq q + 1$ is chosen then we maintain the order of accuracy of the underlying spatial discretization, i.e., we obtain an $O(\Delta x^q)$ truncation error. Conversely, if $p < q + 1$ the error will be dominated by the interpolation error, and we obtain an $O(\Delta x^{p-1})$ truncation error. Notice, in particular, that the method will be inconsistent if linear interpolation is chosen. Section 4.1 compares these analytically derived truncation errors with numerical experiments.

3. Time stepping and matrix properties. The temporal discretization of the semidiscrete problem (2.9) can be performed using implicit linear multistep methods²,

²Implicit Runge–Kutta methods are also straightforwardly accommodated, particularly for the case of diagonally implicit schemes [21, 17].

Δx	$\text{cond}(\mathbf{A}_{\text{BE}})$
0.4	94.1
0.2	281
0.1	664
0.05	1520
0.025	3114
0.0125	6289
0.00625	12703
0.003125	25984

Table 3.1: Condition numbers for the time-stepping matrix $\mathbf{A}_{\text{BE}} = \mathbf{I} - \Delta t \mathbf{M}$ with $\Delta t = \frac{1}{4} \Delta x$ for various Δx . The curve \mathcal{S} is a unit circle in 2D and the interpolation is of degree $p = 4$.

such as the backward differentiation formulas [12]. For each time step, the BDF methods perform a single linear system solve of the form $\mathbf{A}\mathbf{u} = \mathbf{b}$ with $\mathbf{A} = \mathbf{I} - \gamma \Delta t \mathbf{M}$ for some value of γ that depends on the particular method. For example, the BDF2 applied to (2.9) leads to $[\mathbf{I} - \frac{2}{3} \Delta t \mathbf{M}] \mathbf{u}^{n+1} = \frac{4}{3} \mathbf{u}^n - \frac{1}{3} \mathbf{u}^{n-1}$. We denote the time-stepping matrices corresponding to backward Euler and the BDF schemes as \mathbf{A}_{BE} , \mathbf{A}_{BDF2} and \mathbf{A}_{BDF4} . The structure and properties of these matrices are important for understanding the behavior of direct and iterative methods for solving the linear systems.

3.1. Matrix properties. As previously noted in Section 2.2.3 the eigenvalues of \mathbf{M} have negative real parts, leading to stability when the evolution is treated with an appropriate (e.g., A-stable) time-stepping scheme.

Table 3.1 lists the condition number for the time-stepping matrix $\mathbf{A}_{\text{BE}} = \mathbf{I} - \Delta t \mathbf{M}$ for various Δx . Note that with $\Delta t = O(\Delta x)$, the condition number of the time-stepping matrix \mathbf{A}_{BE} is moderate and only doubles when Δx is halved. Similar results are obtained for the time-stepping matrices \mathbf{A}_{BDF2} and \mathbf{A}_{BDF4} . The time-stepping matrices thus appear to be well-conditioned. As discussed in Section 3.2, this is important for solving the linear system at each time step.

Figure 3.1 shows the sparsity structure of the $\mathbf{\Delta}_h$, \mathbf{E} and \mathbf{M} matrices for various surfaces, values of the grid spacing Δx , and degree of interpolation p . The effects of the ghost points G are noticeable on the right of $\mathbf{\Delta}_h$ and the bottom of \mathbf{E} . Note that the composed matrix \mathbf{M} has a band-limited profile; this feature is related to the ordering of the lists L and G chosen by the banding algorithm in Appendix A. Table 3.2 shows some sparsity properties of \mathbf{M} including the number of nonzero elements and the bandwidth. The increasing sparsity as Δx decreases demonstrates the importance of iterative solvers as discussed further the next section.

3.2. Solving the linear systems. Performing time stepping for the implicit Closest Point Method results in large linear systems of the form $\mathbf{A}\mathbf{u} = \mathbf{b}$. As noted above, the time-stepping matrices are well-conditioned for the BDF schemes, and thus direct solves based on variations of LU factorization work well, at least for small systems such as arise in 2D with a moderate or large grid spacing Δx . The MATLAB backslash operator or the sparse solver `scipy.sparse.linalg.spsolve` in SciPy [20] is used in many of the calculations in Section 4.

The time-stepping matrices become increasingly sparse as Δx decreases (e.g., Table 3.2), potentially causing direct methods to become inefficient due to “fill-in” from the LU factorization. For large systems, particularly in 3D, iterative schemes are more efficient both in terms of memory and CPU time. It was shown experimentally in [22, Sections 3.3.3 and 3.3.4] that the backward Euler time-stepping matrix $\mathbf{I} - \Delta t \mathbf{M}$ associated with the implicit Closest Point Method matrix \mathbf{M} is diagonally

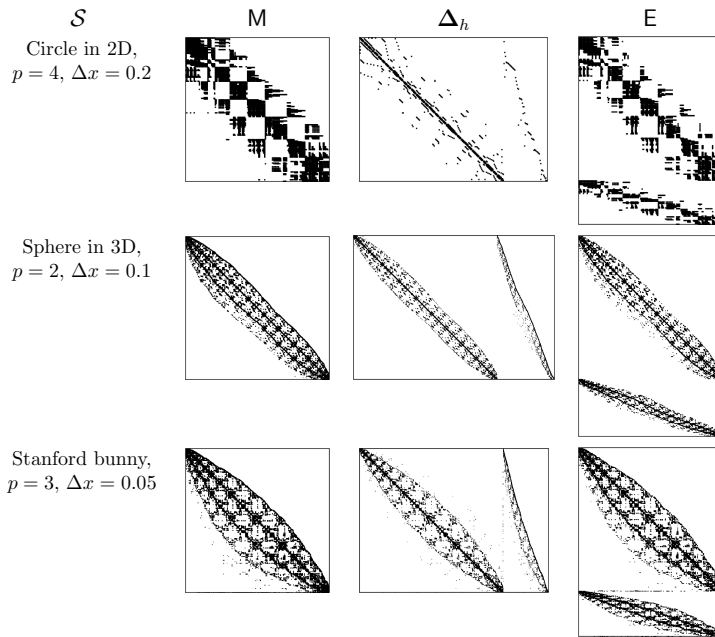


Fig. 3.1: Sparsity structure of the M , Δ_h and E matrices (from left to right) for the Laplace–Beltrami operator on various surfaces using degree- p interpolation.

