Message from the
USENIX ATC '24 Program Co-Chairs

Introduction

Welcome to the 2024 USENIX Annual Technical Conference (USENIX ATC '24), which has set a new USENIX ATC record for the highest number of submissions. Continuing the practice started in 2022, USENIX ATC is co-located with OSDI. We very much look forward to meeting everyone in our systems community whether they are attending USENIX ATC, OSDI, or both. Below, we provide some statistics and insights into the selection process that led to the 77 accepted papers that will be presented at the conference, representing an acceptance rate of 77/488 = 15.8%.

Submissions

As in previous years, USENIX ATC '24 solicited two lengths of papers—full and short—in two tracks—research and deployed systems. The "Deployed Systems Track" papers had different criteria for acceptance from research papers—they were not expected to present new ideas or results to be accepted, but needed to convey practical insights from real-world deployment of systems and networks. The authors submitting to the Deployed Systems Track had the option of indicating the reason why they were submitting to this track. To make the review process more efficient, we continued the practice of allowing authors to indicate on the submission form changes that they have made to a previously rejected submission. This description was made visible to reviewers only after they had submitted their review for that submission.

We received a total of 488 submissions, a surprising 38% increase from USENIX ATC '23. Out of these, 7 were desk rejected for one of two reasons: violating anonymity rules or violating paper formatting rules.

Program Committee Selection Process

We set out to assemble a program committee with several goals in mind: a balance of experience and junior members of our community, groups that are traditionally underrepresented in our community, academia and industry, geographic diversity, and research topic diversity (within the big systems umbrella). The PC was assembled, as is the norm, before the submission deadline. We had 136 members on the committee, the largest in USENIX ATC history in keeping with the record number of submissions. The balance was 66% from academia and 34% from industry with 14% of the PC being female and minority members. To recognize exceptional work in the PC and know how this is a critical and voluntary activity, we decided to recognize "Distinguished Reviewers," who would be announced at the conference.

Once the submission deadline passed and we noted the unexpected surge in the number of submissions, we had to think long and hard about how to handle the submission process. We had two limiting factors: we could not increase the review load promised to the PC members at the time of invitation and we could not add to the PC as authors would not have a chance to flag conflicts of interest (COI). After thinking through the possibilities, and with able guidance from the USENIX ATC Steering Committee, we came up with a solution, workable though not perfect. We instituted an External Review Committee (ERC), which is composed of researchers early in their careers. This meant senior PhD students, post-doctoral scholars, and, in a few cases, early-career faculty or professionals. ERCs were assigned reviews that have already got PC reviews and looked to be borderline in round 1. This limited responsibility ensured the overall quality of the program while allowing ERC members to gain experience servicing the community. Part of the rationale for this choice is that we could identify almost all COI mechanically, by noting only the affiliations of the ERC members.

We had to create the ERC within a short time frame, and we are thankful that three of our energetic and qualified community members, Ram Alagappan (University of Illinois at Urbana-Champaign), Aishwarya Ganesan (University of Illinois at Urbana-Champaign), and Haonan Lu (University at Buffalo), stepped up to organize the ERC. We finally had 70 members in the ERC, each of whom reviewed two or three submissions.

Review Process

Of the 481 submissions being reviewed, 49 were Short (10%), and 27 were Deployed Systems Track (6%) submissions. The most popular topics for submissions, as specified by the authors, were "Systems and Machine Learning," "Cloud Computing," and "Parallel and Distributed Systems." Each submission could choose multiple topics out of eight that we had populated the submission form with.

As always, we adopted a double-blind review process whereby the author names were not visible to the reviewers, and (obviously) the reviewer names were not visible to the authors. The author anonymization was maintained all through the review process until the final decisions were announced to the authors. Conflicts with the PC Chairs were also monitored, and when one of us was conflicted with a submission, the other Chair handled it. In the case of one submission that both of us were conflicted with, Brian Noble from the Steering Committee handled it.

To handle the sharply increased number of submissions and to balance the PC and the ERC responsibilities, we decided to send the submissions out for 2 reviews initially and then, depending on their strength, 1 additional review, from an ERC member, and then finally, those that looked promising, 2 additional reviews. On average, a submission received 3.4 reviews. At each step, we triaged to make sure the committee was spending more of its effort on competitive submissions. This process was more nuanced than the standard one of 3 + 2 reviews, in rounds 1 and 2, respectively, with filtering between the rounds. Our process was necessitated by the above-stated facts of increased submission count and boundary conditions for the PC size and effort.

We had an author rebuttal window of 3 days for the authors to write a recommended 1000-word response to the reviews. Prior to this, the discussion lead for each submission summarized the main comments that the PC wanted addressed by the authors. The reviewers then continued an asynchronous online discussion for 4 days before a synchronous virtual PC meeting held over 2 days. We flagged 65 submissions to be discussed at the PC meeting with the proviso that any PC member could bring up any additional submission for discussion. This set of submissions were ones where there was not a consensus for either accept or reject based on the online discussion. The PC meeting was a vigorous, stimulating, and respectful affair. We thank Brian Noble for helping us, the PC Chairs, during this meeting. At the end of it, we decided on the set of papers to be accepted, 15 of which were with shepherding. This year, we decided to handle the shepherding process through HotCRP with shepherd identity being kept anonymous by default from the authors, while giving shepherds the option to opt out and deanonymize themselves. This had the goal of adding teeth to the shepherding process, which we believe succeeded. After the shepherding process concluded, the shepherds had the option of revealing their identity, and most did. All of the shepherded papers were ultimately accepted.

Artifact Evaluation Process
We continued to run a joint artifact evaluation process with OSDI, led this year by Jianyu Jiang, Ji Qi, and Cesar A. Stuardo. The artifact evaluation committee chairs assembled a committee consisting of 127 members. The authors of all accepted papers were invited to submit an artifact for an evaluation. 46 of the accepted papers (60%) did so. All but one of the submissions received an "Available" badge, 76% of the artifact submissions received a "Functional" badge, and 56% received a "Reproduced" badge. 32% of the total number of accepted papers received all three badges (one artifact was functional and reproduced, but is not publicly available).

Acknowledgements

Saurabh Bagchi, *Purdue University*
Yiying Zhang, *University of California, San Diego*
USENIX ATC '24 Program Co-Chairs