MDS Autumm School Approximation Algorithms for Stochastic Optimization

Problems & References Related to Lectures on Stochastic Scheduling

Marc Uetz (m.uetz@utwente.nl), Applied Mathematics, University of Twente

Problems

(1) Let $t \geq 0$ and for $k \in \mathbb{N}$ consider kt jobs with i.i.d. processing times $P_j \sim \exp(k)$, that is, $\Pr[P_j > t] = e^{-kt}$ (and $\mathbb{E}[P_j] = 1/k$), and weights $w_j = 1/k$. Assume these jobs are scheduled on a single machine from time 0 on. Then for all $\varepsilon > 0$ there exists k large enough so that

$$\mathbb{E}[\sum_j w_j C_j] \leq \int_0^t x\,dx + \varepsilon\,.$$

(From [7].)

(2) Random variables P and Q are (stochastically) independent if $F_{P,Q}(x,y) = F_P(x) \cdot F_Q(y)$ for all x, y, F denoting the respective cumulative distribution functions. For discrete random variables this is equivalent with $\Pr[P = p \wedge Q = q] = \Pr[P = p] \cdot \Pr[Q = q]$ for all p, q. It is well known that, if P and Q are independent, then we have that $\text{Cov}[P,Q] := \mathbb{E}[P \cdot Q] - \mathbb{E}[P] \cdot \mathbb{E}[Q] = 0$[1]. Give an example to show that the reverse in generally not true, i.e., $\text{Cov}[P,Q] = 0$ does not imply independence. (Taken from some probability course.)

(3) Consider the linear program

$$\min \sum_j w_j x_j \quad \text{s.t.} \quad x(W) \geq f(W) \quad \text{for } W \subseteq J$$

for some supermodular set function $f : 2^J \to \mathbb{R}$ (that is, $f(X \cap Y) + f(X \cup Y) \geq f(X) + f(Y)$ for all $X, Y \subseteq J$). Assume $w_1 \geq \cdots \geq w_n$. Show that an optimal solution is given by the greedy solution $x_1^* = f(\{1\})$,

$$x_j^* := f(\{1,\ldots,j\}) - f(\{1,\ldots,j-1\}) \qquad \text{for } j = 2,\ldots,n\,.$$

Here, primal feasibility follows from supermodularity of f and optimality follows by giving a dual solution that fulfills complementary slackness. (From [4].)

(4) Comment on the following argument. Letting W_i be the jobs scheduled on machine i, we have, using the one-machine load inequalities, for any subset of jobs $W \subseteq J$,

$$\sum_{j\in W} p_j S_j = \sum_i \sum_{j \in W_i} p_j S_j \;\; \geq \sum_i \sum_{j,k \in W_i, j\neq k} p_k p_j\,.$$

Here S_j is the start time of a job j. Taking expectations, by independence of P_k and P_j for $j \neq k$, and since policies are non-anticipatory, for any policy Π we get

$$\sum_{j\in W} \mathbb{E}[P_j]\mathbb{E}[S_j^\Pi] \geq \sum_i \sum_{j,k\in W_i, j\neq k} \mathbb{E}[P_k]\mathbb{E}[P_j] = \sum_i \frac{1}{2}\Big(\sum_{j\in W_i}\mathbb{E}[P_j]\Big)^2 - \frac{1}{2}\sum_{j\in W}\mathbb{E}[P_j]^2\,.$$

We now use Cauchy-Schwarz which gives

$$\sum_{j\in W} \mathbb{E}[P_j]\mathbb{E}[S_j^\Pi] \geq \frac{1}{2m}\Big(\sum_{j\in W}\mathbb{E}[P_j]\Big)^2 - \frac{1}{2}\sum_{j\in W}\mathbb{E}[P_j]^2\,.$$

[1] In fact the latter is what we used in deriving the stochastic load inequalities.

Hence,

$$\sum_{j\in W}\mathbb{E}[P_j]\mathbb{E}[C_j^{\Pi}] \geq \frac{1}{2m}\left(\sum_{j\in W}\mathbb{E}[P_j]\right)^2 + \frac{1}{2}\sum_{j\in W}\mathbb{E}[P_j]^2\,. \tag{1}$$

(From <Classified>)

(5) Note that inequalities (1) are indeed valid for stochastic machine scheduling when processing time distributions are NBUE. Show that inequalities (1) are not necessarily valid if processing time distributions P_j are allowed to have a coefficient of variation larger than 1. (From [10].)

(6) As a substantial strengthening of the above claim (why is it a strengthening?), one can show that SEPT —shortest expected processing time first— can have a performance bound as large as $\Omega(\Delta^{1/2})$ for parallel machine scheduling to minimize $\mathbb{E}[\sum_j C_j]$. To that end, take an instance with m^2 deterministic jobs with processing times $1/m$ and $m^3/4$ stochastic jobs with processing times

$$P_j = \begin{cases} 0 & \text{with probability } 1-\frac{1}{m^2}\,, \\ m^{1+\varepsilon} & \text{with probability } \ \frac{1}{m^2}\,. \end{cases}$$

Hint: Use a Chernoff bound to show that the probability for having $m/2$ or more stochastic jobs can be bounded from above by $e^{-m/4}$. (From [2].)