Table 3.2: Sparsity properties of the matrix M for a unit circle in 2D. An increase in Δx or in the degree of interpolation p leads to an increase in the percentage of nonzero entries. Roughly the same relation is reflected in the bandwidth as a percentage of total width.

Δx	nonzero entries (percentage)				bandwidth (percentage)			
	$p = 2$	$p = 3$	$p = 4$	$p = 5$	$p = 2$	$p = 3$	$p = 4$	$p = 5$
0.2	12.8	16.3	18.2	21.3	31.7	55.4	51.5	65.6
0.1	6.5	8.0	9.0	10.5	15.8	21.4	27.2	34.1
0.05	3.3	3.9	4.5	5.2	8.8	11.0	18.2	16.3
0.025	1.6	1.9	2.3	2.6	5.6	6.4	8.6	8.1
0.0125	0.8	1.0	1.1	1.3	2.9	2.8	4.5	3.9

dominant and (unsymmetric) positive definite [13] for time steps of $\Delta t = O(\Delta x^2)$: this guarantees convergence for many iterative schemes. Naturally, in practice we wish to use larger time steps $\Delta t = O(\Delta x)$, and for these step sizes, the backward Euler matrix $I - \Delta t M$ is indefinite and not diagonally dominant. It is interesting to note that in the examples we have considered [22, Section 3.3.1] the spectral radii of the Jacobi and Gauss–Seidel iteration matrices are all strictly less than one, and thus we can expect convergence for those methods. In many of the examples we provide in Section 4, the GMRES algorithm [38] is used instead, as it seems to converge well in practice. In these cases, we chose GMRES subroutines available within either SciPy (`scipy.sparse.linalg.gmres`) or MATLAB (`gmres`).

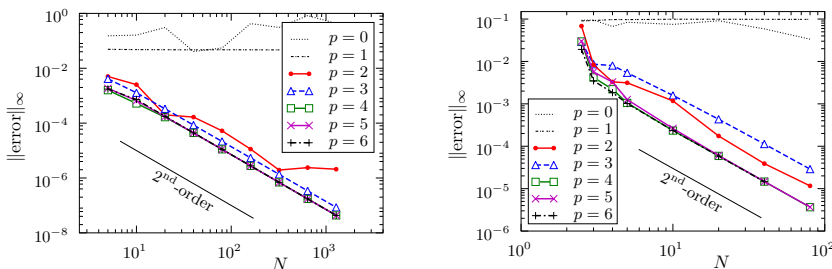


Fig. 4.1: Numerical convergence study results for the in-surface heat equation on a circle (left) and a sphere (right) using second-order finite differences, BDF2 time stepping, and degree p interpolation. Here $N = \frac{1}{\Delta x}$.

We remark that it is not necessary to explicitly construct the matrices M or A in order to use these iterative schemes. For example, GMRES requires only a subroutine that computes the product Au . However, in the numerical results that follow, we have elected to explicitly construct and use the matrices M , A_{BE} , A_{BDF2} and A_{BDF4} .

4. Numerical results.

4.1. Numerical convergence studies. The convergence of the implicit Closest Point Method is tested on two problems.

2D test problem. The problem consists of the in-surface heat equation (2.1) on a unit circle in \mathbb{R}^2 with initial conditions $u(0, \theta) = \cos \theta + \cos 3\theta$. The exact solution is $u(t, \theta) = e^{-t} \cos(\theta) + e^{-9t} \cos(3\theta)$. In the tests below, the problem is solved discretely in time until $t = \frac{1}{2}$ using BDF schemes. The exact solution is used as the starting procedure. Linear systems are solved directly with MATLAB's backslash operator.

3D test problem. The problem consists of the in-surface heat equation on a unit sphere in \mathbb{R}^3 with the exact solution given as a series of spherical harmonics $u(t, \theta, \phi) = \frac{20}{3\pi} \sum_{l=1}^{\infty} e^{-l^2/9} e^{-tl(l+1)} Y_{ll}(\theta, \phi)$, where Y_{lm} represents the degree l order m real spherical harmonic. In practice, the rapid decay of the coefficients means that the series can be truncated after only a few terms. The problem is evolved in time until $t = \frac{1}{2}$ using BDF schemes and the exact solution as a starting procedure. Linear system solves are performed in MATLAB with `gmres` using a tolerance of 1×10^{-10} .

4.1.1. Heat equation, second-order differences in space. Figure 4.1 shows the results of a numerical convergence study on these two problems using second-order finite differences, the second-order BDF2 scheme with $\Delta t = \frac{1}{8} \Delta x$, and various degrees p of interpolating polynomials. We note clear second-order convergence using $p \geq 3$. The results are consistent with the analysis in Section 2.3. The computational results appear identical for $p \geq 4$ because the error is dominated by the Laplacian discretization error term and the interpolation error term is insignificant. Degree $p = 3$ still exhibits second-order convergence but with a larger error constant which suggests that the interpolation error dominates. Note that $p = 2$ behaves surprisingly well (often better than the predicted first-order). As expected, degree $p \leq 1$ interpolation appears inconsistent.

Larger values for the interpolation degree p result in larger interpolation stencils which translates into denser matrices (see Table 3.2). For second-order finite differences, degree $p = 3$ or $p = 4$ are both good choices as they exhibit clear second-order

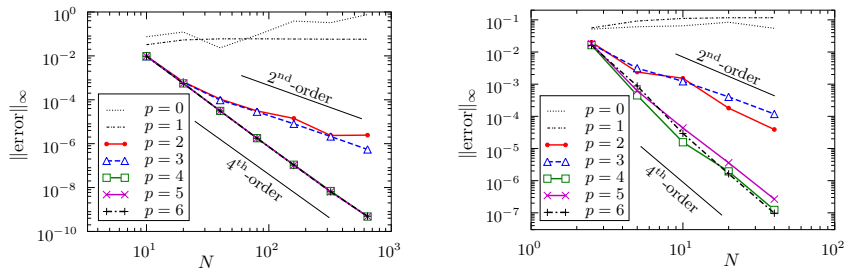


Fig. 4.2: Numerical convergence study results for the in-surface heat equation on a circle (left) and a sphere (right) using fourth-order finite differences, BDF4 time stepping, and degree p interpolation. Here $N = \frac{1}{\Delta x}$.

convergence although the error for $p = 3$ is typically larger than for $p = 4$. Degree $p \geq 5$ would be excessive for second-order calculations as no benefit in accuracy results from the extra work required. Degree $p = 2$ may be appealing for applications where calculation speed is more important than accuracy.

