Abstract

Suppose we are given some fixed (but unknown) set X of binary strings of length n. We would like to learn as much as possible about the elements of X by asking certain binary questions. Each "question" Q is also a binary string, and the "answer" to Q is just the inner product $< Q, x >$ for some x in X. (We view the set of all binary strings of length n as the n-dimensional vector space over the field of two elements.) However, the choice of x is made by a truthful (but possibly malevolent) adversary A, whom we may assume is trying to choose answers so as to yield as little information as possible about X.

We are interested in extracting as much information as possible about X from A's answers. Although A can prevent us from learning the identity of any particular element of X, with appropriate questions we can still learn a lot about X. We determine the maximum amount of information that can be recovered and discuss the optimal strategies for selecting questions. For the case that $|X| = 2$, we give an $O(n^3)$ algorithm for an optimal strategy. However, for the case that $|X| > 2$, we show that no such polynomial-time algorithm can exist, unless $P = NP$.

This is a joint work with Fan Chung and Ronald Graham.