An EPTAS for scheduling jobs on uniform processors

TUE/AE01
15:30–15:50

Klaus Jansen (Univ. Kiel)

We consider the following fundamental problem in scheduling theory. Suppose that we are given a set $\mathscr{J}$ of n independent jobs J_j with processing time p_j and a set $\mathscr{P}$ of m non-identical processors P_i that run at different speeds s_i. If job J_j is executed on processor P_i, the machine needs p_j/s_i time units to complete the job. The problem denoted by $Q||C_{max}$ is to find an assignment $a : \mathscr{J} \to \mathscr{P}$ for the jobs to the processors that minimizes the total execution time, $\max_{i=1,\ldots,m} \sum_{J_j:a(J_j)=P_i} p_j/s_i$.

Hochbaum and Shmoys presented a family of polynomial time approximation algorithms $\{A_\varepsilon|\varepsilon > 0\}$ for scheduling on identical and uniform processors, where each algorithm A_ε generates a schedule of length $(1+\varepsilon)OPT(I)$ for each instance I and has running time polynomial in the input size $|I|$. The running time of the PTAS for uniform processors by Hochbaum and Shmoys is $(n/\varepsilon)^{O(1/\varepsilon^2)}$. If ε is small, then the running time of the algorithm can be very large.

Two restricted classes of approximation schemes were defined that avoid this problem. An efficient polynomial time approximation scheme (EPTAS) is a PTAS with running time $f(1/\varepsilon)poly(|I|)$ (for some function f), while a fully polynomial time approximation scheme (FPTAS) runs in time $poly(1/\varepsilon, |I|)$ (polynomial in $1/\varepsilon$ and the size $|I|$ of the instance). Since the scheduling problem on uniform (and also identical) processors is NP-hard in the strong sense (as it contains bin packing as special case), we cannot hope for an FPTAS. For identical processors, Hochbaum and Shmoys and Alon at el. gave an EPTAS with running time $f(1/\varepsilon) + O(n)$, where f is doubly exponential in $1/\varepsilon$. The existence of an EPTAS for uniform processors is mentioned as an open problem by Epstein and Sgall. Our main result is the following:

THEOREM. *There is an EPTAS (a family of algorithms $\{A_\varepsilon|\varepsilon > 0\}$) which, given an instance I of $Q||C_{max}$ with n jobs and m processors with different speeds and a positive number $\varepsilon > 0$, produces a schedule for the jobs of length $A_\varepsilon(I) \leq (1+\varepsilon)OPT(I)$. The running time of A_ε is $2^{O(1/\varepsilon^2 \log(1/\varepsilon)^3)} poly(n)$.*

Interestingly, the running time of our EPTAS is only singly exponential in $1/\varepsilon$.

[1] N. ALON, Y. AZAR, G.J. WOEGINGER, AND T. YADID: Approximation schemes for scheduling on parallel machines, *Journal on Scheduling* 1, 1998, 55-66.

[2] L. EPSTEIN AND J. SGALL: Approximation schemes for scheduling on uniformly related and identical parallel machines, *Algorithmica*, 39 (2004), 43-57.

[3] D.S. HOCHBAUM AND D.B. SHMOYS: Using dual approximation algorithms for scheduling problems: practical and theoretical results, *Journal of the ACM*, 34 (1987), 144-162.

[4] D.S. HOCHBAUM AND D.B. SHMOYS: A polynomial approximation scheme for scheduling on uniform processors: using the dual approximation approach, *SIAM Journal on Computing* 17 (1988), 539-551.