Distributed Computing Column 36
Distributed Computing: 2009 Edition

Idit Keidar
Dept. of Electrical Engineering, Technion
Haifa, 32000, Israel
idish@ee.technion.ac.il

It's now the season for colorful reviews of 2009. While you can read elsewhere about the year in sports, top-40 pop hit charts, people of the year, and so on, I throw my share into the mix with a (biased) review of this year's major distributed computing events.

Awards First, let's look at awards. This year we learned that two women were recognized with ACM and IEEE prestigious awards for their achievements in, (among other things), distributed computing. Barbara Liskov won the 2008 ACM Turing Award for a range of contributions to practical and theoretical foundations of programming language and system design, some of which are in distributed computing. The award citation mentions her impact on distributed programming, decentralized information flow, replicated storage, and modular upgrading of distributed systems. I include below a short reflection, by Rodrigo Rodrigues, on Barbara Liskov's Turing Award and legacy.

Looking ahead to 2010, it was recently announced that Nancy Lynch is the recipient of the prestigious 2010 IEEE Emanuel R. Piore Award[1] for her contributions to foundations of distributed and concurrent computing. The award is given for outstanding contributions in the field of information processing in relation to computer science. Leslie Lamport has won the same award in 2004 for his seminal contributions to the theory and practice of concurrent programming and fault-tolerant computing.

The 2009 Edsger W. Dijkstra Prize in Distributed Computing was awarded to Joseph Halpern and Yoram Moses for "Knowledge and Common Knowledge in a Distributed Environment". The paper was first published in PODC 1984, and then in JACM in 1990. The same paper won the Gödel prize in 1997. The Dijkstra Prize is jointly awarded by PODC and DISC; Joe and Yoram received it in DISC this year (see review of DISC below). The full Award Citation appears earlier in this issue of SIGACT News (and on the PODC web page[2]), so I do not repeat it here.

[1] http://www.ieee.org/portal/pages/about/awards/sums/piore.html
[2] http://www.podc.org/dijkstra/2009.html

Instead, at the risk of repeating what is already *common knowledge*, let me just say that this paper provides an elegant framework for formally reasoning about what processes "know" or "cannot know", what facts are "common knowledge", and how all of this relates to communication and coordination in distributed systems. The notion of common knowledge can be illustrated in an everyday setting as follows: Alice and Bob love each other. Still, Alice may think to herself "I know that Bob loves me. But does he know that I know that he loves me?" Here the fact that Bob loves Alice is not common knowledge. For it to be common knowledge, Alice would have to know that Bob knows that she knows that he knows . . . *ad infinitum*. Halpern and Moses showed a formal connection between common knowledge and simultaneous coordination in distributed protocols. They proved that common knowledge is not attainable when communication is not reliable. Moreover, they showed that establishing that common knowledge is essential in order to solve Jim Gray's coordinated attack problem, also known as *atomic commit*, and so formalized and generalized the proof that atomic commitment is impossible to achieve with unreliable communication.

The Halpern-Moses paper was influenced by concepts from diverse disciplines, including philosophy, economics, psychology, distributed databases, and distributed systems. One of the upcoming columns will describe what went on "behind the scenes" of this paper, and how all of these influences came together.

A new award was introduced this year– the Prize for Innovation In Distributed Computing[3]. The prize was established by SIROCCO to "recognize individuals whose research contributions on the relationships between information and efficiency in decentralized computing have expanded the collective investigative horizon by formulating new problems, or identifying new research areas, that were at the time of their introduction unorthodox and outside the mainstream." The prize is awarded in SIROCCO, (see review below), and its first recipient was Nicola Santoro, one of the fathers and early organizers of SIROCCO as well as other conferences, such as PODC and DISC (originally called WDAG). The prize was given to Santoro for his contributions to the analysis of labeled graph properties, including introducing the notions of *implicit routing*, *sense of direction*, and *topological awareness*. These notions have been shown to have a significant impact on computability and complexity in systems of communicating entities. The full award citations appears below. Congratulations to all of this year's award winners!

Blog From awards, we move on to cyber-space. The most exciting development in 2009 is the inception of Jared Saia's blog machinations[4]. Jared describes the blog as "containing meandering thoughts on order, disorder and robustness in complex systems, with an eye towards connections with computer science theory, distributed computing and security." The blog has thus far featured topics such as consensus, random sampling in peer-to-peer networks, greedy routing in small world networks, and game theoretic analysis of distributed systems with malevolent users. In short, this is a good place to surf if you're interested in distributed algorithms!

Conferences But blog surfing is no replacement for physically attending conferences. So, we proceed with reviews of this year's distributed computing venues. The first review below is of PODC, the ACM Symposium on Principles of Distributed Computing, which took place in Calgary in August, and was collocated with SPAA, the ACM Symposium on Parallelism in Algorithms and Architectures. (PODC was also covered on a daily basis by Amitabh Trehan in Jared's blog[5].) In continuing the tradition of inviting winners of Best Paper or Best Student Paper Awards to review the conferences where they shined, I solicited a review of PODC by two winners: Keren Censor and Christoph Lenzen.

[3] http://www.imfm.si/sirocco09/prize.htm
[4] http://jsaia.wordpress.com/
[5] http://jsaia.wordpress.com/2009/08/11/day-1-podc-09/

Christoph is a PhD student in ETH Zürich, and he won the Best Paper Award (see picture below) together with Thomas Locher and Roger Wattenhofer for their paper "Tight bounds for clock synchronization". This work deals with *gradient clock synchronization*: the problem of synchronizing clocks of nodes in a network while minimizing the skew between neighboring nodes. The main contribution is in presenting tight bounds that take into account clock drifts and similar parameters, which, by-and-large, were neglected in previous work. More concretely, the focus is on algorithms that bound the minimum and maximum progress rates of the logical clocks. Such algorithms have the practical property that clock values do not change abruptly in a short time or stop changing for too long.

Christoph Lenzen collects the PODC Best Paper Award at the Calgary Zoo. Photo by Marcos Aguilera.

Keren Censor is a PhD student at the Technion, and she won the Best Student Paper Award for her paper "Max registers, counters, and monotone circuits", co-authored with James Aspnes and Hagit Attiya. This work first builds a *max register*– a register that can be written by n processes, and returns the highest previously written value when it is read. These max registers are then used to construct counters. Naïve implementations of counters take $O(n)$ steps, and previous work has shown that there is an $\Omega(n)$ bound on the number of steps that their operations must take. So at first glance, this might not seem like a good choice of a research problem to focus on. But the authors took a closer look at the lower bound proof, and have noticed that the proof requires many operations to be invoked on the counter for these operations to require linear time. So they realized that the key to obtaining more efficient algorithms is limiting the number of operations that can be invoked on the counter. The paper uses this observation in order to "beat" the lower bound: it presents polylogarithmic implementations of max registers, counters, and other data structures that can hold a bounded number of values. When presenting the paper at PODC, Keren concluded that "lower bounds do not always have the final say". Keren has also won the PODC Best Student Presentation Award for this presentation. Kudos!

The second conference review is of DISC– the International Symposium on DIStributed Computing. DISC is reviewed by Andreas Tielmann, who won the Best Paper Award along with his co-authors Carole Delporte-Gallet, Hugues Fauconnier, and Rachid Guerraoui, for their paper "The disagreement power of an adversary" (see picture). This work considers *adversaries* that may crash certain combination of processes and not others. In general, there are 2^{2^n} such adversaries. In order to classify them, the paper introduces the notion of *disagreement power*: the biggest integer k for which the adversary can prevent processes from agreeing on k values. Using this notion, they show that there are in fact n equivalence classes of adversaries.

The winner of DISC's Best Student Paper Award is Henrique Moniz from the University of Lisbon, who won the award for his paper "Randomization Can Be a Healer: Consensus with Dynamic Omission Failures", co-authored with Nuno Neves, Miguel Correia and Paulo Verissimo.

Carole Delporte-Gallet, Rachid Guerraoui, Hugues Fauconnier, and Andreas Tielmann receive the DISC Best Paper Award from Idit Keidar. Photo by Juan Echagüe.

Following the review of DISC, I include summaries of the five workshops co-located with DISC, provided by their organizers: What Theory for Transactional Memory?; BFTW^3: Why? When? Where? (Workshop on Theory and Practice of Byzantine Fault Tolerance); Workshop on Reliability and Security in Wireless Networks; Workshop on Game Theoretic Aspects of Distributed Computing; and Workshop on Theoretical Aspects of Dynamic Distributed Systems (TADDS).

The final review this year is of SIROCCO, a conference that takes pride in promoting crazy ideas in distributed computing, which are often viewed as far-fetched and overly impractical in other venues. SIROCCO is reviewed by Shantanu Das.

In all reviews, you will find fun information about the venues, as well as technical content. Many thanks to Rodrigo, Christoph, Keren, Andreas, the DISC Workshop Organizers (Vincent Gramoli, Petr Kuznetsov, Rodrigo Rodrigues, Seth Gilbert, Darek Kowalski, Roberto Baldoni, Alex Shvartsman, and Chryssis Georgiou), and Shantanu for their colorful contributions!

Introspection Finally, a glance in the mirror, and an eye on the crystal ball. This column has dealt this year with: teaching concurrency in light of the multi-core revolution (in March), a distributed computing research agenda for the new popular platform of cloud computing (in June), and a related perspective on theory and practice in large distributed systems (in September). For future columns, I expect some fascinating contributions on diverse topics– from distributed algorithms and models for radio networks, through round-based models for distributed quantum computing, to random graph transformations, and from recent developments in consensus and Byzantine fault tolerance to game-theoretic approaches for modeling a range user behaviors from selfish to Byzantine.

Best wishes for a fruitful 2010!

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.

Barbara Liskov's Turing Award

Rodrigo Rodrigues
Max Planck Institute for Software Systems (MPI-SWS)
rodrigo@mpi-sws.org

On March 10 this year, Barbara Liskov was awarded the 2008 ACM Turing award, the highest distinction for technical accomplishments in computer science.

