

Foreword

The issue of computer performance evaluation must have concerned computer science and engineering researchers as well as system designers, system operators, and end users ever since computer systems came into existence. But it was not until around 1970 that “performance evaluation”, as a clearly identifiable discipline within the computer science/engineering field, established itself. This was the period when IBM introduced its successful 360 series, and time-shared operating systems were the main focus of the computer community.

The IFIP Working Group 7.3 “Computer system modeling” was initiated by Paul Green (IBM Yorktown), Erol Gelenbe (INRIA, France), and myself (IBM Yorktown). The ACM Sigmetrics group was also established around that time. The 1970s was a really exciting period when the new computer performance community played a major role in advancing the state of the art in queuing network models: a major extension of networks was made with “product-form” solutions by Baskett-Chandy-Muntz-Palacios (1975) and a number of other works on related subjects. Len Kleinrock’s two volumes *Queuing Systems* (1975, 1976) made major contributions by providing both computer and network communities with the state-of-the-art knowledge in queuing theory. Reiser-Sauer-McNair’s efforts at IBM Yorktown in creating QNET and RESQ modeling packages and disseminating them to both industrial and academic communities was a major landmark, and prompted similar efforts undertaken at AT&T, INRIA, and elsewhere. More recently, software packages for the user-friendly specification and automated analysis of stochastic Petri nets like the DSPNexpress package by Christoph Lindemann (1995) made significant impact both in academia and industry.

The MVA algorithm devised by Martin Reiser and Steve Lavenberg (1980) made large scale modeling computationally feasible. Jeff Buzen, who published a seminal paper on the central server model (1973) was among the first successful entrepreneurs to demonstrate a viable business opportunity in performance modeling. The introduction of generalized stochastic Petri nets by Marco Ajmone Marsan, Gianfranco Balbo, and Gianni Conte (1984) allowed the high-level specification of discrete-event systems with exponential events (i.e., Markovian models) and their automated solution process by tool support. A recent landmark in performance modeling of communication networks and the Internet constituted the discovery of self-similar traffic by Walter Willinger and his co-workers at Bellcore (1995). Four personal accounts of these key contributors describing where and how these research results have been discovered are included in this book.

In 1981, an international journal “Performance Evaluation” (North Holland) was created with me as its founding editor-in-chief. Together with the regularly held IFIP WG 7.3’s Performance Symposium and ACM’s Sigmetrics Conferences, the journal established the identity of our performance community, and thanks to the dedicated efforts by succeeding editors (Martin Reiser, and now Werner Bux) the journal has maintained its role as a major forum for the archival

literature of our research community. Through this journal and the aforementioned symposia and through more recent additions such as Performance Tool Conferences regularly organized by Günter Haring and others, we have been quite successful in embracing our outstanding colleagues from related fields, such as operations research and applied probability.

As the paradigm of computing has changed from time-shared mainframe computers to networks of personal computers and workstations, the focus of performance modeling and evaluation has shifted from CPUs to networks and distributed databases. With the rapidly gaining dominance of the Internet, wireless access networks, and optical network backbones, we will witness profound changes and increasing complexities in associated performance issues. For example, as traffic over the Internet expands exponentially, and end user's desire to access the Web servers from anywhere and at any time grows rapidly, interactions among clients and web proxy servers should present major technical challenges to network service providers, and hence great opportunities for our performance community. Tom Leighton of MIT and his colleagues have seized such an opportunity by developing a network of proxy server's cache called "free flow" and have turned their research project into a rapidly expanding business. I expect similar opportunities will abound in the coming decade, as mobile users and a variety of "information appliances" will be interconnected, forming an ever expanding information network.

It is quite timely that this special volume on "Performance Evaluation" contributed by a number of outstanding colleagues is going to be published at this exciting period, when the complexity and sophistication of information networks continue to multiply and challenge our ability to evaluate, predict, and improve system performance. I congratulate Günter Haring, Christoph Lindemann, and Martin Reiser on their success in compiling these important articles into an archival form.

