

A Scientific Story (1985-2025) from WDAG to DISC

Michel Raynal^{*} & Nicola Santoro[◇]

^{*} Univ Rennes (IRISA, Inria, CNRS) France

[◇] Carleton University, Ottawa, Canada

1 What is Distributed Computing?

Distributed computing is the science of cooperation. More precisely a distributed algorithm is composed of a set of computing entities (imposed to the programmers) each providing its own input, that have to cooperate to a common goal depending on the set of inputs (usually in the presence of adversaries such as asynchrony and failures) [2, 3].

. It follows that, in distributed computing, possibility/impossibility results and algorithmic techniques differ from what is encountered in sequential or parallel computing. This characterizes the distributed computing community. DISC and PODC are the very top conferences devoted to distributed computing.

2 Distributed Computing vs Parallel Computing

Distributed computing is sometimes confused with a form of parallel computing. Roughly speaking parallelism is an extension of sequential computing the aim of which aim is to obtain efficient programs. Let us notice that any parallel algorithm could be executed sequentially on a mono-processor(of course in a very inefficient way). This is not the case in distributed computing. Both distributed computing and parallel computing involve concurrency.

3 At the Very Beginning: When DISC was WDAG

3.1 The Start

What is now named DISC began in 1985, a sibling to the 4-years old PODC that, at the time, was the official (and, until then, the only) outlet for the theoretical aspects of the field. Like PODC (and later SIROCCO), the new conference started in Ottawa; created by Nicola Santoro with the collaboration of Eli Gafni, its focus was on the "algorithmic" aspects of distributed computing, and its name at birth was *International Workshop on Distributed Algorithms on Graphs* (hence the acronym *WDAG*¹).

The creation of WDAG was motivated by a strong sense of dissatisfaction felt in regards with what was perceived as an unbalanced official "coverage" of the distributed algorithms research efforts and outputs, and a restricted view of what was important in distributed computing. In this view, the most important concern in the field appeared to be not on general problem solving by the distributed computational entities composing the system but rather on coping with their failures. In other words, unlike practically all the germane computational fields (e.g., sequential computing, parallel computing, computational geometry, etc) the main concern from the start was on fault-tolerance, and the core problems investigated were obviously *consensus* and its many reformulations. Furthermore, the model used in the research almost exclusively assumed the system to be *fully synchronous* and the communication topology of the entities to be the *complete graph*. Most of the algorithmic focus was on the solution of these problems. This in spite of the fact that the majority of problems investigated at the time (especially in the European community and by researchers from combinatorics and discrete mathematics) were about networks and graphs (e.g., spanning-tree construction, routing, etc.), information diffusion (e.g., broadcast, converge-cast, gossip, etc.), control (e.g., symmetry breaking, election, mutual exclusion, etc.), where the concern was to determine the minimal system conditions necessary first of all for solving the problem (solving it in presence of faults was a subsequent concern). Interestingly, most of these investigations were concerned with the more realistic assumptions of *asynchronous* systems and (often unknown) *different graph topologies*.

It was based on this perceived divide that the decision to start WDAG was made, and to make the difference more explicit, the words "Algorithms" and "Graphs" were added to the name. So a new outlet for the DC community with a slightly different view was born (see Photo 1), and its (hard cover!) proceedings, edited by Eli Gafni (EG) and Nicola Santoro (NS), were published in 1986 by

¹A typical mistake has been to assume that the "G" in the acronym was part of the word "algorithm".

Carleton University Press.



Figure 1: Participants to the first WDAG. Among them: Jan van Leeuwen, Christian Lavault, Jan Pachl, Joe Peters, Shmuel Zaks, Yehuda Afek, and ... PhD student Shay Kutten.

3.2 The WDAG Years: 1985-1998

Following the very positive response to the first WDAG, the decision to continue (initially on a biennial basis) was the natural result. The "de facto" steering committee (including now Jan van Leeuwen) decided the location of the workshop to be moved to Europe and, in order to get a wide dissemination, the proceedings to be published by Springer LNCS. So, since its second edition (1987), WDAG became an European event, eventually under the official sponsorship of EATCS (the European Association for Theoretical Computer Science, see below).

Interestingly, with each edition, new members were permanently co-opted into the de-facto steering committee which became a large collaborative collective (a mini distributed system); in particular, starting with Michel Raynal (WDAG 3), we had Sam Toueg and Paul Spirakis (WDAG 5), Shmuel Zaks (WDAG 6), and Paul Vitanyi (WDAG 8) all join the SC collective.

In all these years, the differences between WDAG and PODC were several, from their organizative structure and the choice of geographical locations (Europe for WDAG and Canada at that time for PODC), to the emphasis on content (algorithmic for WDAG and general principles for PODC); a fundamental distinction remained for quite a long time their basic "nature": one was a *workshop*, while the other a *symposium*. The latter difference went beyond the words used; it was



Figure 2: The authors of this paper at WDAG 3.

rather felt and experienced by the participants. Slowly over the years, the differences started to disappear.

3.3 From WDAG to DISC: 1998

Eventually, according to the papers submitted (and accepted) and the scientific interests of the people attending the workshop, it appeared that the meaning of the words *Workshop* and *Graphs* was too restrictive and no longer conveyed the spirit of this scientific meeting, which had become much broader.

So, in 1998, the PC chair Shay Kutten, proposed to replace the name *Workshop*



(a) WDAG 4: NS, Shmuel Zaks, Sergio Rajsbaum, Shlomo Moran.



