

Price of fairness in two-agent single-machine scheduling problems

This is the peer reviewed version of the following article:

Original:

Agnetis, A., Chen, B., Nicosia, G., Pacifici, A. (2019). Price of fairness in two-agent single-machine scheduling problems. EUROPEAN JOURNAL OF OPERATIONAL RESEARCH, 276(1), 79-87 [10.1016/j.ejor.2018.12.048].

Availability:

This version is available <http://hdl.handle.net/11365/1082540> since 2019-10-23T15:23:41Z

Published:

DOI: <http://doi.org/10.1016/j.ejor.2018.12.048>

Terms of use:

Open Access

The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. Works made available under a Creative Commons license can be used according to the terms and conditions of said license.

For all terms of use and more information see the publisher's website.

(Article begins on next page)

Price of Fairness in Two-Agent Single-Machine Scheduling Problems*

Alessandro Agnetis[†] Bo Chen[‡] Gaia Nicosia[§] Andrea Pacifici[¶]

September 13, 2018

Abstract

We investigate the concept of price of fairness in resource allocation and apply it to two-agent single-machine scheduling problems, in which two agents, each having a set of jobs, compete for use of a single machine to execute their jobs. We consider the situation where one agent aims at minimizing the total of the completion times of his jobs, while the other seeks to minimize the maximum tardiness with respect to a common due date for her jobs. We first explore and propose a definition of utility, then we study both max-min and proportionally fair solutions, providing a tight bound on the price of fairness for each notion of fairness. We extend our study further to the problem in which both agents wish to minimize the total of the completion times of their own jobs.

Keywords: two-agent scheduling, max-min fairness, Kalai-Smorodinsky fairness, proportional fairness, price of fairness

1 Introduction

Fairness issues arise in many applications and are studied in different research areas of economics, operations research and mathematics, to name a few. Fairness concepts have been widely studied in the context of fair division problems and in many other application scenarios (see, e.g., Brams and Taylor (1996) for a general overview). In cooperative game theory, classical two-player bargaining is a class of crucial problems that encapsulate the difficulty of fair division of a limited resource between two agents by suitably comparing their utility functions (Chun, 1988; Forsythe et al., 1994; Kalai and Smorodinsky, 1975).

In this paper, we address fairness concepts in the context of classical single-machine scheduling. There are two agents and each owns a set of jobs, which must be scheduled on a common processing resource. Each schedule implies a certain *utility* for each agent. We use weighted sum of the agents' utilities as an index of collective satisfaction (system utility) and we refer to any solution that maximizes system utility as a *system optimum*. Recently, a certain amount of

*This paper is to appear in *European Journal of Operational Research*

[†]Dipartimento di Ingegneria dell'Informazione e Scienze Matematiche, Università degli Studi di Siena, Via Roma 56, 53100 Siena, Italy. agnetis@dii.unisi.it

[‡]Corresponding author: Warwick Business School, University of Warwick, Coventry, CV4 7AL, England, United Kingdom b.chen@warwick.ac.uk

[§]Dipartimento di Ingegneria, Università degli studi Roma Tre, Via della Vasca Navale 79, 00146 Roma, Italy. nicosia@ing.uniroma3.it

[¶]Dipartimento di Ingegneria Civile e Ingegneria Informatica, Università degli Studi di Roma "Tor Vergata", via del Politecnico 1, 00133 Roma, Italy. andrea.pacifici@uniroma2.it

literature has addressed scheduling problems in which jobs are partitioned among two or more agents (Agnetais et al., 2014). In two-agent situations, research mainly focuses on complexity and algorithmic results for the problems of (a) maximizing an agent’s utility under the constraint that the other agent’s utility does not fall below a certain threshold, (b) generating all Pareto optimal solutions, or (c) maximizing system utility. Starting from classical basic single-machine models, research is now addressing these problems in increasingly complex settings (see, e.g., Tang et al., 2017). Now, even if it maximizes system utility, a system optimum may well be highly unfair to (and hence unacceptable by) the worse-off agent. Rather, a solution that incorporates some criterion of fairness may be more acceptable. In this paper, although we also consider complexity issues related to finding fair solutions, our main concern is to investigate how much system utility has to be sacrificed in order to reach a fair solution. The quantity that captures this concept is known as the *price of fairness* (PoF). Given a bargaining problem, and a certain definition of fairness, the PoF is the maximum (over all instances of the given problem) relative loss in overall system utility of a fair solution with respect to the system optimum. We are not aware of any such study in the context of classical single-machine scheduling.

