Abstract

Rescheduling adds to deterministic scheduling the ability to respond to the occurrence of an unexpected event. While scheduling proposes methods to optimize the execution of tasks on machines, rescheduling proposes methods to best *update* known schedules when they become suboptimal for some reason. The unexpected triggering event studied in this dissertation is the arrival of new jobs. The new goal is to optimize the execution of all jobs, but also to avoid excessive deviations from the original schedule on which other work may rely. Methods for rescheduling involve the consideration of an objective function to be minimized and a disruption constraint that limits the deviation from the original solution. The goal of this dissertation is to study, from both a theoretical and algorithmic perspective, single-machine rescheduling problems to minimize classical scheduling objective functions given a disruption constraint measured as a function of the absolute deviation of completion times. Since rescheduling for new orders has only been studied in a fragmented and case-by-case manner in the literature, the work in this dissertation attempts to provide a comprehensive analysis of the nature of rescheduling problems and ways to solve them. First, the analysis considers the formulation of structural properties of the problem and focus on demonstrating the computational complexity of the problems. The structural analysis is followed by a detailed discussion of the timing problem arising from those problems where the problem formulation leads to optimal solutions that may include unforced idle time on the machine. Algorithms for the problems are proposed and their polynomiality is shown. The remainder of the dissertation studies how to approach two computationally hard rescheduling problems. The first problem is solved to optimality in pseudo-polynomial time using dynamic programming, and then using approximation algorithms to obtain solutions with a bounded performance ratio. For the second problem, a branch and memorize algorithm is developed that exploits a generic structure and specific properties. The algorithm is the first exact search tree algorithm for a hard rescheduling problem

that exploits an efficient exploration of the solution space. Computational experiments comparing the algorithm with a commercial solver for optimization problems show that the algorithm does indeed succeed in exploiting problem-specific features for solving the problem in a more efficient way. This concludes the discussion of rescheduling problems analyzed and studied as combinatorial problems, with the aim of providing a starting point for solving them from an application perspective.