Abstract

Scheduling is a decision making process that is used on a regular basis in many real life situations. It takes care of the allocation of resources to tasks over time, and its goal is to optimize one or more objectives. Depending on the problem being solved, tasks and resources can take many forms, and the objectives can also vary. In this dissertation we are interested in manufacturing scheduling, which is an optimization process that allocates limited manufacturing resources over time among parallel and sequential manufacturing activities. Customer orders have to be executed, and each order is composed by a number of operations that have to be processed on the resources or machines available. Each order can have release and due dates associated, and typical objectives functions involve minimizing the tardiness or the makespan.

In real world scheduling problems the environment is so dynamic that all this information is usually not known beforehand. For example, manufacturing scheduling is subject to constant uncertainty, machines break down, orders take longer than expected, and these unexpected events make the original schedule fail. That is the reason why companies prefer to have robust schedules rather than optimal ones. A key issue is to find a proper balance between these two performance measures.

Scheduling problems can be interpreted as a distributed sequential decision-making task, which makes the use of multi-agent reinforcement learning algorithms possible. These algorithms are of high relevance to various real-world problems, as they allow the agents to learn good overall solutions through repeated interactions with the environment.

Our main contribution is a generic multi-agent reinforcement learning approach that can easily be adapted to different scheduling settings, such as the job shop scheduling with parallel machines, online problems or the hybrid flow shop scheduling. It is possible to increase the robustness of the solutions and to look at different objective functions, like the tardiness or the makespan. Furthermore, the proposed approach allows the user to define certain parameters involved in the solution construction process, in order to define the balance between robustness and optimality.

Yailen Martínez

PhD Thesis: A Generic Multi-Agent Reinforcement Learning Approach for Scheduling Problems