Distributed Computing Column 68
Annual Review 2017

Jennifer L. Welch
Department of Computer Science and Engineering
Texas A&M University, College Station, TX 77843-3112, USA
welch@cse.tamu.edu

As with prior December issues, this issue is devoted to a review of notable events related to distributed computing that occurred during the year.

First, congratulations to Elizabeth Borowksy and Eli Gafni, winners of the 2017 Edsger W. Dijkstra Prize in Distributed Computing for their paper "Generalized FLP Impossibility Result for t-Resilient Computations"! Their paper appeared in the 1993 ACM Symposium on Theory of Computing (STOC). The prize is jointly sponsored by ACM and EATCS, and is given alternately at PODC[1] and DISC[2]; this year it was given at DISC. To quote from the award committee's citation,

> "...this paper turned out to be crucial for our understanding of fault-tolerant distributed computing. It proposed a powerful reduction technique, the BG simulation, it introduced the immediate-snapshot model, and it established the fundamental impossibility of k-resilient k-set consensus."

The full citation can be found at http://www.podc.org/dijkstra/2017-dijkstra-prize/.

Congratulations as well to Mohsen Ghaffari, who received the 2017 Principles of Distributed Computing Doctoral Dissertation Award! His thesis is entitled "Improved Distributed Algorithms for Fundamental Graph Algorithms" and was supervised by Professor Nancy Lynch at the Massachusetts Institute of Technology. The award is jointly sponsored by PODC and DISC, and was given at PODC this year. The citation appears at http://www.podc.org/dissertation/2017-dissertation-award/. The thesis includes new and improved results for maximal independent set, vertex- and edge-connectivity, and scheduling mulitple network tasks. And this is just the tip of the iceberg of the results he obtained while in graduate school!

[1] ACM Symposium on Principles of Distributed Computing
[2] EATCS Symposium on Distributed Computing

The first article of the column is a review of SIROCCO[3] 2017 by Christian Konrad. The best paper award was given to Magnus M. Halldorsson, Stephan Holzer, and Evangelia Anna Markatou for their paper "Leader Election in SINR Model with Arbitrary Power Control". Congratulations to Magnus, Stephan, and Evangelia!

The second article of the column is a review of SPAA[4] 2017 by Sepehr Assadi. Sepehr is a coauthor, together with Sanjeev Khanna, of "Randomized Composable Coresets for Matching and Vertex Cover," which received a best paper award at SPAA. The paper "Distributed Partial Clustering" by Sudipto Guha, Yi Li, and Qin Zhang also received a best paper award. Congratulations to Sepehr, Sanjeev, Sudipto, Yi, and Qin!

Next, Peter Davies provides us with a review of PODC 2017. Peter's paper "Exploiting Spontaneous Transmissions for Broadcasting and Leader Election in Radio Networks", coauthored with Artur Czumaj, received the best student paper award. The best paper award went to "A Simple Deterministic Distributed MST Algorithm, with Near-Optimal Time and Message Complexities" by Michael Elkin. Congratulations to Artur, Peter, and Michael!

The column closes with a review of DISC 2017 by Manuela Fischer and Yannic Maus. Manuela was the winner of the best student paper award for "Improved Deterministic Distributed Matching via Rounding". Yannic is an author, together with Mohsen Ghaffari, Juho Hirvonen, Fabian Kuhn, Jukka Suomela and Jara Uitto, of "Improved Distributed Degree Splitting and Edge Coloring" which received the best paper award. Congratulations to Manuela, Mohsen, Juho, Fabian, Yannic, Jukka, and Jara!

Many thanks to Christian, Sepehr, Peter, Manuela and Yannic for their contributions!

Call for contributions: I welcome suggestions for material to include in this column, including news, reviews, open problems, tutorials and surveys, either exposing the community to new and interesting topics, or providing new insight on well-studied topics by organizing them in new ways.

[3]International Colloquium on Structural Information and Communication Complexity
[4]ACM Symposium on Parallelism in Algorithms and Architectures

SIROCCO 2017 Review

Christian Konrad
Department of Computer Science
Centre for Discrete Mathematics and its Applications (DIMAP)
University of Warwick
Coventry, UK
c.konrad@warwick.ac.uk

The 24th International Colloquium on Structural Information and Communication Complexity[1] (SIROCCO 2017) took place from the 19th - 22nd June 2017 on the beautiful island of Porquerolles, which is located 20 minutes by boat off the coast of Hyères in the French Mediterranean Sea. SIROCCO has a tradition of being located in fantastic locations, and this year's edition was no exception. Porquerolles is a small island (roughly 7 x 3 km) with beautiful beaches, impressive cliffs, and a village large enough to host a resort facility (IGESA[2] resort) that offers plenty of accommodation and a small conference venue. This year's edition had 56 participants, which is similar to previous years. I attended the SIROCCO conference for the first time this year. I had heard that the atmosphere at SIROCCO conferences is usually rather colloquial, and I can confirm that this is indeed the case, which is most likely due to the great location, the limited number of participants, and the fantastic organization. The conference received 41 submissions, which is 9 submissions less than last year, and accepted 21 papers. As pointed out in the conference business meeting, the lower number of submissions is most likely due to the fact that the conference deadline was a month earlier than last year and preceded the notification date of other conferences. The 21 paper presentations were scheduled over three days, together with three keynote talks, a SIROCCO award lecture, the business meeting, and the celebration of Andrzej Pelc's 60th birthday. This left enough room for a boat excursion, the conference banquet dinner, and short refreshing dips in the sea to cool down and enjoy the Mediterranean sun for a while.

The scope of SIROCCO is decentralized systems of communicating entities. The focus lies on the interplay of structural knowledge, communication, and computational complexity in such systems. The conference thus encompasses research areas such as distributed and parallel computing, communication complexity, mobile computing, peer to peer networks, social networks and many

[1] https://sirocco2017.lif.univ-mrs.fr

[2] https://www.igesa.fr/vos-vacances/recherche-sejour/reserver-mes-vacances/etablissement/porquerolles/

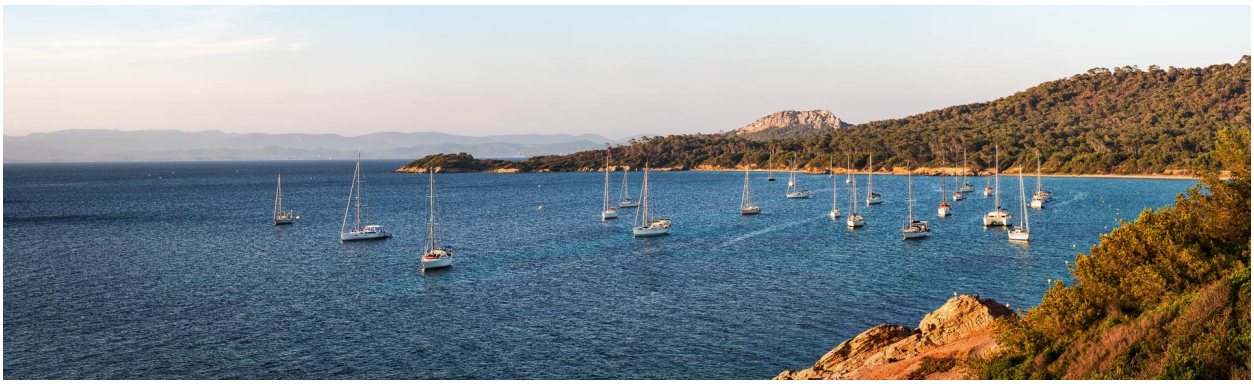

Figure 1: Lovely view of Porquerolles Island. Credit for all photos in this article goes to Sebastien Tixeuil (LIP 6, University Pierre et Marie Curie, Paris).

others. For example, in this year's edition, four papers addressed mobile agents that collectively need to solve a common problem (e.g., exploring the underlying graph or gathering at the same point), three papers considered labelling schemes, two papers dealt with the SINR model for wireless networks, and another two papers addressed the $\mathcal{LOCAL}$ model for distributed computation.