November 1999

Hisashi Kobayashi

Preface

Performance evaluation has been a discipline of computer science for some thirty years. To us, it seemed to be time to take stock of what we - the performance evaluation community - were doing. Towards this end, we decided to organize a workshop on Performance Evaluation of Computer Systems and Communication Networks, which was held at the international conference and research center for computer science, Schloß Dagstuhl, Germany, September 15-19, 1997. The participants discussed, among other things, the following fundamental questions:

What are the scientific contributions of performance evaluation?

What is its relevance in industry and business?

What is its standing in academia?

Where is the field headed?

What are its success stories and failures?

What are its current burning questions?

During this workshop, working groups focused on areas like performance evaluation techniques and tools, communication networks, computer architecture, computer resource management, as well as performance of software systems (see <http://www.ani.univie.ac.at/dagstuhl97/>). The participants summarized the past of performance evaluation and projected future trends and needs in the field. It was identified that - as in many other sciences - at the beginning there was the observation of the behavior of systems, generally by measurements, followed by the development of theories to explain the observed behavior. Especially in system modeling, based on these theories, methodologies have been developed for behavior prediction. At that stage, measurement changed its role from pure phenomenological observation to model-driven parameter estimation. Based on a series of highly successful case studies, tool methodology implemented in versatile software packages has been developed to make the theoretical results amenable to practitioners. Originally focused on computer systems and communications, the scope of performance evaluation broadened to include processor architecture, parallel and distributed architectures, operating systems, database systems, client-server computing, fault tolerant computing, and real-time systems.

With the growing size and complexity of current computer and communication systems, both hardware and software, the existing methodologies and software packages for performance modeling seem to reach their limits. It is interesting to observe that for these complex interoperable systems the methodological cycle starts again with measurements to describe the system behavior, and understand it at a qualitative level. High performance parallel processing systems and the Internet traffic are good examples for this pattern of scientific progress. An example is the discovery of self-similar traffic that spurred a large number of observational and empirical studies.

During three decades of research, performance evaluation has achieved a rich body of knowledge. However, the participants of the Dagstuhl workshop felt, that the results are not as crisply focused as those in other prominent fields of computer science. Therefore, the project of writing a monograph on *Performance Evaluation - Origins and Directions* was begun with the following basic ideas. This monograph should document the history, the key ideas and important success stories of performance evaluation and demonstrate the impact of performance evaluation on different areas through case studies. The book should generate interest in the field and point to future trends. It is aimed at senior undergraduate and graduate students (for seminar reading), faculty members (for course development), and engineers (as orientation about available methods and tools). The contributions should be on an academic level, having an essay-character and they should be highly readable. They should also be as lively and stimulating as possible, where the mathematical topics should be presented on the appropriate level, concentrating on concepts and ideas, avoiding difficult technicalities.

As editors we are proud that we have succeeded in recruiting top researchers in the field to write specific chapters. The spectrum of topical areas is quite broad comprising work from an application as well as from a methods point-of-view. Application topics include software performance, scheduling, parallel and distributed systems, mainframe and storage systems, World Wide Web, wireless networks, availability and reliability, database systems, and polling systems. Methodological chapters treat numerical analysis methods, product-form queueing networks, simulation, mean-value analysis and workload characterization. Altogether the volume contains 19 topical chapters. Additionally, it contains one contribution on the role of performance evaluation in industry, and personal accounts of four key contributors describing the genesis of breakthrough results.

We hope that the book will serve as a reference to the field of performance evaluation. For the first time, it presents a comprehensive view of methods, applications, and case studies expressed by authors who have shaped the field in a major way. It is our expectation, that the volume will be a catalyst for the further evolution of performance evaluation and help to establish this exciting discipline in the curriculum, in industry, and in the IT world at large.

October 1999

Günter Haring
Christoph Lindemann
Martin Reiser