(7) Give an argument to show that on two identical machines, $2\,|\,|\sum_j w_jC_j$, with exponentially distributed processing times, there always exists an optimal scheduling policy that does never leave a machine idle (as long as there are unprocessed jobs). (From [21].)

(8) Give a machine scheduling example to show that the expected value of an optimal stochastic scheduling policy may also overestimate the value of a "corresponding" deterministic instance with $p_j = \mathbb{E}[P_j]$. (From [22].)

(9) If we were to schedule jobs in the order of LP start times S_j^{LP}, instead of LP completion times C_j^{LP}, would we then obtain anything qualitatively different? In terms of the performance bounds? (From [21].)

(10) Consider a precedence constrained problem with stochastic jobs without any machine restrictions (think of $m = n$), and the objective to minimize expected makespan $\mathbb{E}[C_{\max}]$, also known as PERT problem. Show that the problem to compute $\mathbb{E}[C_{\max}]$ is # P-hard, be reducing the reliability problem on a transportation network to it. (From [6].)

(11) Consider, for given $C \in \mathbb{R}^n$ the supermodular maximization problem

$$\max_{W\subseteq J}\left(f(W) - \sum_{j\in W}\mathbb{E}[P_j]C_j\right)$$

with $f(W) = \frac{1}{2m}(\sum_{j\in W}\mathbb{E}[P_j])^2 + \frac{1}{2}\sum_{j\in W}\mathbb{E}[P_j]^2$. Show that this problem can be solved (by sorting) in $\mathrm{O}(\,n\log n\,)$ time. (From [11], showing that if W^* is an inclusion-minimal optimal solution, $j \in W^*$ if and only if C_j fulfills a simple condition.)

(12) Consider any feasible solution $C \in \mathbb{R}^n$ to the linear constraints

$$\begin{aligned} &\textstyle\sum_{j\in W}\mathbb{E}[P_j]\,C_j^{\text{LP}} \geq f(W)\,, && W \subseteq J\,, \\ &C_j^{\text{LP}} \geq \mathbb{E}[P_j]\,, && \forall\, j \in J\,, \end{aligned}$$

with $f(W)$ as above. Assume $C_1^{\text{LP}} \leq \cdots \leq C_n^{\text{LP}}$. Show that $\sum_{k\leq j}\mathbb{E}[P_k] \leq 2C_j^{\text{LP}}$. (From [14].)

(13) Consider the stochastic scheduling problem $\mathrm{P}\,|\,r_j\,|\,\mathbb{E}[\sum_j w_j C_j]$, and assume that processing time distributions are NBUE (so that $\mathbb{CV}[P_j] = \mathbb{V}\mathrm{ar}[P_j]/\mathbb{E}[P_j]^2 \leq 1$ for all jobs j). Show that there is an LP-based $(4-\frac{1}{m})$-approximation algorithm. (From [10].)

An additional problem related to this is to show that the linear program relaxation for the problem with release dates

$$\begin{aligned} \min \quad & \sum\nolimits_j w_j\, C_j^{\mathrm{LP}} \\ & \sum\nolimits_{j\in W} \mathbb{E}[P_j]\, C_j^{\mathrm{LP}} \geq f(W)\,, \quad W \subseteq J\,, \\ & C_j^{\mathrm{LP}} \geq r_j + \mathbb{E}[P_j]\,, \qquad\qquad \forall\, j \in J\,, \end{aligned}$$

with $f(W)$ as above can be solved in time $\mathrm{O}(\,n^2\,)$, by showing that this can be reformulated as a polymatroid optimization problem, too (but then with a slightly more complicated right hand side than f — I guess this needs a hint). (From [10], respectively [5].)

(14) Consider a discrete, integer valued random variable X. Show that

$$\sum\nolimits_{r\geq 0}(r+\frac{1}{2})\,\mathrm{Pr}[X>r] = \frac{1}{2}\mathbb{E}[X^2]\,.$$

(From [17]. This is a discrete version and special case of the (possibly well known) identity $\mathbb{E}[X^\alpha] = \alpha\int_0^\infty t^{\alpha-1}\,\mathrm{Pr}[X>t]dt$ for nonnegative random variable X and $\alpha>0$. It has a simple combinatorial proof.)

(15) Consider a stochastic scheduling instance of $\mathrm{P}\,|\,|\,\mathbb{E}[\sum_j C_j]$ with m identical machines, m^2 jobs with independent, identically distributed processing times

$$P_j = \begin{cases} 1 & \text{with probability } \frac{1-\varepsilon}{m}\,, \\ 0 & \text{with probability } 1-\frac{1-\varepsilon}{m}\,. \end{cases}$$

Show that any fixed assignment policy Π has expected performance $\mathbb{E}[\sum_j C_j^{\Pi}] \geq (1-\varepsilon)\frac{1}{2}m^2$. (From [17].)

