

## Abstract

We consider the problem of online load balancing under  $\ell_p$ -norms: sequential jobs need to be assigned to one of the machines and the goal is to minimize the  $\ell_p$ -norm of the machine loads. This generalizes the classical problem of scheduling for makespan minimization (case  $\ell_\infty$ ) and has been thoroughly studied. However, despite the recent push for beyond worst-case analyses, no such results are known for this problem.

In this paper we provide algorithms with **simultaneous** guarantees for the worst-case model as well as for the random-order (i.e. secretary) model, where an arbitrary set of jobs comes in random order. First, we show that the greedy algorithm (with restart), known to have optimal  $O(p)$  worst-case guarantee, also has a (typically) improved random-order guarantee. However, the behavior of this algorithm in the random-order model degrades with  $p$ . We then propose algorithm SimultaneousLB that has **simultaneously optimal guarantees** (within constants) in both worst-case and random-order models. In particular, the random-order guarantee of SimultaneousLB improves as  $p$  increases.

One of the main components is a new algorithm with improved regret for Online Linear Optimization (OLO) over the non-negative vectors in the  $\ell_q$  ball. Interestingly, this OLO algorithm is also used to prove a purely probabilistic inequality that controls the correlations arising in the random-order model, a common source of difficulty for the analysis. Another important component used in both SimultaneousLB and our OLO algorithm is a smoothing of the  $\ell_p$ -norm that may be of independent interest. This smoothness property allows us to see algorithm SimultaneousLB as essentially a greedy one in the worst-case model and as a primal-dual one in the random-order model, which is instrumental for its simultaneous guarantees.