Distributed Computing Column 64
Annual Review 2016

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 Noga Alon, László Babai, Alon Itai, and Michael Luby who shared the 2016 Dijkstra Prize for their papers on maximal independent set algorithms! Noga, László, and Alon's paper is "A Fast and Simple Randomized Parallel Algorithm for the Maximal Independent Set Problem" and Michael's paper is "A Simple Parallel Algorithm for the Maximal Independent Set Problem". 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 PODC. The full citation can be found at http://www.podc.org/dijkstra/2016-dijkstra-prize/; highlights are:

> "In these seminal works, the authors present, simultaneously and independently, an $O(\log n)$ time randomized distributed/parallel algorithm for the Maximal Independent Set (MIS) problem. MIS is regarded as a crown jewel of distributed symmetry breaking problems, and a central problem in the area of locality in distributed computing. The nominated papers provide a fascinatingly simple, elegant, and efficient randomized solution for this problem."

Congratulations as well to Hsin-Hao Su and Shahar Timnat, who split the Principles of Distributed Computing Doctoral Dissertation Award! Hsin-Hao's dissertation was entitled "Algorithms for Fundamental Problems in Computer Networks", supervised by Seth Pettie at the University of Michigan, Ann Arbor. Shahar's dissertation was entitled "Practical Parallel Data Structures", supervised by Erez Petrank at the Technion. The award is jointly sponsored by PODC and

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

DISC, and was given at DISC this year. The citation appears at
`http://www.podc.org/dissertation/2016-dissertation-award/`; it highlights Hsin-Hao's new and improved algorithms for graph coloring and Shahar's work to make lock-free concurrent data structures wait-free.

The first article of the column is a review of PODC 2016 by Gregory Schwartzman, who won the Best Student Paper award for his paper "A Distributed $(2+\epsilon)$-approximation for Vertex Cover in $O(\log \Delta/\epsilon \log\log \Delta)$ Rounds" with co-authors Reuven Bar-Yehuda and Keren Censor-Hillel. The Best Paper Award was given for the paper "Analysing Snapshot Isolation" by Andrea Cerone and Alexey Gotsman. Congratulations to Gregory, Reuven, Keren, Andrea and Alexey!

The second article is a review of DISC 2016 by Lili Su. Seri Khoury won the Best Student Paper award for the paper "Near-Linear Lower Bounds for Distributed Distance Computations, Even in Sparse Networks," co-authored with Amir Abboud and Keren Censor-Hillel. The Best Paper award went to "Polynomial Lower Bound for Distributed Graph Coloring in a Weak LOCAL Model" by Dan Hefetz, Fabian Kuhn, Yannic Maus and Angelika Steger. Congratulations to Seri, Amir, Keren, Dan, Fabian, Yannic, and Angelika!!

Many thanks to Gregory and Lili 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.

PODC 2016 Review

Gregory Schwarzman
Department of Computer Science
Technion, Israel
gregorys@cs.technion.ac.il

The 35th ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing (PODC 2016) was held on July 25-29, 2016 in Chicago, IL, USA. The conference took place at Loyola University in downtown Chicago. It included three keynote talks by Andrew A. Chien, Faith Ellen and Phillip B. Gibbons. The paper "Analysing Snapshot Isolation" by Andrea Cerone and Alexey Gotsman won the best paper award. The papers "The Greedy Spanner is Existentially Optimal" by Arnold Filtser and Shay Solomon, and "A Distributed $(2+\epsilon)$-approximation for Vertex Cover in $O(\log\Delta/\epsilon\log\log\Delta)$ Rounds" by Reuven Bar-Yehuda, Keren Censor-Hillel and myself, won the best student paper award.

On the evening of the 25th there was a birthday celebration for Faith Ellen. And a banquet was held on the evening of the 27th, where the Dijkstra Prize was awarded jointly to Michael Luby for the paper "A simple parallel algorithm for the maximal independent set problem" and to Noga Alon, László Babai, and Alon Itai for the paper "A Fast and Simple Randomized Parallel Algorithm for the Maximal Independent Set Problem".

Unfortunately this review does not contain all of the papers presented in the conference, but only a few selected talks.