The main purpose of this report is to summarize the technical contributions of the conference. I will also write briefly about the social activities that took place, and hopefully convince the reader that participating in this conference is a very enjoyable experience.

1 Technical Contributions

1.1 Keynote and Award Lectures

Three keynote lectures were given at this year's edition of SIROCCO. On the first day, Faith Ellen (University of Toronto) opened the technical part of the programme with her keynote lecture entitled "Ignorance is Bliss (for Proving Impossibility)". She demonstrated that impossibility results in distributed systems are often obtained by exploiting the lack of knowledge of one entity about the state or input values of others. This talk discussed classic impossibility results for problems such as leader election (á la Angluin), consensus (Fischer, Lynch, Paterson), and coloring the ring (Linial). She also discussed a recent result by Frischknecht, Holzer and Wattenhofer, who proved that if message sizes in a distributed system are bounded by a parameter B, then $\Omega(n/B)$ rounds are necessary for n network nodes to determine the diameter of the input graph, even if the diameter is either 2 or 3. This result is obtained by a reduction to the Set-Disjointness problem in the two-party communication complexity framework.

The second keynote lecture opened the afternoon session of the first day. It was given by Christian Scheideler (University of Paderborn) and concerned programmable matter. Programmable matter is an interdisciplinary research area involving physics, chemistry, bioengeneering, robotics, and, as Christian Scheideler made clear during the presentation, distributed algorithms. It is conceivable that nano-scale programmable devices will be built in the foreseeable future. These devices will require appropriate programming in order to solve specific tasks. Useful applications of programmable matter include shape formation and coating problems (e.g., small particles that coat

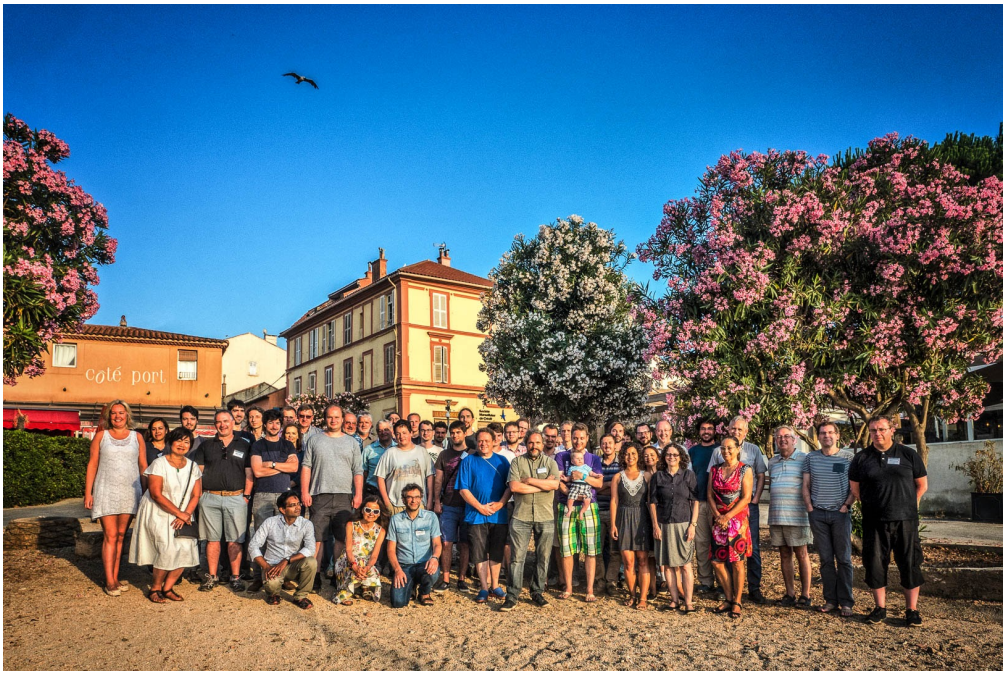

Figure 2: Conference participants before the banquet dinner

buildings or bridges in order to monitor their structural integrity). Christian Scheideler presented the "Amoebot model", where the movements of individual particles share similarities to the movement of amoebas. Leader election is a critical building block for algorithms in this model. He discussed a solution to the leader election problem and how this problem can be used to solve other problems. Finally, extensions of the Amoebot model as well as preliminary results obtained for these extended models were discussed.

No keynote lecture was scheduled on the second day in order to leave time for the SIROCCO 2017 award lecture. Schmuel Zaks (Technion) received the "Prize for Innovation in Distributed Computing 2017". This prize is awarded to outstanding researchers who identified new problems or research areas that were unorthodox at the time of their introduction, but later attracted the interest of the SIROCCO community. Shmuel Zaks received this award for the entirety of his achievements, but especially for his pioneering work on algorithmic aspects of optical networks. He introduced the rigorous study of algorithmic problems related to ATMs and optical networks, which have previously only been tackled by heuristics. His contributions lie in the areas of approximation and online algorithms, complexity, parameterized complexity, and inapproximability. He is the 9th recipient of this award and joins the group of previous recipients consisting of Nicola Santoro, Jean-Claude Bermond, David Peleg, Roger Wattenhofer, Andrzej Pelc, Pierre Fraigniaud, Michel Raynal, and Masafumi Yamashita. His lively lecture was much appreciated and provided a very clear overview of his contributions.

The third and final keynote lecture was given by Christoph Lenzen (Max Planck Institut Saarbrücken) on the morning of the third and final day of the conference. It concerned clock synchronization problems: Even if multiple computational entities start their calculations with synchronized clocks, their clocks will differ from each other after some amount of time, since they

usually tick at slightly different rates. The objective is to keep them synchronized. Christoph Lenzen first argued that clock synchronization problems occur not only in distributed systems but also in hardware chip design, i.e., he justified that chips should in fact be regarded as distributed systems. In his presentation, he addressed various practical and theoretical issues, for example, the presence of uncertainty and faults, and pointed out open problems that are of interest to the distributed computing community.

1.2 Regular Papers

Two of the 21 accepted papers addressed the SINR (Signal-to-interference-plus-noise ratio) model for wireless networks. In this model, wireless network nodes are located on the plane and can transmit with different power levels. A node only receives a transmitted message if the often dreaded SINR condition is fulfilled: a formula that relates the different power levels and distances from the receiver of simultaneously transmitting senders to the power level and distance from the receiver of the sender of interest. This condition ensures that the interference from other senders compared to the signal strength received from the sender of interest is small enough to guarantee error-free reception. Magnús Halldórsson, Stephan Holzer and Evangelia Anna Markatou received the best paper award for their work on the leader election problem in the SINR model. They show that if network nodes can select arbitrarily large transmission powers (at least exponential in the number of network nodes), then the leader election problem can be solved in two rounds. The presentation was given by Magnús Halldórsson, who made the audience illustrate the core of the problem (i.e., interference) by having members of the audience shout out their names simultaneously, which caused reception problems for the listening audience members. The authors' algorithm randomly splits the network nodes into two halves, makes one half transmit their IDs at randomly chosen power levels, and makes the other half listen. If power levels are chosen from a sufficiently large domain, all listeners can distinguish the network node that transmitted with the highest power. This node is chosen to be the leader of the network. In the second round, the listeners of the first round transmit the leader's ID, and transmitting nodes of the first round remain silent and learn the leader's name in this way. Jurdzinski, Różański and Stachowiak presented new results on the token traversal problem in the SINR model. They give a $O(n \log n)$ rounds deterministic algorithm, which is shown to be optimal and improves on previous $O(n \log^2 n)$ rounds algorithms. An interesting aspect of their work is that their result holds in weakly connected SINR networks, while most previous works consider strongly connected networks.

