Maximizing submodular functions via MWU

Chandra Chekuri, University of Illinois, Urbana-Champaign

Abstract
Submodular function maximization subject to independence constraints has received considerable attention in the last decade. One of the key ingredients for theoretical developments has been the multilinear relaxation. Algorithms based on this approach, however, are generally slow due to the cost involved in evaluating the multilinear extension via random sampling. In this talk we discuss the problem of solving the multilinear relaxation when the constraints are given in the form of explicit packing constraints $Ax \leq b$ where A is a non-negative matrix and b is a non-negative vector. Several problems of interest are captured by this constraint class. We show that the multiplicative weight update method, which has been successfully applied for positive linear programs over two decades, can be adapted to approximate the multilinear relaxation both in the sequential and parallel settings, to give faster and also conceptually clean algorithms.

Based on two papers: one with T.S. Jayram and Jan Vondrak on sequential algorithms, and a more recent one with Kent Quanrud on a parallel algorithm.