1 First day

The day started with a keynote lecture by Andrew A. Chien, titled "New Opportunities for PODC? Massive, Volatile, but Highly Predictable Resources". In the talk Andrew addressed the problem of "stranded power" — energy which is generated but has no economic value. Renewable energy sources (such as wind power) can generate massive amounts of stranded power because the amount of power generated is determined by the elements. To harness stranded power, Andrew proposes a new model for cloud computing called Zero-Carbon Cloud. Zero-Carbon Cloud provides massive volatile computational resources, meaning that the amount of computational power available may fluctuate drastically. But this fluctuation can be predicted based on the availability of stranded power. Andrew discusses various use cases for Zero-Carbon Cloud, for example, execution of non-interactive batch tasks.

The first two papers were the winners of the best student paper presented by myself and Shay Solomon. I presented a $(2+\epsilon)$-approximation for distributed vertex cover in $O(\log \Delta/\epsilon \log\log \Delta)$ rounds. The number of communication rounds is tight according to a lower bound by [6]. In the talk I introduced the local ratio technique and several variants which make it a useful tool in the distributed setting. Finally I presented a proof sketch for the special case of $\epsilon = 1/2$. Shay showed that the classic spanner algorithm by [1] is existentially optimal for any graph family that is closed under edge removal. This implies that the greedy construction is just as good as the much more involved construction in [3]. This result also extends to doubling metrics, matching the construction in [5]. Unfortunately his co-author Arnold Filtser could not present at the conference because he was expecting a baby. Congratulations to Arnold!

Mohsen Ghaffari presented an MST construction in the congested clique in $O(\log^* n)$ rounds with high probability. This improves upon the previous best known result of $O(\log\log\log n)$ communication rounds. The winner of the best paper award, Andrea Cerone, presented a method for snapshot isolation. Snapshot isolation is a consistency model used for transaction processing. Andrea presented an alternative specification for snapshot isolation which characterizes it using transactional dependency graphs. This new specification is then used for static analyses of the transactions. Leqi Zhu presented a tight space bound for n-process randomized wait-free and obstruction free consensus protocols. It was known that in the anonymous setting $\Omega(n)$ registers are required to solve the consensus problem. Leqi shows that even if the processors have identifiers, $n-1$ registers are needed.

2 Second day

The second day started with a keynote lecture by Faith Ellen discussing concurrent data structures. In this excellent talk Faith presented the challenges faced when constructing concurrent data structures and various techniques in the design of concurrent data structures. Faith also discussed the evaluation of concurrent data structures, stating that papers presenting an algorithmic result must contain a rigorous proof and not only a high-level description. Experimental results should be reproducible and should measure the right aspect of the data structure.

Yi-Jun Chang presented a brief announcement on an exponential separation between randomized and deterministic complexity in the LOCAL model. Yi-Jun shows that the problem of Δ-coloring of trees, where Δ is the maximal degree in the tree, requires $\Theta(\log_\Delta \log n)$ rounds using random bits and $\Theta(\log_\Delta n)$ deterministically. Another intriguing result in the paper is the connec-

tion between deterministic lower bounds and randomized lower bounds. A connection is shown in both directions: deterministic lower bounds imply randomized lower bounds and randomized lower bounds imply deterministic lower bounds.

Michael Elkin presented a strong network decomposition algorithm with parameters $(O(\log n), O(\log n))$. This expands upon the classical work of Linial and Saks [7] which provides a *weak* diameter network decomposition with the same parameters. Elad Haramaty presented a technique for maintaining an MIS in dynamic graphs. Elad shows that under a single adjustment it is possible to maintain an MIS with a constant expected update time. This is achieved by simulating the greedy sequential algorithm with random ordering of the nodes. Hagit Attiya addressed the problem of collaborative text editing. Although several protocols for collaborative text editing exist, there is no precise specification of their desired behavior. Hagit provides such a specification by defining a replicated list object. This object models the core functionality of collaborative text editing.

The second day ended with a banquet held at the Hotel Palomar. Except for the great food, the main event of the evening was the awarding of the Dijkstra prize. The prize was awarded jointly to Michael Luby for the paper "A simple parallel algorithm for the maximal independent set problem" and to Noga Alon, László Babai, and Alon Itai for the paper "A Fast and Simple Randomized Parallel Algorithm for the Maximal Independent Set Problem". It was a great honor and pleasure to see the faces behind one of the most famous algorithms in the field. The winners shared the story behind the discovery of the algorithm. Upon receiving the award László Babai offered to donate the monetary part of the award in order to allow more veterans of the field to attend the conference in the future.

3 Third day

The third day started with a keynote talk by Phillip B. Gibbons. In the talk Phillip suggests that in the next generation of memories, write operations will be an order of magnitude more expensive than read operations. Thus, we should plan ahead and start developing algorithms which take this into account. We would like to consider the complexity of write operations separately from that of read operations and optimize the running time accordingly.