Four papers addressed problems related to the coordination of mobile agents. Disser et al. presented new results on the collaborative tree exploration problem: Initially, k deterministic agents are located at the same node in a tree. Their common goal is to explore the entire tree. To this end, in each round, they have to choose an incident edge and move along this edge to a different node. They can make use of communication to coordinate their strategies. In other words, the decision as to which edge an agent visits next may depend on the previous actions of all other agents. Their results include a new lower bound that renders the O(1)-competitive algorithm of Dereniowski et al., which works for the $k = n^{1+\epsilon}$ case, in a sense optimal. Brandt et al. study the problem of evacuating k agents that are initially located at the same point on one of m rays. These rays meet at a single junction point. Again, agents can communicate freely with each other. The objective is to evacuate the k agents through a single exit point. Various competitive algorithms are obtained, such as a 4-approximation if $k = m$, and a $2 + \sqrt{3}$-competitive algorithm for the case $k = m = 3$. Czyzowicz et al. study how to evacuate three robots, which can communicate

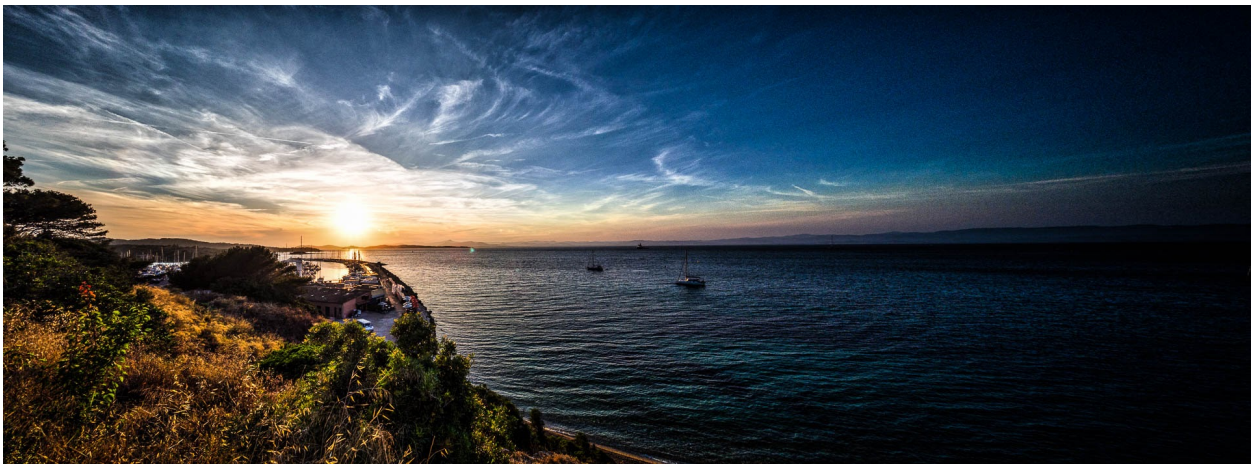

Figure 3: Porquerolles Island at sunset

through a wireless network, from a circle through an exit that is located at an unknown position on the perimeter of the circle. The difficulty stems from the fact that at most one of these robots is faulty. They give algorithms for two different fault models that show that competitive ratios of at most 7 can be achieved, while lower bounds rule out competitive ratios below 5. Di Luna et al. consider the problem of gathering agents that live on an anonymous dynamic ring in a single vertex. In each round, at most one edge of the ring is missing, which keeps the underlying graph structure constantly connected. They consider agents that cannot communicate with each other. Their paper studies the impact of chirality (agents have a common sense of orientation) and cross detection (when traversing an edge, agents can detect whether other agents are crossing the edge in the other direction at the same time) on the feasibility of the gathering problem.

Dynamic or evolving graphs are graphs that change over time such as the constantly connected dynamic ring mentioned above. These graphs have received substantial attention in recent years, and various models for evolving graphs exist today. One such model is characterized by the T-interval connectivity property, which states that, at any moment, there exists a spanning tree in the evolving graph that remains unchanged for at least T iterations. In a previous work, Casteigts et al. gave an algorithm that acts on a dynamic graph and determines the minimal parameter T for which the graph is T-interval connected. In the work that they presented at this conference, they reuse ideas from their previous algorithm and obtain new testing algorithms for other properties, such as the Round-Trip-Temporal-Diameter at a given moment t, i.e., the minimum time d it takes to commute at time t from any node u to any other node v and return back to u.

Feuilloley considered the $\mathcal{LOCAL}$ model of distributed computation and studied average case running times of distributed algorithms. It is known that leader election on a ring requires $\Omega(n)$ rounds (i.e., at least one node runs the algorithm for $\Omega(n)$ rounds). However, if we are interested in the average time required for the nodes to output their results instead of the worst case time, then an algorithm can be designed where the average node requires only $O(\log n)$ rounds. Halldórsson and Konrad studied the minimum vertex coloring problem on interval graphs in the $\mathcal{LOCAL}$ model. In a recent work, they showed that a constant factor approximation can be computed in $O(\log^* n)$ rounds. In their new work, they improved the approximation factor and gave a $(1+\epsilon)$-approximation

algorithm that runs in $O(\frac{1}{\epsilon} \log^* n)$ rounds. Linial's ring coloring lower bound shows that the dependency on $\log^* n$ is optimal. They further prove that the dependency on $\frac{1}{\epsilon}$ is optimal as well.

Three papers addressed labelling schemes. A labelling scheme is an assignment of labels to the vertices of the network graph so that network nodes can make use of these labels and solve a certain problem much faster (i.e., using fewer communication rounds) than without labels. The focus usually lies on the maximum length of a label. Gorain and Pelc studied the topology recognition problem in tree radio networks, where every node is required to learn an isomorphic copy of the input graph and identify its position within the copy. They show that there is a labelling scheme of length $O(\log \log \Delta)$ and a distributed algorithm that makes use of these labels and runs in $O(D\Delta)$ rounds, where D is the diameter of the input graph and Δ the maximum degree. They prove that this is essentially best possible. Ostrovsky, Perry and Rosenbaum addressed proof labelling schemes. In a proof labelling scheme, the distributed algorithm uses the initially assigned labels and usually runs in a constant number of rounds in order to verify a certain property of the input graph. In their paper, the authors give trade-offs between the length of the labels and the number of rounds required by the distributed verification algorithm. They particularly focus on the situation where the verification algorithm runs in $\omega(1)$ rounds. Censor-Hillel, Paz and Perry introduce a new notion for approximate proof labelling schemes, and show that label lengths can be much shorter than for ordinary proof labelling schemes.