Depending on the specific problem setting and also on the agent perception of what a fair solution is, assorted definitions of fairness can be found in the scientific literature. In this paper we focus on two of the most popular fairness notions, namely *Kalai-Smorodinsky fairness* (Kalai and Smorodinsky, 1975) and *proportional fairness* (Kelly, Maulloo, and Tan, 1998).

Caragiannis et al. (2012) are the first to introduce the concept of PoF in the context of fair allocation problems. They compare the total value of agents’ utilities at system optimum with the maximum total of agents’ utilities obtained over all fair solutions based on several notions of fairness, such as proportionality, envy-freeness and equitability. Bertsimas, Farias, and Trichakis (2011) focus on proportional fairness and max-min fairness and provide a tight characterization of PoF for a broad family of allocation problems with compact and convex utility sets of the agents. Nicosia, Pacifici, and Pferschy (2017) establish a number of properties for the PoF, which hold for any general multi-agent problem without any special assumption on the agents’ utilities, focusing on max-min, Kalai-Smorodinsky and proportional fairness. Situations in which the agents pursue the minimization of their costs (rather than maximization of their utilities) have been dealt with by Ertogral and Wu (2000), who derive a measure of fairness among a set of supply chain members. Through an experimental study, they evaluate the degradation of the solution quality with respect to a system optimum, when imposing some fairness level. Another example of fairness in the presence of cost allocations can be found in Bohm and Larsen (1994).

For the scheduling setting addressed in this paper, we can refer the reader to the situations addressed by multi-agent scheduling models (Agnetais et al., 2014), or by decentralized scheduling models (Wellman et al., 2001). Whenever two or more agents compete for the use of a resource over time, the fairness issue arises naturally. As an example, scheduling aircraft on a runway entails fairness considerations, as different flights belong to different companies (Soomer and Franx, 2008). The problem addressed in Section 3 applies to an incoming flight having a nominal arrival time, which has to be inserted within a sequence of flights belonging to another carrier. As another example, one can consider the production problem addressed by Lang and Fink (2015), in which different companies (or departments) jointly hold the production facilities (machines) and compete for their usage, each pursuing its own objective, as is the case of industrial districts (Albino et al., 2005). In order to reach a compromised schedule, a facilitator may impose a negotiation protocol (as proposed by Lang and Fink (2015)), or can enforce a fair schedule if the agents are willing to disclose the information of their jobs, as in the approach adopted in this paper.

Our study of fairness in this paper is organized as follows. In Section 2 we introduce the basic

notation and the concepts of *agent* and *system* utilities, define the notions of fair solutions (such as *Kalai-Smorodinsky* and *proportionally-fair* solutions) and that of *Price of Fairness* (PoF). Sections 3 and 4 are devoted respectively to two scheduling problems with different agents' objectives. In the first scheduling problem, one agent aims to minimize the total completion time of his jobs, and the other the maximum tardiness with respect to a common due date for her jobs. In the second scheduling problem, both agents wish to minimize the total completion time of their jobs. For the above two problems, we investigate whether a specific kind of fair solution exists and, if so, the complexity of finding one. Finally we establish (bounds on) the PoF in the four resulting cases. We provide a further discussion of our results in Section 5.

2 Preliminaries

The problems we address are formally described as follows. There are two agents, namely A and B . Each agent owns a set of jobs and each job has to be performed (for a given period of time) by a single machine that is shared by the two agents. A *feasible* schedule is an assignment of all jobs to their respective starting times such that processing of each job is uninterrupted once started and the machine can only process one job at a time. We use the terms A -jobs and B -jobs to refer to the jobs of the respective agents.

Given a feasible schedule σ , we let $f^A(\sigma)$ and $f^B(\sigma)$ denote the *cost values* for the two agents respectively. If two schedules σ and σ' incur the same costs for the two agents, i.e., $f^i(\sigma) = f^i(\sigma')$ for $i = A, B$, then we are not interested in distinguishing between them and consider σ and σ' as equivalent and regard them as a single solution. A schedule σ^* is said to be *Pareto optimal*, if there is no feasible schedule σ , such that $f^i(\sigma) \leq f^i(\sigma^*)$ for $i = A, B$ and at least one of the two inequalities is strict. In other words, given a Pareto optimal schedule, any other feasible schedule that makes one agent better off will make the other agent worse off. Throughout the paper, we consider as *bargaining set* the set Σ_P of Pareto optimal schedules, as they include all sensible compromise schedules.