Phillip illustrates this clearly with the example of sorting. Due to known lower bounds we cannot have comparison-based sorting algorithms running in time $o(n \log n)$. But we can, however, build a sorting algorithm which has $O(n \log n)$ read operations, but only $O(n)$ write operations. This is achieved simply by inserting all of the element into a self balancing tree and going over the tree in order after we have finished.

Danny Hendler discussed lower bounds for remote memory reference complexity of reader-writer locks in the cache-coherent read/write model. Danny complements the currently known bound of $\Omega(\log m)$, where m is the number of readers, with a matching $\Omega(\log n)$ bound, where n is the number of writers.

Petteri Kaski presented the paper "How Proofs are Prepared at Camelot" whose original description is just too good to not be included here directly: "Picture K Knights seated around the Round Table, distressed. At the center of the table stands the Input. The Knights have been tasked to prepare a Proof about the virtues of the input, and to make extreme haste: a proof must be prepared in time T/K, where T is the fastest time known in all of Britannia for any single soul to reveal such subtle virtues, let alone give uncontestable proof thereof. Nigh impossible is the task in fact, for the Lady Morgana has enchanted many a poor Knight with her cunning dark magic, yet

virtues must be revealed and proof thereof prepared. And not just any proof, but one that any lone soul can check, in time T/K and with a few tosses of a fair coin, and next to never be convinced if foul virtues it claims. However, not all is lost, for the Knights recall the teachings of the wizard Merlin and his powerful proofs that King Arthur so treasures. So the Knights agree to together evaluate in batch, and to individually decode-and-check for signs of the darkest of magic..."

Masahiro Shibata discussed the problem of uniform deployment of robots in a ring. This problem, in a sense, is the opposite of the well known rendezvous problem. Masahiro analyzes the problem under various assumptions on the knowledge of the agents and shows upper bounds on the running time and memory used per agent. If the agents have no knowledge and termination detection is required, then there exists no algorithm that solves the uniform deployment problem. Talal Riaz presented a brief announcement in which he shows that the MIS algorithm of [2] for MIS on a tree can be extended to bounded arboricity graphs with running time $O(poly(\alpha) \cdot \sqrt{\log n \log\log n})$, where α is the arboricity of the graph. Talal achieves this by using a new tail inequality by [4] for read-k families of random variables.

4 Conclusion

This is my first time attending a conference, and it was very enjoyable and enriching. I would like to thank the organizers and the great speakers at the conference.

References

[1] Ingo Althöfer, Gautam Das, David P. Dobkin, Deborah Joseph, and José Soares. On sparse spanners of weighted graphs. *Discrete & Computational Geometry*, 9:81–100, 1993.

[2] Leonid Barenboim, Michael Elkin, Seth Pettie, and Johannes Schneider. The locality of distributed symmetry breaking. In *53rd Annual IEEE Symposium on Foundations of Computer Science, FOCS 2012, New Brunswick, NJ, USA, October 20-23, 2012*, pages 321–330. IEEE Computer Society, 2012.

[3] Shiri Chechik and Christian Wulff-Nilsen. Near-optimal light spanners. In Robert Krauthgamer, editor, *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10-12, 2016*, pages 883–892. SIAM, 2016.

[4] Dmitry Gavinsky, Shachar Lovett, Michael E. Saks, and Srikanth Srinivasan. A tail bound for read-k families of functions. *Random Struct. Algorithms*, 47(1):99–108, 2015.

[5] Lee-Ad Gottlieb. A light metric spanner. In *IEEE 56th Annual Symposium on Foundations of Computer Science, FOCS 2015, Berkeley, CA, USA, 17-20 October, 2015*, pages 759–772, 2015.

[6] Fabian Kuhn, Thomas Moscibroda, and Roger Wattenhofer. Local computation: Lower and upper bounds. *J. ACM*, 63(2):17, 2016.

[7] Nathan Linial and Michael E. Saks. Decomposing graphs into regions of small diameter. In Alok Aggarwal, editor, *Proceedings of the Second Annual ACM/SIGACT-SIAM Symposium on Discrete Algorithms, 28-30 January 1991, San Francisco, California.*, pages 320–330. ACM/SIAM, 1991.