Most of us in this community are familiar with Barbara's work on distributed computing. Her work in this field is impressive in terms of both the breadth of topics it covers, and the impact of her contributions. Examples include the Viewstamped Replication protocol for state machine replication, first presented in PODC in 1988; the work on distributed garbage collection that dates back to the mid-80s, long before environments such as the .NET Framework incorporated such concepts; the Thor project that proposed a transactional storage system that had several innovative techniques in areas like client caching or optimistic concurrency control; and more recently the work on Byzantine fault tolerance that brought new life to a subject that was considered by many to be of limited practical interest.

Barbara's Turing award lecture was given on October 12, 2009 at the opening session of the 22nd ACM Symposium on Operating Systems Principles. Somewhat surprisingly, the lecture focused on her earlier work on programming language design, which influenced several aspects of modern languages such as Java or C#. As she described fundamental concepts such as abstract data types or type hierarchies, and languages she designed such as CLU or Argus, what struck all as fascinating was her immense passion about the process of analyzing a problem, figuring out the limitations of existing approaches, and coming up with solutions that strike the right balance between simplicity, expressiveness, and performance.

Barbara's contributions helped shape the field of computer science in many ways beyond the technical achievements of her work. She supervised over 20 doctoral students, many of whom moved on to academic careers of their own and have now produced a tree of academic descendants with over 80 people that have graduated with a doctoral degree. She has also long been an advocate of the importance of the role of women in computer science. She has fought not only to encourage women to play an active role in the field, but also against gender related barriers to entry and success in computer science in general, and particularly in academic or research careers.

For all these reasons the Turing award is a well-deserved recognition of Barbara's work. I am sure that all of us feel very proud of having her as a colleague, and often case as a mentor, and we are all delighted that she won this award.

Prize for Innovation in Distributed Computing 2009: Nicola Santoro

The Prize for Innovation In Distributed Computing is awarded by the Colloquium on Structural Information and Communication Complexity (SIROCCO). It is established to recognize individuals whose research contributions on the relationships between information and efficiency in decentralized computing have expanded the collective investigative horizon by formulating new problems, or identifying new research areas, that were at the time of their introduction unorthodox and outside the mainstream. The prize recognizes originality, innovation, and creativity. The recipient of the Prize is chosen among the nominated persons for the current year.

The Award Committee has selected *Nicola Santoro* as the recipient of this year's Prize for Innovation In Distributed Computing.

The prize is given to Nicola Santoro for his overall contribution on the analysis of the labeled graph properties which has been shown to have a significant impact on computability and complexity in systems of communicating entities. These contributions are including the notions of "Implicit Routing", "Sense of Direction", and "topological awareness", illustrated, in particular, by the following papers:

- Nicola Santoro, Ramez Khatib: Labeling and Implicit Routing in Networks. Comput. J. 28(1): 5–8 (1985).
- Paola Flocchini, Bernard Mans, Nicola Santoro: Sense of Direction: Formal Definitions and Properties. SIROCCO 1994: 9–34.
- Paola Flocchini, Alessandro Roncato, Nicola Santoro: Backward Consistency and Sense of Direction in Advanced Distributed Systems. SIAM J. Comput. 32(2): 281–306 (2003).

The paper "Labeling and Implicit Routing in Networks" is a pioneering paper investigating the design of compact data structures for distributed routing in networks. In particular, it introduces the Interval Routing technique as one of the very first techniques used for compacting information required for routing. This technique has been implemented by INMOS on its most recent version of the Transputer, and is at the basis of some XML search engines.

The aforementioned papers on Sense of Direction are part of the very few papers that make explicit common implicit notions of knowledge. In particular, the notion of sense of direction formalizes the relationship between the node/edge labeling in a network, the topological structure of this network, and the local view that an entity can get of the entire system. It enables to generalize to any graph results that were known only under certain assumptions such as the ability to distinguish right from left, or to identify the four cardinal directions.

By his results and ideas, Nicola Santoro has enriched Distributed Computing considerably, providing innovative concepts at the source of an extremely large number of current investigations, ranging from compact routing to mobile computing. His pioneering investigations of distributed computing in labeled networks are among the most influential ones, and have opened a vast domains of promising researches, aiming at capturing and understanding the central notion of "local knowledge". He is the author of the book "Design and Analysis of Distributed Algorithms" (Wiley, 2006). Quoting the introduction: "My own experience as well as that of my students leads to the inescapable conclusion that both to teach and to learn distributed algorithms are fun". Those of us who had the opportunity to work with Nicola Santoro are the witnesses that it is not the least of his contributions to have make this conclusion a fact.

The prize has been officially delivered at the Business meeting of the 16th edition of the Colloquium on Structural Information and Communication Complexity (SIROCCO), May 25-27, 2009, Piran, Slovenia.

Award Committee 2009:

Pierre Fraigniaud	CNRS and University Paris Diderot
Leszek Gasieniec	University of Liverpool
David Peleg	Weizmann Institute
Alexander A. Shvartsman	University of Connecticut
Shmuel Zaks	Technion

A Review of PODC 2009

Keren Censor
Computer Science Department
Technion, Haifa, Israel
ckeren@cs.technion.ac.il

Christoph Lenzen
Computer Eng. and Networks Laboratory (TIK)
ETH Zürich, Switzerland
lenzen@tik.ee.ethz.ch

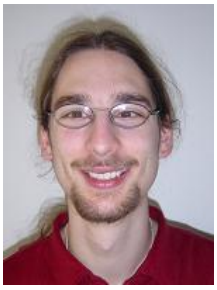

The 28th annual ACM Symposium on Principles of Distributed Computing (PODC), took place on August 10–12, 2009 in Calgary, Canada. It was co-located with the 21st ACM Symposium on Parallelism in Algorithms and Architectures (SPAA). Calgary's cool weather and scattered drizzles were a nice escape from the heat of August. Quite a few rabbits were spotted on the lawn of the University of Calgary, although it was mainly dominated by numerous groups of playful teenagers.[1]

Calgary airport and university campus.

Program The conference featured three invited talks, seven sessions of full paper presentations, and six brief announcements sessions. It started on Sunday evening with a warm welcome reception at the *Nickle Arts Museum*. Presentations took place from Monday until the end of the conference on Wednesday. The

[1] We assume that Canadians reproduce faster than rabbits.

Calgary tower and view of the city.

conference program is available on the web[2]. Papers and extended abstracts are published in [1].

Rather then giving an unbiased overview of the content, we employ a "statistical analysis" to "identify" current streams in the focus of PODC.[3] We compared the relative frequencies of keywords in paper titles from this year, PODC 2007–2009, and PODC 1997–1999.

Entrances to the welcome reception and banquet.

At first glance on the results, two findings are immediately visible. Firstly, the common belief that research is completely random is proved wrong: Four of the top five keywords, namely *distributed*, *systems*, *networks*, and *efficient*, did not change over the last decade,[4] and *distributed* holds first place in all three cases. Secondly, gaining insight into distributed problems seems to become increasingly popular: Before 2000, *distributed* occurred in roughly 4.1% of the titles, until 2007–2009 this doubled to roughly 8.5%, and in Calgary the 10%-barrier has been broken.[5]

So what else has changed? Ten years ago, *consensus* was the third most common keyword in PODC, exceeded only by the much more unspecific words *systems* and *distributed*. It occurred only half as fre-

[2] http://www.podc.org/podc2009/program.shtml

[3] It's always fun doing messy statistics.

[4] Some keywords were tied for the fifth place; being messy, we simply took the topmost five from our output list.

[5] Perhaps the still lasting olympic spirit inspired us.

quently in the last years, three times in 2009. Having a closer look, not all of these papers deal with solving consensus problems. In "Dynamic Storage Without Consensus" by Aguilera et al., it is proved that atomic r/w memory can be emulated in asynchronous message passing systems, avoiding the need for consensus. Intriguingly, the keynote talk "Refining the Way to Consensus" given by Robbert van Renesse also did not address a solution of consensus. Instead, he suggested to teach the topic more efficiently to undergraduates by presenting it through gradual refinement.

Consensus has been replaced by *dynamic* in recent PODCs (∼2.8%), again on third place behind *distributed* and *systems*. Compared to 1997–99 (∼1%, 25th place) it is almost three times as frequent. This might indicate that having solved more and more problems for the static case, research starts to focus more on dynamic models which better resemble many of today's large-scale distributed systems. One such paper examining a problem that is well-understood in the static context is "Parsimonious Flooding in Dynamic Graphs" by Baumann et al. The authors present tight bounds on the complexity of flooding in graphs that evolve by an edge-Markovian process.

The top results of a single year—or at least this year—appear to be less stable than the three-year average. The evergreens *distributed* (∼10.5%) and *networks* (∼8.5%) are followed by *wireless* (∼6%), which means still a comparatively strong polarization towards this issue. It occurred eight times in 2009, four times during the course of the previous two years, and only two times from 97–99. Is this a trend or just a coincidence? There were three regular papers featuring this keyword. In "The Wireless Synchronization Problem", Dolev et al. give almost matching bounds on the task of agreeing on round numbers in a single-hop network where some of the available frequencies are jammed. This is quite different from the work of Avin et al. "SINR Diagrams: Towards Algorithmically Usable SINR Models of Wireless Networks", which strives for basic understanding of the area where a node in the network can be received under the SINR interference model. Assuming uniform power assignments, they prove that these areas are convex and relatively well-rounded. Schneider and Wattenhofer, "Coloring Unstructured Wireless Multi-Hop Networks", proposed an efficient coloring algorithm resilient to asynchronous wake-up and communication, requiring neither collision detection nor initial knowledge on the network topology. This diversity continues throughout the respective brief announcements, showing either the depth of the topic or that it is a popular means to demonstrate raison d'être of theoretical results.

