

2011 Edsger W. Dijkstra Prize in Distributed Computing

We are proud to announce that the 2011 Edsger W. Dijkstra Prize in Distributed Computing is awarded to

Hagit Attiya, Amotz Bar-Noy, and Danny Dolev,

for their paper

Sharing Memory Robustly in Message-Passing Systems

which appeared in the Journal of the ACM (JACM) 42(1):124–142 (1995).

This fundamental paper presents the first asynchronous fault-tolerant simulation of shared memory in a message passing system.

In 1985 Fischer, Lynch, and Paterson (FLP) proved that consensus is not attainable in asynchronous message passing systems with failures. In the following years, Loui and Abu-Amara, and Herlihy proved that solving consensus in an asynchronous shared memory system is harder than implementing a read/write register. However, it was not known whether read/write registers are implementable in asynchronous message-passing systems with failures. Hagit Attiya, Amotz Bar-Noy, and Danny Dolev's paper (ABD) answered this fundamental question affirmatively.

The ABD paper makes an important foundational contribution by allowing one to “view the shared-memory model as a higher-level language for designing algorithms in asynchronous distributed systems.” In a manner similar to that in which Turing Machines were proved equivalent to Random Access Memory, ABD proved that the implications of FLP are not idiosyncratic to a particular model. Impossibility results and lower bounds can be proved in the higher-level shared memory model, and then automatically translated to message passing. Besides FLP, this includes results such as the impossibility of k -set agreement in asynchronous systems as well as various weakest failure detector constructions. On the other hand, shared-memory algorithms, e.g., concurrent time-stamp systems, l -exclusion algorithms, atomic snapshots, consensus objects, and implementations of various other data structures, can all be directly be implemented in a message passing environment. An interesting example is the widely used Paxos algorithm of Lamport: it is significantly easier to understand the algorithm in shared memory, knowing that the actual Paxos protocol in a message passing system can be directly derived using ABD.

The ABD transformation makes use of an arsenal of common-sense solution ingredients, which have since become the bread-and-butter of distributed computing. For example, for fault tolerance, the transformation uses quorums, which are well known tools for high availability. Quorums are used to guarantee and non-empty intersection among the sets of concurrent readers and writers. But the true insight in the paper was the guarantee of intersection among successive readers by writing-back the value in a second round. Unique logical timestamps were used to form a global order, which does not violate causality. Another important component of the paper is to specify a distributed system using a 'service oriented' approach: One specifies the service from the client's perspective, while hiding the internal details of the service implementation and the handling of faults.

Since its publication, the Attiya, Bar-Noy, and Dolev's paper has impacted both the theory and the practice of distributed systems. After ABD, logical distributed storage systems like the Federated Array of Bricks (FAB) were built, which provide "the reliability and performance of enterprise-class disk arrays at a fraction of the cost and with better scalability." The agility of distributed storage was further demonstrated by the RAMBO system, which provided storage over a dynamic set of disks. More recently, it was even shown that this type of dynamic storage might be achieved without synchrony at all. ABD inspired algorithms for sharing data in sensor networks, as well as research exploring extensions of the basic fail-stop fault model to Byzantine failures via Byzantine Quorum Systems. Additionally, the fundamental two-round structure of the ABD solution was analyzed in a series of papers on the round complexity of read/write operations in a variety of fault models.

Finally, we note that over the past 20 years, Distributed Computing Research has become more specialized, with many researchers working either purely on message passing systems or purely on shared memory algorithms. Attiya, Bar-Noy, and Dolev tell us not to worry about this, in the end, what we are doing is pretty much the same.

Award Committee Dijkstra Prize 2011

Shlomi Dolev

Pierre Fraigniaud (Chair)

Dahlia Malkhi

David Peleg

Nir Shavit

Jennifer Welch