The accepted papers covered a wide range of other topics (and I apologize to the authors of these papers for the brevity of my summaries). Spiegelman and Keidar make progress on a dynamic storage problem. Rodeh and Korman consider a parallel search problem that does not allow for coordination. It has been known since the eighties that symmetry breaking is only possible if network nodes are aware of at least "something" about the problem instance (e.g., non-symmetrical topology, random bits, etc.); Peva and Guerraoui focus on identifying the smallest "something" necessary to achieve this. Narayanan and Wu make progress on a problem in viral marketing. Bonnet et al. show how to kill some network nodes in order to avoid virus expansion. A group of twelve Masters students from the University of Paderborn supervised by Meyer auf der Heide give solutions to an interesting monitoring problem in distributed data streams. Biló et al. discuss a way to fix a single-source spanner by swapping edges. The objective of Gotfryd, Klonowski, and Pajak is to hide agents on a graph so that the mutual information between their current and initial positions is minimized. Rabie studies what can be computed in few rounds by probabilistic encounters of agents.

2 Social Activities

The very enjoyable technical programme of the conference was accompanied by great social activities. We celebrated Andrzej Pelc's 60th birthday between the last conference talk of the first day and the business meeting. Members of the audience had the opportunity to congratulate Andrzej, and many of Andrzej's colleagues that could not attend the conference had recorded their birthday wishes on video.

The technical programme on the second day ended before lunch, and we started the afternoon with a boat ride around the island. The views of the island were spectacular, and I refer the reader to the figures placed in this article.

After the boat ride, we were treated to delicious French cuisine and wine at the conference dinner banquet in a restaurant near the harbour. The restaurant provided an excellent atmosphere for

Figure 4: Another lovely view of Porquerolles Island

congratulating Shmuel Zaks on the SIROCCO 2017 award, Magnus Halldórsson and his co-authors on the best paper award, and Andrzej Pelc on his 60th birthday again.

3 Conclusion

SIROCCO 2017 was a fantastic event. The organizational team chaired by Shantanu Das (LIF, University of Aix-Marseille) did a great job and made themselves available throughout the conference. I particularly appreciated that the organizers scheduled transport (bus and boat ride) to and from the conference venue and sorted out accommodation for all participants, which simplified the travel arrangements. Next year's edition of the SIROCCO conference will take place in Israel and will be organized by Boaz Patt-Shamir.

SPAA 2017 Review

Sepehr Assadi
Department of Computer and Information Sciences
University of Pennsylvania
Philadephia, PA, USA
sassadi@cis.upenn.edu

The 29th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA 2017) was held on July 24 - 26, 2017 in Washington D.C., USA. This year, SPAA was co-located with PODC and took place at the Marvin Center at George Washington University. It included three keynote lectures by Piotr Indyk on "Beyond P vs. NP: Quadratic-Time Hardness For Big Data Problems", Guy Blelloch on "Some Sequential Algorithms are Almost Always Parallel", and Maurice Herlihy on "Blockchains and the Future of Distributed Computing" (the latter two were joint with PODC).

The papers "Distributed Partial Clustering" by Sudipto Guha, Yi Li, and Qin Zhang, and "Randomized Composable Coresets for Matching and Vertex Cover" by Sanjeev Khanna and myself won the best paper awards at this SPAA.

Unfortunately this review does not contain all of the papers presented in the conference, but a sample of the talks clearly biased by my personal experience at the conference.

1 First Day

The first day started with a keynote talk by Piotr Indyk on "Beyond P vs. NP: Quadratic-Time Hardness For Big Data Problems" [11]. Piotr talked about the relatively recent area of "hardness in P" in which the goal is to prove lower bounds on the running time of problems solvable in polynomial time (as opposed to NP-hard problems) using popular complexity theoretic conjectures such as Strong Exponential Time Hypothesis (SETH). As obtaining "fast" polynomial time algorithms for some problems, such as edit distance computation or Support Vector Machines have eluded researchers for decades, providing any evidence on "hardness" of these problems is quite interesting on its own and can also lead to the development of new algorithms, say for special cases of these problems. In his talk, Piotr discussed different quadratic-time lower bounds proven in recent years for popular problems such as edit distance computation, regular expression matching, Support Vector Machines, and gradient computation in neural networks.

The first talk of the day was one of the winners of best paper awards presented by myself. I talked about a new approach, namely randomized composable coresets, for solving graph optimization problems on different models of computation on massive graphs such as distributed computing, streaming, and Map-Reduce computation. I then showed an application of this method in solving two prominent graph-theoretic problems, namely the maximum matching and the minimum vertex cover problems [2]. This technique bypasses the strong impossibility results known for these problem using previous approaches such as linear sketching and (deterministic) composable coresets [3].

The morning session continued with a talk by Hossein Esfandiari on streaming algorithms for maximum coverage and set cover problems. These problems have been studied extensively in the streaming literature in the last couple of years. Hossein talked about an improved method of sketching coverage functions that allows for achieving smaller approximation factors for these problems compared to previous work: for example, to improve the previously best $O(p \cdot \log n)$-approximation algorithm for set cover in p passes over the stream by [9] to $O(\log n)$-approximation. Interestingly, a similar sketching approach has been concurrently discovered in [12] as well. Next, Morteza Zadimoghaddam talked about bi-criteria distributed submodular maximization in distributed computing models such as Map-Reduce using the composable coreset approach, and showed how to achieve improved performance by applying this technique over multiple rounds. The session was concluded by the talk of Susanne Albers on designing algorithms for energy conservation in data centers.

The afternoon session started with Hamidreza Jahanjou presenting the first constant-factor approximation algorithm for coflow scheduling. Next, Runtian Ren talked about a lower bound for deterministic online algorithms for flexible job scheduling for minimizing the makespan. Danny Vainstein presented both an improved upper bound and a tight lower bound for the clairvoyant dynamic bin packing problem. This is a variant of the classical bin packing problem in which the items are arriving and departing in an online manner and their departure times are known to the algorithm upon arrival; the goal is to minimize the total accumulated time of all "open" bins.

In the evening, Jason Li talked about his work on solving constraint satisfaction problems in parallel in which he and Ryan O'Donnell studied a variant of the classical multi-prover interactive proof systems where each prover can only communicate 1 bit to the verifier. Elaye Karstadt talked about reducing the leading constant in the O-notation of the $O(n^{\log_2 7})$ time algorithm of Strassen-Winograd for matrix multiplication and bypassing previous lower bounds for this problem by a change of basis. The last talk of the day was a brief announcement by Michael Goodrich on using cuckoo filters for faster set intersection queries and sparse boolean matrix multiplication.

The first day ended with the business meeting in which a heated discussion broke out on whether the process of submission and reviewing papers in SPAA conference (and in fact all TCS conferences) should or should not become double blind.

2 Second Day

The second day at SPAA started with a keynote talk by Guy Blelloch on "Some Sequential Algorithms are Almost Always Parallel" [4]. Guy talked about a meta-theorem stating that many sequential looking algorithms when given their input in a random order become highly parallel in that they have shallow dependency chains of poly-logarithmic length with high probability. He then talked about various examples of this meta-theorem among fundamental algorithms: the Knuth shuffle for random permutations, sorting by insertion into a binary search tree, greedy maximal

independent set (MIS), and more.

