

Improved Bicriteria Existence Theorems for Scheduling

Javed Aslam* April Rasala† Cliff Stein‡ Neal Young§

Dartmouth College

Two common objectives for evaluating a schedule are the *makespan*, or schedule length, and the *average completion time*. In this note, we give improved bounds on the existence of schedules that simultaneously optimize both criteria.

In a scheduling problem, we are given n jobs and m machines. With each job j we associate a non-negative weight w_j . A *schedule* is an assignment of jobs to machines over time, and yields a completion time C_j for each job j . We then define the average completion time as $\sum_{j=1}^n w_j C_j$ and the makespan as $C_{\max} = \max_j C_j$. We use C_{\max}^{opt} and $\sum w_j C_j^*$ to denote the optimal makespan and average completion time.

We will give results which will hold for a wide variety of combinatorial scheduling problems. In particular, we require that valid schedules for the problem satisfy two very general conditions. First, if we take a valid schedule S and remove from it all jobs that complete after time t , the schedule remains a valid schedule for those jobs that remain. Second, given two valid schedules S_1 and S_2 for two sets J_1 and J_2 of jobs (where $J_1 \cap J_2$ is potentially nonempty), the composition of S_1 and S_2 , obtained by appending S_2 to the end of S_1 , and removing from S_2 all jobs that are in $J_1 \cap J_2$, is a valid schedule for $J_1 \cup J_2$.

For the rest of this note we will make claims about “any” scheduling problem, and mean any problem that satisfies the two conditions above. In addition, if a schedule has $C_{\max} \leq \alpha C_{\max}^{\text{opt}}$ and $\sum w_j C_j \leq \beta \sum w_j C_j^*$ we call S an (α, β) -schedule.

Stein and Wein [7] recently gave a powerful but simple theorem on the existence of schedules which are simultaneously good approximations for makespan and for average completion time. They showed that for *any* scheduling problem, there exists a $(2, 2)$ -schedule. The construction is simple. We take an optimal average completion time schedule and replace the subset J' of

jobs that finish after time C_{\max}^{opt} by an optimal makespan schedule for J' . The schedule has length at most $2C_{\max}^{\text{opt}}$, and the completion time of each job at most doubles, thus we obtain a $(2, 2)$ -schedule.

In the $(2, 2)$ -schedule, C_{\max}^{opt} was the *break-point*, the point at which we truncated the average completion time schedule and started the makespan schedule on the remaining jobs. By considering several different break points simultaneously, and taking the best one, Stein and Wein show, via a complicated case analysis, how to achieve improved approximations. In particular, they prove the existence of $(2, 1.735)$ -schedules, $(1.785, 2)$ -schedules and $(1.88, 1.88)$ -schedules.

In this paper, we give improved theorems on the existence of bicriteria schedules. Our first conceptual idea is that an average completion time schedule, appropriately normalized, can be viewed as a continuous probability density function. Even though schedules are actually discrete functions, this mapping to continuous functions facilitates the analysis. We choose, for any average completion time schedule (pdf), the breakpoint that gives the best bicriteria result; this calculation is now expressed as an integral. Choosing the pdf that maximizes this integral provides a worst-case schedule.

We now give an overview of the technical details. Wlog, we can normalize the weights w_j in the optimal average completion time schedule so that $\sum_j w_j C_j^* = 1$. Now let $g(z) = (\sum_{j|C_j^*=z} w_j C_j^*)\delta(0)$, where $\delta(\cdot)$ is Dirac’s delta function. By our normalization assumption, we have that $\int_0^\infty g(z) dz = 1$ and $g(z) \geq 0$. Thus g is a probability density function (pdf). Let L denote the optimal makespan and consider the schedule formed by having a breakpoint at αL . The jobs that complete before time αL have their completion times unaffected, while those that complete at time $z > \alpha L$ have their completion times multiplied by at most $(1 + \alpha)/z$. Thus, the resulting schedule has a makespan of $(1 + \alpha)L$ and an average completion time of

$$\int_0^{\alpha L} g(z) dz + \int_{\alpha L}^\infty \frac{(1 + \alpha)L}{z} g(z) dz \\ = \int_0^\infty g(z) dz + \int_{\alpha L}^\infty \frac{(1 + \alpha)L - z}{z} g(z) dz.$$

*jaa@dartmouth.edu. Research partially supported by Walter and Constance Burke Research Initiation Award 424879.

†apes@dartmouth.edu.

‡cliff@cs.dartmouth.edu. Research partially supported by NSF Career Award CCR-9624828.

§ney@dartmouth.edu. Research partially funded by NSF Career Award CCR-9720664.

