Spatial Adaptivity for Solving PDEs on Manifolds with the Closest Point Method

Nathan King
University of Waterloo
Waterloo, ON, Canada
n5king@uwaterloo.ca

Steven Ruuth
Simon Fraser University
Burnaby, BC, Canada
sruuth@sfu.ca

Christopher Batty
University of Waterloo
Waterloo, ON, Canada
christopher.batty@uwaterloo.ca

ACM Reference Format:
Nathan King, Steven Ruuth, and Christopher Batty. 2025. Spatial Adaptivity for Solving PDEs on Manifolds with the Closest Point Method. In *Special Interest Group on Computer Graphics and Interactive Techniques Conference Posters (SIGGRAPH Posters '25), August 10-14, 2025.* ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3721250.3743012

1 Method

The closest point method (CPM) has recently been used in computer graphics for fluid simulations [Morgenroth et al. 2020] and geometry processing [King et al. 2024a,b]. We propose the first framework to enable spatial adaptivity with CPM, providing a more efficient spatial discretization. To solve manifold PDEs with CPM an *embedding PDE* is constructed whose solution agrees with the solution of the manifold PDE at points $\mathbf{y} \in \mathcal{S}$. The embedding PDE is solved on a tubular neighbourhood of the manifold $\mathcal{S} \subset \mathbb{R}^d$ defined by

$$\mathcal{N}(\mathcal{S}) = \left\{ \mathbf{x} \in \mathbb{R}^d \;\middle|\; \|\mathbf{x} - \mathrm{cp}_{\mathcal{S}}(\mathbf{x})\| \leq r_{\mathcal{N}(\mathcal{S})} \right\}, \tag{1}$$

where $\mathrm{cp}_{\mathcal{S}}(\mathbf{x})$ is the closest point on $\mathcal{S}$ to $\mathbf{x}$ in Euclidean distance and $r_{\mathcal{N}(\mathcal{S})}$ is the *tube radius*. Let $u_{\mathcal{S}}(\mathbf{y})$, for $\mathbf{y} \in \mathcal{S}$, and $u(\mathbf{x})$, for $\mathbf{x} \in \mathcal{N}(\mathcal{S})$, denote the functions defined on the manifold $\mathcal{S}$ and the neighbourhood $\mathcal{N}(\mathcal{S})$, respectively.

The most common discretization of CPM uses a uniformly spaced grid within $\mathcal{N}(\mathcal{S})$ with constant $r_{\mathcal{N}(\mathcal{S})}$. To provide spatial adaptivity for CPM, the tube radius $r_{\mathcal{N}(\mathcal{S})}$ must be allowed to vary over different portions of $\mathcal{S}$. We divide $\mathcal{S}$ into M pieces such that $\mathcal{S} = \mathcal{S}_1 \cup \mathcal{S}_2 \cup \cdots \cup \mathcal{S}_M$. Each subset $\mathcal{S}_m$ of $\mathcal{S}$, for $m = 1, 2, \ldots, M$, is endowed with its own tubular neighbourhood $\mathcal{N}(\mathcal{S}_m)$ that has its own tube radius $r_{\mathcal{N}(\mathcal{S}_m)}$ (see Figure 1 left). Note that each subset $\mathcal{S}_m$ is a disjoint piece of $\mathcal{S}$ except at the boundary between two subsets where $\mathcal{S}$ is duplicated.

The definition of $\mathcal{N}(\mathcal{S}_m)$ given in (1) includes the half-tubular region surrounding the boundary $\partial\mathcal{S}_m$ (darker regions in Figure 1). This half-tubular region, denoted $\mathcal{N}(\partial\mathcal{S}_m)$, is given by $\mathcal{N}(\partial\mathcal{S}_m) = \left\{ \mathbf{x} \in \mathcal{N}(\mathcal{S}_m) \;\middle|\; \mathrm{cp}_{\mathcal{S}_m}(\mathbf{x}) = \mathrm{cp}_{\partial\mathcal{S}_m}(\mathbf{x}) \right\}$. Overlaps of the neighbourhoods $\mathcal{N}(\mathcal{S}_m)$ are provided by these boundary subsets $\mathcal{N}(\partial\mathcal{S}_m)$, which are used to join subsets for the embedding PDE solve globally.

