SIAM Journal of Computing referee report
Online Scheduling of Equal-Length Jobs: Randomization and Restarts Help

0 Preface

The following is a referee report for the paper, *Online Scheduling of Equal-Length Jobs: Randomization and Restarts Help*, by Marek Chrobak, Wojciech Jaword, Jiří Sgall and Tomáš Tichý, submitted to SIAM Journal on Computing.

1 Evaluation

This paper considers models for the online scheduling of a single machine, where jobs have equal processing times, arbitrary release and due dates, and the goal is to maximize the number of jobs completed by their due date (i.e., $1 \,|\, p_j = p \,|\, \sum_j (1 - U_j)$ in standard nomenclature). They present two results:

- In a non-preemptive setting, they present a $\frac{5}{3}$-competitive randomized algorithm which uses a single bit of randomness. This is the first time that randomness has been used to improve upon the known deterministic bound of 2. The best known randomized lower bound is $\frac{4}{3}$. They further examine this concept of a "barely random" algorithm which is based upon only two underlying deterministic algorithms, showing a lower bound of $\frac{3}{2}$ for any such algorithm and $\frac{8}{5}$ for any such algorithm which selects between the two with equal probability.
- They introduce the study of a Preempt-Restart model within this setting, providing a deterministic $\frac{3}{2}$-competitive algorithm and showing that this is the optimal (deterministic) result.

The first of these results is a significant achievement, both for making progress in reducing a gap which has existed in the community for over 6 years, as well as for the novelty and simplicity of their techniques. The proof is not as simple as the algorithm itself, and the technical details are awkward at times, but it is sound and reasonably well presented.

The second of these results is 'cute' though not as significant in terms of history or techniques. The statement of the algorithm is overly complex, as is the proof structure, though it appears technically sound.

Overall, the writing style is conversational in a very enjoyable way, and both the prose and the technical details are generally presented without flaw (with exceptions noted below).

2 Most Significant Recommendations

R1: As a reader, I'm overly distressed by some choices you have made in choosing notations, in particular an inconsistent use of variables j and k in differing parts of the paper. I realize that the scope of notation is localized and so this is not technically a flaw in the proof. But it does place an undue burden upon the reader, and without any just cause. A careful reading of this paper already requires a great deal of cross-reference between various pieces of your proofs.

As a concrete example, in the statement of Lemma 2.2 you use j to denote a member of the domain of $f(\cdot)$. In step, in the proof of Theorem 3.2 on page 10, when defining the partial function $f(\cdot)$, you use j to denote a job from set $\mathcal{X}$ (resp Y), and you denote $k \in \mathcal{Z}$. Yet moments later, when defining the charging scheme, you now let j denote a job in $\mathcal{Z}$ and then you use k in Cases II and III to denote a job in A (resp B). Thus, on the second line of page 11, when you say "If $x_j \geq C_k^{\mathcal{D}}$ then $f^{\mathcal{D}}(k) = j$," I'd like to verify this. So I look back at the definition for the function $f(\cdot)$ on the previous page and see the pre-condition for $f^{\mathcal{A}}(j) = k$ written as $\mathcal{S}_j^{\mathcal{X}} \leq \mathcal{S}_k^{\ddagger} < \mathcal{C}_j^{\mathcal{X}} \leq x_k$. But I now need to remap j to k and k to j because of the changed scope.

As a second example, we look at Section 4. Again, on page 15 you define $f(\cdot)$ by using $j \in \mathcal{X}$ in the domain, and $k \in \mathcal{Z}$ in the range. Yet on page 16, you revert to the notation where $k \in \mathcal{A}$ and $j \in \mathcal{Z}$, both when defining the matching M and in the statement of Lemma 4.3.

If there were some other overwhelming reason why you needed to alter the roles of j and k, then perhaps it is a necessary evil. However I see no such necessity.

R2: Page 14, statement of Algorithm TightRestart
The statement of the algorithm can be greatly simplified in two ways.

First, the phrase "was started as urgent" in (TR2) is unnecessary. It is easy to show that the condition for (TR3) already forbids the preemption of such an urgent job. If $P_s \subseteq P_t^*$, then the urgency of j implies the inflexibility of P_t^*. Alternatively if $P_s \not\subseteq P_t^*$ then some job $j' \in P_s$ must have expired, that is, $x_{j'} < t$. Since j was ED, then $x_j < t$, yet since $j \in P_t^*$, such set cannot be flexible at time t (nor even feasible!).

Secondly, there is no reason in (TR3) to even define P_t^*. It is sufficient to simply state that if P_t is flexible at t, then preempt j; else continue running j. It must be the case that (P_t^* is flexible at t) $\equiv$ (P_t is flexible at t). In short, if you satisfy the existing

condition for preemption, you would have already been willing to preempt upon the release of any earlier preemption candidates. Therefore, no such preemption candidates will be pending in such a case. (there is some relationship between this point and the argument made in the fourth paragraph on page 17)

It is obviously up to you as to whether you wish to make such a simplification, and then how this may alter the remaining proof. I have not looked into that, but presumably you could see where you make use of this condition about urgency in (TR2) or the distinction between P_t^* versus P_t in (TR3).

3 Additional Suggestions

S1: Page 3, line -8
Replace phrase "with resume" by "with resumption"

S2: Page 4, remainder of Section 1.
A citation should be added to, "Maximizing Job Completions Online" by Kalyanasundaram/Pruhs in J.Alg 49(1):63-85, 2003. This paper is closely related to the submission both because of the problem studied (a standard preemptive version of this same problem, with general processing times as well), and because of technique (they use a single-bit, barely-random algorithm to provide a constant competitive randomized result, in a setting where no such deterministic result is possible).

