Message from the USENIX ATC '21 Program Co-Chairs

1. Introduction

Welcome to the 2021 USENIX Annual Technical Conference (ATC).

Running a conference of any size is a significant challenge, one that comes with many opportunities to build upon and perhaps improve what has gone before—and the challenge is all the greater for a conference of the scale and caliber of ATC.

We are fortunate that ATC is sponsored by USENIX, an organization with extensive experience in conference planning. Working with USENIX means joining a well-oiled machine that magically takes care of an almost endless list of tasks that must be completed for a conference to be successful. We are humbled and honored by their trust in asking us to take the helm of USENIX ATC '21 and we are grateful for their invaluable guidance and support throughout the process.

We hope that everyone is as thrilled as we are with the ensuing program. In the rest of this document, we describe the process for USENIX ATC '21, with an emphasis on how we created and managed the program committee.

For USENIX ATC '21, we had to deal with two unusual challenges: first, the ongoing pandemic, and second, the collocation with OSDI '21. The former turned out to be more of a nuisance than a challenge, thanks to the help we received from the USENIX ATC '20 co-chairs, Erez Zadok and Ada Gavrilovska. Ada and Erez had been forced to adapt to the pandemic as it unspooled, switching both their program committee meeting and the conference itself from in-person to virtual. They extensively documented their experience in their own conference message, which we referred to repeatedly, and also offered additional answers to specific questions we had. As a result, we were able to plan a smooth PC meeting (which managed to finish on time!) and are confident that the conference itself will be equally successful.

2. Running a Combined Conference

Navigating the decision to collocate OSDI and ATC was more difficult. In response to requests from the systems research community, USENIX concluded that OSDI needs to be held more often, changing from strict alternation with SOSP to an annual conference. However, for a number of years both SOSP and OSDI have taken place in October; USENIX felt that it would be inappropriate to hold OSDI close to the time when SOSP was scheduled. Instead, after carefully considering the options, the USENIX board decided to experiment with holding OSDI and ATC together.

That presented its own dilemma: since OSDI and ATC are both systems conferences, would they find themselves in competition? How would the the conferences themselves be run? Would either conference experience a significant drop—or a rise—in submissions? Fortunately for us, we had an ace in the hole: the OSDI '21 co-chairs, Angela Demke Brown and Jay Lorch, are two of the nicest and most generous people in systems research. From the beginning, we were able to work closely to ensure a smooth experience.

The first question we had to address, almost immediately after the co-chairs for both conferences were chosen, was submission deadlines. Should the deadlines be the same for both conferences? If not, which conference should "get" to go first? We concluded that OSDI should have an earlier deadline, almost exactly a month before our own. That would allow authors who were unable to make the OSDI deadline to have a significant period of time for completing their papers before they submitted to ATC. It also reduced the burden on the USENIX staff by preventing two simultaneous crunch times.

We then turned to the challenge of selecting program committees for both conferences. OSDI '20 received 400 submissions; USENIX ATC '20 received 350. If those levels continued, we would need a record number of PC members across both conferences. That raised the risk that we would find ourselves competing for PC candidates.

To avoid that possibility, the four co-chairs of the two conferences met several times to discuss PC invitations and create a comprehensive list of candidates. Working from a shared spreadsheet, we created one list for each conference, balancing both lists on a number of axes to ensure that we would have diverse committees that covered a range of viewpoints, areas of expertise, and levels of experience. We gave special attention to including less established researchers, both to advance their careers and to develop new possibilities for future program committees.

Our cooperation in identifying PC members was successful; both conferences were able to assemble strong committees. In the end we wound up with a PC of 120 members: 69 heavy (16-20 reviews expected), 39 light (8-10 reviews), and 12 extended (3-5 reviews).

After considering previous submission levels and growth, we decided not to assume that either conference would experience a drop in submissions; instead we built a program committee that would be able to handle an increase. That proved to be a wise decision.