The first talk of the day was the other winner of the best paper awards presented by Qin Zhang. Qin talked about partial clustering problems, i.e., clustering with outliers, in two rounds of the distributed coordinator model. He presented their approach in [8] for detecting the number of outliers in each machine's input in one round of communication which allows the algorithm to allocate resources efficiently to the machines in the next round that in turn leads to a lower overall communication cost. This results in an improvement of a factor proportional to the number of machines in the total communication needed for solving these problems over previous work.

Dennis Olivetti talked about a distributed model of graph property testing and in particular for testing cycle-freeness of graphs. Property testing is an active area of research in the sub-linear time algorithms literature and has recently been studied in the streaming model [10] and other distributed communication settings [5, 7]. Next, Luca Zanetti talked about an interesting connection between classical results on load balancing and a graph clustering problem in distributed settings. The last talk of the morning session was on scheduling algorithms for distributed transactional memories.

After the excellent lunch break, two more sessions were held. The afternoon session of the day was mainly focused on contention resolution and congestion control in distributed systems and the evening session was mostly about external memory data structures. Unfortunately, I attended the concurrent PODC conference sessions instead and missed the talks given in these session. Hence, I am not able to talk about the outstanding results presented in these two sessions.

The second day ended with a banquet held jointly between SPAA and PODC in which the awards for best papers and best student papers in PODC conference were also given to the winners.

3 Third Day

The third keynote talk of the conference was given by Maurice Herlihy on "Blockchains and the future of distributed computing" in which he talked about blockchain-based distributed ledger systems such as Bitcoin. Maurice pointed out that while much of previous work on these topics originated outside the distributed computing community, the questions raised, such as consensus, replication, fault-tolerance, privacy, and security, and so on, are all issues familiar to this community. His talk surveyed the theory and practice of blockchain-based distributed systems from the point of view of classical distributed computing and pointed to promising future research directions.

Shahbaz Khan talked about a parallel algorithm for maintaining dynamic DFS trees in undirected graphs. In particular, Shahbaz presented a deterministic EREW PRAM algorithm for maintaining a DFS tree of an undirected graph which requires $O(m)$ processors and only polylog(n) update time per each insertion or deletion of an edge in the graph (m and n are, respectively, the number of edges and vertices in the input graph). The previous best deterministic parallel algorithm for this task requires $\widetilde{O}(\sqrt{n})$ update time [1]. Another interesting talk in this session was on the mobile server problem [6] which aims to capture the current trends to move computational tasks from cloud structures to multiple devices close to the end user.

The afternoon session was the final session of SPAA and contained three talks and four brief announcements. Among many interesting talks was a talk on online tree caching which initiated the study of a natural and practically relevant new variant of online caching where the to-be-cached items can have dependencies. Rezaul Chowdhury also talked about wavefront algorithms for evaluating dynamic programming recurrences.

References

[1] A. Aggarwal, R. J. Anderson, and M. Kao. Parallel depth-first search in general directed graphs. *SIAM J. Comput.*, 19(2):397–409, 1990.

[2] S. Assadi and S. Khanna. Randomized composable coresets for matching and vertex cover. In *Proceedings of the 29th ACM Symposium on Parallelism in Algorithms and Architectures, SPAA 2017, Washington DC, USA, July 24-26, 2017*, pages 3–12, 2017.

[3] S. Assadi, S. Khanna, Y. Li, and G. Yaroslavtsev. Maximum matchings in dynamic graph streams and the simultaneous communication model. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10-12, 2016*, pages 1345–1364, 2016.

[4] G. E. Blelloch. Some sequential algorithms are almost always parallel. In *Proceedings of the 29th ACM Symposium on Parallelism in Algorithms and Architectures, SPAA 2017, Washington DC, USA, July 24-26, 2017*, page 141, 2017.

[5] K. Censor-Hillel, E. Fischer, G. Schwartzman, and Y. Vasudev. Fast distributed algorithms for testing graph properties. In *Distributed Computing - 30th International Symposium, DISC 2016, Paris, France, September 27-29, 2016. Proceedings*, pages 43–56, 2016.

[6] B. Feldkord and F. Meyer auf der Heide. The mobile server problem. In *Proceedings of the 29th ACM Symposium on Parallelism in Algorithms and Architectures, SPAA 2017, Washington DC, USA, July 24-26, 2017*, pages 313–319, 2017.

[7] O. Fischer, S. Gershtein, and R. Oshman. On the multiparty communication complexity of testing triangle-freeness. In *Proceedings of the ACM Symposium on Principles of Distributed Computing, PODC 2017, Washington, DC, USA, July 25-27, 2017*, pages 111–120, 2017.

[8] S. Guha, Y. Li, and Q. Zhang. Distributed partial clustering. In *Proceedings of the 29th ACM Symposium on Parallelism in Algorithms and Architectures, SPAA 2017, Washington DC, USA, July 24-26, 2017*, pages 143–152, 2017.

[9] S. Har-Peled, P. Indyk, S. Mahabadi, and A. Vakilian. Towards tight bounds for the streaming set cover problem. In *Proceedings of the 35th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems, PODS 2016, San Francisco, CA, USA, June 26 - July 01, 2016*, pages 371–383, 2016.

[10] Z. Huang and P. Peng. Dynamic graph stream algorithms in o(n) space. In *43rd International Colloquium on Automata, Languages, and Programming, ICALP 2016, July 11-15, 2016, Rome, Italy*, pages 18:1–18:16, 2016.

[11] P. Indyk. Beyond P vs. NP: quadratic-time hardness for big data problems. In *Proceedings of the 29th ACM Symposium on Parallelism in Algorithms and Architectures, SPAA 2017, Washington DC, USA, July 24-26, 2017*, page 1, 2017.

[12] A. McGregor and H. T. Vu. Better streaming algorithms for the maximum coverage problem. In *20th International Conference on Database Theory, ICDT 2017, March 21-24, 2017, Venice, Italy*, pages 22:1–22:18, 2017.

PODC 2017 Review

Peter Davies
Department of Computer Science
Centre for Discrete Mathematics and its Applications (DIMAP)
University of Warwick
Coventry, UK
P.W.Davies@warwick.ac.uk

The 36th ACM Symposium on Principles of Distributed Computing (PODC 2017) was held in Washington D.C. from July 25-27 . The conference venue was the Cloyd Heck Marvin Center, at George Washington University. Alongside PODC, SPAA 2017 was held at the same venue, which was a good opportunity for integration and collaboration between the two closely related communities. A celebration was held on the evening of the 24th for David Peleg's 60th birthday, in which close colleagues gave varied perspectives upon his life and work, followed by the opening reception.

I will now give an overview of a (somewhat arbitrary) selection of the talks; apologies to the equally worthy speakers I do not mention!

Day 1

The program began with a keynote talk by Guy Blelloch entitled "Some Sequential Algorithms are Almost Always Parallel", a fascinating exploration of the potential for parallelism in seemingly very sequential tasks and algorithms. Guy described a method of quantifying the dependencies within algorithms, and applied this concept to a host of simple algorithms for classical tasks, such as sorting, maximal independent set, and graph coloring. His subsequent analysis showed that, on random inputs, these algorithms usually had only a polylogarithmic dependence depth, which translated to simple and efficient parallel algorithms.

The first session of regular talks concerned wireless network communications. I presented my work with Artur Czumaj on exploiting spontaneous transmissions in radio networks, which was awarded best student paper. Building on ideas by Haeupler and Wajc from PODC last year, we showed that allowing communication before reception of a source message admits fast algorithms for both broadcasting and the usually harder task of leader election.