4.1.2. Heat equation, fourth-order differences in space. The convergence of the implicit Closest Point Method is next tested using a fourth-order centered finite difference stencil for Δ_h . Figure 4.2 shows the results of a numerical convergence study on the two test problems above, where time stepping is performed using the fourth-order BDF4 scheme with $\Delta t = \Delta x$ (on the circle) and $\Delta t = \frac{1}{4}\Delta x$ (on the sphere). We note again that $p = 0$ and $p = 1$ degree interpolation appear inconsistent. Degree $p = 2$ or $p = 3$ produce second-order results. Degree $p \geq 4$ produces fourth-order convergence. Note that on this problem, $p = 2$ and $p = 4$ performed roughly one order better than predicted by Section 2.3.

4.2. Heat equation on surfaces. The computations in this section were performed in Python [42] using SciPy [20] and NumPy [26] for numerical calculations. Visualizations use VTK [33].

Throughout this section, closest point representations of surfaces are obtained from complicated triangulated surfaces using the method described in [23]. Briefly, for each triangle in a triangulation, the method determines all nodes that are within the bandwidth of the triangle (for computations using the Closest Point Method, nodes that are further away cannot be influenced by that part of the surface). For each node, this gives a list of triangles. This list is sufficiently small that it is efficient to directly examine each member to determine the closest triangle (and hence the closest point on the surface). Note that this method scales linearly with the number of triangles under the assumption that the closest point function is only needed on a narrow computational band.

4.2.1. Heat equation on Laurent's Hand. Figure 4.3 shows the results of blurring an image on the surface of Laurent's Hand [31] by solving several steps of the in-surface heat equation. In this example, the phone number written on the hand is completely indistinguishable after only a few steps. Note that the Closest Point Method could also be used to apply an in-surface unsharp mask or other image processing technique to attempt to recover the phone number from the blurred image.

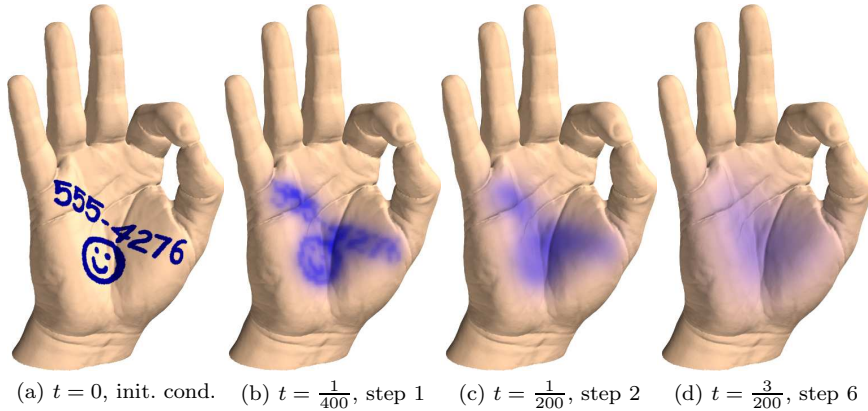


Fig. 4.3: Blurring on the surface of Laurent’s Hand [31] by solving the in-surface heat equation. The phone number becomes increasingly blurred with each time step. Here $\Delta x = \frac{1}{40}$ and the length of the hand is about $80\Delta x$. Time stepping is performed with backward Euler and $\Delta t = \frac{\Delta x}{10}$. The closest point extension uses degree $p = 2$ interpolation.

4.2.2. Heat equation on a complicated domain. Consider a composite domain consisting of Annie Hui’s pig [19] connected by the tail to a filament (an infinitesimally thin wire) which is also connected to a sphere. This surface is therefore comprised of components of mixed codimension: the pig and sphere are codimension-one whereas the filament is codimension-two. The closest point representation is straightforward to obtain for this type of composite domain: simply compute the closest point in each of the three components and then return the closest of the three.

A heat source is applied under the sphere. The temperature u over the surface of the domain is modelled according to

$$u_t = k(\mathbf{x})\Delta_{\mathcal{S}}u + \gamma(T - u)g(\mathbf{x}), \quad (4.1a)$$

where the heat coefficient $k(\mathbf{x})$ is chosen as

$$k(\mathbf{x}) = \begin{cases} 1 & \mathbf{x} \in \text{pig}, \\ 2 & \mathbf{x} \in \text{sphere}, \\ 10 & \mathbf{x} \in \text{filament}, \end{cases} \quad (4.1b)$$

and the forcing term consists of Newton’s law of heating and cooling with constant $\gamma = 30$ and maximum temperature $T = 10$. The function $g(\mathbf{x})$ is a Gaussian blob in \mathbb{R}^3 which localizes the heat source directly under the sphere. Initially the temperature is set to zero everywhere. The problem is solved using $\Delta x = 0.05$ and degree $p = 3$ interpolation. Time stepping is performed to time $t = 25$ using the BDF2 scheme with $\Delta t = 2\Delta x$ and a single step of backward Euler as a starting procedure. Thus, only 250 steps are used in contrast with the roughly 10,000 steps required for explicit time stepping for this problem.

Figure 4.4 shows the initial conditions and heat distribution at various times. This example demonstrates the great flexibility of the method with respect to geometry:

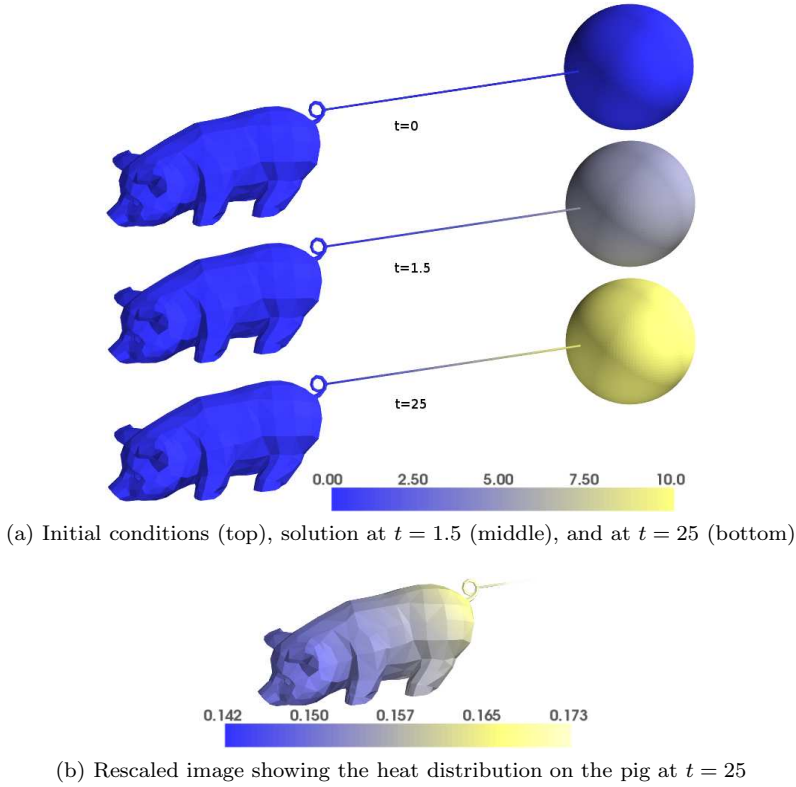


Fig. 4.4: Numerical solution of the in-surface heat equation on the composite pig-wire-sphere domain. At $t = 25$, the temperature is roughly constant at $u = 10$ on the sphere. The pig has warmed up only slightly and there is a temperature gradient in the wire.

heat flow over a complicated geometry of variable codimension is computationally no more complicated than for a sphere.

4.2.3. Pattern formation. The solutions of reaction-diffusion equations can exhibit rich pattern forming behaviour from random initial conditions. The application of reaction-diffusion equations on surfaces has been investigated in previous studies, such as [40] (by computing directly on a surface mesh) and [6] (using a level set representation of the surface). In this section, the implicit Closest Point Method is applied to pattern formation problems.

A possible set of reaction-diffusion equations for two chemicals u and v is

$$u_t = f(u, v) + \nu_u \Delta_S u, \quad u(0, \mathbf{x}) = u_0(\mathbf{x}), \quad (4.2a)$$

$$v_t = g(u, v) + \nu_v \Delta_S v, \quad v(0, \mathbf{x}) = v_0(\mathbf{x}), \quad (4.2b)$$

where f and g are nonlinear reaction terms [39]. The chemicals locally react with one another while at the same time diffusing spatially at different rates. We consider one particular form of (4.2) known as the Brusselator [28, 43] which has

$$f(u, v) = a - (b - 1)u + u^2v, \quad g(u, v) = bu - u^2v.$$

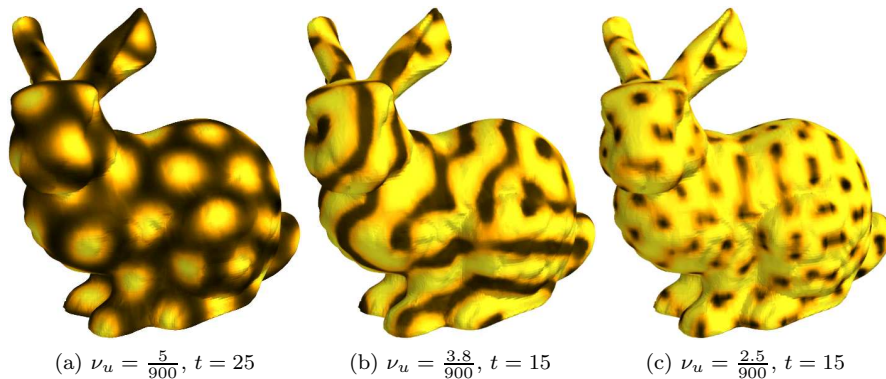


Fig. 4.5: Patterns formed by the Brusselator on the surface of the Stanford Bunny for different values of ν_u with $a = 3$, $b = 10.2$ and $\nu_v = 10/900$. Darker color indicates high concentrations of u . The computation uses $\Delta t = \Delta x = 0.05$ with degree $p = 2$ interpolation and the bunny is scaled to be approximately 2 units long.

The implicit Closest Point Method can be applied to (4.2) using implicit-explicit (IMEX) time-stepping schemes [3, 2, 29]. Using this approach, the nonlinear reaction terms are treated explicitly and the linear diffusion terms are treated implicitly. Notice that this decouples the implicit system solves for u and v and we solve two systems of the form $\mathbf{A}_{\text{BDF2}}\mathbf{u} = \mathbf{b}_1$ and $\mathbf{A}_{\text{BDF2}}\mathbf{v} = \mathbf{b}_2$ to advance to the next time step. The SBDF2 scheme [3] is known to work well for pattern formation problems [29], and it is used here with one step of IMEX Euler as a starting method. For initial conditions, small random perturbations around the zero-diffusion ($\mu_u = \mu_v = 0$) steady state of $u = a$ and $v = \frac{b}{a}$ are used.

Depending on the values of the coefficients a , b , ν_u and ν_v , the Brusselator exhibits different types of pattern forming behaviour [43]. Figure 4.5 shows the results of three simulations of the Brusselator on the surface of the Stanford Bunny [41] with $a = 3$, $b = 10.2$, $\nu_v = 10/900$ and various values of ν_u . This choice of parameters closely matches those used for planar calculations in [43, Figure 2a]³, and indeed we notice the same transition from “honeycomb” (Figure 4.5(a)) to stripes (b) to roughly hexagonal patterns of spots (c).

5. Biharmonic problems. Consider the fourth-order problem

$$u_t = -\nabla_{\mathcal{S}}^4 u, \quad (5.1)$$

where $\nabla_{\mathcal{S}}^4$ is the biharmonic operator intrinsic to the surface. A naïve approach to solving this problem via the Closest Point Method would introduce a single extension matrix into the discretization. Unfortunately, when the gradient and divergence operators are combined into a biharmonic operator using a single extension matrix, the hypotheses for the Principles in Section 1 are violated. We therefore expect that this approach will not work, and indeed, numerical experiments demonstrate the inconsistency of the approach [22].

³We rescale the equations to the size of the bunny which accounts for the additional factors of $\frac{1}{900}$ which do not appear in [43].