Consider the case of $M = 2$ with one boundary between $\mathcal{S}_1$ and $\mathcal{S}_2$ given by $\partial\mathcal{S}_1 = \partial\mathcal{S}_2 = \mathcal{S}_1 \cap \mathcal{S}_2$ (see Figure 1 right). CPM extends

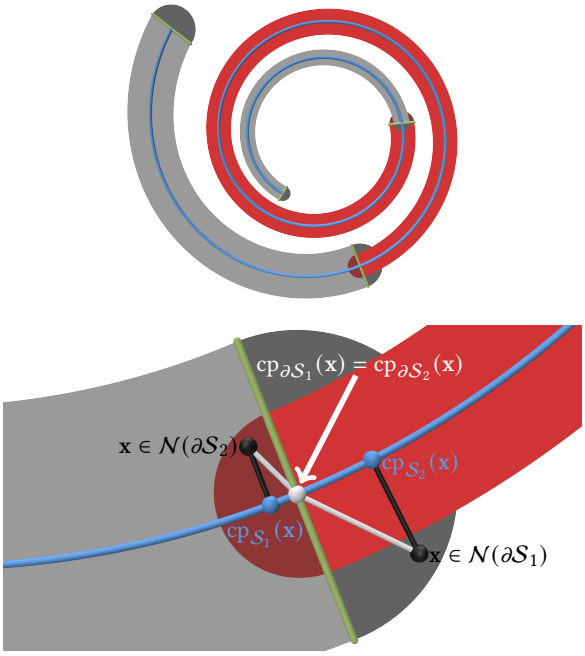

Figure 1: (Top) Overlapping adaptive tubes $\mathcal{N}(\mathcal{S}_m)$ (grey and red) for a spiral curve $\mathcal{S}$ (blue). Boundary subsets $\mathcal{N}(\partial\mathcal{S}_m)$ are coloured darker and are separated by the green lines. (Bottom) Data $u_{\mathcal{S}}$ is CP extended to a point $\mathbf{x} \in \mathcal{N}(\mathcal{S}_2)$ (black) by assigning the value of $u_{\mathcal{S}}$ at $\mathrm{cp}_{\mathcal{S}_1}(\mathbf{x})$ (blue point) instead of the value at $\mathrm{cp}_{\partial\mathcal{S}_2}(\mathbf{x})$ (white) and similarly for $\mathbf{x} \in \mathcal{N}(\partial\mathcal{S}_1)$.

manifold functions $u_{\mathcal{S}}$ onto $\mathcal{N}(\mathcal{S})$ such that the function is constant in the normal direction of $\mathcal{S}$ using the *closest point extension*. The CP extension operator E is defined as $Eu_{\mathcal{S}}(\mathbf{x}) \equiv u_{\mathcal{S}}(\mathrm{cp}_{\mathcal{S}}(\mathbf{x}))$ for $\mathbf{x} \in \mathcal{N}(\mathcal{S})$. Denote the CP extension operator for $\mathcal{N}(\mathcal{S})$, $\mathcal{N}(\mathcal{S}_1)$, and $\mathcal{N}(\mathcal{S}_2)$ by E, E_1, and E_2, respectively.

