Report on the First International Workshop on Teaching Performance Analysis of Computer Systems 2021

Vittoria de Nitto Personé
University of Rome Tor Vergata
denitto@ing.uniroma2.it

Y.C. Tay
National University of Singapore
dcstayyc@nus.edu.sg

1. INTRODUCTION

The teaching of performance analysis started in the early 70's. The importance of such courses is obvious, considering the immense changes in computing systems over the years. However, rarely do academics meet to discuss and take stock of performance modeling and analysis in their teaching. Moreover, in the last two decades, an economic crisis has involved the educational system, and many changes are happening without the awareness of the parties involved.

These reflections point to the need for this Workshop. Some issues to be discussed included: Are we teaching what our students and the industry want or need? Do we need to change the way we teach performance analysis? Can our teaching ride on the contents of popular courses in the curriculum?

The Workshop was held in conjunction with IFIP Performance 2021. Since the Conference was online (because of COVID-19), the Workshop followed suit. This meant the program schedule was severely constrained by time differences.

The 5 talks and 2 discussions were fast-paced (with just two 5-minute breaks) and finished in under 5 hours. The number of participants fluctuated between 20 and 25, and the discussions covered much ground and generated several ideas.

This report summarizes the invited talks, records the discussions and lists some recommendations. (The extended abstracts and slides are available at the conference website.)

2. TALKS

Mor Harchol-Balter (Carnegie Mellon University) gave the first talk, describing her experience as an industry consultant. She provided examples of some basic concepts (utilization vs delay, open vs closed, etc.) that were lacking among computer systems practitioners, and some simple solutions (job size vs importance, pooling, etc.) that she taught them.

Chee Wei Tan (City University of Hong Kong) described how the Chiu and Jain AIMD (Additive Increase Multiplicative Decrease) abstraction and Perron-Frobenius Theory can be used to teach concepts of fairness, utility maximization, transients and convergence for a many-user system. He also extended the AIMD paradigm to the use of polling, quizzes and chatbots to get feedback from students for regulating content delivery to match "content capacity" of a class, and tuning exercises to suit "comprehension capacity" of individual students.

Cathy Xia (The Ohio State University) illustrated how she further incorporated digital technology in using Excel to run live simulations, and Kahoot! for game-based learning; she observed that the students' competitive nature, and the need for speed in the game, motivated them to pay close attention in class. She also recorded short videos for key concepts and follow-along demonstrations for lab sessions.

Jean-Yves Le Boudec (École Polytechnique Federale de Lausanne) explained how his data science students found their experience in implementing simulators instructive in helping them digest probability (sample spaces, etc.) and statistics (random noise, etc.). The subtle differences between samples collected by different observers also become clear when described with simulation variables.

Giuseppe Serazzi (Politecnico di Milano) noted how, over the years, the focus of performance analysis has shifted from theory to applications; in contrast, current textbooks contain much material on theory and algorithms, and do not take advantage of available software and simulators. He recommended the use of applications and case studies in choosing the (minimal) theoretical contents and their ordering.

The highlight of the Workshop was in the two discussions that included audience participation. Vittoria de Nitto Personé (University of Rome "Tor Vergata") prefaced these discussions by surveying the evolution of performance modeling teaching over the last 50 years. She pointed out three important aspects for reasoning about teaching today:

i. The ratio between specialization and basic knowledge, where nowadays specialization seems to replace basic knowledge.

ii. The role of higher education, where Knowledge seems to be increasingly replaced by skills.

iii. The new generation of students have very different habits, having grown up in close contact with digital devices, and always connected. The *Attention Economy* warns us with the consequent distractability that threatens the educational project. In fact, teachers face the difficulty of grabbing the students' attention, as they are increasingly engrossed in digital technologies.

The following record only highlights some of the points discussed, and regroups them (instead of following the discussion order).

3. DISCUSSION1: CURRENT SITUATION

Gianfranco Balbo (University of Torino) observed that university courses on performance are in decline, in terms of enrolment, content, and class time. Evgenia Smirni (College of William & Mary) noted that the students' mathematical preparation has also weakened.

Meanwhile, performance is super-important to the industry, yet practitioners lack basic knowledge on stochastic behavior (Mor, Evgenia). They rely on simulation, rather than models, but using a simulator as a blackbox risks garbage-in-garbage-out (Cathy, Evgenia, Mor).

There were two apparent contradictions in the discussion:

(1) Should students use off-the-shelf simulators (Giuseppe), or should they write their own code, so they understand the simulation model and the data generated? Jean Yves clarified that the students' understanding would benefit from even an imagined implementation as a thought experiment.

(2) Should textbooks drastically cut down on theory, removing out-dated content (Giuseppe, Mor), and risk undermining the students' training in the fundamentals? Gianfranco suggested having two courses, a low-level one on theory and a high-level one on application.

We summarize the main points shared in this first discussion session:

- The importance of performance analysis is unquestionable.
- There is a lot of ignorance in performance analysis (even in the most basic laws), both in industry and among systems engineers from the scientific community.
- A difficulty in understanding even simple stochastic behavior is also generally observed among students.

4. DISCUSSION2: WHAT WE CAN DO

The discussion considered various ways of making the courses attractive to students. In Cathy's courses, the Teaching Assistants played an important role because they know what students want. Moreover, the activity of reviewing the main concepts strengthens learning.

To engage the industry, she suggested having courses or workshops for executives, and stepping out of our comfort zone to publish success stories on performance modeling in business-related publications.

Andre Bondi (Software Performance and Scalability Consulting LLC) pointed out that ICPE has much participation from the industry because the conference is sponsored by SPEC. They now do a lot of performance measurements that can be used for case studies, but they are not willing to provide the examples. Rather, they will spend much money to bring in experts for targeted, proprietary studies (Mor).

To push back against the squeeze on performance-related courses, we will need to bring colleagues on board. However, Diwakar Krishnamurthy (University of Calgary) encountered resistance from colleagues, who pointed out that the curriculum recommended by ACM/IEEE is silent on performance modeling, while students and practitioners find the mathematics too intimidating. He suggested that the community put together a position paper on the need for performance modeling in university curricula.

One possibility in engaging colleagues may be to give performance-related guest lectures in systems courses, but that may not be effective since not much can be taught in one lecture.

Perhaps the community can collaborate and write some wikis to educate our students, industry practitioners and systems colleagues on fundamental concepts and knowledge? That is not effective either, as the knowledge transfer needs to be more "in their face" (Evgenia).

For the same reason, a tutorial at SIGMETRICS or Performance (even if it is done online, post-COVID) does not bring in the systems people who need the tutorial.

Mor recommended reaching out instead, by giving tutorials at systems conferences (e.g. ISCA). Giuseppe also suggested offering courses for practitioners in the industry, like a MOOC on performance modeling and analysis.

We summarize the main points shared in this second discussion session:

- The hours devoted to performance analysis courses are too few — more time is essential, and more classes.
- Teaching should start at the undergraduate level, with a course about probability and statistics that also shows how these concepts can be applied to predict performance; a course at graduate level should be more focused on interesting and actual applications.
- The workshop should continue every year at SIGMETRICS.

5. CONCLUSION

We bravely labeled TeaPACS as a "First International Workshop", without knowing if there will be interest in having a second one. However, the participants did converge on the need to continue and regularly update the discussion on teaching strategies (the role of simulators, using digital technology, etc.), changing the curriculum, reaching out to the systems community, mounting online courses, etc.

This very report of the Workshop is an effort in motivating and mobilizing the community, to collectively educate students, colleagues and practitioners on the basics, relevance and importance of performance analysis for computer systems.