Research Statement

Eric Shaffer, 6/12/2019

Throughout my career, my research focus has been on building computational tools. For me, the most enthralling aspect of computer science has always been the ability to use computation to enable accomplishments and discoveries that would not otherwise have been possible. I feel fortunate to live in a time in which the pervasive nature of modern computing provides opportunities for computer scientists to assist in almost any endeavor.

I completed my PhD dissertation at the University of Illinois working with Michael Garland in 2005. Geometric modeling, the focus of my dissertation, is a critical element of computer graphics applications. Generating a synthetic image of a scene effectively requires simulating the transport of light in some manner. This simulation relies on a geometric model of the surfaces in the scene, typically represented as a mesh of polygons.

A key problem in geometric modeling at the time was that many polygonal models were too large to fit in memory on a typical computer. At the same time there was a paucity of tools available to work with models out-of-core. Our work addressed this issue, as I developed a data structure for managing out-of-core polygonal meshes, as well as out-of-core algorithms for simplifying and smoothing meshes. Looking back on that work, I think the most interesting aspect is how elements of it have ended up being used in other applications. The out-of-core data structure anticipated similar structures now used to build bounding volume hierarchies on modern graphics processing units (GPUs). The mesh simplification technique has been applied to point-set, rather than mesh, surfaces which have become increasingly prevalent with the advent of LiDAR technology for geometry acquisition.

After completing my dissertation, I worked as a post-doc with Michael Heath at the DOE-funded Center for the Simulation of Advanced Rockets (CSAR). My work there still revolved around geometric modeling, but this time within the realm of scientific computation rather than rendering. Many simulations employ geometric models of the physical objects, such as rocket engines, involved in the simulation. The simulated physics will model things like combustion and heat transport in a discretized domain. In the case of a three-dimensional domain, this discretization will take the form of polyhedral elements. Much of my work at CSAR focussed on the problem of optimizing the shape of these elements to limit the numerical error in the simulation, while still accurately representing the shape of the domain. This work contributed to the first fully three-dimensional simulation of solid-propellent rocket combustion, with the software living on through industry partnerships that were developed during the funding of CSAR.

During my time at CSAR, I eventually worked as a research scientist and then technical program manager. I was instrumental in developing two of our industry partnerships, specifically with Caterpillar and Boeing. In both cases, I co-wrote proposals detailing how software and expertise at the University of Illinois could be of benefit. Both Caterpillar and Boeing sponsored research based on those proposals and renewed multiple times. This work further developed our tools for geometry optimization, mesh repair, and feature detection, and integrated it into the production toolchains of the sponsors.

At this same time, I began exploring tools for geological modeling and simulation. I developed a proposal, along with a scientist at the Illinois State Geological Survey (ISGS), to create a freely available tool for oil reservoir visualization. The proposal was funded by the DOE, resulting in the creation of extensions for the Paraview visualization system that have been used both inside the ISGS and at other institutions, principally Wright State University.

During the course of my work on the DOE grant, as well as the work with Boeing, I was able to hire a number of undergraduate research assistants at the University of Illinois. Working with these students was one of the most rewarding aspects of those efforts. It became clear to me that the largest impact of the research really would be in the skills and knowledge gained by the students. Each one of them would move on to other institutions and carry with them and share what they had learned at the University of Illinois. I was fortunate to be given the opportunity to participate in this process on a larger scale by becoming instructional faculty in Computer Science at the University of Illinois.

Since joining the Computer Science Department, my sponsored research has been funded by Exxon-Mobil, developing tools for oil reservoir simulations. The genesis of this work was a visit to campus by representatives of Exxon-Mobil upstream research. Various faculty gave presentations about their research areas and potential collaborations. My proposal was one that was chosen for funding, resulting in a 5 year working relationship and multiple funding renewals. Broadly, the work revolved around developing flow-based similarity measures for reservoir models to ensure that the simulation sample space used in an analysis is not biased. Additional tools for model simplification and the use of machine-learning to predict simulation outcomes where developed as well. This code was integrated into the simulation toolchain at Exxon-Mobil, in an effort to scale up the size and number of simulations run when analyzing potential development sites. Again, this work funded a number of undergraduate hourly positions for University of Illinois students, and provided them with the opportunity to develop skills related to research that are not typically learned inside a classroom.

In my current position as instructional faculty, my research focus has shifted to tool-building for education. Along with my colleague Raluca Ilie, I am working on using virtual reality (VR) to visualize electromagnetic phenomena. These visualizations will be used in virtual reality labs for students taking a course on electromagnetism. Our belief is that being able to see and interact with elements like electromagnetic fields in VR will engender a deeper understanding than gained by non-interactive learning methods. We hope and expect it will prove superior to even existing active, lab-based activities, in that VR will allow the students to see what is unseen in the physical world. We will formally assess the effectiveness of this learning modality in comparison to traditional methods. If the work proves successful, there is a great opportunity to expand beyond just one course and beyond our campus. We have already heard from potential collaborators interested in employing such tools for medical engineering education and high school education.

In addition to the work with VR, I am also interested more generally in active learning and inquiry-based learning methodologies for teaching. I am currently working on developing web-based tools to support inquiry-based learning in the computer graphics courses I teach and plan to pilot their use soon. Again, if assessment shows that this work is effective, it will be publicly released and should generate significant interest.