DISC 2016 Review

Lili Su
University of Illinois at Urbana-Champaign
Urbana, IL, USA
lilisu3@illinois.edu

The 30th International Symposium on Distributed Computing, DISC 2016, was held on September 26 – 30, 2016, in Paris, France. The main conference took place at the Universite Pierre & Marie Curie, in central Paris. Three workshops were co-located with DISC this year.

(1) Monday, 26 September 2016: The 5th Workshop on Advances in Distributed Graph Algorithms (ADGA), chaired by Danupon Nanongkai. The ADGA is a one-day workshop that focuses on the theoretical foundations of distributed graph algorithms. This year, six speakers were invited.

The workshop started with the outstanding talk given by Stephan Holzer, who talked about computing all pairs shortest paths and the problem of computing the diameter in the distributed setting. Optimality of the proposed algorithms in terms of communication complexity was given. Richard Peng, an assistant professor at GaTech, gave an excellent talk on parallel computing. In particular, Prof. Peng surveyed the recent progress of parallel algorithms for graph optimization problems, new algorithmic frameworks and some of their core routines. He also suggested that some of the assumptions made needed to be modified to better approximate certain concrete real-world applications. The four talks after Prof. Peng's talk were also extremely excellent, and collectively comprehensive for the distributed graph algorithms area. Topics included communication complexity, robustness, compact routing, symmetric breaking, etc.

(2) Monday, 26 September 2016: The 6th Workshop on Moving and Computing (MAC), chaired by Prof. Paola Flocchini and Prof. Maria Potop-Butucaru. The MAC workshop is an annual research meeting for people in distributed robot computing to share the cutting-edge techniques, results, ideas, and open problems. Seventeen researchers presented their findings as well as open problems in this workshop, with topics varying from autonomous robots formation in 3D space, constraint localization to biological collaborative systems.

(3) Friday, 30 September 2016: The First Workshop on Dynamic Graphs in Distributed Computing (DGDC). Internet of things brings new challenges to researchers and practitioners of distributed computing and distributed systems – the underlying topology is typically dynamic and changes unpredictably with time. This new workshop focuses on this trendy and promising direction. The program consisted of 8 talks, most of them were invited. Personally, I think the program is very

exciting, and the success of this workshop would motivate more and more researchers to take the possible dynamic properties of the systems into account, and solve the corresponding problems.

The main conference of DISC was held from September 27, 2016 to September 29, 2016. This year, the conference received 132 regular submissions and 13 brief announcements, of which 32 papers and 10 brief announcements were accepted. The resulting technical program covered a wide variety of topics related to distributed computing, including problems in message passing and shared memory systems, radio networks, distributed graph algorithms, fault tolerance, security, and machine learning.

1 First Day

The first day of the conference started with a keynote talk *Verification of Population Protocols* by Prof. Javier Esparza. Prof. Javier Esparza and his colleagues collaboratively solved the decidability of the verification problems which remained open until 2015. In this keynote talk, Prof. Javier Esparza presented the found results and discussed some new developments.

The first session focuses on *shared memory*, and was chaired by Prof. Panagiota Fatourou. The first talk was on *k–Abortable Objects: Progress under High Contention*, authored by Naama Ben-David, David Yu Cheng Chan, Vassos Hadzilacos and Sam Toueg. A new object, termed as "k–abortable object" was defined. This object is the first abortable object that guarantees progress under high contention. Then *Linearizability of Persistent Memory Objects under a Full-System-Crash Failure Model*, authored by Joseph Izraelevitz et al., was presented. This work provided a framework for crash-resilient data structures on a machine with persistent memory. A full-system failure model, in which all transient state is lost on a crash, was considered. This fault model considered is novel, and well captures some particular "real world" systems. Following this, Matthieu Perrin et al.'s work *On Composition and Implementation of Sequential Consistency* was discussed. Existing results say that to implement sequential consistency, linear time is only needed for one of the read or write operations. This paper extended this result to crash-prone asynchronous systems. In particular, the proposed algorithm builds a sequentially consistent shared snapshot memory on top of an asynchronous message-passing system where less than half of the processes may crash. They proved that (1) waiting is needed only when a process invokes a read/snapshot right after a write, and that (2) sequential consistency is composable in some cases commonly encountered. The last talk of this session was on *Upper Bounds for Boundless Tagging with Bounded Objects*, authored by Zahra Aghazadeh and Philipp Woelfel. Tagging is a fundamental technique used in the design of shared memory algorithms. A framework for tagging was proposed.

After the excellent lunch break, two more sessions were held.

