LEVERAGING RANDOMIZATION: FROM SCHEDULING TO ADWORDS

At the very center of algorithmic computer science lie *assignment problems*, where we assign objects to pairs or larger sets often satisfying some additional constraints. The quintessential representative of this class of problems is BIPARTITE MATCHING, where objects (think two groups of students) are to be matched into as many pairs (of lab partners) as possible while respecting the preferences in the form of a graph.

The theory of *scheduling problems* tackles those assignment problems where a specific set of *jobs* is to be assigned into *timeslots* on available machines, which are to execute and finish said jobs. Scheduling problems are one of the central and most applicable problems, studied since the inception of optimization theory in the 1940s. Their applications include production and distribution planning, workflow optimization and much more.

All assignment problems (matching and scheduling) have also been intensely studied under the lens of *online optimization*, which is important for applications in network load balancing and routing, CPU scheduling and real-time computation. In the online model, we consider the input to arrive sequentially (think one job after the other), and the *online algorithm* must then make an immediate and permanent decision about the assignment before the next piece of input is revealed. The goal is to design algorithms that have a strong performance guarantee (called *competitive ratio*) compared to the best solution in hindsight.

There is an important connection between online optimization and game theory. We can see this through the fact that the algorithm makes immediate decision, but its performance is measured on the worst case – thus, we can imagine a second player, called *adversary* in the literature, that presents the next item based on the previous decisions of the algorithm, and chooses such an item as to hinder the algorithm's performance.

One fundamental result for online assignment problems is that ONLINE BIPARTITE MATCHING has been resolved even in the randomized setting, where the best possible randomized algorithm reaches a competitive ratio of $\frac{e}{e-1}$. (In ONLINE BIPARTITE MATCHING, the algorithm is given vertices of the "offline partition" in advance, and needs to match as many as possible to vertices of the "online partition", which arrive one by one, revealing their possible matches only upon arrival.)

In contrast to ONLINE BIPARTITE MATCHING, several central problems in online scheduling still have an unsolved optimal competitive ratio. One problem that was until recently in this category is ONLINE PACKET SCHEDULING, where unit-size jobs arrive online with an integer *deadline* for their last possible slot as well as positive *weight* corresponding to the priority or profit that is generated when the job is to be scheduled before the deadline. The task of the ONLINE PACKET SCHEDULING is then to select one job for the current time slot and maximize the overall weight of the scheduled jobs.

It was long suspected that the optimal competitive ratio in ONLINE PACKET SCHEDULING is the golden ratio $\phi \approx 1.618$, but only in 2019 was this confirmed by a research group including the P.I. of this grant.

In this research proposal we wish to employ the power of *randomization* for several central problems in online scheduling revolving around ONLINE PACKET SCHEDULING. Randomization is a strong technique in online computation that often improves results on two fronts: First, randomized algorithms often perform better than their deterministic algorithm – this is the case for ONLINE PACKET SCHEDULING, but is not true for some other similar problems. In either case, quantifying the benefit of randomization is precisely our research focus.

Second, randomization often allows for some constructions in both the algorithm design and the analysis to be simplified, which is equally important. For example, the recent analysis of the ϕ-competitive algorithm has over 40 pages, which contrasts with the fact that an extremely simple randomized algorithm, whose definition and analysis fit on a single page with a superior ratio of $\frac{e}{e-1}$ for the problem was designed early on.

Online algorithms for assignment problems have connections to *mechanism design* of auction theory, which can be seen through their common game-theoretical roots and through their common applications in networking and routing, even though the fields differ in assumptions on the individual choices of objects being routed.

Indeed, it both ONLINE PACKET SCHEDULING and ONLINE MATCHING can be seen as a special case of the ADWORDS PROBLEM, in which edges have weights and offline vertices can be matched to multiple online vertices, subject to their individual capacity for total weight of chosen incident edges. It has received a lot of attention as a clean generalization, but also due to the application for Internet companies such as Google, Microsoft, and Facebook. The ADWORDS PROBLEM captures their aim of maximizing revenue from ads being displayed with search results or a feed — think of offline vertices as advertisers willing to pay an amount they specify for each display (in specified circumstances), subject to a total budget constraint. This problem has been resolved with the optimal ratio $\frac{e}{e-1}$ in some special cases with additional constraints on the bidders' behavior, but the general case is still not fully explored.

We are hopeful that using the new insights into ONLINE PACKET SCHEDULING along our original research into the problem will suffice to resolve the optimal ratio in a broader set of instances of the ADWORDS PROBLEM.