The CP extension of manifold data $u_{\mathcal{S}}$ for $\mathbf{x} \in \mathcal{N}(\partial\mathcal{S}_1)$ would ordinarily be $E_1u_{\mathcal{S}}(\mathbf{x}) = u_{\mathcal{S}}(\mathrm{cp}_{\partial\mathcal{S}_1}(\mathbf{x}))$. However, this choice is incorrect for the global problem since globally $Eu_{\mathcal{S}} = u_{\mathcal{S}}(\mathrm{cp}_{\mathcal{S}})$ and $\mathrm{cp}_{\partial\mathcal{S}_1} \neq \mathrm{cp}_{\mathcal{S}}$ on $\mathcal{N}(\partial\mathcal{S}_1)$. Since $\mathrm{cp}_{\mathcal{S}_2} = \mathrm{cp}_{\mathcal{S}}$ on $\mathcal{N}(\partial\mathcal{S}_1)$, we instead perform the CP extension for $\mathbf{x} \in \mathcal{N}(\partial\mathcal{S}_1)$ using $\mathrm{cp}_{\mathcal{S}_2}$:

$$E_1u_{\mathcal{S}}(\mathbf{x}) \equiv E_2u_{\mathcal{S}}(\mathbf{x}) = u_{\mathcal{S}}(\mathrm{cp}_{\mathcal{S}_2}(\mathbf{x})), \quad \text{for} \quad \mathbf{x} \in \mathcal{N}(\partial\mathcal{S}_1), \tag{2}$$

and vice versa for $\mathbf{x} \in \mathcal{N}(\partial\mathcal{S}_2)$. In short, points in the boundary subset for one region simply use the CP extension for the neighbouring region. CPM can then be applied on $\mathcal{N}(\mathcal{S}_1)$ and $\mathcal{N}(\mathcal{S}_2)$ with no other changes since (2) provides a sufficient condition to couple the problem globally. In general, a subset $\mathcal{S}_m$ can be bordered by

SIGGRAPH Posters '25, August 10-14, 2025, Vancouver, BC, Canada

ACM ISBN 979-8-4007-1549-5/2025/08
https://doi.org/10.1145/3721250.3743012

multiple subsets $\mathcal{S}_j$, $j \neq m$. Therefore, $E_m u_{\mathcal{S}}(\mathbf{x})$ for $\mathbf{x} \in \mathcal{N}(\partial\mathcal{S}_m)$ can be composed of CP extensions from the multiple bordering $\mathcal{S}_j$.

2 Numerical Results

In the discrete setting, we construct uniformly spaced grids $\Omega(\mathcal{S}_m)$ within each $\mathcal{N}(\mathcal{S}_m)$ for $m = 1, 2, \ldots, M$ with grid spacing h_m. Our approach adapts that of King et al. [2024b] for the construction of the computational tubes $\Omega(\mathcal{S}_m)$. Adaptivity allows efficient implementations of highly realistic simulations of phenomena in computer graphics. Manteaux et al. [2017] review these methods, highlighting the large number of criteria used to invoke adaptivity. We explore adaptivity based on solution gradients or local feature size. Further examples and details of our approach are given in [King 2025, Section 3.4].

2.1 Solution Gradients

Fine grid resolutions are required to capture large variations in the solution of a PDE. When the large variations are localized to a small region of the computational domain, spatial adaptivity is an effective approach to improve efficiency.

Consider the Poisson equation

$$\begin{aligned} \Delta_{\mathcal{S}} u_{\mathcal{S}} &= -\frac{80000}{\pi^2} \frac{\tanh\left(200\,\theta/\pi + 50\right)}{\cosh^2\left(200\,\theta/\pi + 50\right)}, \\ u_{\mathcal{S}}\left(-3\pi/4\right) &= -\tanh(100) \approx -1, \\ u_{\mathcal{S}}\left(\pi/4\right) &= \tanh(100) \approx 1, \end{aligned} \tag{3}$$

on the arc of the unit circle with $\theta \in [-\frac{3\pi}{4}, \frac{\pi}{4}]$. The exact solution is $u_{\mathcal{S}}(\theta) = \tanh\left(200\,\theta/\pi + 50\right)$. The solution rapidly transitions from -1 to 1 over a small region centred at $\theta = -\frac{\pi}{4}$.

