Learning Performance Engineering through Applications

Giuseppe Serazzi
Politecnico di Milano

ABSTRACT

In 1986 D. Ferrari, in a position paper published on IEEE Transactions on Software Engineering, emphasized that the field of Computer Systems Performance Evaluation (CSPE) since its birth had grown in isolation from the other areas of computer science. It was pointed out that the drawbacks of such insularity outweigh its advantages. Some urgent actions to improve the situation were listed. Among them, on the educational side, the integration of performance evaluation concepts in the computer systems basic courses and a more application-oriented approach of the teaching have been advocated. Another problem highlighted was the existence of a time lag between the available performance evaluation techniques and the type of problems created by the new system architectures. In 2003 Ferrari revisited and updated the concepts described in the first paper. Some limited positive results have been perceived. But, what has happened during the last 14 years, about these and other problems related to the education and practice of performance engineering? The old problems, at least partially, have been solved? What are the new problems that arose?

Keywords

Education; Performamce Evaluation Courses; Applications

1. BEFORE 2003

In the papers published in 1986 [1] and 2003 [2], Domenico Ferrari described the situation of performance evaluation discipline in two periods of its existence, prior to 1986 and from 1986 to 2003. The characteristic that dominate the first 20 years of its life, was the insularity of this new discipline with respect to the basic ones of computer science such as computer architectures, operating systems, programming languages, and software engineering. This situation has created a gap between researchers expert in the area and other professional profiles as system architects, designers, installation and applications managers, capacity planning specialists. In some cases, researchers have lost any contact with the reality. The same problem involved also the teaching side. The performance evaluation courses are often seen as special courses that have a limited impact on the completeness of the curricula of computer science engineering. The author pointed out that the drawbacks of such a situation of insularity outweigh its advantages. Some urgent actions to improve the situation were listed. Among them, on the educational side, the integration of performance evaluation concepts in the computer systems basic courses and a more application oriented approach of the teaching have been advocated. In the 2003 paper, Ferrari reconsidered the goals that has been set in the previous paper as *desirable to reach* and assessed whether they have been, at least partly, achieved. The overall feeling was that the insularity has been somewhat reduced both on the educational and, most noticeably, on the research side. He substantiated the conclusions with quantitative data concerning some of the ACM SIGs computer systems engineering conference series: SIGMETRICS, SIGOPS, SIGCOMM, SIGARCH. The number of *applied* performance evaluation papers with respect to *non-applied* (e.g., theoretical, methodological) presented in the annual meetings show a continuous increase, with different rates as a function of the subject of the conference. This aspect has been considered a positive signal that the distance between the performance researchers and the practical applications was slowly decreasing. Also on the educational side a positive trend seemed to appear, but not very noticeable and could not be quantified because of the difficulty of collecting data having a good statistical validity. The most important objectives desirable to reach, described in the first paper, and recognized, at least partially, still valid in the second paper are:

a. Reduction of the distance between the performance evaluation discipline and the other disciplines involved in the computer science

b. Reduction of the time lag between the solutions offered by performance evaluation techniques and the problems raised by new developments in computer systems technologies and applications.

c. Integration of the performance evaluation concepts into the courses of computer systems that apply the related techniques to motivate the solutions adopted

2. THE LAST 14 YEARS

In the most recent years the trends of the technologies and of the applications made the current situation very different

ICPE '17 Companion, April 22 - 26, 2017, L'Aquila, Italy

ISBN 978-1-4503-4899-7/17/04...$15.00

DOI: http://dx.doi.org/10.1145/3053600.3053639

from that which was present at the end of 90's. Some of them had positive effects on the education and practice of performance engineering, while others have had a negative influence. Among the former, we may consider the evolution of the concept of performance and the proliferation of the conferences/workshops on computer applications. Over the years the importance of performance has increased and now is no longer confined to the concept of speed of computation, as it was originally, but it is directly connected to the notion of business. For example, who still uses a web service that is too slow, or who accesses again a website that takes a long time to answer the requests? This tendency has boosted the spread of applied performance evaluation papers in hundreds of conferences/workshops that are proliferating in recent years focusing on all the computer applications. Case studies describing the applications of performance evaluation techniques to solve real-world problems can be found easily in the large literature available. As a consequence, the distance between the performance evaluation and the other disciplines involved in the computer science is continuously reducing (as described by objective a). Another trend that may be considered positive is the maturity level reached by several tools available in the open source market. Some of them are very popular, even with a hundred thousand downloads or more. Almost all implement the new performance evaluation techniques appeared in the last years, and have a good level of reliability. The application of these tools allows researchers to reduce the time lag between the occurrence of performance problems raised by new architectures and their solution (as described by objective b). On the educational side, the situation we see today is very different from fourteen years ago and in some ways controversial. In the last decade, the extremely fast pace of technologies of computer architectures and applications has forced academia to update the computer science curricula with high frequency. The number of courses on the new (and sometimes exoteric) topics increased exponentially. Each university has many master's programs with trendy titles most of them devoted to very focused applications. The consequence is that the space (in term of credits) for courses devoted to performance evaluation concepts and techniques is continuously decreasing. A problem, already pointed out by Ferrari [1], that seems to affect chronically the education of performance engineering is related to the teachers characteristics. Basically, they can be subdivided into two groups: the applied and the theoreticians. In the first group there are professors with a good background on all (or many) topics involved in the performance discipline, who teach courses (referred to as applied) on how to use the most appropriate techniques, methodologies, and tools to solve the current performance evaluation problems. The second group consists of applied mathematicians that are involved in the development of the theory of new analytical techniques or in the advancement of known techniques to cover new problems raised by the technological progress of computer science. Typically, the *applied* courses attract more students than the *theoretical* ones, but they are far less numerous than the latter. Indeed, historically it seems that the community of theoretical professors is more populated than that of the applied. The risk, in this case, it may be that the performance engineering courses may be perceived too far from the reality, and that the approached problems are irrelevant compared to those that must be solved. It should be pointed out that several excellent textbooks have appeared in recent years covering all theoretical concepts and techniques used in performance evaluation. These books clearly represent a significant support for the teachers of theoretical courses. The same phenomenon has not occurred (with a few exceptions) for application-oriented texts. In summary, a more application-oriented approach in the performance engineering courses and the integration of performance concepts in some of the computer science courses seem the right directions to follow.

3. REFERENCES

[1] D. Ferrari. Considerations on the insularity of performance evaluation. *IEEE Transactions on Software Engineering*, (6):678–683, 1986.

[2] D. Ferrari. Insularity revisited. *G. Kotsis Ed., Performance Evaluation – Stories and Perspectives, Oesterreichische Computer Gesellschaft*, pages 1–10, 2003.