In [30], it was argued that operators which are composed of multiple Laplace–Beltrami $\Delta_{\mathcal{S}}$ operators (or other high-order derivative terms) could be implemented using multiple extension steps. Recalling that $\nabla_{\mathcal{S}}^4 u = \Delta_{\mathcal{S}}(\Delta_{\mathcal{S}}u)$ and applying the principles in Section 1

$$\nabla_{\mathcal{S}}^4 u(\mathbf{x}) = \Delta_{\mathcal{S}}(\underbrace{\Delta(u(\text{cp}(\mathbf{x})))}_v) = \Delta(v(\text{cp}(\mathbf{x}))) = \Delta(\underbrace{[\Delta(u(\text{cp}(\mathbf{x})))]}_v)(\text{cp}(\mathbf{x}))$$

for points \mathbf{x} on the surface, that is, a closest point extension of the inner Laplacian v is performed before computing the outer Laplacian. With Δ_h as the discrete diffusion operator from Section 2.2.2, we could consistently approximate the previous expression using the matrix

$$\widetilde{M}_{\text{bi}} = \Delta_h \mathbf{E} \Delta_h \mathbf{E}.$$

Similar to the case of the Laplace–Beltrami operator, the matrix $-\widetilde{M}_{\text{bi}}$ has eigenvalues with positive real components, leading to instability. A stable variant arises naturally by using the procedure from Section 2.2.3. Specifically, this leads to the matrix

$$M_{\text{bi}} = \text{stab}(\Delta_h, \mathbf{E}) \text{stab}(\Delta_h, \mathbf{E}) = \mathbf{M}\mathbf{M},$$

which in turn gives rise to the ODE system

$$\frac{\partial}{\partial t} \mathbf{u} = -M_{\text{bi}} \mathbf{u}. \quad (5.2)$$

Note that this is a “one-line” change in computer code for the in-surface heat equation, requiring only the squaring of the \mathbf{M} matrix.

5.1. Numerical convergence studies. The convergence of the implicit Closest Point Method for biharmonic problems is tested on the two problems of Section 4.1, where here the exact solutions are $u(t, \theta) = e^{-t} \cos(\theta) + e^{-81t} \cos(3\theta)$ and $u(t, \theta, \phi) = \frac{20}{3\pi} \sum_{l=1}^{\infty} e^{-l^2/9} e^{-tl^2(l+1)^2} Y_{ll}(\theta, \phi)$. The 2D problem is solved until $t = 0.5$ and the 3D problem is solved until $t = 0.15$. Figure 5.1 demonstrates that the Closest Point Method achieves the expected accuracy using both second-order and fourth-order finite differences, although the error appears to stagnate at around 10^{-7} . Note that the entries in the M_{bi} matrix are each multiplied by $1/\Delta x^4$. Heuristically, the stagnating error appears to be consistent with rounding errors that are magnified by about $1/\Delta x^4$ with $\Delta x \approx 10^{-2}$.

Our experiments demonstrate that the implicit Closest Point Method gives a viable and extremely simple way to treat problems with biharmonic terms, although for highly accurate results some attention should be paid to the control of roundoff error. In the following section, we apply the technique to a fourth-order problem of interface motion.

5.2. “Surface diffusion” on a surface. In the “surface diffusion” problem, “atoms [on an interface] diffuse from areas of high mean curvature to low mean curvature” [34]. In the standard settings of \mathbb{R}^2 and \mathbb{R}^3 , the interface moves at a speed proportional to the intrinsic Laplacian of its curvature. Here, we are interested in the corresponding motion of an interface (defined by the zero contour of the function u) on a curved surface \mathcal{S} : the term “surface diffusion” is somewhat misleading in this case because the diffusion is occurring on the interface (which in turn lies on the surface \mathcal{S}).




Fig. 5.1: Biharmonic numerical convergence studies using the operator M_{bi} with second-order finite differences (and BDF2 time stepping) and fourth-order finite differences (and BDF4 time stepping). The calculations on circles use $\Delta t = \frac{1}{4}\Delta x$ and the calculations on the sphere use $\Delta t = \frac{1}{8}\Delta x$. Note $N = \frac{1}{\Delta x}$.

A level set formulation of this problem is given by the fourth-order nonlinear surface PDE

$$u_t = -K_{\mathcal{S}}(u) = -\Delta_I \kappa |\nabla_{\mathcal{S}} u|,$$

where $\kappa = \nabla_{\mathcal{S}} \cdot (\nabla_{\mathcal{S}} u / |\nabla_{\mathcal{S}} u|)$ is the in-surface curvature of the local level set of u . The Laplacian Δ_I operator is the in-surface analogue of the “surface Laplacian” operator in [34]. Smereka [34] presented a simple and effective splitting

$$u_t = -\beta \nabla_{\mathcal{S}}^4 u + \beta \nabla_{\mathcal{S}}^4 u - K_{\mathcal{S}}(u),$$

from which a semi-discretization in time using IMEX Euler is

$$\frac{u^{n+1} - u^n}{\Delta t} = -\beta \nabla_{\mathcal{S}}^4 u^{n+1} + \beta \nabla_{\mathcal{S}}^4 u^n - K_{\mathcal{S}}(u^n) + O(\Delta t).$$

Notice this is linear in u^{n+1} because the fourth-order $K_{\mathcal{S}}(u)$ is treated explicitly. A full discretization follows using the discrete M_{bi} operator

$$[I + \Delta t \beta M_{\text{bi}}] (\mathbf{u}^{n+1} - \mathbf{u}^n) = -\Delta t K^h(\mathbf{u}^n),$$

where K^h is a discrete form of $K_{\mathcal{S}}$ obtained using the explicit Closest Point Method. That is, the term $K^h(\mathbf{u}^n)$ is formed by first computing κ , the mean curvature of \mathbf{u}^n , in the embedding space \mathbb{R}^d . This computation is followed by a closest point extension of κ . Finally, $\Delta_I \kappa$ is computed by following the procedure in [34, 32]. This explicit computation is performed only locally near the interface. We use $\beta = \frac{1}{2}$ as is recommended in [34].

Figure 5.2 shows two example computations. Our goal here is to demonstrate the effectiveness of the implicit Closest Point Method on this problem and we have not implemented the full algorithm of [34]. Specifically, some care must be taken when dealing with merging of interfaces. Finally, for simplicity, we have not implemented the local (to the interface) level set approach of [32].

6. Odd-order terms. So far, this paper has concentrated on Laplacian–Beltrami and in-surface biharmonic operators. However, the Closest Point Method can also be applied to odd-order terms. We outline a brief description of a possible approach, based on discretizing the surface divergence of a vector field intrinsic to the surface.