Given a particular schedule, $g(z)$, we choose the α that minimizes the above expression to find the minimum average completion time. If we wish to restrict ourselves to finding the best schedule of makespan no more than $1 + \rho$, then we allow α to range from 0 to ρ , and choose the worst possible schedule $g(z)$. This corresponds to evaluating

$$\max_g \min_{0 \leq \alpha \leq \rho} \int_{\alpha L}^{\infty} \frac{(1 + \alpha)L - z}{z} g(z) dz,$$

where g is a probability distribution over $[0, \infty)$. This can be shown to be equivalent to the expression

$$(1) \quad \max_f \min_{0 \leq \alpha \leq \rho} \int_{\alpha}^{\infty} \frac{1 + \alpha - x}{x} f(x) dx,$$

where f now ranges over all distributions.

We can show that the maximum of (1) is achieved by the following function

$$f_{\text{opt}}(x) = \begin{cases} \frac{e^{\rho}}{e^{\rho}-1} x e^{-x} & 0 \leq x < \rho \\ \frac{\rho}{e^{\rho}-1} \delta(0) & x = \rho \\ 0 & x > \rho. \end{cases}$$

which yields the following bound.

THEOREM 1. *For any $\rho \in [0, 1]$, for any scheduling problem, there exists a $(1 + \rho, e^{\rho}/(e^{\rho} - 1))$ approximation.*

We omit the proof, but note that this theorem can be verified by viewing the integral as a continuous infinite-dimensional linear program and computing the dual, which is

$$\min_h \max_{x \geq 0} \int_0^{\min \rho, 1} \frac{(1 + \alpha - x)}{x} h(\alpha) d\alpha,$$

where h is a pdf over the interval $[0, 1]$. This dual is optimized by choosing $h(x) = e^x/(e^{\delta} - 1)$ for $x \in [0, 1]$ and 0 otherwise.

COROLLARY 1. *For any scheduling problem, there exists a $(2, 1.582)$ -schedule, a $(1.695, 2)$ -schedule and a $(1.806, 1.806)$ -schedule.*

These results, for some scheduling models, provide better bicriteria algorithms than can be achieved by Chakrabarti et. al.[2] or Stein and Wein[7]. Consider the case of $w_j = 1$ for all j . For the scheduling of jobs on parallel machines of different speeds, there is a polynomial-approximation scheme for makespan [3], and a polynomial-time algorithm for average completion time[4, 1]; resulting in a $(1.806 + \epsilon, 1.806)$ -algorithm. When considering the problem of scheduling jobs on *unrelated* parallel machines there is a 2-approximation

algorithm for makespan [6] and a polynomial-time algorithm for average completion time [4, 1]; thus a $(3.612, 1.806)$ -algorithm exists. Finally, for the scheduling of jobs, now with general weights, on parallel identical machines, for average weighted completion time there is a $(\frac{\sqrt{2}+1}{2})$ -approximation algorithm [5]; together with the polynomial-approximation scheme for makespan[3] we achieve a $(1.806 + \epsilon, 2.180)$ -algorithm.

Our results also apply to bicriteria optimization of the travelling salesman and travelling repairman problems. In the travelling repairman problem, we have a start vertex v , and define c_i to be the distance in the tour from vertex c to vertex i . Associated with vertex i is a nonnegative weight w_i and the goal is to find a tour that minimizes $\sum_i w_i c_i$. Combining results from [7] with the techniques in this paper we see that the existence of an (α, β) schedule implies the existence of a tour that is simultaneously a $1 + \alpha$ approximation for the travelling salesman problem and a β -approximation for the travelling repairman problems.

This model also bounds the completion time of *each job* by a factor of β times its completion in an optimal schedule. Therefore our results have consequences for minsum criteria other than $\sum w_j C_j$, such as $\sum w_j C_j^2$.

References

- [1] J.L. Bruno, E.G. Coffman, and R. Sethi. Scheduling independent tasks to reduce mean finishing time. *Communications of the ACM*, 17:382–387, 1974.
- [2] S. Chakrabarti, C. A. Phillips, A. S. Schulz, D. B. Shmoys, C. Stein, and J. Wein. Improved scheduling algorithms for minsum criteria. In F. Meyer auf der Heide and B. Monien, editors, *Automata, Languages and Programming*, number 1099 in Lecture Notes in Computer Science. Springer, Berlin, 1996. Proceedings of the 23rd International Colloquium (ICALP’96).
- [3] D.S. Hochbaum and D.B. Shmoys. A polynomial approximation scheme for machine scheduling on uniform processors: using the dual approximation approach. *SIAM Journal on Computing*, 17:539–551, 1988.
- [4] W. Horn. Minimizing average flow time with parallel machines. *Operations Research*, 21:846–847, 1973.
- [5] T. Kawaguchi and S. Kyan. Worst case bound of an LRF schedule for the mean weighted flow-time problem. *SIAM Journal on Computing*, 15:1119–1129, 1986.
- [6] J.K. Lenstra, D.B. Shmoys, and É. Tardos. Approximation algorithms for scheduling unrelated parallel machines. *Mathematical Programming*, 46:259–271, 1990.
- [7] C. Stein and J. Wein. On the existence of schedules that are near-optimal for both makespan and total weighted completion time. *Operations Research Letters*, 21, 1997.