I am applying to graduate school to explore how best to physically implement computation and how best to pass on what I learn to others. As physical devices get pushed to their limits, new design paradigms have begun to emerge across the spectrum of hardware architectures. This spectrum spans from prominent software-programmable devices, such as many-core CPUs and modern GPUs, to hardware-programmable technologies such as FPGAs. I am fascinated with the design challenges inherent in these and other types of devices and want to contribute to meeting these challenges.

Implementing computation involves design decisions from the application to the hardware architecture level. As non-parallel CPU performance is no longer scaling with time, these devices are evolving to support higher core counts and more elaborate system interconnects. I want to learn how these next-generation compute engines can be designed to support next-generation applications. Furthermore, I would like to study how novel architectures can support different applications sharing similar computational needs. I believe that with parallelism becoming ubiquitous, applications will again drive design at the hardware architecture level. Given different applications, I believe that the optimal hardware solution will find common ground between general-purpose compute engines and fully custom accelerators. To match physical computation to applications, one must consider the entire hardware architecture design space.

During my undergraduate career, research experience has given me the background to explore these issues in more depth as a graduate student. This year, I have been working with Professors John Wawrzynek (U.C. Berkeley) and Garry Nolan (Stanford) to create an FPGA-based accelerator for learning cell signaling network structure. The purpose of this approach is to provide a map of how cancer cells process information and resist drug treatment. Mapping this problem to hardware is desirable because the learning algorithm's complexity is super-exponential with the number of proteins in a cell. One of my contributions to the project was to optimize a computationally expensive part of the algorithm, built on an FPGA. Given that we were designing to a specific algorithm, our design achieved speedup through micro-architectural optimizations, such as parallelism and multi-threading, and characteristics of the FPGA fabric itself. By designing with both the application and device in mind, we were able to develop a highly optimized solution to an important problem in systems biology.

After I complete my graduate study, I want to conduct interdisciplinary research and share what I learn with others, as a professor. I feel that discovery and education are equally important at a time when computer architecture is becoming an integral component in other fields. As a case in point, the signaling networks project demonstrates how advances across computer architecture, statistics, and biology are coming together to solve an interdisciplinary problem. I recently gave a talk on the project at the Gigascale Systems Research Center. (Talk slides can be found at http://cwfletcher.net/GSRC2009.pdf.) Giving the talk, and seeing how the audience responded to it, inspired me. Working on this project has shown me that research is about answering questions that we care about as a civilization.

In summary, my goals as a graduate student are to research novel hardware architectures to solve next-generation problems and to learn how best to communicate these challenges and proposed solutions to others. Looking back as a teacher, one challenge continues to be emphasizing the connection between the application and device. As computation evolves to match different applications' needs, only in respecting both domains

will we be able to continue the trend of high-performance over time. Looking forward, I am excited to have the opportunity to contribute to addressing this challenge.