For each schedule $\sigma \in \Sigma_P$, we want to define utility values $u^A(\sigma)$ and $u^B(\sigma)$, so that, for $i = A, B$, $u^i(\sigma) \geq 0$ and $u^i(\sigma)$ increases as $f^i(\sigma)$ decreases. To this end, we propose the following definition of utility (an alternative definition is briefly discussed in Section 5). Let $f_\infty^i = \max\{f^i(\sigma) : \sigma \in \Sigma_P\}$. For *regular functions* $f^i(\sigma)$ (i.e., functions that are non-decreasing in the completion times of agent i 's jobs), this is the minimum cost agent i bears if his jobs are scheduled after all the jobs of the other agent. Then we define agent utilities as follows:

$$u^i(\sigma) = f_\infty^i - f^i(\sigma), \quad i = A, B. \quad (1)$$

In other words, the utility of an agent is represented by the *saving* achieved with respect to the worst schedule for the agent.

Let us now consider the concept of *system utility*. Perhaps the most common definition, mostly adopted in economics, is to define system utility as the *sum* of the individuals' utilities. This is also the definition given by Bertsimas, Farias, and Trichakis (2011). Such a definition is especially appropriate if the agents hold similar objectives. However, this may not be the case in general under our scheduling setting, in which, e.g., one agent may be interested in minimizing the makespan of her jobs, while the other agent the total of his job completion times. For this reason, we consider a more general system utility, obtained as a (positively) weighted sum of the agents' individual utilities. Therefore, without loss of generality, given a schedule σ , for any exogenously given and fixed value $\alpha > 0$, we define the system utility $U(\sigma)$ as follows:

$$U(\sigma) = u^A(\sigma) + \alpha u^B(\sigma). \quad (2)$$

We also let σ^* denote the schedule that maximizes system utility (*system optimum*), i.e.,

$$U(\sigma^*) = \max_{\sigma \in \Sigma_P} \{U(\sigma)\}.$$

Denote $f^{i*} = \min\{f^i(\sigma) : \sigma \in \Sigma_P\}$. The main purpose of our study concerns the comparison in terms of system utility between a fair schedule and a system optimal schedule. We consider the following two fairness notions.

Kalai-Smorodinsky fairness. Given $\sigma \in \Sigma_P$, let

$$\bar{u}^i(\sigma) = \frac{u^i(\sigma)}{f_\infty^i - f^{i*}} \quad (3)$$

be the *normalized utility* of σ for agent i . Consider the set $\Sigma_{KS} \subseteq \Sigma_P$ of schedules that maximize the normalized utility of the agent who is worse-off, i.e.,

$$\Sigma_{KS} = \left\{ \sigma_{KS} : \sigma_{KS} = \arg \max_{\sigma} \min_{i=A,B} \{ \bar{u}^i(\sigma) \} \right\}. \quad (4)$$

We say that all the schedules in Σ_{KS} are *KS-fair*. Notice that Σ_{KS} is always nonempty. In bargaining games for which the bargaining set is compact and convex, $|\Sigma_{KS}| = 1$ (Kalai and Smorodinsky, 1975). Note that in our scheduling setting, it may be indeed $|\Sigma_{KS}| > 1$. This implies that some care has to be taken when defining the price of fairness.

Another classical notion of fairness, known as *max-min* fairness, arises if we use in Equation (4) the agents utility as defined in Equation (1), instead of the normalized utility (3). Indeed, the two concepts of Kalai-Smorodinsky and max-min solution coincide if $f_\infty^A - f^{A*} = f_\infty^B - f^{B*}$. Max-min fairness makes sense in contexts where the range of values of the agent objectives are comparable. This is not the case in our study, at least when the two agents pursue different objectives that take values which may be far apart from each other.

Proportional fairness. A schedule σ_{PF} is proportionally fair if, for any other Pareto optimal schedule σ , it holds that

$$\frac{u^A(\sigma) - u^A(\sigma_{PF})}{u^A(\sigma_{PF})} + \frac{u^B(\sigma) - u^B(\sigma_{PF})}{u^B(\sigma_{PF})} \leq 0. \quad (5)$$

In other words, when moving from schedule σ_{PF} to any other schedule σ , the relative benefit any one agent may obtain is at the cost of a no-smaller relative utility decrease of the other agent. This is actually the same rationale behind the concept of Nash Bargaining Solution (NBS) (Nash, 1950). In fact, the two concepts coincide when the bargaining set is compact and convex, in which case it is the (unique) allocation maximizing the product of the two agents' utilities. If this is not the case and the feasible set of alternatives is finite—as it is in the present setting—the problem of extending the concept of NBS is not trivial: A comprehensive discussion can be found in, e.g., Mariotti (1998). Indeed, one can always define the NBS as the one maximizing such product (see, e.g., Agnetis, de Pascale, and Pranzo (2009)), yet the existence of an NBS does not guarantee that there is a proportionally fair solution. On the other hand, one can show that if a proportionally fair solution exists, then it is unique (Nicosia, Pacifici, and Pferschy, 2017). Clearly, the latter general property holds true also in our scheduling setting.