Self and *algorithms* share the fourth and fifth position on this year's hitlist. At the latest by now weaknesses of our approach become apparent: Though the absolute count of *self* is five, two occurrences are in the same title. Moreover, it does not refer to the same topic in all cases: Aaron Sterling took the opportunity of a brief announcement to advertise results from the application of distributed computing theory to *self-assembly*, where predefined tilesets perform computations by assembling larger structures starting from a "seed" resembling the input. The other sightings of *self* are in the form of *self-stabilization*, the single regular paper carrying this word being "A Distributed Polylogarithmic Time Algorithm for Self-Stabilizing Skip Graphs" from Jacob et al. The authors present an augmented skip graph and a polylog-time self-stabilizing asynchronous local algorithm that recovers from any weakly connected initial topology. *Algorithms* on the other hand does not give any particular information, rather it lines in with the abundance of general terms inhabiting the top ranks: PODC is all about *efficient algorithms* for *networks* and *distributed systems*.

At this point we conclude our wordy trip through careless PODC statistics, apologizing for the wide gaps in the coverage of content. We did not even touch some fascinating results on game theory and randomized topology control, just to name two examples. Perhaps techniques proposed by Das Sarma and Nanongkai in "Fast Distributed Random Walks" permit a more efficient exploration of future PODC's landscapes than our simplistic approach. To remedy this deficiency at least in parts, we next give a short overview of the keynote talks.

Keynote Talks As mentioned above, the first keynote talk, given by Robert van Renesse, was titled "Refining the Way to Consensus". He addressed the issue of teaching undergraduate students about programming complex software systems by considering first a basic algorithm, and then gradually refining it to satisfy more specifications. This was demonstrated on consensus protocols, which are an important example since they lie in the base of all replication techniques, crucial for web-based services. The next keynote talk was "Memory Models: A Case for Rethinking Parallel Languages and Hardware", given by Sarita Adve. She talked about the gap between popular parallel environments and concepts of structured and safe programming, specifically. She discussed research directions for parallel programming languages that prevent, for example, the occurrence of data races. The last invited talk was given by Bruce Hendrikson, "Emerging Challenges and Opportunities in Parallel Computing: The Cretaceous Redux?" Addressing the gap between theory and practice, he described the recent changes in the parallel computing landscape, such as the rise of multi-core processors, and the observation that "Silicon is now essentially free". With amusing dinosaur animations he described how these changes would reshape the world of parallel computing.

Industry Session & Lunch In addition, there was an invited session on industrial applications of algorithms, chaired by Marcus Aguilera from Microsoft Research. This was something new for PODC, and showed the interesting design difficulties that arise when moving from theory to practice. The session included three talks. One was about Yahoo!'s *ZooKeeper*, which is a system used for coordination in distributed applications. ZooKeeper provides synchronization primitives, such as leader election and group membership, and can be used by application developers instead of re-implementing them per project. Another talk was about *Cassandra*, which is a distributed storage system developed by Facebook for managing large scale data. It is designed to provide reliability while achieving high performance and availability. The last talk was about *Pregel*, Google's system used for processing of large scale graphs. These appear in many computing problems, such as social networks, and network graphs. The session came after a sponsored lunch for students, postdocs, and employees of industrial sponsors, where representatives of the companies talked about their experience in doing research in industry research labs.

Finally, this review would not be complete without mentioning the wonderful organization of the conference due to Lorenzo Alvisi and the PC, Conference, and Steering Committee members. This is particularly remarkable since everything ran smooth even though Lorenzo, being caught in a legal dispute, could not attend the conference.[6] Another success story are the local arrangements, well-managed by Lisa Higham and Philipp Woelfel. And finally some news from the business meeting: We wish to thank the great contributions of Faith Ellen, who completed her term as Steering Committee Chair, and Michel Raynal, the corresponding Member-at-Large. Their roles are now taken by Andrjez Pelc and Jennifer Welch. We are looking forward to PODC 2010, which will take place at ETH Zurich in Switzerland!

Acknowledgements Photos are courtesy of Dmitri Perelman, Johannes Schneider, and Amitabh Trehan. Michael Kuhn kindly provided the tools to generate the keyword statistics. Thanks!

References

[1] *PODC '09: Proceedings of the 28th ACM Symposium on Principles of Distributed Computing*, 2009. General Chair: Srikanta Tirthapura; Program Chair: Lorenzo Alvisi.

[6] At the business meeting it was announced that he had been engaged as an expert and thus became an important witness. Badly informed sources speculate, however, that the judge was sulky because he has not been invited to PODC as a keynote speaker.

A Review of DISC 2009

Andreas Tielmann
LIAFA, University Paris Diderot, France
tielmann@liafa.jussieu.fr

The 23rd International Symposium on DIStributed Computing (DISC 2009) took place in Elche/Elx, Spain on September 23-25, 2009. It is an international symposium on the theory, design, analysis, implementation and application of distributed systems and networks that is organized in cooperation with the European Association for Theoretical Computer Science (EATCS). The symposium was established in 1985 as a biannual International Workshop on Distributed Algorithms on Graphs (WDAG). The scope was soon extended to cover all aspects of distributed algorithms and is nowadays an annual symposium (since 1989) under the name DISC (since 1998) to reflect the exciting and rapid developments in this field.

This year's DISC featured 3 keynotes, 1 tutorial, 33 regular presentations, 15 brief announcements, 5 co-located workshops and one double-birthday celebration session. Furthermore, in this edition of DISC, also the 2009 Edsger W. Dijkstra Prize in Distributed Computing was awarded.

The conference was located in the beautiful city of Elche/Elx, the spanish footwear capital and third largest city in the Comunidad Valenciana. The conference hotel was directly located in the historical centre, next to Europe's largest palm grove with 300,000 palm trees. Some of them can be seen in Figure 1(a), but it seems that some pomegranates attracted the attention of some of the conference attendees even more (Figure 1(b)). Elche has two UNESCO World Heritage items: the palm grove and the "Mystery Play" (Misteri), a religious lyrical drama from the middle ages. Fortunately, the conference program arranged visits to these sights[1]. More photos of the conference can be found on the web[2].

The *Program Committee* was chaired by Idit Keidar (Technion) and selected 48 great papers out of 121 submissions (116 for regular papers and 5 for brief announcements). Thirty-three of these were accepted for regular presentations and 15 as brief announcements. To make these decisions, 402 reviews were made. Of the submitted papers, 21 received only positive scores by the reviewers and more papers got an "accept" average score than could possibly be accepted for the conference.

The *Organizing Committee* was co-chaired by Antonio Fernández Anta (Universidad Rey Juan Carlos) and Vicent Cholvi Juan (Universitat Jaume I) who both did an excellent job with the local arrangements (especially the excellent spanish food at the conference and the banquet (Figure 2) should be highlighted here). The *Steering Committee* was chaired by Rachid Guerraoui (EPFL).

[1] As it is appropriate for a forefront computer science conference and since the "Mystery Play" takes place only in August, it was presented *virtually* using modern audiovisual technologies.

[2] http://disc2009.gsyc.es/photos

Figure 1: (a) Discovering Elche; (b) The forbidden fruit(s). Photos by Jukka Suomela.

This year, more than 116 persons registered for DISC, among which 93 were regular registrants and 23 registered as students. Apart from that, 3 keynote speakers, 2 Dijkstra-Award winners, 9 local organizers and 5 sponsored students attended the conference.

For the first time, this year's DISC featured a special *"Award Nominees"* session, in which 5 papers that were selected by the program committee competed for the "Best Paper Award" and the "Best Student Paper Award". The award winners were then chosen by the votes of the attendees of the award session.

As Idit Keidar pointed out at the conference business meeting, this year's DISC was also remarkable in other ways. Several records were broken: e.g. in the number of provided meals, the number of awards, the number of sponsorships and the number of papers about "failure detectors" to mention only a few.

Workshops and co-located events. Before and after DISC, 5 wokshops, 1 tutorial and 1 double-birthday celebration took place. Detailed reviews of the 5 workshops can also be found in this column. I will give here only a short list:

- *"What Theory for Transactional Memory?"* organized by Rachid Guerraoui (EPFL) and Vincent Gramoli (EPFL and Univ. of Neuchâtel).
- *"Workshop on Reliability and Security in Wireless Networks"* organized by Seth Gilbert (EPFL) and Dariusz Kowalski (Univ. of Liverpool).
- *"BFTW3: Why? When? Where? (Workshop on Theory and Practice of Byzantine Fault Tolerance)"* organized by Petr Kuznetsov (TU Berlin/ Deutsche Telekom Laboratories) and Rodrigo Rodrigues (Max Planck Institute for Software Systems MPI-SWS)
- *"Workshop on Theoretical Aspects of Dynamic Distributed Systems"* organized by Roberto Baldoni (Univ. of Rome "La Sapienza") and Alexander A. Shvartsman (Univ. of Connecticut).
- *"Workshop on Game Theoretic Aspects of Distributed Computing"* organized by Chryssis Georgiou (Univ. of Cyprus) and Paul Spirakis (CTI and Univ. of Patras).
- Half day tutorial on *"Cloud Computing Architecture and Application Programming"* organized by Roger Barga, Jose Bernabeu-Auban, Dennis Gannon and Christophe Poulain from Microsoft Corporation.

Figure 2: The banquet. Photo by Juan Echagüe.

Figure 3: Shmuel Zaks and Michel Raynal discussing. Photo by Dmitri Perelman.

- Birthday celebration session: *"Michel Raynal and Shmuel Zaks: Sixty and Beyond"*. In celebration of Michel Raynal and Shmuel Zaks's sixtieth birthday, DISC 2009 featured a series of invited lectures illustrating and celebrating the influence of these colleagues and their work. The birthday boys can be seen in Figure 3.

Dijkstra Award. The 2009 Edsger W. Dijkstra Prize in Distributed Computing was awarded at the conference banquet. It was given to Joseph Halpern and Yoram Moses for their outstanding paper *"Knowledge and Common Knowledge in a Distributed Environment", Journal of the ACM, Vol. 37, No. 3, January, 1990*. The prize is sponsored jointly by the ACM Symposium on Principles of Distributed Computing (PODC) and DISC.

Conference Sessions. The conference consisted of 11 sessions, including the 3 keynote talks by Nir Shavit, Lorenzo Alvisi and Willy Zwaenepoel and the award nominees session.

