Abstract

Given an underlying graph, we consider the following *dynamics*: Initially, each node locally chooses a value in $\{-1, 1\}$, uniformly at random and independently of other nodes. Then, in each consecutive round, every node updates its local value to the average of the values held by its neighbors, at the same time applying an elementary, local clustering rule that only depends on the current and the previous values held by the node. We prove that the process resulting from this dynamics produces a clustering that exactly or approximately (depending on the graph) reflects the underlying cut in logarithmic time, under various graph models that exhibit a sparse balanced cut, including the stochastic block model. We also prove that a natural extension of this dynamics performs community detection on a regularized version of the stochastic block model with multiple communities. Rather surprisingly, our results provide rigorous evidence for the ability of an extremely simple and natural dynamics to address a computational problem that is non-trivial even in a centralized setting.