The Distributed Graph Algorithm session was chaired by Christoph Lenzen. Four regular papers were presented. The first paper presented was *Fast Distributed Algorithms for Testing Graph Properties*, authored by Keren Censor-Hillel et al. A thorough study of distributed property testing was provided. Among other fundamental results, a logarithmic lower bound for testing bipartiteness and cycle-freeness, which holds even in the LOCAL model, was derived. This talk was followed by the second paper on testing with different focus – *Distributed Testing of Excluded Subgraphs*, authored by Pierre Fraigniaud, Ivan Rapaport, Ville Salo and Ioan Todinca. *Deterministic Leader Election in $O(D + \log n)$ Time with Messages of Size $O(1)$*, authored by Arnaud Casteigts, Yves Métivier, John Michael Robson and Akka Zemmari, presented a very interesting deterministic algorithm construction, which elects deterministically a leader in an arbitrary network. The proposed

algorithm has a bit round complexity of $O(D+\log n)$, where D is the diameter of the network. This bit round complexity was proved to be optimal. *Further Algebraic Algorithms in the Congested Clique Model and Applications to Graph-Theoretic Problems* was presented by Francois Le Gall, who is the single author of this paper. Additional algebraic techniques for designing algorithms in congested clique model were developed. Both deterministic and randomized algorithms were proposed for efficiently computing multiple independent instances of matrix products, computing the determinant, the rank and the inverse of a matrix, and solving systems of linear equations.

The session on "Information Spreading" contained three regular papers and three brief announcements. This session was chaired by Prof. Swan Dubois. The focuses of these papers were evacuation, routing, fault-tolerance, missing information, communication complexity, and maximal independent set. Due to space limit, I am not able to talk more details about these six outstanding works.

2 Second Day

The second day of the conference started with a keynote talk *Personal Information Management Systems and Knowledge Integration* by Prof. Serge Abiteboul. Personal data is constantly collected, either voluntarily or passively by various applications. Thus those personal data is more or less accessed and controlled by the commercial companies. And this not only threatens the privacy of users, and may even be illegal. Observing this, Prof. Abiteboul introduced the concept of Personal Information Management Systems (PIMS), which enables user to manage his/her personal information directly without having to revealing it to others. Research issues raised by PIMS, either new or that acquire a new favor in a PIMS context, were discussed.

The "Asynchrony" session, chaired by Prof. Panagiota Fatourou, started with the talk *Anonymity-Preserving Failure Detectors*, authored by Zohir Bouzid and Corentin Travers. The paper investigated the consensus problem in anonymous, failure-prone and asynchronous shared memory systems. A new class of failure detectors (that preserve anonymity) were proposed. Following this talk, *Asynchronous Computability Theorems for t-Resilient Systems*, authored by Vikram Saraph, Maurice Herlihy and Eli Gafni, was presented. The t-resilient asynchronous computability theorem stated characterized the tasks that have t-resilient protocols in a shared-memory model. This result generalized the prior (wait-free) asynchronous computability theorem of Herlihy and Shavit to a broader class of failure models. The priority mutual exclusion problem was revisited in the paper *Priority Mutual Exclusion: Specification and Algorithm*, authored by Chien-Chung Huang and Prasad Jayanti. This paper showed that despite extensive efforts that have been made on the problem of priority mutual exclusion, none of the existing algorithms meet the reasonable expectations we would have of an algorithm that claims to respect process priorities. Observing this, this paper suggested more specification of the priority mutual exclusion problem and designed an efficient algorithm accordingly. In the next paper, *Opacity vs TMS2: Expectations and Reality*, authored by Sandeep Hans et al., the theoretical and practical implications of relaxing the existing algorithms were investigated.

Prof. Cyril Gavoille chaired the "Local Model" session. Dan Hefetz, Fabian Kuhn, Yannic Maus and Angelika Steger received the Best Paper Award for the paper *Polynomial Lower Bound for Distributed Graph Coloring in a Weak LOCAL Model*. An $\Omega(\Delta^{\frac{1}{3}-\frac{\zeta}{3}})$ lower bound is given on the runtime of any deterministic $O(\Delta^{1+\zeta})$-graph coloring algorithm in a weak variant of the LOCAL model. In the weak LOCAL model, the nodes have no identifiers, but it is required that the com-

putation starts with an initial valid vertex coloring. *Near-Optimal Low-Congestion Shortcuts on Bounded Parameter Graphs*, authored by Bernhard Haeupler et al., generalized the results obtained by Ghaffari, Haeupler, Izumi and Zuzic [SODA'16, PODC'16]. The complexity of construction of purely additive spanners in the CONGEST model was studied in *Distributed Construction of Purely Additive Spanners*, authored by Keren Censor-Hillel et al. Three outstanding brief announcements were presented subsequently, with topics ranging from MDS approximation, local distributed verification, to MST sparsification.

