

The 2014 Edsger W. Dijkstra Prize in Distributed Computing

The Dijkstra Prize Committee has selected **Kanianthra Mani Chandy** and **Leslie Lamport** as the recipients of this year's Edsger W. Dijkstra Prize in Distributed Computing. The prize is given to them for their outstanding paper "Distributed Snapshots: Determining Global States of Distributed Systems", published in ACM Transactions on Computer Systems, Vol. 3, No. 1, 1985, pages 63–75.

The ACM-EATCS 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 prize is sponsored jointly by the ACM Symposium on Principles of Distributed Computing (PODC) and the EATCS Symposium on Distributed Computing (DISC).

In their paper, Chandy and Lamport describe a distributed algorithm to record a consistent global state of an asynchronous distributed computation. Determining such a distributed snapshot is essential for several fundamental tasks and it is for example at the basis of solutions for termination or deadlock detection or to verify whether some other stable global property holds. Note that the inherent lack of a common clock or a global observer means that global states that are asynchronously recorded can be potentially inconsistent. The solution provided by Chandy and Lamport is extremely elegant and also remarkably simple, a fact that certainly also contributed to its success. The correctness proof is based on swapping of prerecording and postrecording events without affecting the correctness of the observed state. It is very significant that the recorded global state may not have been any of the global states that the system passed through in a real-time view of its execution.

The paper provides the first clear understanding of the definition of consistent global states in distributed systems. Consistent global states form the cornerstone of asynchronous distributed execution observation, and subsequent concepts such as those of the hugely popular vector clocks and concurrent common knowledge are based on consistent global states. Further, the paper demonstrates that the recorded state in the global snapshot is a valid equivalent state in the sense that an equivalent execution may very well pass through that state, even though the actual execution may not have. In fact, the possibility of the occurrence of the recorded state is indistinguishable from the possibility of the occurrence of the actual states that occurred in the distributed execution. This property led to the definition of global predicates on entire executions, as opposed to the definition of global predicates on individual observations of executions. Similarly, this spawned the important concept of isomorphism of executions.

In conclusion, the paper by Chandy and Lamport has had a profound and lasting impact on the theory and implementation of distributed algorithms and systems. It has led to concepts such as vector time, isomorphism of executions, global predicate detection, and concurrent common knowledge. Applications of the results of observing the system in consistent states include the development of vector clocks, checkpointing and message logging protocols, correct protocols for detecting stable properties such as distributed deadlocks and termination, mutual exclusion algorithms, garbage collection protocols, cache coherency and file coherency protocols in distributed replicated file systems, distributed debugging protocols, protocols for total message order and causal message order in group communication systems, global virtual time algorithms used particularly in parallel and distributed simulations of discrete event systems, and collaborative sessions and editing protocols in wide area systems and on the grid.

Dijkstra Prize Committee 2014

Lorenzo Alvisi, Shlomi Dolev, Rachid Guerraoui,
Idit Keidar, Fabian Kuhn (chair), Shay Kutten