- The first session of the conference was about *"Transactional Memory"* and chaired by Idit Keidar. It consisted of 4 talks, including a keynote lecture by Nir Shavit (Tel-Aviv University) with the title *"Software Transactional Memory: Where do we come from? What are we? Where are we going?*.
- The second session was about *"Shared Memory"* with 6 talks, chaired by Eric Ruppert.
- The third session was about *"Distributed and Local Graph Algorithms"* with 6 talks, chaired by Fabian Kuhn.
- The first day of the conference ended with the session *"Modeling Issues"* with 5 talks, chaired by Pascal Felber.
- The second day of the conference started with a session about *"Game Theory"*, chaired by Idit Keidar. It consisted of 3 talks, including a keynote lecture by Lorenzo Alvisi (University of Texas at Austin) with the title *"BFT we can believe in"*.
- After that, a special session with the *"Award Nominees"* started (see Figure 4). This session was chaired by Yehuda Afek and consisted of 5 talks that were elected by the program committee. Every

Figure 4: Award Nominees Session: (a) Eli Gafni heckling a speaker; (b) Henrique Moniz giving his talk. Photos by Jukka Suomela.

conference attendee was given a ballot paper to vote for the talks in order to elect the "Best Paper" and the "Best Student Paper" of the conference. These 5 nominated talks were

- *"The Disagreement Power of an Adversary"* by Carole Delporte-Gallet, Hugues Fauconnier, Rachid Guerraoui and Andreas Tielmann
- *"New Bounds for the Controller Problem"* by Yuval Emek and Amos Korman
- *"On Set Consensus Numbers"* by Eli Gafni and Petr Kuznetsov
- *"The Abstract MAC Layer"* by Fabian Kuhn, Nancy Lynch and Calvin Newport
- *"Randomization Can Be a Healer: Consensus with Dynamic Omission Failures"* by Henrique Moniz, Nuno Neves, Miguel Correia and Paulo Verissimo.

Among these papers, *"The Disagreement Power of an Adversary"* was given the best paper award and *"Randomization Can Be a Healer: Consensus with Dynamic Omission Failures"* was given the best student paper award.

- The third session of the second day was entitled *"Failure Detectors"* and consisted of 6 talks and was chaired by Carole Delporte-Gallet.
- The third day of the conference started with the session *"From Theory to Practice"* and was chaired by Idit Keidar. It consisted of 2 talks, including a keynote lecture by Willy Zwaenepoel (École Polytechnique Fédérale Lausanne (EPFL)) with the title *"P2P, DSM, and Other Products of the Complexity Factory"*. Zwaenepoel, who is Dean of Computer Science at EPFL, made the rather discouraging observation that solutions that tend to succeed in academia do not do well in industry, and vice versa. In order to publish papers in research conferences, including systems-oriented ones, one needs to propose complex solutions, whereas in the real world, simple solutions are preferred.
- The next session was about *"Graph Algorithms and Routing"* and consisted of 5 talks and was chaired by Antonio Fernández.
- The third session of the third day was entitled *"Consensus and Byzantine Agreement"*, consisted of 4 talks and was chaired by Christian Cachin.

- The last session of the conference with 5 talks was about *"Radio Networks"* and chaired by Seth Gilbert.

DISC 2010. In 2010, DISC will be held in Cambridge, Massachusetts, USA. This will be the first time it is not held in Europe. The program committee will be co-chaired by Nancy Lynch and Alex Shvartsman and the conference will take place on September 13-15, 2010.

What Theory for Transactional Memory?

Vincent Gramoli
EPFL and University of Neuchâtel
Switzerland
vincent.gramoli@epfl.ch

Introduction

The Transaction Memory (TM) paradigm is argued to be as easy to use as coarse-grained locking and nearly as efficient on multi-core systems as hand-crafted fine-grained locking. It is thus not surprising to see a large body of work dedicated to implementing the paradigm and exploring its limitations. Very little work has however been devoted to explore its theoretical ramifications. This workshop, organized with Rachid Guerraoui, asks what theory is needed for TM. The aim of this document is to summarize a subset of the numerous ideas that have been debated during this workshop.

Towards a Unified Transactional Memory Model

Modeling TM is the first step towards unifying the understanding of what TMs can achieve. In its simplest form a TM abstraction is a service that provides `read`, `write` and `commit` operations to programmers. This basic interface provides very little control on the execution. As a result it is impossible for a TM to ensure, in asynchronous system with crash failures, that every thread that keeps executing transactions eventually commits a transaction [6]. This property called wait-freedom is the strongest liveness property and does not assume any particular schedule [7]. Conversely, in a different model where the TM not only

implements those accesses but is able to map some emulated threads to the native threads and to change the execution of a thread at runtime, the transactional memory can ensure wait-freedom as shown in [3].

It is also noteworthy that recent consistency criteria require TM to be isolated [8, 5] whereas previous TM designs suggested isolation to be explicitly requested on demand or to be provided by an external sandboxing service [4]. Isolation prevents all transactions from reaching an uncommitted and possibly inconsistent state thus preventing pointer arithmetic issues or division-by-zero errors. A TM requires incremental validation to ensure isolation, which may reduce significantly its performance. To some extent, this observation raises the question whether a TM should be fault-tolerant.

Making a Simple Abstraction Efficient

In an attempt to make the programming abstraction as general as possible, TM often hampers performance. In contrast with TMs, synchronization techniques at the core of a concurrent program that targets a single specific application can be genuinely tuned to enhance concurrency and to limit contention. It is thus crucial to better understand what properties make a speculative execution efficient and how these properties can be applied to a large set of scenarios. Scalability is a very appealing property as it indicates that performance increases as parallelism grows. This property is, however, less interesting when considered alone as the performance obtained from a sequential execution of a non-thread-safe code can be even faster [2]. Hence, it is important to consider scalability with some complementary properties.

A first property that allows TM to be scalable is *disjoint-access-parallelism* (DAP). Informally speaking, the definition recently proposed for the TM context is that two transactions that do not access the same shared locations access also disjoint metadata. We know that no TMs with wait-free read-only transactions can be DAP if its read operations are invisible from other concurrent transactions [1]. A second appealing property, called *conflict-free time complexity*, captures the time complexity in the absence of contention as a function of the number of read operations and the number of write operations this transaction executes. The conflict-free time complexity is generally quadratic unless all transactions use a shared global counter or a perfectly synchronized hardware clock [11]. Interestingly, both techniques violate DAP as soon as two transactions access disjoint sets of locations.

Another idea to enhance performance is to minimize the number of *unnecessary aborts*, an abort that occurs while a commit would have been safe instead. A typical example is when a transaction reads a location that is overwritten before it commits. Most existing TMs abort in this scenario in an over-conservative way. One solution to minimize the number of unnecessary aborts is to handle multiple versions at each location: in the previous example, the transaction reading the location could safely return the version that has been overwritten [9]. Unfortunately, keeping track of multiple versions is costly as some memory locations cannot be garbage collected and there is a clear tradeoff between the cost of maintaining a large number of versions for each location and the cost of aborting unnecessarily.

Other ideas include thread synchronizers [10] for scattering the operations of a single transaction into concurrent transactions that synchronize either to commit or to abort together.

Acknowledgments

I wish to thank Rachid Guerraoui for helping with the organization and the contributors to the workshop: Hagit Attiya, Aleksandar Dragojević, Pascal Felber, Christof Fetzer, Tim Harris, Maurice Herlihy, Damien Imbs, Michał Kapałka, Eric Koskinen, Victor Luchangco, Maged Michael, Alessia Milani, Mark Moir, Dmitri Perelman, Torvald Riegel, Nir Shavit and Vasu Singh.

References

[1] Hagit Attiya, Eshcar Hillel, and Alessia Milani. Inherent limitations on disjoint-access parallel implementations of transactional memory. In *SPAA*, pages 69–78, 2009.

[2] Aleksandar Dragojević, Pascal Felber, Vincent Gramoli, and Rachid Guerraoui. Why STM can be more than a research toy. Technical Report LPD-REPORT-2009-003, EPFL, 2009.

[3] Christof Fetzer. Robust transactions, 2009. Personal communication.

[4] Keir Fraser and Tim Harris. Concurrent programming without locks. *ACM TOCS*, 2007.

[5] Rachid Guerraoui, Thomas A. Henzinger, Michałl Kapałka, and Vasu Singh. Generalizing the correctness of transactional memory. In *Workshop on Exploiting Concurrency Efficiently and Correctly*, 2009.

[6] Rachid Guerraoui and Michał Kapałka. How live can a transactional memory be? Technical Report LPD-REPORT-2009-001, EPFL, 2009.

[7] Maurice Herlihy and Nir Shavit. On the nature of progress. Unpublished work.

[8] Damien Imbs, Jos Ramn Gonzlez de Mendvil, and Michel Raynal. Virtual world consistency: a new condition for STM systems. In *PODC*, pages 280–281, 2009.

[9] Idit Keidar and Dmitri Perelman. On avoiding spare aborts in transactional memory. In *SPAA*, pages 59–68, 2009.

[10] Victor Luchangco and Virendra Marathe. Transaction synchronizers. In *SCOOL*, 2005.

[11] Maged M. Michael. Conflict-free time complexity. Technical report, IBM Research, November 2009.

BFTW^3: Why? When? Where? Workshop on the Theory and Practice of Byzantine Fault Tolerance

Petr Kuznetsov
TU Berlin/Deutsche Telekom Laboratories

Rodrigo Rodrigues
Max Planck Institute for Software Systems

Abstract

The first BFTW3 workshop took place in Elche, Spain, on the 22nd of September, just before the 23rd International Symposium on Distributed Computing (DISC).

The workshop gathered researchers from both the theory and systems communities and aimed at discussing recent advances and understanding why the impressive research activity in the area of Byzantine fault-tolerance is not yet instantiated in practice. Has the moment for a wide deployment of BFT systems arrived, and if so, where BFT systems should be deployed in the first place?

