

2009 Edsger W. Dijkstra Prize in Distributed Computing

The 2009 Edsger W. Dijkstra Prize in Distributed Computing was awarded to the paper “Knowledge and Common Knowledge in a Distributed Environment” by Joseph Halpern and Yoram Moses. The presentation of the award was made at DISC’09.

The Edsger W. Dijkstra Prize in Distributed Computing is awarded for an outstanding paper 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 Dijkstra Award Committee has selected Joseph Halpern and Yoram Moses as the recipients of this year’s Edsger W. Dijkstra Prize in Distributed Computing. The prize is given to them for their outstanding paper: “Knowledge and Common Knowledge in a Distributed Environment” published in *Proceedings of the Third Annual ACM Symposium on Principles of Distributed Computing (PODC’84)* pp. 50-61, 1984 ([pdf](#)), and in *Journal of the ACM (JACM)*, 37(3):549-587, 1990 ([pdf](#)).

The “Knowledge and Common Knowledge in a Distributed Environment” paper presented by Halpern and Moses in PODC 1984 provided an effective new way of reasoning about distributed systems, which has proven incredibly influential in ensuing years; its influence continues to be felt today. This influence extends far beyond the distributed systems community and can be seen in current work in AI, security, and game theory, demonstrating how research in distributed computing is relevant to and applicable in a variety of settings involving multi-agent interaction.

The paper presented a novel rigorous and elegant framework supporting the intuition that the most fundamental characteristic of distributed algorithms is the fact that they must cope with uncertainty (i.e., lack of knowledge). When reasoning informally about distributed protocols, researchers naturally think (and speak) in terms of agents “knowing” certain facts about the global system state. The key insight of Halpern and Moses was that this informal notion of knowledge could be given a rigorous mathematical formulation. The resulting new “knowledge framework” shed useful, new light on old results and enabled the discovery of new ones.

The paper is seminal in many respects. First, the paper introduced a model of knowledge, which is now essentially standard in the distributed systems, formal methods, and multi-agent systems communities. The second key aspect of the paper is perhaps its most famous and possibly most cited result: common knowledge cannot be achieved in systems where the receipt of

messages is not guaranteed. This result, and the role of common knowledge in the associated “coordinated attack problem”, are now part of the computing folklore; they are known and cited by many, many researchers, including many who no doubt have never read the original paper, and are not even part of the distributed systems community. Although common knowledge had been introduced before, this paper for the first time identified its role in distributed systems, and particularly coordination problems. Third, the “hierarchies” of knowledge identified in the paper have been very influential. The different levels of knowledge (distributed/implicit knowledge, through to common knowledge) and their relationships have subsequently been extended, challenged, championed, and refined by many others.

For a long while, this paper provided the foundation for many negative and impossibility results, showing that certain tasks cannot be performed in a distributed environment, such as the impossibility of gaining common knowledge where none was initially present. However, subsequent developments have shown that the terminology and notation of the logic of knowledge introduced in this seminal paper can also be used in a positive way to derive new distributed algorithms and re-derive existing ones. Thus, the framework developed in this paper can serve as a high-level language for the natural and intuitive development of new distributed algorithms. With the recent renewed interest in synthesis of code from logical specifications we can expect that this foundation will find new and exciting applications.

In the context of security, since at least the late 1970s, it has been recognized that matters of belief and knowledge are central to the design and to the understanding of systems for security-critical tasks. However, the reasoning used in this context was informal and, although fruitful, also error-prone. Clearly, more rigorous ways of analyzing belief and knowledge are of great value in this context. Later research on security protocols relied on more precise definitions and on more systematic procedures for protocol design and analysis, which were influenced by the work of Halpern and Moses.

In an explicit attempt to address the common knowledge paradox raised by the Halpern-Moses paper, the notion of internal knowledge consistency was later considered, not only by the distributed computing community, but also by Nobel laureates in the economics community.

More generally, this paper sparked a considerable effort in the study of logics of knowledge involving multi-agent settings (most earlier work in philosophy dealt with the knowledge of a single agent in isolation). Indeed, the excitement and activity generated by this work had a central role in bringing about the biennial TARK conference (Theoretical Aspects of Reasoning about Knowledge, recently renamed Theoretical Aspects of Rationality and Knowledge), with its community consisting of theoretical computer scientists, AI researchers, economists and philosophers.