We divide $\mathcal{S}$ into three subsets $\mathcal{S}_1, \mathcal{S}_2, \mathcal{S}_3$ with $\theta \in [-\frac{3\pi}{4}, -\frac{3\pi}{10}]$, $\theta \in [-\frac{3\pi}{10}, -\frac{2\pi}{10}]$, $\theta \in [-\frac{2\pi}{10}, \frac{\pi}{4}]$, respectively (see Figure 1 on the accompanying poster). CPM is used to solve (3) with cubic interpolation, second-order centred finite-differences, and second-order BCs (see [King et al. 2024b, Section 3.3]) on the adaptive computational tube $\Omega(\mathcal{S}_1) \cup \Omega(\mathcal{S}_2) \cup \Omega(\mathcal{S}_3)$. The tube radii of $\Omega(\mathcal{S}_1)$ and $\Omega(\mathcal{S}_3)$ are held constant with $r_{\Omega(\mathcal{S}_1)} = r_{\Omega(\mathcal{S}_3)}$ with $h_1 = h_3 = 0.025$. To investigate the effect of the grading between subsets we vary $r_{\Omega(\mathcal{S}_2)}$ using $h_2 = 0.025 \times 2^{-i}$ for $i = 1, 2, \ldots, 8$ and some arbitrarily chosen values $h_2 = 0.008$, 0.0045, 0.00229, 0.00113, 0.0005, 0.000283, 0.0001426. Note that these arbitrarily chosen h_2 give a $\Omega(\mathcal{S}_2)$ that is not aligned with the other two subsets.

Eigen's SparseLU is used to solve the linear system for results in Figure 2, on the poster, which presents a convergence study (left) for the max-norm error using both an adaptive computational tube and a uniform computational tube with $h = h_2$. The errors using the adaptive and uniform computational tubes are nearly identical and show the expected second-order convergence. Note that the error on $\mathcal{S}$ is dominated by the error on $\mathcal{S}_2$, which is why the max-norm error of the adaptive and uniform solutions match. The max-norm error is computed from 100 equally spaced points in the θ parameter.

Figure 2 (right) on the poster illustrates the improvement in efficiency of the adaptive approach by comparing the corresponding values for the uniform approach with those of the adaptive approach, specifically the LU decomposition time, forward/backward solve time, and number of DOFs. There is a runtime improvement of up to 10× and 13× for the LU decomposition and solve time, respectively. The timings are the average values of 20 trials.

2.2 Local Feature Size

Manifolds with small local feature size (LFS) can drastically increase CPM's computational cost if a uniform computational tube is used. LFS is defined as the minimum Euclidean distance from $\mathbf{y} \in \mathcal{S}$ to the medial axis of $\mathcal{S}$. With a uniform computational tube, the tube radius must satisfy $r_{\Omega(\mathcal{S})} < \text{reach}(\mathcal{S}) = \min_{\mathbf{y}\in\mathcal{S}} \text{LFS}(\mathbf{y})$. This leads to a large number of unnecessary DOFs for subsets $\mathcal{S}_m$ of $\mathcal{S}$ where $\text{LFS}(\mathbf{y}) > \text{reach}(\mathcal{S})$. Therefore, an adaptive computational tube can improve the computational cost of CPM by allowing larger $r_{\Omega(\mathcal{S}_m)}$ for subsets $\mathcal{S}_m$ with LFS > reach($\mathcal{S}$).

Consider a spiral sheet surface constructed by extending the hyperbolic spiral, $\mathbf{r}(\theta) = \frac{1}{\theta}[\cos\theta, \sin\theta]^T$, in the z-direction, creating a surface akin to a rolled-up sheet of paper with width W (see Figure 3, bottom, on the poster). The example below takes $\theta \in [0.3, 12.5]$ and $z \in [0, 4]$. We solve the screened-Poisson equation

$$-\Delta_{\mathcal{S}} u_{\mathcal{S}} + u_{\mathcal{S}} = \left(1 + \frac{4\pi^2(L^2 + W^2)}{L^2 W^2}\right) \cos\left(\frac{2\pi s}{L}\right) \cos\left(\frac{2\pi(z - z_0)}{W}\right), \tag{4}$$

on this surface with zero-Neumann BCs, which has the exact solution

$$u_{\mathcal{S}}(s, z) = \cos\left(\frac{2\pi s}{L}\right) \cos\left(\frac{2\pi(z - z_0)}{W}\right). \tag{5}$$

