

# **Conference Reports Column**

by

Laurent Feuilleley

Laboratoire d'informatique  
en image et systèmes d'information (LIRIS)

CNRS and Université Lyon 1

43 Bd du 11 novembre 1918, 69622 Villeurbanne, France

[laurent.feuilleley@cnrs.fr](mailto:laurent.feuilleley@cnrs.fr)

# Conference report on STOC 2026

Matthias Bentert (Technische Universität Berlin)

The 58th Annual ACM Symposium on Theory of Computing (STOC'26), sponsored by the ACM Special Interest Group on Algorithms and Computation Theory (SIGACT), took place from June 22 to 27 in Salt Lake City, Utah. The Hilton Salt Lake City Center hosted six days of paper presentations, workshops, invited talks, social events, and—thanks to a late addition on Saturday—a full-day discussion of whether AI can do theory.

For context, I am a postdoc in Germany working on algorithms. I was at STOC to present joint work on a polynomial-time algorithm for fault-tolerant packet forwarding in communication networks. Overall, I had an amazing experience and I would like to share some of my memories with you. But first, here is a riddle to guess if you can: why did I return from a very hot week in Utah with only one regret: not having brought warmer clothes?<sup>1</sup>

I cannot possibly do justice to the technical program in a conference report as STOC featured 211 accepted papers often in four parallel sessions. Luckily, I do not have to. All the talks will be available on the youtube channel of the SIGACT: <https://www.youtube.com/@sigactec8649>. At the time of writing, only some of the talks are online, but hopefully, all talks are available at the time of reading. Let me just highlight one presentation:

A large online collaboration set out to investigate the fifth busy beaver number and, with the help of the Coq proof assistant, they proved that it is 47176870.<sup>2</sup> The proof classifies more than 181 million five-state Turing machines and determines whether each of them halts. It is the first new Busy Beaver value determined in more than forty years and the first one to be formally verified. Tristan Stérin presented the result on behalf of the collaboration. Besides settling a wonderfully concrete question about an uncomputable function, the project was also a strong demonstration of what a large, distributed group of contributors can accomplish.

---

<sup>1</sup>Hint: this was only the second time I visiting the United States.

<sup>2</sup>Justin Blanchard, Daniel Briggs, Konrad Deka, Nathan Fenner, Yannick Forster, Georgi Georgiev (Skelet), Matthew L. House, Maja Kadziolka, Pavel Kropitz, Shawn Ligocki, mx dys, Mateusz Nasciszewski, Tristan Stérin, Chris Xu, Jason Yuen, Théo Zimmermann. *Determination of the Fifth Busy Beaver Value. STOC 2026: 565-574*




Figure 1: A figure from the Busy beaver paper (see footnote) that describes the "maximum" five-state Turing machine, and its space-time diagram at different scales.

Tuesday evening's banquet was held at the Ken Garff University Club. Unfortunately, jet lag overcome me this evening, so the following is second-hand information: both the food and the evening's Luca Trevisan Award Talk by Scott Aaronson were wonderful, so I am sad to have missed it. At least I will get the chance to listen to the talk once it is uploaded.

Thursday brought another exceptional event. Charles H. Bennett, co-recipient with Gilles Brassard of the 2025 ACM A. M. Turing Award, chose to give his award lecture at STOC. You may know Bennett for many different contributions. The first results I encountered bearing his name was the Bennett–Gill theorem: Relative to a random oracle  $O$ , it holds that  $P^O \neq NP^O$  with probability one. His Turing Award lecture was titled "Quantizing Shannon, and Freeing Cosmology from the Need Define Consciousness or Count Observers". Bennett went "slightly" over time, but please, see for yourself. The recording is already up on youtube.

On Friday evening, Jane Street organized what has by now become a welcome tradition: the Estimathon. Teams had to estimate quantities that nobody could reasonably be expected to know. One memorable example question: What is the smallest integer  $n$  such that the  $n^{\text{th}}$  Fibonacci number contains the substring "2026"?

I was unsure whether to include a report on the served food or not. In the end, I decided to add some information, but please be aware that the sample size was very small. Apart from small refreshments, only three meals were provided. Monday's lunch boxes were not especially well received. The Tuesday banquet was, by all reports, excellent, and Saturday's buffet was also very good. The resulting empirical picture is therefore mixed but improves

substantially if one conditions on the food not being handed over in a box.

Let me now report some of the numbers from the business meeting: This year saw 740 submissions, continuing the upwards trend of the last years (for reference, 721 submission in 2025, 580 in 2024, and 479 in 2023). Of these, 211 papers were accepted, giving an acceptance rate of roughly 29%. The reported acceptance rates were 28% for papers without a program-committee author and 34% for papers with at least one PC member as an author. The program committee consisted of 80 members, each assigned between 27 and 33 papers. There were 376 registered participants, most of them from North America. This was below the 490 participants in 2025, which appears to have been an outlier, and broadly in line with earlier years. Of 41 applications for student travel support, 27 were funded. Finally, STOC 2027 will take place in Atlanta.