Given a certain scheduling problem, let \mathcal{I} denote the set of all its instances. Given an instance $I \in \mathcal{I}$, let $\sigma^*(I)$ and $\Sigma_F(I)$ denote a system optimum and the set of fair schedules for I . We

next introduce a quantity that indicates how much utility one has to give up (with respect to the system optimum) in order to get a fair solution. Note that there may be more than one fair solution (i.e., one may have $|\Sigma_F(I)| > 1$), differing in terms of global utility. So, one might question whether the loss in system efficiency should be measured with respect to the best or to the worst fair solution. The concept of fairness typically arises as something that has to be enforced by a *super-partes* facilitator, who trades some system efficiency for fairness. In this respect, it seems reasonable that the facilitator chooses the *best* fair solution. This is also the point of view adopted in other papers, such as Karsu and Morton (2015) and Naldi et al. (2016), where a fair solution is selected by a third party or central decision maker who takes into account the utility functions of the two agents (hence their degree of satisfaction), *and* the overall system utility.

Hence, by suitably modifying the definition by Bertsimas, Farias, and Trichakis (2011), we define the *price of fairness* as

$$\text{PoF} = \sup_{I \in \mathcal{I}} \min_{\sigma_F \in \Sigma_F} \left\{ \frac{U(\sigma^*(I)) - U(\sigma_F(I))}{U(\sigma^*(I))} \right\}. \quad (6)$$

Note that this is a definition similar to the well-known *Price of Stability* (Anshelevich et al., 2004), replacing the role of Nash equilibrium with fairness.

Hereafter, we indicate with PoF_{KS} and PoF_{PF} the price of Kalai-Smorodinsky fairness and proportional fairness, respectively. In this paper we investigate the values of PoF_{KS} and PoF_{PF} in the following scheduling problems:

1. Agent A pursues the minimization of the sum of completion times of his jobs, while agent B is interested in minimization of the maximum tardiness of her jobs with respect to a common due date d , where the tardiness of a job with completion time C is defined as $\max\{0, C - d\}$. We denote such a problem as

$$1|d_j^B = d|(\sum C_j^A, T_{\max}^B).$$

Note that this problem includes as a special case ($d = 0$) the situation where agent B wishes to minimize the makespan of her jobs, i.e., $1||(\sum C_j^A, C_{\max}^B)$.

2. Both agents pursue the minimization of the sum of the completion times of their jobs. We denote the problem as

$$1||(\sum C_j^A, \sum C_j^B).$$

For problem $1|d_j^B = d|(\sum C_j^A, T_{\max}^B)$, we show that $\text{PoF}_{\text{KS}} = 2/3$ and $\text{PoF}_{\text{PF}} = 1/2$. Moreover, we show that, if the jobs are given in SPT order, i.e., in order of non-decreasing processing times, in logarithmic time a proportionally fair solution can be computed or proved that it does not exist. For problem $1||(\sum C_j^A, \sum C_j^B)$, we show that $\text{PoF}_{\text{KS}} \geq 2/3$ and $\text{PoF}_{\text{PF}} = 1/2$. Finally, we show that it is NP-hard to compute a KS or a proportionally fair schedule even if such a schedule exists.

3 Problem $1|d_j^B = d|(\sum C_j^A, T_{\max}^B)$

Let m denote the number of jobs of agent A , and p_1, p_2, \dots, p_m their processing times. Denote $P = \sum_{j=1}^m p_j$. We assume that A -jobs are numbered in shortest processing time (SPT) order. Using a standard pairwise exchange argument it is easy to prove that in any Pareto optimal schedule A -jobs are scheduled in SPT order.

For what concerns agent B , we observe that, in any schedule σ , only the maximum completion time of a B -job is relevant to the value of $f^B(\sigma)$. Therefore, in each schedule $\sigma \in \Sigma_P$ all the B -jobs are scheduled consecutively in one block, and as a consequence, with no loss of generality, in this section we assume that agent B has a *single* job of length K . (For the sake of simplicity, and with a small abuse of notation, we refer sometimes to such a job as *job K* .)

Since in any Pareto optimal schedule the A -jobs are scheduled in SPT order, a Pareto optimal schedule is completely characterized by the number of A -jobs following the job K . Let σ_ℓ denote the schedule in which ℓ jobs follow job K , and let P_ℓ denote the total processing time of the jobs scheduled after K , i.e.,

$$P_\ell = \sum_{j=m-\ell+1}^m p_j.$$

Hence $P_m = P$.