Context

The Byzantine fault model assumes that a faulty system component is allowed to deviate arbitrarily from its expected behavior, modulo the ability of breaking asymmetric cryptography [6]. Starting from the original "Byzantine generals" paper by Lamport et al. [13], systems tolerating Byzantine faults were extensively studied.

Besides intellectual curiosity, there seem to be practical reasons to look at Byzantine fault-tolerant systems. Multiple cases of system outages caused by sophisticated non-crash faults (e.g., [1, 7, 14]) suggest that BFT may be a very attractive feature in modern computing systems. In recent years, multiple researchers focused on improving the performance and practicality of BFT systems, which also increases their chances of adoption (e.g., [3, 15, 18, 4, 12, 17]). Also, as an indication of perceived importance of this work, the last Turing award went to Barbara Liskov partially for her work on BFT.

However, there are few visible instantiations of these results in practical systems. Industrial software tends to ignore the BFT-related research and heads for less consistent but (apparently) simpler and more efficient solutions (e.g., [5, 16, 11]).

In this workshop, we discussed the state of the art in BFT systems, and tried to understand why BFT systems have not seen a widespread adoption, and what we could do to increase the chances of deploying BFT systems.

Talks

The workshop featured seven talks. Below we give a very short overview of the program (more detailed abstracts and slides of the talks can be found at http://www.net.t-labs.tu-berlin.de/~petr/BFTW3/program.html):

- *BFT for the skeptics*. Yee Jiun Song (Cornell University), Flavio Junqueira (Yahoo! Research), Benjamin Reed (Yahoo! Research)

  In this somewhat provocative talk, Flavio Junqueira (Yahoo! Research) highlighted several critical assumptions that are usually made in the BFT-related research. The talk also described typical difficulties that one faces trying to convince managers, developers, and maintenance engineers that BFT makes sense.

- *Practical Intrusion-tolerance in the Cloud*. Rüdiger Kapitza (Informatik 4, University of Erlangen-Nürnberg), Hans P. Reiser (LaSIGE, University of Lisboa), and Tobias Distler (Informatik 4, University of Erlangen-Nürnberg)

  This talk, presented by Rüdiger Kapitza (Univ. of Erlangen-Nürnberg), described *Spare*, a BFT system that uses $f+1$ replicas in the normal case, and rapidly activates up to f more replicas in case of failures or intolerable delays using the virtualization features provided by modern cloud-computing platforms.

- *Towards Recoverable Hybrid Byzantine Consensus*. Hans P. Reiser (LaSIGE, University of Lisboa) and Rüdiger Kapitza (Informatik 4, University of Erlangen-Nürnberg)

  Hans P. Reiser (Univ. of Lisbon) addressed the problem of *proactive recovery*, a mechanism that periodically "revitalizes" a BFT system by refreshing the servers and eliminating faulty components. The talk argued that, due to intrinsic relations between a BFT system and recovery variables, proactive recovery must be implemented as an integral feature of a BFT algorithm, and not as an orthogonal complementary mechanism.

- *From Byzantine-Tolerant to Intrusion-Safe Services*. Christian Cachin (IBM Research - Zurich), Idit Keidar (Technion), and Alexander Shraer (Technion)

  The talk, presented by Christian Cachin (IBM Research - Zurich), proposed to combine BFT systems with the notion of fail-awareness [2], which ensures correct (wait-free and linearizable) operation as long as the service as whole is not compromised, and guarantees (weak) fork consistency and eventual fault detection otherwise.

- *Abstractions for Maintainable Byzantine Fault Tolerance*. Marko Vukolic (IBM Research - Zurich), Rachid Guerraoui (EPFL), and Vivien Quéma (INRIA)

  In this talk, Marko Vukolic (IBM Research - Zurich) described *ABortable STate mAChine Replication (ABSTRACT)* [9, 8], an abstraction that aims at facilitating the design of BFT protocols, and enables the use of existing protocols in the same computation and "on-the-fly" switching between the protocols, depending on the environment.

- *The Fault Detection Problem*. Andreas Haeberlen (Max Planck Institute for Software Systems) and Petr Kuznetsov (TU Berlin/Deutsche Telekom Laboratories)

In this talk, Andreas Haeberlen (MPI-SWS) described a formalism for *intermediate* fault models, more general than crash faults and more restrictive than the most general Byzantine faults. The formalism helped defining the *fault detection problem* for a given fault class, and determining lower bounds on message complexity of fault detection (find more in [10]).

- *Intrusion Tolerance: The killer application for BFT Protocols (?).* Alysson Neves Bessani, Miguel Correia, Paulo Sousa, Nuno Ferreira Neves, Paulo Verissimo (University of Lisbon, Faculty of Sciences, Portugal)

  The talk, presented by Alysson Neves Bessani (U Lisbon), argued that the killer application for BFT is (finally!) found, and it is intrusion-tolerance. However, in order to use BFT to resist attacks and intrusions we need to augment replication with diversity, recovery, accountability and confidentiality schemes. The challenge is, how to integrate all these mechanisms without increasing the complexity and the possible number of vulnerabilities of the system?

The announced talk by Allen Clement, "Experiences with the UpRight Library," did not, unfortunately, take place because the speaker had to cancel his presence at the last minute.

Summary

Our workshop contributed to understanding some of the difficulties in deploying techniques related to BFT. We believe that this kind of discussions help in bridging the gap between theory and practice, and may eventually lead to a common understanding that tolerating sophisticated non-crash faults in software systems is practically important, which may result in a wide application of BFT.

The workshop also contributed to discussing some of the limitations of current BFT-related research. Some of them, such as unrealistic assumptions and intractable complexity featuring existing BFT systems, were discussed in our workshop, and potential ways out were proposed. Also, we still do not have a clear indication of the very necessity of working in the (most pessimistic) Byzantine fault model. Would intermediate fault models that capture "relevant" non-crash faults be enough? How to define such fault models and how to devise fault-tolerant protocols in them are open questions, and we expect these questions to be of high practical importance.

Overall, the workshop was quite successful: we had a good collection of talks and interesting discussions. There was a short and inconclusive discussion about making this a regular workshop. Ironically, we would prefer BFTW3 not to become regular, since we hope that the "three W" questions raised in the title will eventually be answered, and there will be no need to revisit them. We hope that one day this research field will become not only intellectually satisfying but also practically relevant.

References

[1] Amazon S3 Availability Event: July 20, 2008. http://status.aws.amazon.com/s3-20080720.html, 2008.

[2] Christian Cachin, Idit Keidar, and Alexander Shraer. Fail-aware untrusted storage. In *Proc. Intl. Conference on Dependable Systems and Networks (DSN)*, June 2009.

[3] Miguel Castro and Barbara Liskov. Practical Byzantine Fault Tolerance and Proactive Recovery. *ACM Transactions on Computer Systems*, 20(4), 2002.

[4] Byung-Gon Chun, Petros Maniatis, Scott Shenker, and John Kubiatowicz. Attested Append-Only Memory: Making Adversaries Stick to their Word. In *Proceedings of ACM Symposium on Operating System Principles (SOSP)*, Stevenson, WA, USA, 2007.

[5] Giuseppe DeCandia, Deniz Hastorun, Madan Jampani, Gunavardhan Kakulapati, Avinash Lakshman, Alex Pilchin, Swaminathan Sivasubramanian, Peter Vosshall, and Werner Vogels. Dynamo: Amazon's highly available key-value store. In *Proceedings of ACM Symposium on Operating System Principles (SOSP)*, Stevenson, WA, USA, 2007.

[6] Danny Dolev and Andrew Chi-Chih Yao. On the security of public key protocols. *IEEE Transactions on Information Theory*, 29(2):198–207, 1983.

[7] Google. App Engine Outage today. http://groups.google.com/group/google-appengine/browse_thread/thread/f7ce559b3b8b303b?pli=1, 2008.

[8] Rachid Guerraoui. The next 700 bft protocols. In *OPODIS '08: Proceedings of the 12th International Conference on Principles of Distributed Systems*, pages 1–1, Berlin, Heidelberg, 2008. Springer-Verlag.

[9] Rachid Guerraoui, Vivien Quéma, and Marko Vukolić. The Next 700 BFT Protocols. Technical Report LPD-REPORT-2008-008, 2008.

[10] Andreas Haeberlen and Petr Kuznetsov. The fault detection problem. In *OPODIS*, 2009.

[11] James Hamilton. Internet-Scale Service Efficiency. In *Proceedings of 2nd Large-Scale Distributed Systems and Middleware Workshop (LADIS)*, New York, USA, 2008.

[12] Ramakrishna Kotla, Lorenzo Alvisi, Mike Dahlin, Allen Clement, and Edmund Wong. Zyzzyva: Speculative Byzantine Fault Tolerance. In *Proceedings of ACM Symposium on Operating System Principles (SOSP)*, Stevenson, WA, USA, 2007.

[13] Leslie Lamport, Robert Shostak, and Marshall Pease. The Byzantine Generals Problem. *ACM Transactions on Programming Languages and Systems*, 4(3):382–401, 1982.

[14] Netflix Blog. Shipping Delay. http://blog.netflix.com/2008/08/shipping-delay-recap.html, 2008.

[15] Rodrigo Rodrigues, Miguel Castro, and Barbara Liskov. BASE: Using abstraction to improve fault tolerance. In *Proceedings of ACM Symposium on Operating System Principles (SOSP)*, Banff, Canada, 2001.

[16] Franco Travostino and Randy Shoup. eBay's Scalability Odyssey: Growing and Evolving a Large eCommerce Site. In *Proceedings of 2nd Large-Scale Distributed Systems and Middleware Workshop (LADIS)*, New York, USA, 2008.

[17] Benjamin Vandiver, Hari Balakrishnan, Barbara Liskov, and Sam Madden. Tolerating Byzantine Faults in Database Systems using Commit Barrier Scheduling. In *Proceedings of ACM Symposium on Operating System Principles (SOSP)*, Stevenson, WA, USA, 2007.

[18] Jian Yin, Jean-Philippe Martin, Arun Venkataramani, Lorenzo Alvisi, and Mike Dahlin. Separating Agreement from Execution for Byzantine Fault Tolerant Services. In *Proceedings of ACM Symposium on Operating System Principles (SOSP)*, Bolton Landing, NY, USA, 2003.