Note that the arc-length s of the hyperbolic spiral is used above, which can be computed from θ via

$$s = \int_{\theta_0}^{\theta} \frac{\sqrt{1 + \theta^2}}{\theta^2} d\theta = \left[\ln\left(\theta + \sqrt{1 + \theta^2}\right) - \frac{\sqrt{1 + \theta^2}}{\theta}\right]_{\theta_0}^{\theta}. \tag{6}$$

Figure 3 on the poster shows the adaptive computational tube $\Omega(\mathcal{S}_1) \cup \Omega(\mathcal{S}_2) \cup \cdots \cup \Omega(\mathcal{S}_5)$ and solution when solving (4) on the spiral sheet. The subsets $\mathcal{S}_1, \ldots, \mathcal{S}_5$ are taken to be constant in the z parameter and divided at $\theta = 1.7, 3.1, 4.7, 11$. The resolutions used on $\Omega(\mathcal{S}_1), \ldots, \Omega(\mathcal{S}_5)$ are $h_m = 0.064, 0.032, 0.016, 0.012, 0.008$, respectively. The same interpolation and finite-difference methods that were used for the arc example above are used here, but only first-order BCs are imposed. The BiCGSTAB solver of King et al. [2024b] is used for this problem.

A fast, but incorrect, solution can be computed with $h = 0.064$. The max and average errors are 25× and 178× larger, respectively, on the uniform tube with $h = 0.064$ compared to the adaptive tube (see Table 1 on the poster). To obtain an accurate solution on a uniform $\Omega(\mathcal{S})$ we need $r_{\Omega(\mathcal{S})} < \text{reach}(\mathcal{S})$, which requires $h = 0.008$. This results in a $\Omega(\mathcal{S})$ with ~2.9 million DOFs and requires about 22 minutes to compute the solution (see Table 1 on the poster). The adaptive computational tube can produce a solution that is visually indistinguishable from the uniform solution with $h = 0.008$ but with 6.5× fewer DOFs and in only 76 seconds (around 17× faster). The max error is about 7× larger for our adaptive solution but only 2.5× larger on average.

Acknowledgments

Nathan King was supported in part by the QEII Graduate Scholarship in Science & Technology. Steven Ruuth was supported in

part by an NSERC Discovery grant (RGPIN-2022-03302). Christopher Batty was supported in part by an NSERC Discovery grant (RGPIN-2021-02524) and the CFI-JELF program (Grant 40132).

References

Nathan King. 2025. *Closest Point Geometry Processing: Extensions and Applications of the Closest Point Method for Geometric Problems in Computer Graphics.* PhD thesis. University of Waterloo. https://hdl.handle.net/10012/21853

Nathan King, Steven Ruuth, and Christopher Batty. 2024a. A Simple Heat Method for Computing Geodesic Paths on General Manifold Representations. In *SIGGRAPH Asia 2024 Posters (SA '24).* ACM, New York, NY, USA, Article 69, 2 pages. https://doi.org/10.1145/3681756.3697920

Nathan King, Haozhe Su, Mridul Aanjaneya, Steven Ruuth, and Christopher Batty. 2024b. A Closest Point Method for PDEs on Manifolds with Interior Boundary Conditions for Geometry Processing. *ACM Transactions on Graphics* (2024). https://doi.org/10.1145/3673652

P.-L. Manteaux, C. Wojtan, R. Narain, S. Redon, F. Faure, and M.-P. Cani. 2017. Adaptive Physically Based Models in Computer Graphics. *Computer Graphics Forum* 36, 6 (2017), 312–337. https://doi.org/10.1111/cgf.12941

D. Morgenroth, S. Reinhardt, D. Weiskopf, and B. Eberhardt. 2020. Efficient 2D Simulation on Moving 3D Surfaces. *Computer Graphics Forum* 39, 8 (2020), 27–38. https://doi.org/10.1111/cgf.14098