If $K \leq d$, then there exists at least a value of $\ell \in \{0, 1, \dots, m\}$ such that the completion time of job K in σ_ℓ does not exceed the due date d , in which case we say job K is *early*. Let ℓ_d be the smallest among such values of ℓ , noticing that, if $K > d$, job K cannot be early in any schedule σ_ℓ :

$$\ell_d = \begin{cases} \min_{0 \leq \ell \leq m} \{ \ell : \sum_{j=1}^{m-\ell} p_j + K \leq d \}, & \text{if } K \leq d; \\ m, & \text{if } K > d. \end{cases} \quad (7)$$

The above definition implies that job K is early in σ_ℓ if and only if $K \leq d$ and $\ell \geq \ell_d$, i.e.,

$$P_\ell \geq P + K - d \xleftrightarrow{K \leq d} \ell \geq \ell_d. \quad (8)$$

In fact, the number of A -jobs processed before job K in the best Pareto optimal solution for agent B is $m - \ell_d$.

Example 1. Consider the following instance of problem $1|d_j^B = d|(\sum C_j^A, T_{\max}^B)$. Agent A has $m = 9$ jobs with processing times $p_j = 3, 4, 5, 7, 8, 16, 26, 88, 130$, so that $P = 287$. The B -job has a processing time $K = 33$ with common due date $d = 43$. We have $\ell_d = 7$. In fact $p_1 + p_2 + K = 40 \leq d$ and $p_1 + p_2 + p_3 + K = 45 > d$. Therefore, $P_{\ell_d} = \sum_{\ell=3}^9 p_\ell = 280$, $P + K - d = 277$. See the upper pane in Figure 1 for an illustration. Clearly, among Pareto optimal schedules, σ_0 is the best schedule for agent A (and the worst for agent B) and σ_{ℓ_d} the best schedule for agent B (and the worst for agent A), i.e., $\bar{u}^A(\sigma_0) = \bar{u}^B(\sigma_{\ell_d}) = 1$, while $\bar{u}^A(\sigma_{\ell_d}) = \bar{u}^B(\sigma_0) = 0$. Observe that schedules σ_ℓ for $\ell > \ell_d$ are not considered, since they are not Pareto optimal.

The lower pane in Figure 1 corresponds to the same instance of the problem except that $d = 30$.

3.1 Utility values of Pareto optimal schedules

Note that if $\ell_d = 0$, then schedule σ_0 provides maximum utility to each agent and hence it is trivially fair, giving $\text{PoF} = 0$ under any fairness measure. Therefore, we assume without loss of generality that $1 \leq \ell_d \leq m$, which implies that $P + K - d > 0$.

To simplify notation, in the remainder of this section we write $f^i(\ell)$, $u^i(\ell)$, $\bar{u}^i(\ell)$ and $U(\ell)$, instead of $f^i(\sigma_\ell)$, $u^i(\sigma_\ell)$, $\bar{u}^i(\sigma_\ell)$ and $U(\sigma_\ell)$ for $i = A, B$. In Table 1 the values of cost, utility and normalized utility for each Pareto optimal solution σ_ℓ are reported for both $K \leq d$ and $K > d$. Note that, in particular, for the case $K \leq d$, $f^B(\ell) = P - P_\ell + K - d > 0$ when $\ell < \ell_d$, and $f^B(\ell) = 0$ when $\ell = \ell_d$. Finally we obtain

$$U(\ell) = K(\ell_d - \ell) + \alpha \min\{P_\ell, P + K - d\}. \quad (9)$$




Figure 1: An example with schedule σ_{ℓ_d} when $K < d$ and $K > d$

	Agent A	Agent B	
		$K \leq d$	$K > d$
f^{i*}	$\sum_{j=1}^m \sum_{h=1}^j p_h$	0	$K - d$
f_{∞}^i	$f^{A*} + K\ell_d$	$P + K - d$	$P + K - d$
$f^i(\ell)$	$f^{A*} + K\ell$	$\max\{0, P - P_{\ell} + K - d\}$	$P - P_{\ell} + K - d$
$u^i(\ell)$	$K(\ell_d - \ell)$	$\min\left\{\underbrace{P_{\ell}}_{\ell < \ell_d}, \underbrace{P + K - d}_{\ell = \ell_d}\right\}$	P_{ℓ}
$\bar{u}^i(\ell)$	$(\ell_d - \ell)/\ell_d$	$\min\{P_{\ell}, P + K - d\}/(P + K - d)$	P_{ℓ}/P