3. Program Committee

The most important task that we did as USENIX ATC '21 program chairs was selecting and inviting the program committee. The PC had 120 members, with expertise in a wide range of topics. The areas with the highest expertise among the PC members were Distributed Systems, Storage, Cloud, Architecture, OS, Analytics and Machine Learning. Those fields allowed us to cover most of the submissions well, but we could have used more expertise in Quantum Computing. We also were surprised to discover that we had multiple submissions in the area of Low-Earth-Orbit Satellite Networks; we had to seek additional help on that topic.

In addition to covering a wide range of research topics, we also wanted to ensure diverse representation in the PC across several dimensions:

- Junior vs. Senior members of the community. We believe it is important to invite new community members to take part in the review process. Our PC had 34% junior members.
- Academia vs. Industry. USENIX ATC has broad industry impact, so we wanted to have this duality reflected in our PC membership as well. Industry and Industrial Research Labs members constituted 34% of the PC.
- Women and Under-represented Minorities. The PC representation was 22%.
- Geography. 63% of the PC members are located in North America, 24% in Europe, 8% in Asia, and the remaining 5% in other areas.

The PC was split into three categories, depending on the number of reviews the PC members were able to provide, as well as whether they could attend the PC meeting or not: heavy (up to 17 reviews), light (up to 9 reviews), and extended (up to 5 reviews). The heavy PC consisted of 68 members, the light PC 39 members, and the extended PC 12 members.

We found it important for each PC member to provide us information before the reviewing process started. In particular, it simplified the conflict verification process if each PC member updated their own lists of collaborators and other conflicts. Doing this before the submission deadline is ideal, as it provides the authors the opportunity to double-check the identified conflicts through HotCRP. In addition, the PC members rank their topics of interest from the conference topics. While this step is optional for PC members who spent the time to go through the submissions after the deadline and submit their paper preferences, we asked the PC members who did not to make sure they at least had their topics preference filled in. It would have been impossible for us to assign papers to PC members who had expressed neither a paper nor a topic preference. In retrospect, the PC members who actually took the time to bid on (a larger number of) papers also got a better assignment. In addition, it was easier for us to assign papers when extra reviews were needed.

3.1. Submissions Chairs

Following the prior two editions of USENIX ATC, we adopted the practice of having two submissions chairs. Alex Conway (VMware Research) and Chris Stone (Harvey Mudd College) provided invaluable support during the USENIX ATC evaluation period. In particular, the submissions chairs helped with the pre-review checks, verified the authors' reported conflicts, helped with checking the quality of the reviews and took notes during the PC meeting. We would not have been able to put the conference together without their able and tireless assistance.

4. Review Process

Perhaps the most important part of putting a conference program together is the lengthy period of peer review. As is common for most large conferences, we divided the review into several steps.

4.1. Pre-Review Quick Checks

After the deadline passed, we had a total of 341 completed submissions, excluding spurious ones. We divided the papers into four equal groups assigned to the co-chairs and submissions co-chairs for a quick validity check. Each person scanned their pile for several flaws:

- Anonymity violations, both obvious (authors on the title page) and more subtle (self-citation, mention of authors' institutions, etc.).

- Formatting problems not caught by the automated format checker (e.g. unreadable figures).
- Violations of the Call for Papers, such as appendices following the bibliography.
- Other random submission errors, such as accidentally submitting a version of the paper as supplemental material.

During the quick checks, we also identified an oversight in the Call for Papers. A number of authors had included links to online copies of their source code in their papers. Such links had not been prohibited by the CFP, but they created the risk of an anonymity violations. In addition, since online resources can be modified at any time, there was no way to ensure that they were representative of work done before the deadline. For those reasons, we decided to ask authors to remove those links.

