

Research Statement

Dee A. B. Weikle

The art of scientific inquiry and invention has always fascinated me. When I entered college, computers were the latest invention and I wanted to be able to both build and program them. As a result, I majored in Electrical Engineering and fulfilled the requirements for a degree in Computer Science. In graduate school I received a PhD in Computer Science, specializing in computer architecture. Math was a central tool to both of these degrees and I have a long term interest in applying mathematics to problems.

My dissertation research focused on analytical models of caching systems. The idea behind this work was that current black box simulation methodologies, while frequently useful, often fall short in promoting the understanding of why certain systems succeed or fail. The model I developed in my dissertation attempted to distill the essential operation of a cache, giving designers insight into the interaction of memory access patterns and cache design. The framework consisted of three components. The first, TSpec, was a language (or domain) for representing memory accesses issued by a processor. The second was a filtering function that operated on memory access traces in TSpec, filtering those accesses as they would be by a caching system. The final component of the framework was two new measures, instantaneous locality and instantaneous hit-rate. These measures gave different insight into the operation of a cache than traditional hit-rate and helped illustrate why certain combinations of caches worked better (or not) on particular memory reference patterns. This work was described in “Caches As Filters: A New Approach to Cache Analysis” in MASCOTS’98, which won a best paper award.

In my position as a Research Scientist at the University of Virginia, I expanded my work to include characterization of workloads and the experimental methods used by the computer architecture community, in addition to mathematical modeling. Historically, computer architecture has made great strides using a generally ad-hoc approach to research. The most common methodology uses standard benchmark suites and a simulator (often written or modified by the researcher) to run the benchmark suite. A comparison is made between execution times for a current design and a slightly modified future design. If the modified design executes the benchmark suite faster on average in the simulator, it is considered the better design. This approach fails to recognize the impact of the selection of the workload represented by the benchmark suite, the sampling of the benchmark suite, the error introduced by the simulator, and the application of appropriate statistics. As the field of computer architecture has advanced, recognition of how tenuous the methodology has become is growing. Faster computers, the reality of multiple processor cores, and the wide range of computer applications are compromising this methodology even more. In our workshop paper titled, *Exploring the Impact of Normality and Significance Tests in Architecture Experiments*, we show the impact of statistical methods and benchmark suites in particular, and show how to use the appropriate mean and determine statistical significance for an experiment. I have also continued my modeling research by applying it to power performance tradeoffs in branch predictors for embedded processors. By modeling these tradeoffs it is possible to perform initial design space exploration without a detailed power simulation for the new

processor. These results are described in *A Break-Even Formulation for Evaluating Branch Predictor Energy Efficiency*, published in the Workshop on Complexity-Effective Design held with the 2005 International Symposium in Computer Architecture.

The thrust of this combined research has led me to characterizing workloads as a way to see what applications need from a new computer architecture. Characterization has been used by several researchers to reduce simulation time, predict performance, and create synthetic benchmarks for single core processors. My future work is to expand this methodology, set of metrics, and models to incorporate multiple processors (parallel characterization). The goal is to use these characterizations to determine the granularity of parallelism that can be detected in current applications and drive the level and type of parallelism provided by new multi-core chips. Understanding workloads and the inherent parallelism within them requires understanding not just current applications, but the real-world problems they are trying to solve and how those problems are and can be expressed by the people trying to solve them.

Making things work well and efficiently has always appealed to me, and I see solid scientific method applied to real world problems as a long term solution to making things work. This philosophy applied to teaching has led me into research in education. I want to teach effectively and retain a diverse student population. My collaboration with the Curry School has included efforts on retaining minorities at the college level as well as our current project developing elementary and middle school differentiated, problem-based curriculum. Research shows that retention of minorities must begin in elementary and middle school. It also shows that minorities (especially women) are motivated by seeing the societal benefit of solving problems (problem based learning) and that gearing lessons to the individual students' learning style, ability, and interest (differentiation) increases student success. Problem based learning and differentiation can be the solution at the college level as well. I see service-learning as an excellent example of problem based learning and I am motivated to introduce this and differentiation into my college classes with the appropriate methodology to publish results on my experience. My contacts with the Curry School will enable me to do this effectively.

Guiding students to perform successful research projects requires a different skill set than performing research on ones own. While a research scientist, I supervised two undergraduate theses in education, assisted in the supervision of several others, and co-advised two PhD students, one through graduation. Consequently, I know what is necessary to enable undergraduates to perform successful research projects and what is expected of graduate students at research universities. They need to learn the fundamentals of the field, how to read and analyze technical papers, and how to explain their thoughts and ideas through both presentations and the written word. All of my recent papers have been published with student authors, including one undergraduate author. I also funded students' attendance at conferences, even when not presenting papers so that they could develop a vision of what they need to accomplish. I look forward to bringing this experience to a research group of my own.