Table 1: Cost and utility values for schedule σ_{ℓ}

It is convenient to plot normalized utilities of the two agents on a chart with ℓ on the horizontal axis and $\bar{u}^i(\ell)$ on the vertical axis (see Figure 2 for an illustration). Recalling that $\ell_d = m$ when $K > d$, as ℓ varies from 0 to ℓ_d , clearly $\bar{u}^A(\sigma)$ linearly decreases from 1 to 0. The linear function $\bar{u}^A(x) = 1 - (x/\ell_d)$, for $x \in [0, \ell_d]$, coincides with $\bar{u}^A(\sigma)$ for integer values of x . On the contrary, as ℓ varies from 0 to ℓ_d , $\bar{u}^B(\sigma)$ goes from 0 to 1. If we consider the piecewise linear curve $\tilde{u}^B(x)$ obtained from joining consecutive points on the chart, we observe that such a curve is concave, since $p_{m-\ell+1} \geq p_{m-\ell}$, for all $0 \leq \ell \leq \ell_d$. As a consequence, the curve $\tilde{u}^B(x)$ is entirely *on or above* the line x/ℓ_d . Actually, it coincides with such a line if the ℓ_d longest jobs of agent B have the same length.

In order to better understand the properties of the schedules σ_{ℓ} , we analyze the trends of $f^i(\ell)$ and $\bar{u}^i(\ell)$ when moving from $\sigma_{\ell-1}$ to σ_{ℓ} . To begin with, it is easy to observe that f^A increases by K , so \bar{u}^A decreases by $K/(K\ell_d) = 1/\ell_d$ (or $1/m$ if $K > d$). For what concerns the objective of agent B, we need to distinguish the cases of $K \leq d$ and $K > d$.

If $K \leq d$ and $\ell \leq \ell_d - 1$, f^B decreases by $p_{m-\ell+1}$, and therefore \bar{u}^B increases by $p_{m-\ell+1}/(P + K - d)$. In the schedule σ_{ℓ_d} , K is not tardy, and we denote its slack time as $\Delta = d - (P - P_{\ell_d} + K) \geq 0$ (see Figure 1, where $\Delta = 3$), with $\Delta = 0$ if and only if K completes exactly at d . Hence, moving from σ_{ℓ_d-1} to σ_{ℓ_d} , $f^B(\sigma)$ decreases by $(p_{m-\ell_d+1} - \Delta)/(P + K - d)$, so that $\bar{u}^A(\ell_d) = 1$.

If $K > d$, $\ell_d = m$ and in the best Pareto optimal schedule σ_m for agent B, job K is already tardy. As above, f^B decreases by $p_{m-\ell+1}$ while $\bar{u}^B(\sigma)$ increases by $p_{m-\ell+1}/P$ for all values $\ell = 1, \dots, m$.




Figure 2: The graphics of normalized utility for Example 1

3.2 KS fairness

Let us characterize σ_{KS} among all schedules of Σ_P . First of all, we note that the case $\ell_d = 1$ can be ruled out. In fact, in this case Σ_P only consists of σ_0 and σ_1 . These two schedules give maximum (i.e., 1) normalized utility to one agent and zero (normalized) utility to the other, and at least one of them is also a system optimum. Such a system optimum can be therefore taken as σ_{KS} , and also in this case $\text{PoF}_{KS} = 0$. In this case, the KS fairness concept fails to capture any fairness. Hence, from now on we assume $\ell_d \geq 2$ in our consideration of KS fairness.

In what follows, we let ℓ_{KS} denote the index such that $\sigma_{\ell_{KS}} \equiv \sigma_{KS}$. Consider the point $x = x_{KS}$ where the two curves $\tilde{u}^A(x)$ and $\tilde{u}^B(x)$ intersect in Figure 2. From the definition of KS fairness, one has that if x_{KS} is integer, then $x_{KS} = \ell_{KS}$. Otherwise, either $\ell_{KS} = \lfloor x_{KS} \rfloor$ or $\ell_{KS} = \lceil x_{KS} \rceil$, depending on whether or not $\min\{\bar{u}^A(\lfloor x_{KS} \rfloor), \bar{u}^B(\lfloor x_{KS} \rfloor)\} \leq \min\{\bar{u}^A(\lceil x_{KS} \rceil), \bar{u}^B(\lceil x_{KS} \rceil)\}$. Due to the concavity of $\tilde{u}^B(x)$, the intersection certainly takes place for $x \leq \ell_d/2$. In conclusion, if we denote by $\bar{\ell}$ the smallest ℓ such that $\bar{u}^A(\ell) < \bar{u}^B(\ell)$, then $\ell_{KS} \in \{\bar{\ell}-1, \bar{\ell}\}$. From the monotonicity of $\bar{u}^A(\cdot)$ and $\bar{u}^B(\cdot)$, the value of $\bar{\ell}$ (and therefore ℓ_{KS}) can be found through a binary search between 2 and $\lfloor \ell_d/2 \rfloor$. Hence, the following theorem holds.