S3: Page 4, first line of Section 2
I'm very confused by your use of the phrase "instance on input" here, so much so that I am not even sure what to suggest as an intended alternative (presumably just remove "instance on")

S4: Page 6, line 20
The phrase "if there are more several choices" is awkward. Presumably should remove the word "more"

S5: Page 7, line 13
Replace word "an" with "a" preceding symbol $\subseteq$

S6: Page 8, lines 15–16
The statement "Since $f(j) \in \mathcal{Q}_{t+p}$ for $t = \mathcal{S}_j^{\mathcal{X}}$," is not technically true as $f(j)$ may have already been completed by that time. Of course the end result of Lemma 2.2(1) is clearly true in this case as well.

S7: Pages 9–13
In general, I'm a bit dissatisfied with the nesting of the pieces leading to Theorem 3.2. You have two lemmas with proofs nested within the proof of the theorem, and even an additional claim with proof nested within the proof of one of the lemmas. Would

it be possible to describe in narrative the charging scheme, and then save the formal statement of Theorem 3.2 until *after* all of the necessary lemmas have been established?

S8: Page 9, Lemma 3.1
Given the above decision, why is Lemma 3.1 then left *outside* of the main theorem. This is particularly confusing because the definitions of $\mathcal{D}$ and $\overline{\mathcal{D}}$ differ within the Lemma (relative to $\mathcal{X}$ and $\mathcal{Y}$) versus the application of the lemma (relative to $\mathcal{A}$ and $\mathcal{B}$). As it happens, the equivalences as per Lemma 2.2 allow this lemma to carry over to $\mathcal{A}$ and $\mathcal{B}$, but in that case, why not state the lemma directly for those schedules (in which case you cannot do so until after defining those schedules, which currently is done within the proof of the main theorem).

S9: Page 11, line 4
Replace "an urgent job" with "a tight job" as urgency is only defined upon the actual start time of a job, not based upon its pending status.

S10: Page 11, lines 6–7 of Case (IIIa)
Please explicitly cite Lemma 3.1 as part of the final sentence as an aid to the reader (i.e., why you may conclude that j is executed in $\overline{\mathcal{E}}$ as flexible before time t'.)

S11: Page 11, statement of Lemma 3.3
Should capitalize "Case" in expression "if case (IV)"

S12: Page 11, last line
A bad page break in this case between "Case" and "(III)". In general, you may want to uniformly protect against these line and page breaks by making sure to use `Case~(III)` in latex source.

S13: Page 12, line 4 of caption to figure 3
Expression "$j = k$" should instead read "$j \neq k$"

S14: Page 12, line 7–8 of prose
As for Case (IIIa), it is indeed true, but it is not clear what "same" argument applies as the presumption of Case (IIIa) does not seem to imply that $\mathcal{S}^{\overline{\mathcal{D}}}_{k_{\overline{\mathcal{D}}}} \leq \mathcal{S}^{\mathcal{Z}}_{k_{\overline{\mathcal{D}}}}$

S15: Page 12, line 8 of prose
I'm not sure what to suggest, but was quite puzzled for a while when interpreting, "and in the subcase of (IIIb) when $\mathcal{S}^{\mathcal{Z}}_{\overline{\mathcal{D}}} > t$. I think part of my trouble was in correctly interpreting which instantiation of 't' was intended, as there is the t that was defined in regard to the application of case IV which began Lemma 3.3, but also the t based upon the presumed application of case (IIIb) which generates such a self-charge.

S16: Page 14, line -5
suggest replacing "by P_t we denote" with "we denote as P_t" for a smoother flow to the sentence

S17: Page 15, line 8
Since you explicitly assume "flexible" j here, you should probably quickly note that an urgent job is never preempted (as per your existing definition of (TR2), or by lemma if accepting the recommendation R2 suggested in this report)

S18: Page 15, line 13
Replace "preemption preemption" with "preemption"

S19: Page 15, first paragraph in Proof of Theorem 4.2
Despite your efforts to draw a distinction between TightRestart versus $\mathcal{X}$, you are confusing them with each other here. You say "If $\mathcal{X}$ is running a job k which is later preempted by l" yet the definition of $\mathcal{X}$ already presumes the removal of such preempted jobs. Presumably you mean to say TightRestart here, and in fact for all three current occurrences of X in this paragraph.

S20: Page 15, Figure 4 and caption
The caption claims this is schedule A, yet the figure itself labels the top schedule as X. Which is it? (or are they the same based upon this particular application of Lemma 2.2.

S21: Page 17, line 15
You state that $P_S \cup K - \{j\} = P^*_{r_l} \cup K$ yet this is not true. In particular $j \in P^*_{r_l}$ by definition, so you must write $P_S \cup K - \{j\} = P^*_{r_l} \cup K - \{j\}$

S22: Page 17, lines 17–18
The statement "TightRestart would then execute urgent jobs from s' until at least the time x_j" is not immediately obvious and deserved justification. In particular, Lemma 2.1 does not suffice, as that would only guarantee urgency up to the expiration of the minimal element of the set U, which may be from K. Presumably an argument based upon Jackson's rule could be provided.

S23: Page 17, line 20
Two significant problems here. The first clause of this sentence should read "by the feasibility of $\underline{U}$ at s'" (note U rather than K). That, together with Lemma 2.2(3) seems to be sufficient to infer that A completes all jobs in U.

It is not clear why the statement "the flexibility of U at s and Lemma 2.2(2)" is included here. Of note, set U may include items which arrived after time s, so it does not make sense to discuss that set's flexibility at time s.

S24: Page 19, lines 7–8 following Figure 6
The phrase "a randomly chosen one job" should either read "a randomly chosen job" or "one randomly chosen job."