The Stochastic Knapsack Problem

February 4, 2005

In this talk I will discuss a generalization of the classic knapsack problem in which the sizes of the items are not known in advance. Instead for each of the n items we are given a deterministic value and a random variable for the size with a known arbitrary distribution. The goal is to maximize the expected value of items placed in the knapsack, but the size of an item is not known until it is placed in the knapsack (and it cannot be removed). Many questions associated with the stochastic knapsack problem are PSPACE-hard so we can attempt to find an approximate solution by developing either a non-adaptive policy which performs the same sequence of insertions no matter what the sizes of the items turn out to be or an adaptive policy which makes decisions based on the actual sizes of the items it has already placed. A greedy non-adaptive algorithm that approximates the optimal adaptive policy within a factor of 7 is given in [BD04]. The authors also present an adaptive algorithm that computes the next choice in polynomial time but adaptivity only gives a constant factor improvement, yielding a $(5+\epsilon)$-approximation.

References

[BD04] J. Vondrak B. Dean, M. Goemans. Approximating the stochastic knapsack problem: The benefit of adaptivity. *Symposium on Foundations of Computer Science*, pages 208–217, 2004.