Function evaluation in the priced information framework

Ferdinando Cicalese

We study the fundamental problem of evaluating a function by sequentially selecting a subset of variables whose values uniquely identify the function's value. This basic problem arises in several domains of computer science, e.g., Automatic diagnosis, AI, applied Game Theory, Data Base Query optimization, just to menation a few.

A function f over a set of variables V must be computed and, for many inputs of the domain, not all the variables need to be read in order to determine the value of f on those inputs. A deterministic algorithm for this problem adaptively reads the value of the variables of f until the values read so far uniquely determine the value of f. Classically, each read operation is assumed to incur a unit-cost and the number of variables read (for the worst case input setting) is the measure used to analyze the efficiency of the algorithms. However, it is well known that a large class of functions of interest enjoy the evasiveness property, i.e., in the worst case any deterministic algorithm must read all the variables. Such classes show that the worst case analysis is not generally able to distinguish among the performances of different algorithms for the function evaluation problem. Other metrics that employ probabilistic and competitive analysis have been investigated in the literature.

Following Charikar et. al., here we address the variant of the function evaluation problem where different variables can incur different reading costs and competitive analysis is employed to measure the performance of the evaluation algorithm.