Reliability and Security in Wireless Networks

Seth Gilbert
EPFL, Switzerland
seth.gilbert@epfl.ch

Dariusz R. Kowalski
University of Liverpool, UK
darek@liv.ac.uk

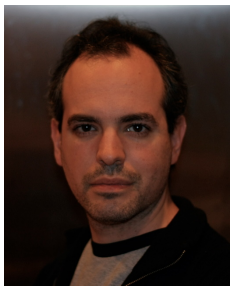

The Workshop on Reliability and Security in Wireless Networks was held on September 22, 2009, immediately prior to the 23rd Symposium on Distributed Computing in Elche, Spain.

Reliability and security are becoming increasingly important in wireless networks. The 21st century heralds the coming of the age of open airwaves. Over the last ten years, a plethora of new wireless applications have been developed and deployed to facilitate communication and to help people in their everyday lives. To a large extent, these applications rely on open airwaves, the unlicensed bands of spectrum, to accomplish their work.

However, the fact that the airwaves are open to all poses its own set of challenges that risk derailing the wireless revolution. The rapid growth of competing applications leads to a crowding of the airwaves, resulting in inadvertent interference among users. And the shared, open spectrum verily invites malcontents and attackers that may wish to disrupt applications and to corrupt data. Thus, it is critical that we address the problems of reliability and security in wireless networks.

The workshop began with a keynote talk by **Jean-Pierre Hubaux** (EPFL) on *Security and Non-Cooperative Behavior in Wireless Networks*. He surveyed the challenges caused by malicious and selfish users in wireless networks, focusing on "Chimp-Bonobo cross-layer schizophrenia," i.e., the interaction between cooperative and non-cooperative aspects of a wireless protocol stack. Hubaux emphasized the importance of game theory in understanding the behavior of users in wireless networks, and proposed several interesting directions. His talk was followed by a talk by **Stefan Schmid** (University of Paderborn) on *Dealing with Bad Vibes in Open Airwaves*, addressing the problem of jamming in wireless networks, and a talk by **Marek Klonowski** (Wrocław University of Technology) taking us *Towards Fair Leader Election in Wireless Networks*, addressing the problem of how to choose an honest leader, despite the presence of malicious participants.

The second session focused on the problems caused by contention on the wireless airwaves, beginning with a keynote talk by **Bogdan Chlebus** (University of Denver, Colorado) on *Conflict for Access Resolution*

and Adversity. He reviewed several different types of adversarial behavior, including both benign interference and malicious attacks. He presented a comprehensive overview of existing algorithmic techniques for coping with these challenges. He was followed by two talks: **Miguel Mosteiro** (LADyR-URJC and Rutgers) talked about *Unbounded Contention Resolution: k-Selection in Radio Networks*, addressing the problem of resolving conflicts among an unknown number of devices who want access to a single channel; and **Mariusz Rokicki** (University of Liverpool) talked about the *Complexity of Broadcast in Single-hop Radio Networks*, focusing on resolving conflicts in the dynamic case of continual online injection of new packets.

The third session began with a keynote talk by **Nitin Vaidya** (University of Illinois at Urbana-Champaign) on *Byzantine Faults in Wireless Networks*. He discussed how to design algorithms that could cope with malicious (i.e., Byzantine) participants, focusing on both the opportunities and challenges posed by the broadcast nature of wireless networks. He concluded his talk with some new results that apply network coding techniques to provide improved throughput (and improved capacity) in wireless networks subject to Byzantine failures. He was followed by **Henrique Moniz** (University of Lisbon) who presented his work on *Modeling Wireless Ad-hoc Networks for Efficient Fault- and Intrusion-Tolerant Protocols*. The workshop concluded with a talk by **Mauro Conti** (Vrije Universiteit) on his new project *S-MOBILE: Security of software and services for mobile systems*, which aims to simplify the development of secure software for mobile devices.

Overall, the workshop was quite successful in bringing together researchers with different perspectives on the challenges of reliability and security. The insightful talks led to many productive discussions, and we anticipate significant progress in the years ahead. *For more information on the workshop, see the web page*[1].

Theoretical Aspects of Dynamic Distributed Systems:

Report on the Workshop, Elche, Spain, September 26, 2009

Roberto Baldoni
Dip. di Informatica e Sistemistica "Antonio Ruberti"
Univ. degli Studi di Roma "La Sapienza", Roma, Italy

Alexander A. Shvartsman
Computer Science and Engineering
Univ. of Connecticut, Storrs, USA

This is a report on the Workshop on Theoretical Aspects of Dynamic Distributed Systems, TADDS, that was held in Elche on September 26, 2009. The workshop was co-located with the International Symposium on Distributed Computing, DISC 2009. The meeting had its focus on the dynamic aspects of distributed

[1] http://lpd.epfl.ch/sgilbert/workshop/

systems, encompassing systems in existence today and looking into the future development and deployment of dynamic distributed systems on the basis of sound theoretical foundations.

Design, development, and deployment of distributed systems are ubiquitous in modern computing. The nature of distributed systems is evolving in response to specific application requirements, and the advent of new classes of applications and technologies, such as VANET, Airborne Networks, Smart Environments, P2P, Broad Area Supercomputing, and Cloud Computing are radically changing the way we think about them. In these applications the initial set of processes is not known in advance, and the number of processes may change substantially over time as the system evolves due to changing needs of the applications and due to planned and unexpected perturbations in the computing medium. It is imperative for applications to accommodate such dynamic changes while maintaining their correctness. In order to manage such systems one cannot rely on a centralized management schemes that can either become a bottleneck or simply be unfeasible due to the size of the system. Instead, the processes need to behave autonomously, employing some form of self-management. In contrast with externally managed systems, in a dynamic distributed system there is no manager that guarantees a suitable environment for a period of time sufficient for the distributed system to accomplish its goal, while behaving correctly with respect to its system model assumptions. For example, in dynamic setting, it is hard or impossible to guarantee needed, sufficient, or appropriate entities that the application can assume in ensuring its correctness. Likewise, it is difficult to provide operational guarantees of QoS and the necessary degree of synchrony in the underlying distributed platform. In a dynamic distributed system each process must autonomously decides to locally run a component of a distributed application (i.e., when joining and leaving the system). Here the system and/or its components do not start with a known and pre-defined setting, and a "nice" manageable system model assumptions either cannot be guaranteed or do not last for long. It is also important to have measurable metrics and notions of efficiency that would allow one to analyze formally the performance and fault tolerance of dynamic systems when one cannot rely on the initial parameterization and a priori knowledge to characterize their behavior. Understanding the fundamentals of how to master this dynamic dimension is of primary importance to design of robust, dependable, and predictable distributed systems.

The workshop proceedings consisted of the introductory motivation presented by the organizers, four featured presentations, and it concluded with a panel discussion. The main presentations were given by worldwide recognized experts of distributed systems: Peter Druschel (Max Plank Institute), Leslie Lamport (Microsoft Research), Michel Raynal (University of Rennes and IRISA), and Maarten van Steen (VU Amsterdam). Each expert offered a view of the state-of-the-art achievements and discussed challenges and problems. The concluding panel with the presenters, led by the organizers, considered the following topics and questions. Dynamic distributed system: is this something really new or are we reinventing the wheel? Dynamic distributed systems vs. large scale systems: are these orthogonal concepts? What are the key attributes of dynamic distributed systems that are distinct from static systems? Do dynamic distributed systems need new specific theoretical foundations, models, complexity?

Michel Raynal focused on the differences between implementing fundamental abstractions (e.g., regular register) in a static distributed systems and in a dynamic distributed systems. He stressed the importance of providing rigorously correct solutions in a system prone to churn. Also he pointed out the importance of the notion of "adaptation" and he saw the resiliency to churn of a system as the way to adapt to system size dynamics, while scalability being a way to adapt to large scale systems. Leslie Lamport presented a grand tour of the approach using replicated state machines, and he offered a view of how to implement dynamic systems using a replicated state machine running on tens of machines within a cloud computing environment formed by thousands of machines. He underlined the importance of a DNS-like mechanism to locate the right state machine within a cloud. Despite large data center needs of some form of device autonomy

to self manage the system, state machines always need a certain degree of human management. Maarten van Steen introduced the case of extreme (ultra large) distributed systems consisting of increasingly miniaturized devices, pointing out how the application context impacts the system model. For example, extreme distributed systems on wired networks (such as IPTV distribution, p2p telephony, etc.) have negligible node churn while having heavy user churn. This difference could be used in the design phase of the application. Contrarily, node churn becomes a key notion in extreme wireless distributed systems where churn is generated by node mobility. Maarten concluded that more emphasis is needed on finding (validated) models that capture the behavior of our systems. Peter Druschel introduced the notion of accountability to detect and expose node faults in very large wired distributed systems consisting of full-featured components. Here accountability means maintaining a tamper-evident record in the system that contains a verifiable (digest of) of each node's actions. Based on this record, a faulty node whose observable behavior deviates from that of a correct node can be detected eventually. At the same time, a correct node can defend itself against any false accusations. Complete accountability scales to moderately large systems.

During the panel discussion concluding the workshop, there was general agreement that the study of sound theoretical foundation for dynamic distributed systems is one of the central upcoming research areas. Presenters also observed that in dynamic distributed systems one may need to relax strong deterministic guarantees in favor of probabilistic guarantees for system models and applications. Finally, the panel broadly concluded that subsequent meetings should be organized to encompass practitioners, system scientists, and theoreticians to foster collaboration that can, on one hand, launch theoreticians to create foundations based on real application scenarios and, on the other, let systems people better understand the scientific bases that may govern the behavior of their systems.

The presentations from the TADDS workshop can be found at the workshop web page[1].