The conference ended on Saturday with the workshop “Can AI Do Theory?”. Scott Aaronson suggested that current AI systems are particularly strong at two things: trying many examples very quickly and recalling and combining facts from different areas. Prabhakar Raghavan presented compelling evidence that these abilities can contribute at the highest research level: an AI system discovered a new gadget that improves the hardness-of-approximation bound for metric TSP from  $117/116$  to  $111/110$ . Most memorable for me was the concluding panel and I want to highlight some ideas that stuck with me. Yael Kalai observed that imperfect AI reviews should not be compared with an imaginary standard of perfect human reviewing; human reviews have serious weaknesses too. At the same time, she warned that asking AI at every opportunity is tempting but may erode our own capabilities in the long run. Jon Kleinberg emphasized that, while accelerating research is desirable, the community must preserve pathways by which young researchers can learn to do research and enter the field. Adding to Scott Aaronson’s list of strong sides of AI, the panelists agreed that personalized learning is already a particularly effective application. Kleinberg also gave some encouraging words, especially for younger researchers who might feel threatened by the rapidly changing AI world. He argued that predicting the situation ten years from now is impossible, so remaining adaptive is a better strategy for the future than giving up on shaping it.

To conclude this report, I want to thank the organizers. In particular, Program Committee Chair Artur Czumaj, Local Arrangements Chair Aditya Bhaskara, and TheoryFest Chair Andrew McGregor, but also the many committee members, volunteers, and organizers who made the week possible.

Is there anything else to report? Oh yeah, the opening riddle. Salt Lake City was indeed hot, but as this was only my second trip to the United States, I was unprepared for the formidable power of American conference-venue air

conditioning. Outside, Utah supplied the summer, but inside, the Hilton felt closer to a German winter. My warmest recommendation for future STOC attendees is therefore a literal one: bring a sweater!

## Conference report on PODC 2026

Dominick Banasik and Nithin Salevemula

This year's ACM Symposium on Principles of Distributed Computing (PODC) was held from July 6–10 in Egham, England, at the Royal Holloway. It was co-located with ACM Symposium on Parallelism in Algorithms and Architectures (SPAA) and EATCS International Colloquium on Automata, Languages and Programming (ICALP). Attendees of any conference were allowed to attend all three, which gave us all the opportunity to see other keynote talks and ongoing research (although the program was a lot to keep track of).



Figure 2: The Founder's Building at night.

Below are a few highlights:

- The Founder's Building, where we all enjoyed our English breakfast, is beautiful, and it glows in the dark!

- Introducing the Student Spotlight was a great idea. It gave students an opportunity to present their research interests and receive valuable suggestions and feedback from experts. It would have been even better if the poster presentations had been officially scheduled for at least 1 hour during one of the main conference days, as many participants could not attend the Student Spotlight on the first day.
- LLMs were acknowledged for their contributions to multiple papers. *Large language models* were acknowledged for developing a new lower bound technique, and *lukewarm liquid machines* were acknowledged for, presumably, helping to stay awake.

## 1 Business Meeting

The business meeting ran long. Below are a few key takeaways:

- Co-location with SPAA and ICALP was popular, although it made organization more difficult.
- This year’s Call for Papers was the first to feature an AI policy for authors. AI policies for authors, reviewers, and the program committee will be further fleshed out for next year.
- There were long discussions about having two submission windows. The idea was not so popular, but we will see how it works out for SPAA.
- Another main topic was the location of future conferences. One proposal for the next next PODC was Crete, Greece. Extended discussions followed about hosting PODC in more central locations more frequently.

## 2 Interesting Papers (personally)

The following are a few personally interesting talks.

- The talk “*Distributed Approximate Maximum Matching and Minimum Vertex Cover via Generalized Graph Decomposition*” by Peter Davies-Peck shed some light on the necessity of both  $\Delta$ - and  $n$ -dependent terms in the well-known distributed lower bounds for these problems.

- The talk “*Ranking Opinions with Few States in Population Protocols*” by Antoine El-Hayek introduced a neat application of stable matchings for finding the plurality opinion in population protocols.
- The Dijkstra Prize Talk, “*The  $\Omega(D + \sqrt{n})$  Lower Bound Story of Distributed Algorithms*” by Gopal Pandurangan, was particularly interesting. The talk provided an overview of the origins of the  $\Omega(D + \sqrt{n})$  lower bound and highlighted the wide range of distributed computing problems in which this bound appears.
- The talk “*2-Coloring Cycles in One Round*” by Jukka Suomela was interesting. It presented a one-round randomized distributed algorithm that achieves an expected fraction of monochromatic edges below 0.24118, while showing that no one-round algorithm can achieve a fraction below 0.23879. It was also interesting that the proofs were largely developed using large language models and formally verified in Lean 4.