Research Statement

Bardia Sadri

sadri@cs.duke.edu
http://www.cs.duke.edu/~sadri/

My research mostly focuses on computational geometry problems with topological twists. These are algorithmic problems concerning topological properties of geometric objects. Although mathematically rigorous, many of these problems stem from fundamental questions in computer graphics, geometric molecular biology, geographical information systems, and other applied domains, and therefore have remarkable practical applications. In what follows, I present a brief overview of my results[†]. Following that, I mention a few research ideas on which I am currently working or expect to work in near future. A listing of cited works can be found under "Publications" in my CV.

PhD Thesis Research. Surfaces, or more precisely, codimension one smooth submanifolds of Euclidean spaces, are the main subjects of interest in my PhD research. Such surfaces famously emerge in inexhaustibly many theoretical and applied problems in many of which the topology of the surface in question plays a central role. Exact representation of most non-trivial surfaces is rather costly and often beyond our reach. In fact, in many applications a surface is known only through a discrete sample of it. Consequently, topological properties of the surface have to be extracted from this sample. A natural approach is to try to interpolate a simple surface from the given point cloud and study its topology. This is known as the *surface reconstruction* problem which, put more accurately, calls for approximating (reconstructing) a surface from a given sample point set while guaranteeing the resulting surface to be topologically equivalent and geometrically close to the original one, provided that the given sample satisfies certain density requirements.

There is a quite rich literature on surface reconstruction but the majority of the proposed solutions are essentially heuristics that work well in practice but come with no theoretical guarantees. Among the relatively few that do offer guarantees, the idea common to most is to locally approximate a neighborhood of every sample point by one or more surface patches and at the end extract a manifold from the collected set of patches. Despite being clever, these solutions give little insight into the nature of the problem and topological correctness is generally achieved through building a homeomorphism between the original and reconstructed surfaces explicitly. The goal in my thesis is to address this somewhat old problem (and some of its relatives) systematically using some more solid mathematical tools.

The idea is to closely examine the *distance function* induced by the surface and the one induced by the given sample. Consider a surface S and let $h(x)$ be the distance from a point x to S. It is well known that h encodes a great deal of information about the surface S. This is of course no surprise considering that S is trivially the zero-set of h. But more interestingly, the medial axis M of S[‡] consists exactly of those points in space at which ∇h is undefined. It has been recently shown that the medial axis M of a surface S has close topological ties to S, namely, the volume enclosed by S has the same homotopy type as the piece of M contained in it. In this manner, M, or just as well h, characterize the geometry of S and its embedding in space. For example a regular torus and a knotted one can be told apart by comparing the homotopy types of their medial axes.

When a good sample of the surface is at hand, the distance to the surface can be approximated by the distance to the sample. It is a natural question whether the rich topological information encoded in h can also be extracted from the distance function induced by the sample. Critical point theory of distance functions is a well-developed field and I employ several of known properties of such functions in developing my topological tools. Among the most important of these properties is the existence of a unique direction of steepest ascent (even at places where the gradient is undefined) that only vanishes in critical points. Although it is not smooth, this vector field can be integrated to result a continuous flow map that can be characterized in terms of the Voronoi and Delaunay complexes of the sample. A close study of stable and unstable manifolds of the critical points under this flow map, and their corresponding stable and unstable

[†] I should hereby acknowledge my coauthors as listed in my curriculum vitæ.

[‡] The medial axis of a surface is the locus of the points in the space with two or more closest points on the surface.

flow complexes, is at the heart of my work. My first result in this direction [3, 7] is that the critical points of distance functions induced by dense samples of surfaces are sharply separated into two groups: those close to the surface and those close to its medial axis. Under Amenta-Bern style ε-sampling assumptions, I showed that critical points can be classified algorithmically. A second result shows that the boundary of the unions of stable manifolds of the critical points near the medial axis does indeed approximate the original surface faithfully in both geometric and topological senses in 3D [3]. Recently, I proved a stronger variant of this result for arbitrary submanifolds of Euclidean spaces of arbitrary dimensions which in addition to the topology of arbitrary dimensional submanifolds asserts the topology of their complements (hence for example guaranteeing that a simple loop is not reconstructed as a knotted cycle) [14]. I then studied the unstable manifolds of medial axis critical points and proved that their union approximates a subset of the medial axis and captures its topology. This topological "core" can be easily extended freely to provide better geometric approximation of the medial axis without sacrificing the achieved topological guarantee [2, 6].

The developed techniques can further be extended to allow the analysis of Edelsbrunner's famous WRAP algorithm. This algorithm, which was proposed in mid-nineties, is successfully implemented for commercial use[§] and has been one of the most successful practical reconstruction tools. However, although WRAP received much attention particularly due to the solid mathematical intuitions behind its introduction, it was never analyzed under any sampling assumptions. In particular, a major obstacle in doing so was that the main variant of this algorithm per se always generated a surface which was the boundary of a contractible volume. Using the separation of critical points, I came up with a slightly modified version of WRAP that can guarantee the topological type of its output. In [1] I analyzed this algorithm and showed that under ε-sampling conditions, the topology of its output can be guaranteed to agree with the sampled surface.