Game-Theoretic Aspects of Distributed Computing

Chryssis Georgiou
Dept. of Computer Science, University of Cyprus
1678 Nicosia, Cyprus
chryssis@cs.ucy.ac.cy

The Workshop on Game-Theoretic Aspects of Distributed Computing took place on September 26th and it was co-located with DISC 2009[1]. The workshop was organized by Chryssis Georgiou (University of

[1] http://disc2009.gsyc.es/workshops
[1] http://disc2009.gsyc.es

Figure 1: Some of the workshop participants. Picture courtesy of Antonio Fernandez Anta.

Cyprus, Cyprus) and Paul Spirakis (RACTI and University of Patras, Greece) and it was sponsored by the University of Cyprus[2] and the Comunidad de Madrid[3]. Further details can be obtained from the Workshop's website[4].

Workshop's Objective

In traditional Distributed Computing, the behavior of the system components (i.e., processors/processes/nodes/agents) is characterized *a priori* as either "good" or "bad", depending on whether they follow the prescribed protocol or not. In Game Theory, players are assumed to be *rational* or *selfish*, that is, they act on their own self-interest and they do not have an *a priori* established behavior. In other words, the players decide on how to act in an attempt to increase their own benefit (a quantified measure).

Game Theory has long been considered in many fields, ranging from Economics to Law Enforcement and Voting Decision. With the evolution of the Internet, Game Theory has found many applications in Networks, and in Distributed Computing in general. Such examples include Internet Routing, Resource/Facility Location and Sharing, Containment of Viruses Spreading, Secret Sharing, and Web-based Task Computations.

The purpose of the workshop was to present recent works that consider Distributed Computing issues from a Game-Theoretic view and approach. The *main objective* of the workshop was to enable members of the Distributed Computing community (especially students and junior researchers) to realize the potential of what it can be called *Game-Theoretic Distributed Computing*.

The workshop included five invited talks delivered by prominent researchers working on Game-Theoretic Distributed Computing issues: *Ioannis Caragiannis* (University of Patras, Greece), *Stefano Leonardi* (Sapienza University of Rome, Greece), *Giuseppe (Pino) Persiano* (Universita di Salerno, Italy), *Maria Serna* (Universitat Politecnica de Catalunya, Spain), and *Paul Spirakis* (RACTI and University of Patras, Greece). Figure 1 shows the workshop's invited speakers, organizers and some of the participants.

Overview of the Invited Talks

We give an overview of the five invited talks in the order they were presented.

[2] http://www.ucy.ac.cy
[3] http://www.madrid.org
[4] http://sites.google.com/site/wgtadc

***Strategic Games to Analyze Unreliable Systems*, by Maria Serna.** Being the first speaker, Maria first presented some fundamental terminology of Game Theory (including strategic games, pure and mixed strategies, utilities and cost functions, and Nash equilibria). Then she proceeded to the main topic of the talk, which is a work performed jointly with J. Gabarro, A. Garcia, P. Kilpatricka and A. Stewart. They formulated and analyzed, from a game-theoretic view, two problems arising in web or grid environments with unreliable behavior: service failures during grid orchestrations and resource allocation in grids.

Maria first presented the idea behind the problems' formulation. A service failure may be catastrophic in that it causes an entire grid application to fail. Alternatively, a grid manager may utilize alternative services in the case of a failure, allowing an orchestration to recover. The proposed approach attempts to provide an alternative for the case of bounded number of failures in between optimistic and pessimistic situations and encapsulates the idea of *limited failures* and *destructive* versus *not destructive* behavior. The talk concentrated in dealing with failures prior to the execution of the tasks; such failures can be anticipated and modeled by what they call *risk profile*. A risk profile is the mean of modeling situations in a way that is neither overly optimistic nor overly pessimistic. Risk profiles provide the opportunity of defining a class of games that they call *Angel and Daemon* games, in which the angel controls non-malicious failures while the daemon controls malicious failures. Then Maria presented the two problems under investigation and showed how they can be modeled as Angel-Daemon games. She concluded the talk with the presentation of some results on the existence of pure Nash equilibria and the complexity of the corresponding computational problems.

***Efficient Coordination Mechanisms for Selfish Scheduling*, by Ioannis Caragiannis.** Ioannis presented *coordination mechanisms* for scheduling selfish jobs on unrelated machines. As Ioannis explained, a coordination mechanism aims to mitigate the impact of selfishness of jobs on the efficiency of schedules by defining a local scheduling policy on each machine. These scheduling policies induce a game among the jobs and each job prefers to be scheduled on a machine so that its completion time is minimum given the assignments of the other jobs; the maximum completion time among all jobs is considered as the measure of the efficiency of schedules. The approximation ratio of a coordination mechanism quantifies the efficiency of pure Nash equilibria (price of anarchy) of the induced game.

The talk covered simple strongly local and non-preemptive coordination mechanisms (such as Shortest-First, LongerstFirst, and Randomized) as well as recent preemptive ones (such as ACOORD, BCOORD, and CCOORD) that achieve logarithmic or polylogarithmic (in terms of the number of the machines) approximation ratio. These preemptive coordination mechanisms make use of m*-efficient assignments*. In such assignments, a job is never assigned to a machine where its inefficiency is more than m. The talk concluded with a discussion on open questions. For example, is it possible to obtain constant approximation ratio? Is the knowledge of m really necessary? How about mixed Nash equilibria?

***Routing Traffic Through Selfish Capacitated Links*, by Giuseppe (Pino) Persiano.** Pino's talk considered the problem of *routing traffic* over communication links that are owned by selfish agents; agents are selfish in the sense that might misreport link characteristics to the routing algorithm so to induce a better (for the agent) allocation of traffic to links. As Pino pointed out, this is a classical problem in *mechanism design* where one seeks payments for the agents so to make it profitable to report the true characteristics of the links (that is, make the agents *truthful*), and thus the routing algorithm can work on real data.

In the talk, Pino presented a recent work, performed jointly with V. Auletta and P. Penna, where they study the novel scenario in which links have traffic capacities. As Pino explained in detail, the obtained results are in stark contrast with the known results for another important case of agents with capacities

(bidders with budget constraints). Several reseach questions on this topic remain open, for some of which Pino gave insight on possible directions in addressing them.

***Utilitarian Mechanism Design for Multi-Objective Optimization*, by Stefano Leonardi.** Stefano presented recent work, performed jointly with F. Grandoni, P. Krysta and C. Ventre, on algorithmic mechanism design for *multi-objective optimization* problems. In a classic optimization problem the complete input data is assumed to be known to the algorithm. As Stefano pointed out, this assumption may not be true anymore in optimization problems motivated by the Internet where part of the input data is private knowledge of independent selfish agents. The goal of *algorithmic mechanism design* is to provide (in polynomial time) a solution to the optimization problem and a set of incentives for the agents such that disclosing the input data is a dominant strategy for the agents. In case of NP-hard problems, the solution computed should also be a good approximation of the optimum.

In the talk, Stefano focused on mechanism design for multi-objective optimization problems. In this setting a main objective function is given, along with a set of secondary objectives which are modeled via budget constraints. Multi-objective optimization is a natural setting for mechanism design as many economical choices ask for a compromise between different, partially conflicting, goals. The main contribution of this work (which Stefano explained rigorously) is showing that two of the main tools for the design of approximation algorithms for multiobjective optimization problems, namely *approximate Pareto curves* and *Lagrangian relaxation*, can lead to truthful approximation schemes.

***Learning and Enforcement in Competitive Environments*, by Paul Spirakis.** Paul's talk considered the problem of efficiently *learning and enforcing* equilibria strategies in competitive environments. Paul first presented basic notions of Machine Learning and discussed what can be learned in competitive environments consisting of selfish agents. In particular, he showed that *correlated equilibria strategies* can be learned efficiently in such environments. Furthermore, he demonstrated that Nash equilibria can be learned in the special case of zero sum bimatrix games. Then the talk focused on *repeated games* of many agents and discussed the issue of collaborative enforcement of individual behaviours that lead to equilibria play. The notions of *regret*, of a *threat point* and of a *correlated threat point* were defined and explained through several examples. The talk concluded by demonstrating how the notion of a correlated threat point can help in efficient enforcement of equilibria strategies.

Acknowledgements

Paul Spirakis and I would like to thank the DISC'09 Organization Committee chairs, *Antonio Fernández Anta* (Universidad Rey Juan Carlos, Spain) and *Vicent Cholvi Juan* (Universitat Jaume I, Spain) for their cordial help and support. We especially appreciate their involvement in obtaining sponsorship from the Comunidad de Madrid. Also, I would like to thank *Idit Keidar* (Technion, Israel) for inviting me to write this report.

Review of SIROCCO 2009

Shantanu Das
ETH Zürich, Switzerland
shantanu.das@inf.ethz.ch

The 16th edition of the *International Colloquium on Structural Information and Communication Complexity*, SIROCCO 2009, took place during May 25-27, 2009 in Piran, Slovenia. SIROCCO is a conference on distributed computing in graphs and networks with a special focus on the interplay between structural knowledge, local computation and the efficiency of global communication in distributed environments. On one hand, the focus on a smaller domain within the area of distributed computing, makes this conference a venue for close interactions between researchers specializing in this particular domain. On the other hand, this also leaves enough room to accommodate new ideas and unconventional results which expand the frontiers of the field of distributed computing. This edition of the conference was a lively event consisting of two invited talks, 22 regular presentations and four poster presentations, covering a range of topics including wireless computing, autonomous robots, self stabilization, population protocols, software transactional memory, graph explorations and graph coloring, to name a few. Collocated with the conference was the one-day workshop **IMAGINE 2009** (*International Workshop on Mobility, Algorithms, Graph Theory in Dynamic Networks*) which is a yearly event showcasing the work of junior researchers mainly in the area of distributed algorithms for dynamic networks.

The 2009 edition of SIROCCO was special due to two reasons. First, this year marked the institution of the SIROCCO *Prize for Innovation in Distributed Computing*, an award that recognizes original and innovative research that expands the frontier of conventional research specifically in the areas related to SIROCCO (more about this in Section 3). Another novel initiative started this year at SIROCCO was the presentation of posters in addition to the regular program of talks.

