

**Position Paper**  
**NSF Workshop on Computer Performance Evaluation Techniques**  
**Gurindar Sohi**  
**Computer Sciences Department, University of Wisconsin**

## **1. Background**

Before discussing what performance tools are going to be needed in the future, it is important to discuss the role of these tools in computer research and design. There are several aspects: (i) the objectives of the users of performance evaluation tools, (ii) the varieties of performance evaluation tools and their abilities, (iii) putting together the objectives of the users and the capabilities of the tools.

I will consider two broad categories of tool users: researchers and developers of computer systems. These two categories of users have different objectives. In my view, computer systems researchers play two important roles. First, they develop a better understanding of aspects of computer systems behavior. For example, a better understanding of aspects of program behavior, or a better understanding of the aspects of performance of a commercial design. Second, they propose new concepts for systems design, possibly with an initial evaluation of the concept in a limited context.

Developers generally have a different objective from researchers: they are mainly concerned with a thorough understanding of a proposed computer system design. The purpose is to not only get an understanding of the likely performance of the proposed design, but also to make sure that there are no performance "unknowns". They are less concerned with proposing new concepts and their potential.

We have with us a suite of techniques performance evaluation techniques to pick from, including analytical models, simulators, and measurement. The tools required in the two different roles are, of course, going to be different.

Building detailed performance evaluation tools take a lot of effort. Effort spent in building tools takes away from effort that can be spent on other research using the tools. Even more difficult are tools that are "general-purpose", i.e., that are not customized for a particular design. These tools have to be designed to allow easy customization, and yet be sufficiently accurate to achieve their purpose. Such tools get abused despite the best intentions of both the tool developers and the users.

There is also no convincing way to determine whether a performance tool accurately models a desired system. One can validate a model with an existing system by comparing summary and intermediate results from the two. However, such a "validation" is only slightly more convincing than a more constrained "validation" for that system, and there is no reason to believe that a "validated" performance model will accurately model another system whose hardware parameters are different.

## **2. Position**

Where we should go with performance evaluation tools depends upon who is going to be using them, and for what purpose. Is the purpose to better understand an existing system? To evaluate a new architectural concept? To help come up with a new architectural concept? Different tools are likely to be required, or certainly adequate, for different purposes.

Developers of computer systems will continue to be concerned with tools that allow them to predict the performance of a design, and to learn about potential performance bottlenecks. Simulation techniques similar to what are in use today, coupled with analytical models of network/wire delays will continue to be the tool of choice.

For researchers mainly interested in proposing new concepts, simulation techniques that allow an initial evaluation of the concept should be the technique of choice. The important point here is that proposing new ideas does not require detailed evaluation methods for an initial evaluation. In fact

emphasizing a detailed evaluation is likely to detract from "idea-getting".

The choice of workloads is another issue that needs discussion. A decade ago the emphasis on detailed quantitative evaluation, with a consequent sacrificing of the "idea-getting" aspect of research. Today the analogous emphasis is on full-system simulation and commercial workloads. For most researchers, I believe that this emphasis is as misplaced as the emphasis on building detailed simulators that can accurately model an existing computer system.

### **3. Summary**

Building accurate, widely-usable performance tools that serve the objectives of the spectrum of users of performance tools is not only a very difficult proposition, but perhaps even the wrong approach. For most researchers, a performance evaluation tool is a means to an end -- novel observations and concepts in computer systems design -- and not the end itself. Overemphasizing quantitative analysis and performance evaluation tools is likely to detract from the "idea-getting" aspect of research.