In the end, we rejected one submission immediately because it was too flawed to be rescued and because it violated the anonymity requirement by including the authors' names. We gave the authors of 67 (!) others a brief window in which to correct the problems; all accepted the offer. Although the CFP clearly stated that format violations would result in instant rejection, we chose to be gentle because the violations appeared to be unintentional mistakes rather than an attempt to evade the rules.

Finally, we also reviewed the papers for undeclared conflicts with PC members. That task was eased because the submission form had included a request for DBLP links for all co-authors. Nevertheless, that task was a major one, and we hugely appreciate Alex and Chris for carrying it out. We added a number of new conflicts into HotCRP, and notified the authors so that they would be more careful in the future.

4.2. Review Timeline

The review process was scheduled over two rounds, with ample discussion time after each round. The discussions were very useful not only because PC members had time to read each other's reviews and reach a common decision, but also because it gave us time to supplement reviews where necessary by asking other PC members or external reviewers. For example, it was necessary to supplement reviews when a paper had lower reviewer expertise, or when the reviews were late or missing. We advise future co-chairs to allow for these buffer periods as well, while still enforcing the round deadlines.

4.3. Review Sequence

Once we had 340 conforming submissions, we assigned reviews to the members of the program committee. In the first round, heavy PC members were assigned 10-12 papers, with proportionally fewer going to light and extended members. Each paper received three first-round reviews. An unexpected issue arose when one member was forced to withdraw from the PC due to serious health issues. We are hugely grateful to Reto Achermann, David Cock, Dave Dice, and Scott Stoller, who were willing to accept an additional assignment on insufficient notice to help out.

After the round-1 reviews arrived, we held brief online discussions to decide which papers would advance to round 2. We took a cautious approach, leaning toward advancement for papers that had received at least one weak accept and papers where no reviewer had rated themselves as having high expertise in the area. At the same time, though, we were cognizant of the fact that we were planning to send round-1 reject notices right away (rather than waiting until after the PC meeting), since rejected authors generally prefer to receive the news early so that they can begin work on a revision. In the end, 178 papers advanced to the second round.[1]

Papers rejected in round 1 did not have the option to provide an author response, due to the early notification. We found that a small number of the authors of the papers rejected in round 1 would have preferred to take the option of a later notification, so that they could have the chance to submit a response to the reviews. Thus, future chairs might consider incorporating an option to choose between early notification and response at submission time, or adding a response period to round-1 rejected papers.

For the papers that advanced, we solicited brief (500-word) responses to the reviews from the authors. Unfortunately HotCRP does not enforce the word limit, and some authors went far over. Nevertheless, the program committee found the responses useful in a number of cases, and some papers were accepted at least partly based on the clarifications given in the responses.

Each round-2 paper received at least two additional reviews, so that in the end every paper had at least five[2]. After those reviews were in, we again held online discussions to divide the papers into three categories: clear accepts, clear rejects, and those that needed in-person discussion at the online PC meeting. In several cases we asked for additional reviews from experts in the area, either on the PC or external.

[1] A few were advanced despite low ratings because some reviews were still missing at the deadline.

[2] The papers that would have been rejected in R1 but for their late reviews were an exception, having only 3 or 4 total reviews.

4.4. PC Meeting

In the end, we wound up with 52 papers that needed discussion by the program committee. We had planned from the first to hold the PC meeting online over a period of two days. Inspired by USENIX's conference-scheduling practices, we decided to limit the meeting to five-hour slots, running from 7:00 am to noon Pacific Time; those times are acceptable on the West Coast of the U.S., pleasant on the East Coast, and not horrible in Europe. Unfortunately they work less well in some other locales (e.g., 7:30 pm to past midnight in India, midnight to 5:00 am in Korea, and 1:00 am–6:00 am in Eastern Australia).

To deal with the time-zone problem, we polled the PC members to ask where they would be on the days of the meeting. We then hand-scheduled the papers for discussion on a rolling basis, beginning with those that had reviewers with the most challenging time zones so that those people could leave the meeting as soon as possible. The approach worked effectively, but only because of the heroic efforts of our distant PC members to stay up far beyond what most of us would consider a reasonable time.