Talks on related models followed: one on communication in cognitive radio networks, a popular model in the practical community which allows devices to choose one of many possible frequencies in which to communicate, and another on noisy radio networks, showing how to adapt the known-topology broadcasting algorithm of Gasieniec et al. (2007) to make it robust against random edge failures. Calvin Newport introduced a new model aiming to modernize the telephone model, and capture the peer-to-peer capability of mobile phones. Shared memory was the theme of the second session, beginning with Naama Ben-David's talk. She described a model which aimed to isolate delays caused by contention, and then gave an adaptive algorithm for memory access which performs better than the standard exponential backoff in this model. This was followed by work on fault-tolerant storage, with a talk about using erasure codes in distributed storage systems, and another about emulating fault-tolerant registers on fault-prone ones. On the subject of emulating registers, the sequence of brief announcements which ended the session included Eric Ruppert's interesting observation that in any wait-free implementation of unbounded registers atop bounded ones, the simulated read operations must include writes to the underlying bounded registers.

The final session of the day concerned the locality and communication requirements of graph problems. The first talk showed a complexity trichotomy for LCL problems (those whose solutions are checkable within constant-radius neighborhoods) in the LOCAL model on 2-dimensional grids, and showed how algorithm design could be automated for the most interesting class of such problems: those requiring $\Theta(\log^\star n)$ time. Work on the related CONGEST model was presented by Hsin-Hao Su, who showed an algorithm for MST whose running time is parameterized by the mixing time of the graph, beating the general graph lower bound for some graph families. Mohsen Ghaffari presented (remotely, since he was unfortunately unable to attend) another algorithm for a classical graph problem in a distributed setting, demonstrating that maximal matching can be solved faster in the CONGESTED-CLIQUE model (where nodes are not restricted to communicating with only their neighbors) than in CONGEST.

Day 2

Maurice Herlihy began the second day's proceedings with his brief, and colorful, history of blockchains, and made the case for attempting to bring them into the distributed computing fold. His contention was that, perhaps despite appearances, many of the issues now facing blockchain-based systems would prove familiar to distributed computing researchers.

After the break, the regular talks picked up where they had left off the previous day: with algorithms for distributed graph problems. Winner of the best paper award, Michael Elkin presented his remarkably simple 'best of both worlds' algorithm for deterministic MST in the CONGEST model, achieving near-optimal message and time complexity. Based on some clever adaptations of previous methods, in particular the randomized algorithm of Pandurangan et al., the result closed the long-standing open problem of whether good message and time complexity could be achieved simultaneously (and deterministically).

Talks in the session after lunch mostly concerned concurrent data structures, and included Wojciech Golab and Danny Hendler's paper on a sub-logarithmic algorithm for recoverable mutual exclusion. This task is an extension of the standard mutual exclusion problem (a synchronization problem of ensuring that at most one process at a time is executing the 'critical section' of its instructions) with the twist that processes are allowed to recover from crashes and resume executing at a later time.

A variety of topics were covered in the last session of the day, including Seth Gilbert's exploration of 'computational noise', the intriguing concept of modelling noise within a computational entity itself, rather than the environment in which it operates. He gave a model motivated by biological applications, and some algorithmic results for typical distributed computing tasks. Paul Gazzillo closed out the day with a throwback to Maurice Herlihy's keynote that morning: he demonstrated how smart contracts in cryptocurrency systems can be executed concurrently, providing significant performance improvements. This gave a nice example of the potential of Maurice's call to unite distributed computing concepts with blockchain technology.

Day 3

Thursday's keynote talk was given by Rosario Gennaro, who addressed the question of how one can trust computation which has been outsourced to the cloud. He gave an overview of the effort to revisit the classical complexity-theoretic results on interactive proofs and similar paradigms, and adapt and apply them to practical cloud-computing systems.

Talks after the break included Coordination Without Prior Agreement, an examination of what can be done by processes using shared memory, if they have no prior agreement on labels for the memory locations. The author Gadi Taubenfeld presented a nice selection of algorithms, lower bounds and impossibility results for the classical problems of consensus and mutual exclusion in the setting, challenging the often-overlooked assumption of common address labels.

Among the talks in the third session of the day was Ralph Keusch's exploration of the 'small-world' phenomenon, the well-known observation that in real world networks, two agents can usually reach each other in a surprisingly small number of hops. He discussed the recently introduced model of Geometric Inhomogeneous Random Graphs (GIRGs), which aim to capture the properties of such networks, and showed that greedy routing on these graphs (i.e. forwarding to the neighbour maximizing a simple objective function) is indeed surprisingly effective, succeeding with constant probability and following an almost-shortest path.

The final session of PODC was on population protocols and dynamics, models to study the computational power of systems of individually weak devices. Particularly exotic was James Aspnes's result that providing a clock that ticks on limit ordinals (to abstract any mechanism which conveys to the population that enough time has passed for it to converge or loop) grants population protocols power equivalent to the NL complexity class. This was a very clean and general addition to the recent line of work on strengthening population protocols.

Conclusion

This was my first year attending PODC, and I found it a very welcoming and engaging conference. I'd like to thank the organizers, and of course all of the speakers, for their work.

DISC 2017 Review

Manuela Fischer
ETH Zurich
manuela.fischer@inf.ethz.ch

Yannic Maus[1]
University of Freiburg
yannic.maus@cs.uni-freiburg.de

The *31st International Symposium on Distributed Computing, DISC 2017*, was held on October 16-20, 2017, in Vienna, Austria. The main conference took place at the *Austria Trend Hotel Park Royal Palace* close to the Schönbrunn Castle. Many of the attendees claimed that they had not been to Vienna before, and, besides the conference events, also enjoyed visiting the city. However, attendance at all talks and the six (!) colocated workshops was very high. Many will remember this conference also for the great sweets and cakes that were served; there was probably no single person who did not go through some kind of a sugar shock.

This review contains a few selected events and talks which constitute our personal highlights of the conference.

Welcome Reception. The conference itself began on Monday evening with a reception in the *Vienna Technical Museum*. Besides strolling around through the exhibition, the attendees in particular enjoyed the unexpectedly interesting high voltage show.

Business Meeting. In the business meeting on Tuesday evening it was pointed out that the quality of submitted papers (37 accepted) was very high this year, and also the number of submissions was an all time high for DISC (171 reviewed papers versus, e.g., 130 in the year before).

Birthday Celebration for Yoram Moses

One of the highlights was the celebration of the 60th birthday of Yoram Moses on Tuesday afternoon. Besides the gigantic cake the crowd enjoyed the talks by Joe Halpern (Cornell University, USA), Moshe Tennenholtz (Technion, Israel, Yoram's first Ph.D. student), Moshe Vardi (Rice University, USA), Shafi Goldwasser (MIT, USA), and Sergio Rajsbaum (UNAM, Mexico), all circling around the topic of *reasoning about knowledge* (Yoram is famous for his contributions in this area). Sergio

[1]Supported by ERC Grant No. 336495 (ACDC).

convinced us that indistinguishability is everywhere in distributed computing. He had recently picked up the term from Yoram's talk at *BIRS, Oaxaca* where Yoram advocated that the term *knowledge* should be replaced by *indistinguishability*.

Conference Banquet and Awards

On Wednesday evening, the conference banquet took place at the beautiful Schönbrunn Castle. After an aperitif, a delicious 3 course menu was served. This festive occasion was also used for the award ceremony.