This year, a special session "Best Student Paper Election" was held, wherein all three nominated papers' student authors presented their corresponding work. The paper "Non-Bayesian Learning in the Presence of Byzantine Agents" was presented by Lili Su (co-authored with Prof. Nitin H. Vaidya). The paper provided a solution for the Bayesian learning problem using the computation and memory light non-Bayesian learning rule. The second paper *Near-Linear Lower Bounds for Distributed Distance Computations, Even in Sparse Networks* was presented by Seri Khoury (co-authored with Amir Abboud and Prof. Keren Censor-Hillel). This work developed a new technique for constructing sparse graphs that enable the authors to prove near-linear lower bounds on the round complexity of computing distances in the CONGEST model. In particular, an $\tilde{\Omega}(n)$ lower bound for computing the diameter in sparse networks was given. Then Ohad Ben-Baruch (co-authored with Prof. Hagit Attiya and Prof. Danny Hendler) presented the paper *Lower Bound on the Step Complexity of Anonymous Binary Consensus*. This paper presented a lower bound of $\Omega(\log n)$ on the solo step complexity of obstruction-free binary anonymous consensus.

Right after these three presentations, all the audience voted for the winner of the Best Student Paper Award. Based on the voting results, this year, the award goes to paper *Near-Linear Lower Bounds for Distributed Distance Computations, Even in Sparse Networks*.

This year, the conference banquet was held on boat. The cruise started and finished at the Pont Neuf (near Saint Chapelle and just down from the Louvre). This dinner cruise enabled the DISC participants to enjoy French food, French wine and Champagne, and the Eiffel Tower, Notre Dame, the Louvre and more top Paris attractions at the same time.

3 Third Day

The third day of the conference started with a keynote talk *Matching and Covering in Streaming Graphs* by Prof. Graham Cormode. Matching and covering problems both have a long history. In the big data era, due to the memory constraints of machines, researchers have been seeking

solutions in the streaming setting, where each edge is seen once only. Prof. Graham Cormode presented several new results that improved the state of the art fundamentally.

The "Distributed Data-Structures and Algorithms" session was chaired by Prof. Philipp Woelfel. *Optimal Consistent Network Updates in Polynomial Time* presented a polynomial-time algorithm for computing optimal consistent order updates. Several interesting findings and techniques were presented in the paper *Depth of a Random Binary Search Tree with Concurrent Insertions.* A distance labeling scheme is an assignment of bit-labels to the vertices of an undirected, unweighted graph such that the distance between any pair of vertices can be decoded solely from their labels. *Sublinear-Space Distance Labeling using Hubs*, authored by Pawel Gawrychowski, Adrian Kosowski and Przemyslaw Uznanski, presented a hub labeling which allows us to decode exact distances in sparse graphs using labels of size sublinear in the number of nodes. A new problem—Balanced RePartitioning (BRP) problem—was introduced in *Online Balanced Repartitioning* by Chen Avin, Andreas Loukas, Maciej Pacut and Stefan Schmid.

The last two sessions were "Distributed Models and Computation" and "Robots and Shared Memory", chaired by Prof. Emmanuel Godard and Prof. David Ilcinkas, respectively.

The paper *Are Byzantine Failures Really Different from Crash Failures?*, authored by Damien Imbs, Michel Raynal and Julien Stainer, presented a new distributed simulation in asynchronous systems, where up to t processes can fail, and communication is by read/write registers or reliable message-passing. It is shown that (from a computability point of view) Byzantine failures are **NOT** "different" from crash failures. Jingjing Wang (co-authored with Prof. Rachid Guerraoui) then talked about their work on *Optimal Fair Computation.* They proved, for the first time, a tight lower-bound on the message complexity of optimistic fair computation for n parties among which $n-1$ can be malicious in an asynchronous network. A new information spreading model was introduced in *How to Discreetly Spread a Rumor in a Crowd*, by Mohsen Ghaffari and Calvin Newport. This model can guarantee some notion of privacy. They argued that this model matched the trends of future peer-to-peer communication standards for mobile devices.