ASYMPTOTICALLY OPTIMAL ONLINE STOCHASTIC BIN PACKING

Varun Gupta, University of Chicago, USA, guptav@uchicago.edu
Ana Radovanović, Google Research, New York, USA, anaradovanovic@google.com

Motivated by the problem of packing Virtual Machines on physical servers in the cloud, we study the problem of online stochastic bin packing. Items with sizes i.i.d. from an unknown distribution with integral support arrive as a stream and must be packed on arrival in bins of size B, also an integer, with the goal of minimizing the number of bins used. All currently known heuristics for online stochastic bin packing are either optimal for only certain classes of item size distributions, or rely on learning the distribution. The state-of-the-art Sum of Squares heuristic (Csirik et al.) obtains sublinear (in number of items seen) waste for distributions where the expected waste for the optimal offline algorithm is sublinear, but has a constant factor larger waste for distributions with linear waste under OPT.

We present the first distribution agnostic bin packing heuristic that achieves additive $O(\sqrt{n})$ waste compared to OPT for all distributions. Our algorithm is essentially gradient descent on a suitably defined Lagrangian relaxation of the bin packing Linear Program, and extends to multi-dimensional packing constraints as well.

Next, we consider the more general problem of online stochastic bin packing with item departures where the time requirement of an item is unknown at arrival. Our algorithm extends as is to the case of item departures. We also briefly revisit the Best Fit heuristic which has not been studied in the scenario of item departures yet.