

Leader-election protocol

The following protocol is due to [Dolev, Klawe, Rodeh \(1982\)](#).

The protocol consists of n participants (where n is a parameter). The participants are connected by a ring of unidirectional message channels. Communication is asynchronous, and the channels are reliable. Each participant has a unique ID (e.g., some random number).

Goal: The participants communicate to elect a “leader” (i.e., some distinguished participant). The protocol shown here ensures low communication overhead ($\mathcal{O}(n \log n)$ messages; most naïve protocols have quadratic message overhead).

Leader-election protocol


Participants are either **active** or **inactive**. Initially, all participants are *active*.

The protocol proceeds in **rounds**. In each round, at least half of the participants will become inactive. (As a consequence, there are at most $\mathcal{O}(\log n)$ rounds.


In each round every active participant receives the numbers of the two nearest active participants (in incoming direction). A participant remains active only if the value of the nearest neighbour is the largest of the three. (In the following slides, the participant adopts this largest number as its own; this is optional.)

The last remaining active participant is declared the leader.


Leader Election: Example




Leader Election: First round




Leader Election: Result of the first round



Leader Election: Second round



Leader Election: Third round



Leader Election: Final result