This year, the chosen venue was Slovenia, the 11th country that SIROCCO has been to, in the sixteen years of its existence. The peaceful and picturesque town of Piran on the Slovenian coast proved to be an ideal venue for enjoying the series of interesting talks and informal discussions in a relaxed atmosphere. The conference was held in Hotel Barbara Fiesa, located directly on the coast, and all the participants were accommodated at the same hotel, so that the casual discussions continued from the dinner table to the breakfast buffet. At the end of the three days of the conference, each of us returned with fond memories of the Slovenian coast, some interesting facts about the process of salt-making (see Section 2) and a bottle of pure Slovenian honey in our backpacks.

In the following, I discuss the excellent program of talks at the conference, the fun-filled social events and finally, the SIROCCO prize that was awarded for the first time this year.

1 The Technical Program

The technical program started on the morning of 25th, with a warm welcome note by the organizing chair Janez Žerovnik, followed by a brief opening speech by Shay Kutten, the Program Committee chair. The first invited speaker was Israel Cidon from the Technion - Israel Institute of Technology, who presented an interesting overview of the emerging research field of Network-on-Chip (NoC) in his talk titled "*Zooming in on Network-on-Chip Architectures*". The NoC paradigm advocates the design of VLSI chips using techniques from the field of networking to implement the communication between submodules. Prof. Cidon pointed out that many of the research problems for telecommunication networks, such as routing, congestion control, topology design, etc. appear also in the field of NoC; However, in light of the special constraints imposed by the characteristics of silicon chips, new solutions and techniques are required for these traditional networking problems. This opens up an interesting avenue of research for the distributed computing community and the speaker left us convinced that new innovations from the community would shape the future of this emerging area of research.

The second day of the conference began with the invited talk "*Efficient gossiping methods in wireless/radio networks*" presented by Leszek Gasieniec from the University of Liverpool. The speaker gave a very entertaining talk on broadcasting and gossiping in wireless radio networks. In radio networks, communication is achieved by each node broadcasting a message to all its neighbors. Due to interference between transmissions by nearby nodes, there can be collisions, leading to retransmissions. Therefore clever algorithms are needed that achieve global dissemination of information within a small number of transmission rounds. The talk surveyed the recent results on efficient communication in radio networks and provided some ideas of the key techniques used for time and energy efficient communication in such networks. It was highly motivating to learn about the active collaborative research being conducted in this area, as scientists compete with each-other to come up with more and more efficient algorithms. The talk also sketched problems that are still left open and where further improvements are possible. The presentation ended on a light note with a joke[1] about wireless communication which everyone appreciated.

Figure 1: (a) The participants of SIROCCO 2009. (b) Discussions over coffee.

The regular presentations took place over several small sessions of three to four talks, with enough time between sessions for discussions over coffee. During the coffee breaks, the participants could also examine

[1] http://thinkmorethunk.typepad.com/think_more_thunk_ideas_in/2007/04/a_wireless_joke.html

the posters on display at the lobby just outside the conference hall. There were eight sessions in total and from among the many interesting results that were presented, I will attempt to sketch ideas from some of them. The paper titled "*An Optimal Bit Complexity Randomized Distributed MIS Algorithm*" by Métivier, Robson, Saheb-Djahromi and Zemmari, presented an algorithm for distributed *Maximum Independent Set* construction that is optimal in terms of bit complexity. This randomized algorithm uses a technique of local exchange of bits among processors, thus reducing the bit complexity to $O(\log n)$ instead of $O(\log^2 n)$ (as in previous algorithms), without changing the asymptotically optimal time complexity of $O(\log n)$.

The talk by Moti Medina on "*Revisiting Randomized Parallel Load Balancing Algorithms*" was an insightful study of the well-known problem of parallel load balancing (or allocating balls to bins). The authors presented a nice analysis of the previous algorithms by Adler *et al.* and finally provided a simpler algorithm for the problem using some clever techniques.

There were several presentations on network routing. Karol Suchan presented schemes for computing short routes in k-chordal graphs, the class of graphs containing no chordless cycle of length more than k. For computing routes in such topologies, the presentation discussed the tradeoffs between the length of the routes and the time required to compute them. The paper titled "*Designing Hypergraph Layouts to GMPLS Routing Strategies*" presented by Ignasi Sau, considered the problem of routing in optical networks using label switching, with the objective of minimizing the number of labels used. Some hardness results were shown and approximation algorithms were presented both for general graphs and for the special case of a path. In a related presentation, Prudence Wong showed how to minimize the number of ADM switches for routing in optical networks (they considered the star topology in particular), by relating the problem to that of maximum matching in multipartite graphs.

Another topic which featured prominently in the conference was the problem of graph (or network) exploration. Jesper Jansson presented a fascinating talk on the periodic traversal of anonymous graphs. The problem considered was the assignment of port-labellings in an undirected graph such that a finite automaton can explore the graph with the smallest possible period of exploration. This problem was initially presented at the same conference four years ago where a period of $10n$ was achieved for one-state automata exploring arbitrary graphs having n vertices. The current paper significantly improved that result to obtain a period slightly over $4n$ and at the same time presented a lower bound of $2.8n$. The exploration of grid graphs was considered in the presentation by Adrian Kosowski. Frank Petit presented a probabilistic algorithm for exploration of ring networks with multiple agents. Stefan Dobrev's presentation on exploring and searching for a black hole in directed networks, received a lot of attention from the audience.

There were two papers on autonomous mobile robots. The first result presented by Kenta Yakamoto, studied the convergence problem for robots having inaccurate sensors when there is a known margin of observation-error. This robot model is based on practical considerations and was initially proposed in a paper by Cohen and Peleg, where the impossibility of convergence was shown for measurement-errors greater than $\pi/3$. The paper by Yakamoto and others showed that under the uniform error model (i.e. when the measurement error in each round is uniform across all the active robots), it possible to achieve convergence with an error margin of anything less than $\pi/2$. The second talk on mobile robots presented by Yann Disser, was very entertaining and considered robots having limited visual capabilities. The paper showed which capabilities are sufficient for a robot moving on a polygon to construct the visibility graph of the polygon.

There was also a lot of interest on *population protocols* during this edition of the conference with three presentations on the subject. The model of population protocol, initially introduced by Angluin and others, is a computational model based on pairwise interactions within a set of finite-state agents. The paper titled "*Space Complexity of Self-Stabilizing Leader Election in Passively-Mobile Anonymous Agents*" by Cai,

Izumi and Wada, showed that more than $n - 1$ states are necessary for any self-stabilizing leader election protocol for n agents in the proposed model, while n states are sufficient for achieving the same. In a related paper by Sudo, Nakamura, Yamauchi, Ooshita, Kakugawa and Masuzawa, the same problem was considered under a relaxed condition which the authors introduce as the notion of *loose stabilzation*. The main result shown in the paper is that the knowledge of the network size is unnecessary for a leader election protocol under the modified requirement. The paper "*On Gossip and Populations*" presented by Yann Busnel, provided an interesting discussion on the relationship between population protocols and gossiping in networks.

These were by no means the only interesting results presented at the conference and there were several other, equally fascinating presentations which are omitted for the sake of keeping this review concise.

2 The Social Event and the Banquet

Figure 2: Excursions: (a) The Saltpans of Sečovlje and (b) The Church of the Holy Trinity.

The conference program included several fun events to complement the interesting series of talks. The afternoon of the second day was reserved for some exiting social events starting with a trip to the *Sečovlje Salina Nature Park*. We were told that this is one of the few rare saltpans where salt is still produced in the traditional way. After a tour of the facility, we were shown a documentary film on the salt-making process and the unique ecosystem of plants, animals and birds which has developed over 700 years on these saltpans. The trip to the saltpans was followed by a visit to one of the traditional Slovenian vineyards. The hostess welcomed us and explained how wine is produced and matured in the large wooden containers kept in the cellars. We then sat down to taste some of the nicest wines produced in Slovenia.

The afternoon was still not over and the organizers had yet another excursion planned for us. This time we were introduced to the cultural and religious history of Slovenia as we visited a fifteenth century church located just above the village of *Hrastovlje*. This medieval church, named *The Church of the Holy Trinity*, is famous for the frescos decorating its inner walls and ceilings which have been carefully preserved over the years. A short audio presentation provided some interesting facts regarding the significance of these medieval paintings. The evening ended with a special dinner banquet at a classic Slovenian restaurant. During the banquet, the SIROCCO prize was presented and a vote of thanks was given to the organizers and contributors to the conference.

Figure 3: (a) The Dinner Banquet. (b) Leszek Gasieniec congratulates Nicola Santoro on receiving the prize.

3 The SIROCCO Prize

The *Prize for Innovation In Distributed Computing* is awarded at the SIROCCO conference to honor significant research contributions on topics of interest to SIROCCO, which have "expanded the collective investigative horizon by formulating new problems, or identifying new research areas, that were at the time of their introduction unorthodox and outside the mainstream". This year the award was presented to **Prof. Nicola Santoro** from Carleton University, Canada, for his work on investigating labelled graph properties which led to the formalization of some key concepts such as *implicit routing*, *sense of direction* and *topological awareness*. The full award citation appears earlier in this column.

The award was presented by Pierre Fraigniaud, the Steering Committee chair of SIROCCO, who praised the efforts of Prof. Santoro in enriching the field of Distributed Computing with several innovative concepts and ideas. In his brief acceptance speech, Prof. Santoro encouraged the junior researchers to keep pursuing research on topics that interest them the most, even if these ideas are not immediately accepted or adopted by the community. He reminded us that SIROCCO has always been an avenue for such unorthodox and innovative research ideas which have the potential of making a significant impact on the future of distributed computing.

4 The Next Edition

As announced in the business meeting of the conference, the next edition of SIROCCO will be held during June 7–11, 2010 in the *Nesin Mathematics Village* situated on a beautiful hillside about 1 km away from Sirince in Turkey[2]. See you in Turkey next year!

Acknowledgements

Many thanks to the Organizing Committee members Janez Žerovnik, Igor Pesek and Petra Šparl for making this a memorable event and a special thanks to Igor for the pictures from the conference.

[2] http://www.sirocco2010.boun.edu.tr/