Fig. 5.2: “Surface diffusion” of interfaces on surfaces. Top row (a): an initial interface on a unit sphere evolves toward a circular steady state (from left-to-right: $t = 0$, $t = 3\Delta t$, $t = 9\Delta t$, $t = 18\Delta t$ with $p = 3$, $\Delta x = 0.1$ and $\Delta t = 100\Delta x^4$). Bottom row (b): interface evolving on a bumpy sphere [1] (left-to-right: $t = 0$, $t = 3\Delta t$, $t = 6\Delta t$ and $t = 50\Delta t$ with $p = 3$, $\Delta x = 0.05$, $\Delta t = 50\Delta x^4$).

6.1. Surface divergence terms. Consider a vector field $\mathbf{w}(\mathbf{x}) : \mathcal{S} \rightarrow \mathbb{R}^d$ intrinsic to the surface \mathcal{S} , that is, any vector in \mathbf{w} is tangent to \mathcal{S} . Now define the vector field $\mathbf{v} : \mathbb{R}^d \rightarrow \mathbb{R}^d$ as the closest point extension of the vector field \mathbf{w} , i.e., $\mathbf{v}(\mathbf{x}) := \mathbf{w}(\text{cp}(\mathbf{x}))$. It follows that $\mathbf{v}(\mathbf{x})$ satisfies the requirements of Principle 2 and thus

$$\nabla_{\mathcal{S}} \cdot \mathbf{w} = \nabla \cdot [\mathbf{w}(\text{cp}(\mathbf{x}))], \quad \text{for points } \mathbf{x} \text{ on the surface } \mathcal{S}.$$

We discretize in a similar fashion to Section 2.2. Let \mathbf{w}_i be a list of length m consisting of the i^{th} component of the vector field \mathbf{w} evaluated at each of the grid points in the list L . The closest point extension of each component of the vector field can then be approximated using the extension matrix \mathbf{E} . Combining this with a discretization of the divergence operator gives

$$\nabla_{\mathcal{S}} \cdot (\mathbf{w}) \approx D_{x_1} \mathbf{E} \mathbf{w}_1 + D_{x_2} \mathbf{E} \mathbf{w}_2 + \dots + D_{x_d} \mathbf{E} \mathbf{w}_d, \quad (6.1)$$

where D_{x_i} is an $m \times (m + m_g)$ matrix representing a finite difference approximation to the derivative in the x_i direction (the matrices D_{x_i} are analogous to the matrix $\mathbf{\Delta}_h$ in Section 2.2.2).

Another consistent (and potentially more stable) discretization is obtained using the splitting function $\text{stab}(\cdot, \cdot)$ given in Section 2.2.3. This yields

$$\begin{aligned} \nabla_{\mathcal{S}} \cdot \mathbf{w} &\approx \text{stab}(D_{x_1}, \mathbf{E}) \mathbf{w}_1 + \text{stab}(D_{x_2}, \mathbf{E}) \mathbf{w}_2 + \dots + \text{stab}(D_{x_d}, \mathbf{E}) \mathbf{w}_d, \\ &= \mathbf{M}_{x_1} \mathbf{w}_1 + \mathbf{M}_{x_2} \mathbf{w}_2 + \dots + \mathbf{M}_{x_d} \mathbf{w}_d, \end{aligned} \quad (6.2)$$

where the $m \times m$ matrices $\mathbf{M}_{x_i} = \text{stab}(D_{x_i}, \mathbf{E})$ approximate the differential terms of the surface divergence operator. Note that the forms (6.1) and (6.2) are identical if a central difference discretization is chosen since D_{x_i} has zeros along the main diagonal.

6.2. Third-order terms. For a given vector field \mathbf{w} , suppose we want to form

$$\Delta_{\mathcal{S}}(\nabla_{\mathcal{S}} \cdot \mathbf{w}),$$

where $\mathbf{w} : \mathcal{S} \rightarrow \mathbb{R}^d$ is a vector field defined on the surface and tangent to the surface. As in the biharmonic case in Section 5, we perform two closest point extensions. However, here we make use of Principles 1 and 2 to obtain

$$\Delta_{\mathcal{S}}(\nabla_{\mathcal{S}} \cdot \mathbf{w}) = \Delta_{\mathcal{S}}(\underbrace{\nabla \cdot (\mathbf{w}(\text{cp}(\mathbf{x})))}_u) = \Delta(u(\text{cp}(\mathbf{x}))) = \Delta(\underbrace{[\nabla \cdot (\mathbf{w}(\text{cp}(\mathbf{x})))]}_u)(\text{cp}(\mathbf{x})),$$

for points \mathbf{x} on the surface \mathcal{S} . Applying the discretizations (2.8) and (6.2) yields

$$\begin{aligned} \Delta_{\mathcal{S}}(\nabla_{\mathcal{S}} \cdot \mathbf{w}) &\approx \text{stab}(\mathbf{\Delta}_h, \mathbf{E}) \left(\text{stab}(\mathbf{D}_{x_1}, \mathbf{E}) \mathbf{w}_1 + \dots + \text{stab}(\mathbf{D}_{x_d}, \mathbf{E}) \mathbf{w}_d \right), \\ &= (\mathbf{M} \mathbf{M}_{x_1}) \mathbf{w}_1 + (\mathbf{M} \mathbf{M}_{x_2}) \mathbf{w}_2 + \dots + (\mathbf{M} \mathbf{M}_{x_d}) \mathbf{w}_d. \end{aligned} \quad (6.3)$$

Preliminary tests on the linear advection equation show the approach (6.2) is both consistent and stable. We leave to future work the testing and application of the method to third-order surface PDEs.

7. Conclusions and future work. This paper presents an implicit Closest Point method for solving surface PDEs. The method introduces one or more matrices into the discretization to carry out the interpolation steps that are central to the Closest Point Method. For the important Laplace–Beltrami and in-surface biharmonic operators we have shown that a stable variant of the algorithm is achieved by introducing the closest point operator into off-diagonal terms only. The corresponding discretization is simple and has the property that computations can be carried out in a small band near the surface with no loss of accuracy whatsoever. The discretization of the in-surface biharmonic operator is obtained trivially, by squaring the discretization of the Laplace–Beltrami operator. The power and flexibility of the implicit Closest Point Method is illustrated for a number of examples including in-surface heat flow, reaction-diffusion equations, and higher-order motions. Our experiments are carried out on a variety of interesting surfaces, including triangulated surfaces and surfaces of mixed codimension.

