Research Statement

Bardia Sadri

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

My research is mostly focused on computational geometry problems with topological sides. 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, or even verification of properties of concurrent programs and other applied domains, and therefore have remarkable practical applications. In what follows, I present a brief overview of my results[†] and 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. These are fundamental mathematical objects that emerge in inexhaustibly many theoretical and applied problems. In many of these problems the topology of the surface in question plays a central role. Exact representation of non-trivial surfaces is often costly. In fact, in many applications a surface is known only through a discrete sample of it. Consequently, topological properties of the surface, such as Homology groups, have to be extracted from this sample. A natural approach is to try to interpolate a simple surface through this sample and study its topology. This is known as the *surface reconstruction* problem. A satisfactory solution to this problem is an algorithm that approximates (reconstructs) a surface from the given sample and guarantees that the output is topologically equivalent and geometrically close to the original surface, 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 smooth 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) characterizes the geometry of S and its embedding in space. For example a regular torus and a knotted one, although topologically equivalent (homeomorphic), are distinguished through 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 if the sample is dense enough. 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 flow complexes, is at the heart of my work. My first result in this direction [7, 11] 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

[†] 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.

[7]. 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 one) [2]. 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 and can be used generically to make virtually any medial axis approximation method topologically accurate [6, 10].

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 popular 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 this algorithm as originally described always generated a surface which was the boundary of a contractible volume and thus provided no topological fidelity. Using the separation of critical points, I came up with a slightly modified version of WRAP that fully addresses this major shortcoming. This very natural modification comes in the form of a sensible filtering of the critical points on which the original WRAP algorithm acts indiscriminately. In [5] 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 and future directions. Here I briefly describe the research that I have conducted at Duke University and later at the University of Toronto since the completion of my PhD. The main project at Duke in which I was involved was the STREAM project[¶] that aimed 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 therefore most classical solutions to these problems are far too expensive. In [3], I worked on the problem of generating contour maps of such terrains. A proper output must report segments in any contour in cyclic order and this proves to be rather challenging to be done optimally under the I/O model. Interestingly, the solution exploits topological properties of the terrain to generate a total ordering of the triangles (faces) of the terrain in such a way that the triangles intersecting each connected component of every level set appear around it in sorted order and triangles from different components do not interleave. Such an ordering is called a *level-ordering*. The fact that every terrain admits a level-ordering is rather nontrivial which together with the supplied I/O-efficient algorithm for the generation of this ordering are the main contributions of this paper. Currently, I am working on extending this result to triangulated 2-manifolds of arbitrary genus endowed with a "tame" function (corresponding to the height function in terrains) defined on their vertex set. Moreover, I am studying the problem of defining and computing a *joint level-ordering* of triangles in a sequence of terrains that represent a *multi-resolution* Dobkin-Kirkpatrick (or similarly Heckbert-Garland) type simplification of the terrain, thus allowing to efficiently patch together contour maps of various resolutions from various parts of a terrain. In my opinion, the importance of answers to these questions compares with that of their applications in construction of I/O-efficient algorithms; in essence, what we like to find out is whether it is possible to embed a discretized 2-dimensional manifold (or a hierarchy of them) into a one-dimensional space (such as a sequential medium) in such a way that level sets are embedded almost contiguously.

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 [4]. 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. There is much more to do on untangling. For example, the *crease edges* in a tangled mesh (edges at which the *orientation* 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. In this sense, untangling corresponds to a higher order variant of topological simplification through cancellation of critical points. It is a very deep mathematical problem to formulate

[§]Raindrop Geomagic WRAP (`http://www.geomagic.com`).

[¶]`http://terrain.cs.duke.edu/`

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.

My work on terrains [3] has increased my interest in 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 MADALGO institute at the Aarhus University in Denmark. I intend to stay close to the trunk of this development effort, which has seen a lot of interest from the industry, 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. My recent work in [1] is an example of this with regards to simplification (denoising) of terrains. The commonly used method in practice is to "flood" pits of *low persistence*. However, this results "flat" regions over the flooded areas and may amount to massive geometric alterations to the terrain. One idea is to formulate the problem as an optimization problem in which one intends to minimize some measure of movement for the terrain while meeting some *topological constraints*. For example, one may try to minimize the modifications to the terrains as measured by a norm such as least squares distance. This family of problems generalize (one-dimensional) shape-constrained regression problems commonly studied in statistics and are interestingly related to network flows and circulation problems extensively studied in combinatorial optimization, but their topological aspects are widely unexplored. In [1], we developed algorithms for modifying a function defined on a tree so that the resulting function has a unique local minimum while minimizing the ℓ_2 distance to the original function. This can be used as a tool for topological simplification of terrains through using the least amount of *carving*.

To recap, I am very interested in a host of computational problems involving topology of combinatorial objects. To a great extent, perhaps due to the remoteness of topology from the mainstream of theoretical computing, computational topology has attracted less attention than it deserves and to date remains an active field with much left to explore. It is my intention to stay involved in the development of this branch of theoretical computer science.