Codes, Lower Bounds and Phase Transitions in the Symmetric Rendezvous Problem

Abstract

In the Rendezvous problem on the complete graph, two parties are trying to meet at some vertex at the same time, despite starting out with independent random labelings of the vertices. It is well known that the optimal strategy is for one player to wait at any vertex, while the other visits all n vertices in consecutive steps, which guarantees a rendezvous within n steps and takes $(n+1)/2$ steps on average. This strategy is very far from being symmetric, however. Anderson and Weber **??** presented a symmetric algorithm that achieves an expected meeting time of $0.829n$, which has been conjectured to be optimal in the symmetric setting.

We change perspective slightly: instead of trying to minimize the expected meeting time, we try to maximize the probability of successfully meeting within a specified number of timesteps. In this setting, for all time horizons that are linear in n, the Anderson-Weber strategy has a constant probability of failure. Surprisingly, we show that this is not optimal: there exists a different symmetric strategy that almost surely guarantees meeting within $4n$ timesteps. This bound is tight, in that the factor 4 cannot be replaced by any smaller constant. Our strategy depends on the construction of a new kind of combinatorial object that we dub "rendezvous code."

On the positive side, for $T < n$, we show that the probability of meeting within T steps is indeed (approximately) maximized by the Anderson-Weber strategy. Our results imply new lower bounds on the expected meeting time for any symmetric strategy, which establishes an asymptotic difference between the best symmetric and asymmetric strategies.

Finally, we examine the symmetric rendezvous problem on other vertex-transitive graphs.