Our approach generalizes in a straightforward fashion to other operators, including odd-order operators and nonlinear operators [22], and we intend to investigate this further. Efficient methods for iteratively solving the implicit equations will be another focus of our studies.

Appendix A. Constructing the computational band. This appendix gives one possible algorithm to construct the band of grid points L enveloping the surface \mathcal{S} and the band of ghost points G . The algorithm is related to the “stencil set approach” in [23] and differs mainly in that for the implicit Closest Point Method we need to store the indices and weights associated with differentiation and extension in order to build the matrices $\mathbf{\Delta}_h$ and \mathbf{E} .

The algorithm proceeds in two passes. In the first pass, we build the list L of all grid points needed for the closest point extension; these are exactly the points on which the approximate solution will be propagated. This pass also identifies special *base points* $\mathbf{b} \in L$ which are grid points that appear in the lower-left corner of an interpolation stencil (e.g., the lower-left grid points in the shaded regions in Figure 2.1). Each of these base points \mathbf{b} stores the corresponding interpolation stencil as a list $S_{\text{interp}(\mathbf{b})}$ of indices pointing into the list L .

The second pass adds the list G of ghost points; these grid points appear in an differentiation stencil but not in any of the interpolation stencils. Each point \mathbf{x} in $L \cup G$ stores $\text{cp}(\mathbf{x})$, its closest point on \mathcal{S} , and the index (pointing into L) of the base point \mathbf{b} of its interpolation stencil. Points \mathbf{x} in L also store their differentiation stencil as a list $S_{\text{diff}(\mathbf{x})}$ of indices pointing into L and G .

Algorithm 1 implements the two passes. Given the computational band, Algorithm 2 constructs the Δ_h and \mathbf{E} matrices.

Algorithm 1: Building L and G

```

// First Pass
Add any one grid point near  $\mathcal{S}$  to  $L$ ;
repeat
  Get next element  $\mathbf{x}$  from  $L$ ;
  foreach point  $\mathbf{y}$  in the differentiation stencil for the point  $\mathbf{x}$  do
    Calculate  $\text{cp}(\mathbf{y})$ ;
    Find  $\mathbf{b}$  the base point of the interpolation stencil for  $\text{cp}(\mathbf{y})$ ;
    if  $\mathbf{b}$  not yet processed then
      Let  $S_{\text{interp}(\mathbf{b})} = \emptyset$ ;
      foreach  $\mathbf{z}$  in the interpolation stencil for  $\text{cp}(\mathbf{y})$  do
        Look for  $\mathbf{z}$  in  $L$ , note its index if found;
        If  $\mathbf{z}$  is not found, add it to  $L$  and note its index;
        Add the index of  $\mathbf{z}$  to  $S_{\text{interp}(\mathbf{b})}$ ;
until no more elements in  $L$  ;

// Second Pass
foreach  $\mathbf{x}$  in  $L$  do
  Set  $S_{\text{diff}(\mathbf{x})} = \emptyset$ ;
  foreach  $\mathbf{y}$  in the differentiation stencil for the point  $\mathbf{x}$  do
    Find the index in  $L$  of the base point  $\mathbf{b}$  corresponding to  $\text{cp}(\mathbf{y})$ . Associate this
    index and the value of  $\text{cp}(\mathbf{y})$  with  $\mathbf{y}$ ;
    Look for  $\mathbf{y}$  in  $L$  and  $G$ , note its index if found;
    If not found, add it to  $G$  and note its index;
    Add the index of  $\mathbf{y}$  (and which band it is in) to  $S_{\text{diff}(\mathbf{x})}$ ;

```

Algorithm 2: Forming the Δ_h and \mathbf{E} matrices

```

foreach index  $i$  where  $\mathbf{x}_i \in L$  do
  foreach  $j \in S_{\text{diff}(\mathbf{x}_i)}$ , the differentiation stencil for  $\mathbf{x}_i$  do
    Set  $[\Delta_h]_{i,j}$  to the weight for the  $j^{\text{th}}$  component of the differentiation stencil;
foreach index  $i$  where  $\mathbf{x}_i \in L \cup G$  do
  Retrieve the base point  $\mathbf{b}$  for  $\mathbf{x}_i$ , and  $\text{cp}(\mathbf{x}_i)$ ;
  Compute weights  $w^{\text{cp}(\mathbf{x}_i)}$  with barycentric Lagrange interpolation;
  foreach index  $j$  in the interpolation  $S_{\text{interp}(\mathbf{b})}$  do
    Set  $[\mathbf{E}]_{i,j} = w_j^{\text{cp}(\mathbf{x}_i)}$  (the weight corresponding to the  $j^{\text{th}}$  component of the
    interpolation stencil  $S_{\text{interp}(\mathbf{b})}$ );

```

REFERENCES

- [1] AIM@SHAPE, “Bumpy Sphere” from the AIM@SHAPE shape repository. <http://shapes.aimatshape.net>, 2008. Accessed 2008-08-13.
- [2] U. M. ASCHER, S. J. RUUTH, AND R. J. SPITERI, *Implicit-explicit Runge–Kutta methods for time-dependent partial differential equations*, Appl. Numer. Math., 25 (1997), pp. 151–167.
- [3] U. M. ASCHER, S. J. RUUTH, AND B. T. R. WETTON, *Implicit-explicit methods for time-dependent partial differential equations*, SIAM J. Numer. Anal., 32 (1995), pp. 797–823.