(16) Give an argument to show the —somewhat surprising— fact that, for parallel machine scheduling, the LP relaxation based on expected completion time variables C_j^{LP} can be stronger than the formulation using time-indexed variables x_{ijt}. (From [17].)

References

[1] C. Chekuri, R. Motwani, B. Natarajan, and C. Stein. Approximation techniques for average completion time scheduling. *SIAM Journal on Computing*, 31:146–166, 2001.

[2] W. C. Cheung, F. Fischer, J. Matuschke, N. Megow. A $\Omega(\Delta^{1/2})$ gap example for the WSEPT policy. Personal Communication, 2014.

[3] C.F. Chou, H. Liu, M. Queyranne, and D. Simchi-Levi. On the asymptotic optimality of a simple online algorithm for the stochastic single machine weighted completion time problem and its extensions. *Operations Research*, 54:464– 474, 2006.

[4] J. Edmonds. Submodular functions, matroids and certain polyhedra. In R. K. Guy and H. Hanani (Eds.), *Combinatorial Structures and their Applications*, pp. 69–87. New York: Gordon & Breach, 1970.

[5] S. Fujishige. *Submodular functions and optimization*, Volume 47 of *Annals of Discrete Mathematics*. Amsterdam: North-Holland, 1991.

[6] J. N. Hagstrom. Computational complexity of PERT problems. *Networks 18*, 139–147, 1988.

[7] C. Jagtenberg, U. Schwiegelshohn, and M. Uetz. Analysis of Smith's Rule in Stochastic Machine Scheduling, *Operations Research Letters*, 41(6): 570-575, 2013.

[8] N. Megow, M. Uetz, and T. Vredeveld. Models and algorithms for stochastic online scheduling. *Mathematics of Operations Research*, 31(3):513–525, 2006.

[9] R. H. Möhring, F. J. Radermacher, and G. Weiss. Stochastic scheduling problems I: General strategies. *ZOR - Zeitschrift für Operations Research*, 28:193–260, 1984.

[10] R. H. Möhring, A. S. Schulz, and M. Uetz. Approximation in stochastic scheduling: The power of LP-based priority policies. *Journal of the ACM*, 46:924–942, 1999.

[11] M. Queyranne. Structure of a simple scheduling polyhedron. *Mathematical Programming*, 58:263–285, 1993.

[12] M. H. Rothkopf. Scheduling with random service times. *Management Science*, 12:703–713, 1966.

[13] M. Scharbrodt, T. Schickinger, and A. Steger. A new average case analysis for completion time scheduling. *Journal of the ACM*, 53:121–146, 2006.

[14] A. S. Schulz. Scheduling to minimize total weighted completion time: Performance guarantees of LP-based heuristics and lower bounds. In W. H. Cunningham, S. T. McCormick, and M. Queyranne, editors, *Integer Programming and Combinatorial Optimization*, volume 1084 of *Lecture Notes in Computer Science*, pages 301–315. Springer, 1996.

[15] A. S. Schulz. Stochastic online scheduling revisited. In B. Yang, D.-Z. Du, and C. Wang, editors, *Combinatorial Optimization and Applications*, volume 5165 of *Lecture Notes in Computer Science*, pages 448–457. Springer, 2008.

[16] A. S. Schulz and M. Skutella. Scheduling unrelated machines by randomized rounding. *SIAM Journal on Discrete Mathematics*, 15:450–469, 2002.

[17] M. Skutella, M. Sviridenko, and M. Uetz. Stochastic scheduling on unrelated machines. N. Portier E. W. Mayr, ed., *Proc. 31st Symp. Theoretical Aspects of Computer Science*, *Leibniz International Proceedings in Informatics*, vol. 25. Dagstuhl Publishing, Germany, 639–650, 2014.

[18] M. Skutella and M. Uetz. Stochastic machine scheduling with precedence constraints. *SIAM Journal on Computing*, 34:788–802, 2005.

[19] W. E. Smith. Various optimizers for single-stage production. *Naval Research and Logistics Quarterly*, 3:59–66, 1956.

[20] A. Souza and A. Steger. The expected competitive ratio for weighted completion time scheduling. *Theory of Computing Systems*, 39:121–136, 2006.

[21] M. Uetz. *Algorithms for Deterministic and Stochastic Scheduling* (PhD thesis), Cuvillier Verlag Göttingen, Germany, 2002

[22] M. Uetz. When greediness fails: Examples from stochastic scheduling. *Operations Research Letters*, 31:413–419, 2003.

[23] Gideon Weiss. Approximation results in parallel machines stochastic scheduling. *Annals of Operations Research*, 26:195–242, 1990.

A Chernoff Bound:

Theorem (Chernoff). Given n i.i.d. random variables X_i that take values $\{0,1\}$, with $\mu := \mathbb{E}[X_i]$, then for $0 < \varepsilon < 1$,

$$\Pr[\tfrac{1}{n}\sum_i X_i \geq (1+\varepsilon)\mu] \leq e^{-\varepsilon^2 n\mu/3}\,.$$

(I apologize for all errors!)