More Recent Results. Here I briefly describe my research from the past year at Duke. The main project I am involved with is the STREAM project[¶] that aims at designing I/O-efficient algorithms for various GIS (Geographical Information System) problems on massive terrains, where the size of the input data (at the order of hundreds of gigabytes) forces the efficiency to be measured under the I/O model (where the number of disk I/Os is the dominant factor) and this renders most classical solutions to these problems far too expensive. In [12], I worked on the problem of generating contour maps of such terrains. A proper output must generate segments in any contour cycle sorted in cyclic order and this proves to be rather challenging to be done optimally under the I/O model. They key is to approach the problem from a topological point of view by first simplifying the terrain topologically, eliminating all but two critical points through surgeries and then solve the problem for this topologically simplified terrain. The solution explores some elegant properties of the modified terrain which become crucially useful in achieving optimal performance.

A different problem I have worked on is that of *untangling* a triangulation of a two dimensional domain that has become invalid because of the movement of vertices [13]. The untangling attempts to change the topology of the triangulation, locally, only near the "affected" areas. The very characterization of these areas is a nontrivial problem and their structural properties span a considerable fraction of the work. Using these properties, we demonstrate output-sensitive algorithms that untangle the triangulation in time and space proportional to a certain measure of the complexity of the overlay in the tangled triangulation, as opposed to the size of the mesh in question. Significantly, the method allows the local fixing of the triangulation even if only a seed triangle in the affected areas is given for input. This allows the algorithm to be applicable in cases where due to numerical inaccuracy, tangling events cannot be computed exactly and the algorithm is only aware of parts of the invalidities in the tangled triangulation.

Future Directions. There are a number of interesting lines of research I intend to pursue. Below I mention some of the more promising ones. For one, my thesis research on distance functions and their induced flows has lead to various open problems. I mention two here.

- My recent work in [14] puts even stronger emphasis on the importance of the flow complex in manifold reconstruction. This complex, although can potentially coincide with Delaunay, is much more globally defined. Due to the fact that the complexity of this complex can be quadratically larger than that of the Delaunay and because of the numerical instability in computing it, it is a highly practical problem

[§]Raindrop Geomagic WRAP (`http://www.geomagic.com`).
[¶]`http://terrain.cs.duke.edu/`

whether this complex can be approximated by a simpler one. Effectively, the WRAP algorithm suggests that this should at least be doable by approximating each cell of the flow complex by a subcomplex of the Delaunay. However, a rigorous solution to the question demands a more general approach.

- From a more practical standpoint, flow complex methods have recently drawn considerable attention as powerful tools in geometric biology. For example, J. Giesen* has interesting results in finding protein pockets using fairly simple properties of flow complexes. This is not a surprise since flow complex offers a coarse segmentation of geometric objects which has been used earlier in similar contexts although all such applications were all heuristics. I am very keen on evaluating the theoretical power of this methods in similar geometric problems in molecular biology. This of course is to some extent in dependent on collaboration with biologists who can supply real-world problems and can determine the effectiveness of the approach from a practical point of view.

Aside from pursuing my PhD research, the past year in Duke has given me several fresh research directions. I restrict myself to research ideas directly following my recent work mentioned above. My work on terrains [12] has closely familiarized me with the literature and importance of I/O-efficient computation and its existing challenges. Working on the STREAM project has also taken me to the implementation level where I am now directly involved with development of a major project jointly pursued both in Duke and Aarhus University in Denmark. I intend to stay close to the trunk of this project and continue the collaboration by hopefully involving interested students in my future place of affiliation. These problems particularly interest me because many of them must be approached with an eye on topology. For example, currently I am working on a number of potential schemes for simplification (denoising) of terrains, as suggested by topologically persistent pairing of critical points. The commonly used method in practice is to "flood" pits of *low persistence*. However, this results "flat" regions over the flooded areas and this affects further processing of the simplified terrains. Our idea is to formulate the problem as an optimization problem in which one intends to minimize some measure of movement for the terrain while satisfying some topological constraint. For example, one may try to minimize the sum of the squares of the height differences throughout a fragment of the terrain. Alternatively, one may wish to keep the volume under the terrain unchanged and minimizes the "earth movers distance" needed to cancel critical points.

On the other hand, many of the questions asked in the GIS domain about terrains, which are simply topological disks, can be generalized to arbitrary surfaces. For example, it is an immediate question whether the approach used in [12] for I/O-efficient generation of contour lines applies to surfaces of arbitrary genus.

Many aspects of the triangulation untangling problem [13] are also wide open. An immediate question is to what extent and at what cost can these ideas be extended to higher dimensions. Especially in 3 dimensions, mesh untangling, has numerous direct practical applications. Theoretically also there is much to explore. The *crease edges* in a tangled mesh (edges at which the "orienation" of neighboring triangles in the tangled mesh disagree) are critical sets (Jacobi sets) of the motion map that maps the untangled triangulation to the tangled one, and therefore generalize the notion of critical points in the case of a real-valued function. Untangling calls for a higher order variant of "cancellation" of critical points for these sets. It is a very deep topological problem to formulate these cancellations as a generalization of pairing of critical points in Morse theory or to explore their "persistence" as is measured through persistent homology for the case of critical points.

At the end, I should emphasize that like most other researchers working on "algorithms", I constantly search for new open problems. In addition to being interested in problems in computational geometry and topology and the related areas in general, I am specially also keen on finding grounds for interdisciplinary research. I value such research philosophically and weigh the accessibility of researchers in other fields, from mathematics and physics to geography or biologists, who are interested in collaboration rather positively in choosing my future research environment.

*my coauthor in [2, 3]