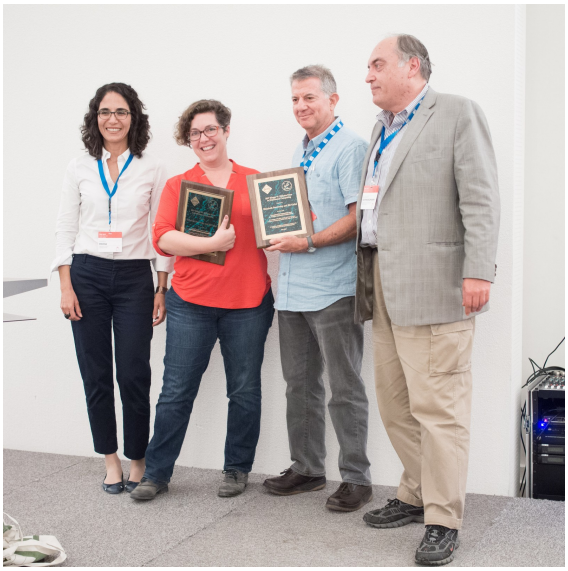

Figure 1: Andréa Richa, Elizabeth Borowsky, Eli Gafni, and Alex Schwarzmann.

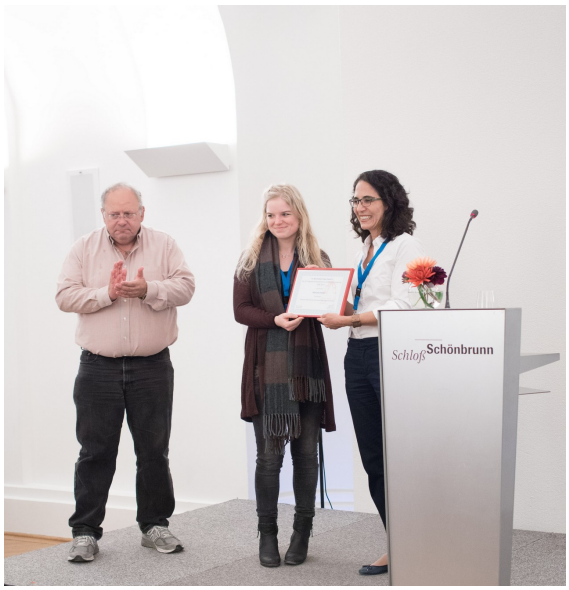

Figure 2: Shlomi Dolev, Manuela Fischer, and Andréa Richa.

The 2017 **Edsger W. Dijkstra Prize in Distributed Computing** was awarded to Elizabeth Borowsky and Eli Gafni for their 1993 work *Generalized FLP Impossibility Result for t-Resilient Asynchronous Computations*, containing several important results about *consensus*. Their paper has had—with hundreds of citations—and still has an enormous impact on the field of distributed computing. During his speech, Eli told the anecdote how they missed out on the Goedel Prize, for lack of a journal version, and that most of his publications still do not have one.

The **Best Paper Award** went to *Improved Distributed Degree Splitting and Edge Coloring* by Mohsen Ghaffari, Juho Hirvonen, Fabian Kuhn, Yannic Maus, Jukka Suomela, and Jara Uitto.

The **Best Student Paper Award** was given to *Improved Deterministic Distributed Matching via Rounding* by Manuela Fischer.

Congratulations to the award winners!

Selected Talks

Best Paper: Improved Distributed Degree Splitting and Edge Coloring
Mohsen Ghaffari, Juho Hirvonen, Fabian Kuhn, Yannic Maus, Jukka Suomela, Jara Uitto
Yannic explained that studying *degree splitting*—which is a coloring of the edges with red and blue such that every node has almost the same number of incident edges in each color—is useful in the distributed setting, and particularly important for the LOCAL model. He presented their newly devised algorithm for degree splitting, which is simpler, faster and achieves smaller discrepancy, that is, a smaller maximum difference in the number of red and blue incident edges of a node. This, among others, gives rise to a faster deterministic algorithm for $(2+o(1))\Delta$-edge-coloring, improving on that of Ghaffari and Su [SODA'17].

Best Student Paper: Improved Deterministic Distributed Matching via Rounding
Manuela Fischer
Manuela showed how a deterministic distributed rounding method can be used to devise simpler and faster algorithms for a number of well-studied variants of the matching problem in the LOCAL model. In particular, this rounding-based approach leads to an $O(\log^2 \Delta \cdot \log n)$-round deterministic algorithm for Maximal Matching, the first improvement in about 20 years over the celebrated $O(\log^4 n)$-algorithm by Hańćkowiak, Karoński, and Panconesi [SODA'98, PODC'99].

Self-stabilizing Byzantine Clock Synchronization is Almost as Easy as Consensus
Christoph Lenzen, Joel Rybicki
The talk was given by Joel. The main objective of the paper is to give fault-tolerant algorithms for establishing synchrony in distributed systems in which each of the nodes has its own clock; the protocol tolerates up to $f < n/3$ Byzantine nodes. In the deterministic setting the paper achieves a stabilisation time of $\Theta(f)$ (this has been achieved before) with $\Theta(\log f)$ bits (an exponential improvement on previous work). In the randomized setting they obtain an algorithm that stabilizes in time $O(\text{poly}\log f)$ while each node broadcasts $O(\text{poly}\log f)$ bits. The title stems from the fact that the authors reduce the stabilization time (with polylog overhead only) to non-self-stabilising binary consensus in the synchronous model. A reduction in the other direction as well as any lower bounds for the problem are not known.

Some Lower Bounds in Dynamic Networks with Oblivious Adversaries
Irvan Jahja, Haifeng Yu, Yuda Zhao
The main result of the paper (for several problems such as binary consensus, leader election, aggregation problems) is to show a $\Omega(\text{Diameter} + \text{poly}(n))$ lower bound in the CONGEST model with an oblivious adversary. Previously such bounds were only known for adaptive adversaries.

In the talk Irvan pointed out two convincing novelties in their reduction to a variant of the two-party communication problem *gap disjointness* which allows to prove a lower bound for oblivious adversaries. The first novelty is the concept of a leaker, i.e., a third party that randomly reveals (unimportant) parts of the input to both parties. The second introduced novelty is the concept of a sanitized adversary, i.e., a wisely constructed adversary that is adaptive but turns out to not be more powerful than an oblivious adversary.

Dynamic Analysis of the Arrow Distributed Directory Protocol in General Networks
Abdolhamid Ghodselahi, Fabian Kuhn
The presentation was given by Abdholhamid. In the model there is a single shared object that can be requested by nodes of a given network at any time. The well-known and heavily studied *Arrow*

protocol is a distributed queuing protocol for tree networks, i.e., it orders the requests. The cost for a request is the time until it finds out who its predecessor in the order is.

The main result of the paper shows that *Arrow* is $O(1)$-competitive on so called hierarchically ordered trees (HSTs). One way to transform the Arrow protocol into a protocol for general networks is to embed the network graph into a tree. Using the distributed algorithm by Ghaffari and Lenzen [DISC'14] that embeds any network graph into an HST with expected $O(\log n)$ stretch, one can show that *Arrow* is $O(\log n)$-competitive (expected) for general networks. These results improve significantly on previous results and show that *Arrow* even beats the competitiveness of other protocols such as *Spiral*. Furthermore, the analysis also works for the asynchronous setting.