In past face-to-face PC meetings, a significant amount of time was lost moving conflicted members in and out of the meeting room; HotCRP offers an option to automatically generate a schedule that minimizes such motion. We discarded that option in favor of accommodating time zones; we created a single Zoom breakout room named "Hallway" and used it for conflicts. That worked well; we were usually able to move people into and out of the room in 10–30 seconds so little time was lost.

Given necessary overheads such as introductions and breaks, we wound up with about 10 minutes to discuss each paper. As is common in PC meetings, discussions of the early papers often ran over, but the committee gradually developed a sense of timing so that we finished the first day's work on schedule. As is also common, the reviewers of a few papers came to an accept/reject decision overnight, and in the end we were able to finish slightly early.

5. Statistics

It has become traditional to offer a number of statistics about the submissions and the review process. Some of the numbers are impressive:

- 341 submitted papers (25 short)
- 64 accepted (1 short), 277 rejected (23.1% acceptance rate)
- 119 PC members wrote 1357 reviews
- 750,218 total words in reviews

As has been common in recent years (and not just in ATC), short papers faced something of an uphill battle, with only a 4% acceptance rate. We encourage future chairs to consider whether the short-paper category is a good idea.

6. Conference Plans

Although we are writing this message well in advance of the conference, we are looking forward to a few changes that we hope will make it run smoothly and improve the virtual conference experience for attendees:

- Before the pandemic, several conferences had run successful experiments with "preview sessions" aimed at new attendees. The goal of these sessions was to provide enough background for the uninitiated to understand a group of papers being presented during the conference. This year, we have asked some PC members to create similar previews that will be posted in advance of the conference itself. The preview session will cover common topics from both USENIX ATC and OSDI.
- One of the most important things that was lost in the shift to virtual conferences was the so-called "hallway track"—the break times during which attendees mingled, met new people, and discussed research ideas. There have been several attempts to find substitutes, including extended Zoom sessions, Zoom breakout rooms, and Slack channels; however, none of these options have been entirely satisfactory. At USENIX ATC and OSDI, we will be experimenting with a new option, ohyay.co, which we hope will provide an attractive and fun alternative where small groups can interact with each other in a productive fashion. We want to provide a more interactive environment for attendees to replicate the informal discussions that would normally happen at the conference, as well as provide additional time for authors to answer questions and discuss their work in a smaller group setting.
- Several editions of OSDI have implemented a successful mentoring program—matching students with senior researchers and facilitating 1-1 discussions outside the conference program. A few USENIX ATC '21 PC members are collaborating with OSDI '21 PC members to create a common mentoring program for both conferences.

7. Final Thoughts

Co-chairing USENIX ATC ’21 has been a challenging, instructive, and rewarding experience. In the end, we are pleased to have been able to create a strong program featuring many new insights and creative ideas. We are grateful to everyone who contributed to creating the program and to all who are still working behind the scenes to ensure that USENIX ATC ’21 runs smoothly and provides value to all attendees; to the authors who worked hard on their submissions and final papers; to the program committee and external reviewers, who spent countless hours reading, discussing and selecting the best papers; to the shepherds, who ensured that each paper appears in the program in its best shape; and to the volunteers who are organizing mentoring, networking, and preview sessions jointly with OSDI ’21. At last, but most importantly, we are grateful to our submissions chairs, who jumped in at a moment’s notice to help out with all the tasks where we felt overwhelmed, and to the wonderful USENIX staff, who directed us and cleared any roadblocks.

We thank everyone for their hard work and dedication. We are humbled by your willingness to go above and beyond to make USENIX ATC ’21 a success!

We hope that you enjoy the conference!

USENIX ATC ’21 Program Co-Chairs
Irina Calciu, *VMware Research*
Geoff Kuenning, *Harvey Mudd College*