Scheduling to Minimize Total Flow Time

Stefano Leonardi, University of Roma "La Sapienza", Roma, Italy

We consider the classical problem of scheduling jobs in a multiprocessor setting in order to minimize the flow time (total time in the system). The performance of the algorithm, both in offline and online settings, can be significantly improved if we allow preemption: i.e., interrupt a job and later continue its execution. The best known algorithm for the non-preemptive case has approximation ratio $\sqrt{n/m}$ where n is the number of jobs and m the number of parallel machines [4, 5]. Preemption is inherent to make a scheduling algorithm efficient [4, 5]. However, minimizing the total flow time on parallel machines with preemption is known to be NP-hard on $m \geq 2$ machines.

Leonardi and Raz [5] showed that the well known heuristic *shortest remaining processing time* (SRPT) performs within a logarithmic factor of the optimal offline algorithm on parallel machines. It is not known if better approximation factors can be reached in the off-line setting.

In this first seminar we present basic techniques to devise approximation algorithms for minimizing total flow time for the non-preemptive and the preemptive case. In particular we will present the non-preemptive algorithm for parallel machines [4, 5] and a simple proof of the approximation ratio of SRPT never appeared in literature.

We will also discuss special cases of the problem, for instance the case in which job preemption but not job migration is allowed [2] and the case in which the objective function is the average stretch, i.e. the average ratio of the flow time to the processing time of a job [1]

References

[1] J. E. Gehrke, S. Muthukrishnan, R. Rajaraman, and A. Shaheen. Scheduling to Minimize Average Stretch Online. *To appear in Proc. Annual Symposium on Foundations of Computer Science*, 1999.

[2] B. Awerbuch, Y. Azar, S. Leonardi and O. Regev. Minimizing the flow time without migration. *Proc. of the 31st annual ACM Symposim on Theory of Computing*, pp. 198-205, 1999.

[3] L. Becchetti and S. Leonardi. Non-Clairvoyant scheduling to minimize the average flow time on single and parallel machines. *Proc. of the 33rd annual ACM Symposim on Theory of Computing*, pp. 94-103, 2001.

[4] H. Kellerer, T. Tautenhahn and G.J. Woeginger. Approximability and nonapproximability results for minimizing total flow time on a single machine. *SIAM Journal on Computing*, Vol. 28, Number 4, pp. 1155-1166, 1999.

[5] S. Leonardi and D. Raz. Approximating total flow time on parallel machines. In *Proceedings of the Twenty-Ninth Annual ACM Symposium on Theory of Computing*, pages 110–119, El Paso, Texas, 1997. To appear in *Journal of Computer and ystem Sciences – special issue for STOC '97.*