

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 its 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