Theorem 1. *If the A-jobs are already sorted in SPT order, then it is possible to find a Kalai-Smorodinsky fair schedule in $O(\log m)$ time.*

Example 1 (continued). Since $\ell_d = 7$ is odd, the curves $\tilde{u}^B(x)$ in Figure 2 intersect at the fractional point $1 < x < 2$ while $\ell_{KS} = 2$ (see shaded rectangle R). In addition, the curve $\tilde{u}'^B(\ell)$ illustrates the case corresponding to a different set of processing time values for agent A: $p'_j = 1, 6, 32, 36, 39, 40, 41, 42, 50$ with $P = 287$. Everything remains the same for agent B and for the quantities $\ell_d = 7$, $P_{\ell_d} = 280$, $P + K - d = 277$ and $\Delta = 3$. In this case, because $\tilde{u}'^B(x)$ is only

slightly concave, the intersection x' lies between $\lfloor \ell_d/2 \rfloor$ and $\lceil \ell_d/2 \rceil$. We are in the case $\ell_d = \lfloor \ell_d/2 \rfloor$, and $\tilde{u}^{tB}(\ell_{\text{KS}}) < 1/2$ (see rectangle R').

Lemma 1. *If ℓ_d is even, then $\bar{u}^B(\ell_{\text{KS}}) \geq 1/2$.*

Proof. Since $x_{\text{KS}} \leq \ell_d/2$ and $\ell_d/2$ is integer, $\lceil x_{\text{KS}} \rceil \leq \ell_d/2$. Because of the concavity of $\tilde{u}^B(x)$, if $\bar{u}^B(\lfloor x_{\text{KS}} \rfloor) \geq 1/2$ we are done. Otherwise, if $\bar{u}^B(\lfloor x_{\text{KS}} \rfloor) < 1/2$, then certainly $\bar{u}^B(\lceil x_{\text{KS}} \rceil) \geq 1/2$, and also in this case $\bar{u}^B(\ell_{\text{KS}}) \geq 1/2$. \square

Lemma 2. *For any $\ell_d \geq 2$, it holds that $\ell_{\text{KS}} \leq \ell_d/2$.*

Proof. Let us distinguish the cases $\lceil x_{\text{KS}} \rceil \leq \ell_d/2$ and $\lceil x_{\text{KS}} \rceil > \ell_d/2$. In the former case, ℓ_{KS} can be either $\lfloor x_{\text{KS}} \rfloor$ or $\lceil x_{\text{KS}} \rceil$, and we are done. Let us therefore consider the case $\lceil x_{\text{KS}} \rceil > \ell_d/2$ (according to Lemma 1 this can occur only if ℓ_d is odd). In this case (see Figure 2 for an illustration),

$$\bar{u}^A(\lfloor \ell_d/2 \rfloor) > \bar{u}^B(\lfloor \ell_d/2 \rfloor) > \bar{u}^A(\lceil \ell_d/2 \rceil)$$

and

$$\bar{u}^B(\lceil \ell_d/2 \rceil) > \bar{u}^A(\lceil \ell_d/2 \rceil).$$

Hence $\ell_{\text{KS}} = \lfloor x_{\text{KS}} \rfloor$. \square

Lemma 3. *If $\bar{u}^B(\ell_{\text{KS}}) < 1/2$, then $\ell_{\text{KS}} = \lfloor \ell_d/2 \rfloor$.*

Proof. Suppose to the contrary that $\ell_{\text{KS}} < \lfloor \ell_d/2 \rfloor$ (recall that ℓ_{KS} cannot exceed $\ell_d/2$ due to Lemma 2). In this case, $\bar{u}^A(\ell_{\text{KS}}) \geq 1/2$ and hence $\bar{u}^A(\ell_{\text{KS}}) \geq \bar{u}^B(\ell_{\text{KS}})$. We consider two cases.

Case (a): $\bar{u}^B(\lfloor \ell_d/2 \rfloor) \geq \bar{u}^A(\lfloor \ell_d/2 \rfloor)$. Since $\bar{u}^A(\lfloor \ell_d/2 \rfloor) \geq 1/2$ and $\bar{u}^B(\ell_{\text{KS}}) < 1/2$, one has that $\min\{\bar{u}^A(\lfloor \ell_d/2 \rfloor), \bar{u}^B(\lfloor \ell_d/2 \rfloor)\} > \bar{u}^B(\ell_{\text{KS}})$, which contradicts the fact that $\sigma_{\ell_{\text{KS}}}$ is KS fair.

