

Abstract

Population protocols are a popular model of distributed computing, in which randomly-interacting agents with little computational power cooperate to jointly perform computational tasks. Inspired by developments in molecular computation, and in particular DNA computing, recent algorithmic work has focused on the complexity of solving simple yet fundamental tasks in the population model, such as *leader election* (which requires convergence to a single agent in a special “leader” state), and *majority* (in which agents must converge to a decision as to which of two possible initial states had higher initial count). Known results point towards an inherent trade-off between the *time complexity* of such algorithms, and the *space complexity*, i.e. size of the memory available to each agent. In this paper, we explore this trade-off and provide new upper and lower bounds for majority and leader election. First, we prove a unified lower bound, which relates the space available per node with the time complexity achievable by a protocol: for instance, our result implies that any protocol solving either of these tasks for n agents using $O(\log \log n)$ states must take $\Omega(n/\text{polylog} n)$ expected time. This is the first result to characterize time complexity for protocols which employ super-constant number of states per node, and proves that fast, poly-logarithmic running times require protocols to have relatively large space costs.

On the positive side, we give algorithms showing that fast, poly-logarithmic convergence time can be achieved using $O(\log^2 n)$ space per node, in the case of both tasks. Overall, our results highlight a time complexity separation between $O(\log \log n)$ and $\Theta(\log^2 n)$ state space size for both majority and leader election in population protocols, and introduce new techniques, which should be applicable more broadly.