

Time—It's Time for a Change

(Invited Paper)

Boudewijn R. Haverkort

Embedded Systems Institute & University of Twente

P.O. Box 513, 5600 MB Eindhoven, Netherlands

<http://www.esi.nl/> <http://dacs.ewi.utwente.nl/>

Abstract—Since the 1970's, the scientific field of model-based performance and dependability evaluation has been flourishing. Starting with breakthroughs in the area of closed queueing networks in the 1970's, the 1980's brought new results on state-based methods, such as those for stochastic Petri nets and matrix-geometric methods, whereas the 1990's introduced process algebra-type models. Since the turn of the century, techniques for stochastic model checking are being introduced, to name just a few major developments. The applicability of all these techniques has been boosted enormously through Moore's law; these days, stochastic models with tens of millions of states can easily be dealt with on a standard desktop or laptop computer. A dozen or so dedicated conferences serve the scientific field, as well as a number of scientific journals.

However, for the field as a whole to make progress, it is important to step back, and to consider how all these as-such important developments have really changed the way computer and communication systems are being designed and operated. The answer to this question is most probable rather disappointing. I do observe a rather strong discrepancy between what is being published in top conferences and journals, and what is being used in real practice. Blaming industry for this would be too easy a way out. Currently, we do not see model-based performance and dependability evaluation as key step in the design process for new computer and communication systems.

Moreover, in the exceptional cases that we do see performance and dependability evaluation being part of a design practice, the employed techniques are not the ones referred to above, but instead, depending on the application area, techniques like discrete-event simulation on the basis of hand-crafted simulation programs (communication protocols), or techniques based on (non-stochastic) timed-automata or timeless behavioral models (embedded systems). In all these cases, however, the scalability of the employed methods, also for discrete-event simulation, forms a limiting factor. Still, industry is serving the world with ever better, faster and more impressive computing machinery and software! What went wrong? When and why did "our field" land on a side track?

In this presentation I will argue that it is probably time for a change, for a change toward a new way of looking at performance and dependability models and evaluation of computer and communication systems, a way that is, if you like, closer to the way physicists deal with very large scale systems, by applying different type of abstractions. In particular, I will argue that computer scientist should "stop counting things". Instead, a more fluid way of thinking about system behavior is deemed to be necessary to be able to evaluate the performance and dependability of the next generation of very large scale omnipresent systems. First successes of such new approaches have recently been reported. Will we witness a paradigm shift in the years to come?