Three Notes on Distributed Property Testing
Guy Even, Orr Fischer, Pierre Fraigniaud, Tzlil Gonen, Reut Levi, Moti Medina, Pedro Montealegre, Dennis Olivetti, Rotem Oshman, Ivan Rapaport, Ioan Todinca
This was probably the first time that DISC had a paper consisting of three merged papers now having eleven authors! Three of the authors presented different results for distributed property-testing in the CONGEST model, with an emphasis on testing subgraph-freeness.

Quadratic and Near-Quadratic Lower Bounds for the CONGEST Model
Keren Censor-Hillel, Seri Khoury, Ami Paz
Seri presented the first super-linear lower bounds for natural graph problems such as minimum vertex cover, maximum independent set, and even some problems in P in the CONGEST model, answering long-standing open questions. The paper also gives a linear lower bound for the problem of weighted all-pairs-shortest-paths (APSP). All these lower bounds follow by reducing a 2-party communication problem to the corresponding problem in the CONGEST model. The authors show that for the APSP problem, this technique cannot lead to any stronger lower bound, and hence completely new ideas would be required to further improve to a superlinear lower bound.

Sublogarithmic Distributed Algorithms for Lovász Local Lemma, and the Complexity Hierarchy
Manuela Fischer, Mohsen Ghaffari
Manuela presented an $2^{O(\sqrt{\log \log n})}$-round algorithm for the Lovász Local Lemma (LLL) in the LOCAL model, improving over the $O(\log n)$-algorithm by Chung, Pettie, and Su [PODC'14], and making the first step of progress towards the conjectured $O(\log \log n)$-complexity by Chang and Pettie [STOC'17]. Combining this algorithm with the enlightening revelation by Chang and Pettie that any algorithm for a Locally-Checkable Labeling (LCL) problem (on bounded degree graphs) with round complexity $o(\log n)$ can be automatically sped up to the round complexity of LLL, this shows that not only LLL but in fact any sublogarithmic LCL problem (on bounded degree graphs) can be solved in $2^{O(\sqrt{\log \log n})}$ rounds.

Derandomizing Local Distributed Algorithms under Bandwidth Restrictions
Keren Censor-Hillel, Merav Parter, Gregory Schwartzman
Merav presented general tools for derandomizing algorithms under bandwidth restrictions. More concretely, the authors show how to use the technique of conditional expectations (and other methods) to derandomize algorithms that use random variables with bounded dependency. This leads to improved deterministic algorithms for many problems, such as Maximal Independent Set in the CONGEST and the Congested Clique model. This sheds some light on the curious gap between randomized and deterministic algorithms under bandwidth restrictions.

Keynote Talks

Each day of the conference started with a one-hour keynote talk by invited speakers.

From the Lab to the Wild - Scalable Recommenders
Anne-Marie Kermarrec (INRIA Rennes)
Tuesday started with a keynote by Anne-Marie about scalable recommenders as we all know them from Amazon recommending goods, Netflix recommending movies and so on. She talked about her experiences with her startup *Mediego*, which offers services such as personalized newsletters. Anne-Marie explained collaborative filtering (one of the most popular approaches for online recommenders), focusing on the high level algorithmic ideas, the difficulties that arise during implementation, and the complications with customers.

Blockchain Consensus Protocols in the Wild
Christian Cachin (IBM Research Zurich)
On Wednesday morning, Christian talked about blockchain consensus protocols—protocols to ensure agreement of all participating nodes on a unique transaction order—in the presence of adversarial nodes and faults. He reviewed known consensus protocols with respect to their resilience against attacks and explained how the established practice in cryptography and computer security can be applied to blockchains to arrive at an agreement.

Phase Transitions and Emergent Phenomena in Random Structures and Algorithms
Dana Randall (Georgia Tech)
In an insightful keynote on Thursday morning, Dana showed that many processes have so called *phase transitions*, i.e., the behavior of the process changes drastically when a parameter is changed above or below a problem dependent threshold. Dana presented many processes, e.g., from statistical physics or demographics, which do have phase transitions. She convinced the audience, which was curious about the topic, that studying these processes might be useful to gain knowledge about random structures.

Workshops

There were six workshops scheduled around the conference. All had full day programs consisting of several (in fact, up to 14) invited talks, aimed at the general DISC audience. On Monday, **FRIDA** (*Formal Reasoning in Distributed Algorithms*), **BTT** (*Workshop on Blockchain Technology and Theory*), and **TPC** (*Workshop on the Theory and Practice of Concurrency*) took place. On Friday there were the workshops **ADGA** (*Workshop on Advances in Distributed Graph Algorithms*), **CoDyn** (*Workshop on Computing in Dynamic Network*), and **HDT** (*Workshop on Hardware Design and Theory*).

A big thanks to the organizers and the speakers for the insightful workshops!

FRIDA had an interesting mix of well-established techniques, such as consensus lower bounds through the topological approach, and recent developments in distributed computing, such as property testing/distributed decision.

In the BTT workshop, which had the highest number of registrations (namely 58), it was pointed out by Ittai Abraham that the well-studied Byzantine Fault Tolerant protocols are not so different

from what happens in Blockchains. Sara Tucci-Piergiovanni sensitized the audience for studying the fairness for users instead of miners in Blockchains.

The TPC workshop enjoyed high popularity as well. Some people even claimed that this workshop was the only reason for them to attend the conference.

ADGA Workshop

For us, a special highlight was the ADGA workshop, whose purpose is to cover general aspects, main remaining objectives, and major technical obstacles of current topics in the research area of Distributed Graph Algorithms.

Moti Medina, in his talk titled *A (Centralized) Local Guide*, tried (and succeeded) to motivate the distributed algorithms community to study centralized local algorithms—called Local Computation Algorithms (LCA) or CentLocal Algorithms—where only a tiny part of the input can be read. He outlined connections between LCAs and distributed local computation, property testing, as well as sublinear approximation algorithms, and provided an overview of the state of the art and of the known techniques for these models.

In the talk *LOCAL Algorithms: The Chasm Between Deterministic and Randomized*, Mohsen Ghaffari addressed the wide gap between deterministic and randomized algorithms in the LOCAL model: While we know very simple $O(\log n)$-algorithms for many of the classic local graph problems, devising deterministic $\operatorname{poly}\log n$-round algorithms remains one of the central open questions of the area. Mohsen elaborated on the centrality of deterministic rounding methods, due to recent findings by Ghaffari, Kuhn, Maus [STOC'17]. Roughly speaking, rounding fractional to integral variables, while approximately preserving some linear constraints, is the only obstacle for obtaining efficient deterministic algorithms for all the classic graph problems. In the second part of the talk, Mohsen then explained concrete rounding methods for some special cases, leading to improved matching algorithms as well as the first efficient deterministic $(2\Delta - 1)$-edge-coloring algorithm, solving an almost 30 years old open problem.

Juho Hirvonen talked about *Proving Lower Bounds in the LOCAL Model*. He provided a great survey about core ideas and limitations of recent lower bound techniques, focusing on the two approaches *algorithmic simulation* of lower bounds and the *lifting* of lower bounds from weaker to stronger models, and elaborated on the role of unique node identifiers for lower bounds. In the end of his talk, Juho mentioned some central open questions, one of them being an $\Omega(\Delta)$ lower bound for Bipartite Maximal Matching.

Further, and by no means less interesting, talks were given by Andreas Karrenbauer about *Continuous Optimization Methods for Network Problems*, Magnus Halldorsson about *Graph Abstractions in Wireless Networking*, and Adrian Kosowski about *Distance Labeling: Understanding the Source of Hardness*.