

Scheduling under Non-Uniform Job and Machine Delays

Thursday, July 13, 2023 12:10 PM (20 minutes)

David Stalfa, Rajmohan Rajaraman and Sheng Yang

Abstract: We study the problem of scheduling precedence-constrained jobs on heterogeneous machines in the presence of non-uniform job and machine communication delays. We are given a set of n unit size precedence-ordered jobs, and a set of m related machines each with size m_i (machine i can execute at most m_i jobs at any time). Each machine i has an associated in-delay ρ_i^{in} and out-delay ρ_i^{out} . Each job v also has an associated in-delay ρ_v^{in} and out-delay ρ_v^{out} .

In a schedule, job v may be executed on machine i at time t if each predecessor u of v is completed on i before time t or on any machine j before time $t - (\rho_i^{\text{in}} + \rho_j^{\text{out}} + \rho_u^{\text{out}} + \rho_v^{\text{in}})$. The objective is to construct a schedule that minimizes makespan, which is the maximum completion time over all jobs.

We consider schedules which allow duplication of jobs as well as schedules which do not. When duplication is allowed, we provide an asymptotic $\text{polylog}(n)$ -approximation algorithm. This approximation is further improved in the setting with uniform machine speeds and sizes. Our best approximation for non-uniform delays is provided for the setting with uniform speeds, uniform sizes, and no job delays. For schedules with no duplication, we obtain an asymptotic $\text{polylog}(n)$ -approximation for the above model, and a true $\text{polylog}(n)$ -approximation for symmetric machine and job delays. These results represent the first polylogarithmic approximation algorithms for scheduling with non-uniform communication delays.

Finally, we consider a more general model, where the delay can be an arbitrary function of the job and the machine executing it: job v can be executed on machine i at time t if all of v 's predecessors are executed on i by time $t - 1$ or on any machine by time $t - \rho_{v,i}$. We present an approximation-preserving reduction from the Unique Machines Precedence-constrained Scheduling (\textsc{umps}) problem, first defined in [13], to this job-machine delay model. The reduction entails logarithmic hardness for this delay setting, as well as polynomial hardness if the conjectured hardness of \textsc{umps} holds.

This set of results is among the first steps toward cataloging the rich landscape of problems in non-uniform delay scheduling.

Presenter: STALFA, David

Session Classification: Track A-3