- [4] J.-P. BERRUT AND L. N. TREFETHEN, *Barycentric Lagrange interpolation*, SIAM Rev., 46 (2004), pp. 501–517.
- [5] M. BERTALMÍO, A. BERTOZZI, AND G. SAPIRO, *Navier–Stokes, fluid dynamics, and image and video inpainting*, Proc. of IEEE-CVPR, (2001), pp. 355–362.
- [6] M. BERTALMÍO, L.-T. CHENG, S. OSHER, AND G. SAPIRO, *Variational problems and partial differential equations on implicit surfaces*, J. Comput. Phys., 174 (2001), pp. 759–780.
- [7] M. BURGER, *Finite element approximation of elliptic partial differential equations on implicit surfaces*, CAM Report 05-46, UCLA, 2005.
- [8] L.-T. CHENG, P. BURCHARD, B. MERRIMAN, AND S. OSHER, *Motion of curves constrained on surfaces using a level-set approach*, J. Comput. Phys., 175 (2002), pp. 604–644.
- [9] G. DZIUK AND C. ELLIOTT, *Surface finite elements for parabolic equations*, J. Comp. Math., 25 (2007), pp. 385–407.
- [10] ———, *Eulerian finite element method for parabolic PDEs on implicit surfaces*, Interf. Free Bound., 10 (2008), pp. 119–138.
- [11] M. FLOATER AND K. HORMANN, *Surface parameterization: a tutorial and survey*, in Advances in Multiresolution for Geometric Modelling, Springer, 2005.
- [12] C. GEAR, *Numerical Initial Value Problems in Ordinary Differential Equations*, Prentice Hall, 1971.
- [13] G. GOLUB AND C. VAN LOAN, *Matrix Computations*, Johns Hopkins University Press, 1996.
- [14] J. B. GREER, *An improvement of a recent Eulerian method for solving PDEs on general geometries*, J. Sci. Comput., 29 (2006), pp. 321–352.
- [15] J. B. GREER, A. L. BERTOZZI, AND G. SAPIRO, *Fourth order partial differential equations on general geometries*, J. Comput. Phys., 216 (2006), pp. 216–246.
- [16] X. GU, Y. WANG, T. CHAN, P. THOMPSON, AND S. YAU, *Genus zero surface conformal mapping and its application to brain surface mapping*, Medical Imaging, IEEE Transactions on, 23 (2004), pp. 949–958.
- [17] E. HAIRER AND G. WANNER, *Solving ordinary differential equations. II: Stiff and differential-algebraic problems*, Springer-Verlag, second ed., 1996.
- [18] M. HOLST, *Adaptive numerical treatment of elliptic systems on manifolds*, in Advances in Computational Mathematics, 2001, pp. 139–191.
- [19] A. HUI, “Annie Hui’s pig”, the AIM@SHAPE shape repository. <http://shapes.aimatshape.net>, 2008. Accessed 2008-04-09.
- [20] E. JONES, T. OLIPHANT, P. PETERSON, ET AL., *SciPy: Open source scientific tools for Python*. <http://www.scipy.org>, 2001. Accessed 2008-07-15.
- [21] D. I. KETCHESON, C. B. MACDONALD, AND S. GOTTLIEB, *Optimal implicit strong stability preserving Runge–Kutta methods*, Appl. Numer. Math., 59 (2009), pp. 373–392. doi:10.1016/j.apnum.2008.03.034.
- [22] C. B. MACDONALD, *The Closest Point Method for time-dependent processes on surfaces*, PhD thesis, Simon Fraser University, August 2008.
- [23] C. B. MACDONALD AND S. J. RUUTH, *Level set equations on surfaces via the Closest Point Method*, J. Sci. Comput., 35 (2008), pp. 219–240. doi:10.1007/s10915-008-9196-6.
- [24] J. MURRAY, *Mathematical Biology II: Spatial Models and Biomedical Applications*, Springer, third ed., 2003.
- [25] O. NEMITZ, M. NIELSEN, M. RUMPF, AND R. WHITAKER, *Finite element methods on very large, dynamic tubular grid encoded implicit surfaces*, SIAM J. Sci. Comput., (submitted).
- [26] T. E. OLIPHANT, *Python for scientific computing*, Computing in Science & Engineering, 9 (2007), pp. 10–20.
- [27] L. OLSEN, P. MAINI, AND J. SHERRATT, *Spatially varying equilibria of mechanical models: Application to dermal wound contraction*, Math. Biosci., 147 (1998), pp. 113–129.
- [28] I. PRIGOGINE AND R. LEFEVER, *Symmetry breaking instabilities in dissipative systems. II*, The Journal of Chemical Physics, 48 (1968), pp. 1695–1700.
- [29] S. J. RUUTH, *Implicit-explicit methods for reaction-diffusion problems in pattern formation*, J. Math. Bio., 34 (1995), pp. 148–176.
- [30] S. J. RUUTH AND B. MERRIMAN, *A simple embedding method for solving partial differential equations on surfaces*, J. Comput. Phys., 227 (2008), pp. 1943–1961.
- [31] L. SABORET, M. ATTENE, AND P. ALLIEZ, “Laurent’s Hand”, the AIM@SHAPE shape repository. <http://shapes.aimatshape.net>, 2007. Accessed 2007-02-16.
- [32] D. SALAC AND W. LU, *A local semi-implicit level-set method for interface motion*, J. Sci. Comput., 35 (2008), pp. 330–349.
- [33] W. SCHROEDER, K. MARTIN, AND B. LORENSEN, *The Visualization Toolkit: an object-oriented approach to 3D graphics*, Prentice Hall, 1998.
- [34] P. SMEREKA, *Semi-implicit level set methods for curvature and surface diffusion motion*, J.

- Sci. Comput., 19 (2003), pp. 439–455.
- [35] J. STAM, *Flows on surfaces of arbitrary topology*, ACM Trans. Graph., 22 (2003), pp. 724–731.
 - [36] P. TANG, F. QIU, H. ZHANG, AND Y. YANG, *Phase separation patterns for diblock copolymers on spherical surfaces: A finite volume method*, Phys. Rev. E, 72 (2005), p. 016710.
 - [37] L. TIAN, C. B. MACDONALD, AND S. J. RUUTH, *Segmentation on surfaces with the Closest Point Method*, in Proc. ICIP09, International Conference on Image Processing (to appear), Cairo, Egypt, 2009.
 - [38] L. TREFETHEN AND D. BAU, *Numerical Linear Algebra*, SIAM, 1997.
 - [39] A. TURING, *The chemical basis of morphogenesis*, Phil. Trans. Roy. Soc. Lond., B237 (1952), pp. 37–72.
 - [40] G. TURK, *Generating textures on arbitrary surfaces using reaction-diffusion*, Computer Graphics, 25 (1991), pp. 289–298.
 - [41] G. TURK AND M. LEVOY, *The Stanford Bunny, the Stanford 3D scanning repository*. <http://www-graphics.stanford.edu/data/3Dscanrep>, 1994. Accessed 2008-06-18.
 - [42] G. VAN ROSSUM ET AL., *The Python programming language*. <http://www.python.org>, 1991. Accessed 2008-07-15.
 - [43] L. YANG, A. M. ZHABOTINSKY, AND I. R. EPSTEIN, *Stable squares and other oscillatory Turing patterns in a reaction-diffusion model*, Phys. Rev. Lett., 92 (2004), pp. 198–303.