(b) WDAG 9: Hagit Attiya, Shmuel Zaks, NS.

on Distributed Algorithms on Graphs by International Symposium on Distributed Computing, in short DISC. The Business meeting fully agreed on this proposal, and WDAG 12 became DISC 12.



Figure 4: WDAG 12 = DISC 12.

The reader will find the full list of WDAG workshops and DISC symposia (with years, places and programs committee chairs) at “dblp DISC”.



Figure 5: A historical piece: the very first DISC T-shirt

4 DISC 20th Anniversary

This venue of DISC was Stockholm. To celebrate this 20th anniversary of the DISC Symposia three famous speakers were invited to give a talk, namely Leslie Lamport, Nancy Lynch, and Michael Rabin.

5 Sharing Dijkstra Award with PODC

The Edsger W. Dijkstra Prize in Distributed Computing is named for Edsger Wybe Dijkstra (1930-2002), a pioneer in the area of distributed computing. His foundational work on concurrency, semaphores, mutual exclusion, deadlock, finding shortest paths in graphs, fault-tolerance, self-stabilization, among many other contributions comprises one of the most important supports upon which the field of distributed computing is built. No other individual has had a larger influence on research in principles of distributed computing. The prize is given for outstanding papers on the principles of distributed computing, whose significance and impact on the theory and/or practice of distributed computing has been evident for at least a decade.

The Prize was initially sponsored by the ACM Symposium on Principles of Distributed Computing (PODC). In 2003 the chair and the vice-chair (Alex Shvartsman) of the steering committee of DISC initiated a procedure for the prize to be awarded by both PODC and DISC. The current chair of the PODC steering committee (Mark Tuttle) agreed, but as PODC depends on ACM, the head of ACM required DISC to be part of an official organisation. After a short discussion the head of EATCS agreed to support DISC and since then the award is presented annually, with the presentation taking place alternately at ACM PODC and EATCS DISC.

6 From Springer LNCS to LZI LIPIcs

An important change appeared in 2017 (31st DISC) where the steering committee decided to have the proceedings produced by LIPIcs (Leibniz International Proceedings in Informatics), that is a series of high-quality conference proceedings across all fields in informatics established in cooperation with Schloss Dagstuhl - Leibniz Center for Informatics. Among other points, this favoured free open access to all the DISC articles.

7 Which Future?

Prediction is very difficult, especially about the future.
Niels Bohr (1885–1962).

8 The Guardians of the DISC Temple

(Past Chairs of the DISC Steering Committee)

Sam Toueg, Cornell, USA, 1996-1998
Shmuel Zaks, Technion, Israel, 1998-2000
André Schiper, EPFL, Switzerland, 2000-2002
Michel Raynal, Irisa, France, 2002-2004
Alex Shvartsman, University of Connecticut, 2004-2007
Rachid Guerraoui, EPFL, Switzerland, 2007-20098
Nicola Santoro, Carleton University, Canada, 2009-2011
Sergio Rajsbaum, UNAM, Mexico, 2011-20138
Antonio Fernandez Anta , IMDEA Netw., Spain, 2013-2015
Shlomi Dolev, Ben Gurion University, Israel, 2015-2017
Roberto Baldoni, Sapienza Universita di Roma, Italy, 2017-2018
Yoram Moses, Technion, Israel, 2018-2020
Andréa Richa, Arizona State University, USA, 2020-2022
Jukka Suomela, Aalto University, Finland, 2022-2024
Hagit Attiya, Technion, Israel 2024-2026.

9 What is a Good Conference?

Of course this is an opinion and not a definition. A good conference is located in a nice place, and is the product of good submissions and good referees. It is also a place where people discuss on their reseach topics (or any other topic!) and where we can meet scientific people among which some of them will become good friends² with which we have discussions that can last several years, improving our view of the topics we are working on, etc. In some sense, a good conference is like a (scientific) family meeting composed of researchers interested in common topics.

²This notion of friendship has nothing to do the notion of “friend” as imposed by Facebook.

10 What is a Good Paper?

This is an important question. The answer to this question given below is from [2]. At the very beginning (when I (MR) was younger, i.e. in the previous millennium!) my answer was mainly based on an objective numerical criterion, namely, a good paper is a paper with numerous citations”. Later, I was saying “a paper that won the best paper award in a top conference”. Still later I was saying “a paper that won a prize devoted to more than ten years old papers”, etc. But over time, none of these integer-based definitions fully satisfied me, and I started thinking about the papers that I myself consider as very important papers ... and I discovered that those were papers I was a little bit kindly jealous ... not to be a co-author! This was because those are papers I like to read (and reread) because they are nicely written, their ideas go beyond their technical content, they introduce new ideas in a simple and efficient way, and have a very strong impact on the community. This is the effect of good papers: everyone makes them ”theirs”, assimilating them and passing their essence to students.

Finally, to conclude, considering that the pair research/teaching defines the golden coin that lies at the heart of Universities all over the world, and as research feeds teaching (and vice-versa) it is important to consider that “Teaching is not an accumulation of facts” as stated by L. Lamport in [1].

References

- [1] Lamport L., Teaching concurrency. *ACM Sigact NEWS*, 40(1):58-62 (2009)
- [2] Raynal M., About informatics, distributed Computing, and our Job: a personal view. *Proc. 30th Int’l Colloquium on Structural Information and Communication Complexity (SIROCCO’23)*, Springer LNCS 13892, pp. 44-56 (2023)
- [3] Raynal M., Mutual exclusion vs consensus: both sides of the same coin? *Bulletin of the European Association of Theoretical Computer Science (EATCS)*, Vol.140, 14 pages (2023)