Case (b): $\bar{u}^B(\lfloor \ell_d/2 \rfloor) \leq \bar{u}^A(\lfloor \ell_d/2 \rfloor)$. Then $\min\{\bar{u}^A(\lfloor \ell_d/2 \rfloor), \bar{u}^B(\lfloor \ell_d/2 \rfloor)\} = \bar{u}^B(\lfloor \ell_d/2 \rfloor)$, and in turn, since $\bar{u}^B(\cdot)$ is an increasing function, $\bar{u}^B(\lfloor \ell_d/2 \rfloor) > \bar{u}^B(\ell_{\text{KS}})$. In this case, $\sigma_{\ell_{\text{KS}}}$ cannot be KS fair. \square

Now let us establish bounds on PoF_{KS} and their tightness.

Theorem 2. *If ℓ_d is even, then $(U(\sigma^*) - U(\sigma_{\text{KS}}))/U(\sigma^*) \leq 1/2$; if $\ell_d \geq 3$ is odd, then $(U(\sigma^*) - U(\sigma_{\text{KS}}))/U(\sigma^*) \leq (\ell_d + 1)/(2\ell_d)$.*

Proof. It suffices to prove that, for all $0 \leq \ell \leq \ell_d$ with $2 \leq \ell_d \leq m$, the following hold:

$$\frac{U(\ell_{\text{KS}})}{U(\ell)} \geq \begin{cases} 1/2, & \ell_d \text{ even;} \\ (\ell_d - 1)/(2\ell_d), & \ell_d \text{ odd.} \end{cases} \quad (10)$$

Let us first focus on the situation of $K \leq d$ with two steps of proving (10). Since $1/2 \geq (\ell_d - 1)/(2\ell_d)$, our first step is to show the first inequality of (10) for even ℓ_d and for odd ℓ_d with

$$\bar{u}^B(\ell_{\text{KS}}) \geq 1/2. \quad (11)$$

Note that inequality (11) is automatically satisfied when ℓ_d is even according to Lemma 1. Also note that when $\ell_d \geq 2$ we have $\ell_{\text{KS}} < \ell_d$, which gives $u^B(\ell_{\text{KS}}) = P_{\ell_{\text{KS}}}$ and hence inequality (11) becomes

$$2P_{\ell_{\text{KS}}} \geq P + K - d. \quad (12)$$

Recalling the definition of system utility (9), the first inequality in (10) is equivalent to the following, according to Table 1:

$$(\ell_d - 2\ell_{\text{KS}} + \ell)K \geq \alpha(P_\ell - 2P_{\ell_{\text{KS}}}), \quad \text{if } \ell < \ell_d; \quad (13a)$$

$$2(\ell_d - \ell_{\text{KS}})K \geq \alpha(P + K - d - 2P_{\ell_{\text{KS}}}), \quad \text{if } \ell = \ell_d. \quad (13b)$$

Now consider inequality (13a). Since in this case $\ell < \ell_d$, from (8) one has that $P_\ell < P + K - d$, which, along with (12), implies $P_\ell < 2P_{\ell_{\text{KS}}}$ and hence the right-hand side of (13a) is negative. Similarly, when $\ell = \ell_d$, the right-hand side of (13b) is non-positive due to (12). On the other hand, the left-hand sides of inequalities in (13) are both non-negative according to Lemma 2. Therefore, both inequalities in (13) hold indeed.

Our second step of proving (10) is to show its second inequality for odd ℓ_d with $\bar{u}^B(\ell_{\text{KS}}) < 1/2$. From Lemma 3, we get $\ell_{\text{KS}} = \lfloor \ell_d/2 \rfloor = (\ell_d - 1)/2$. The definition of KS schedule implies

$$\bar{u}^B(\ell_{\text{KS}}) \geq \bar{u}^A(\lceil \ell_d/2 \rceil),$$

which together with the fact that $\bar{u}^A(\lceil \ell_d/2 \rceil) = (\ell_d - 1)/(2\ell_d)$ (see Figure 2) implies that

$$2\ell_d P_{\ell_{\text{KS}}} \geq (\ell_d - 1)(P + K - d). \quad (14)$$

With $\ell_{\text{KS}} = (\ell_d - 1)/2$ the second inequality in (10) is equivalent to the following, according to Table 1:

$$(2\ell_d + (\ell_d - 1)\ell)K \geq \alpha((\ell_d - 1)P_\ell - 2\ell_d P_{\ell_{\text{KS}}}), \quad \text{if } \ell < \ell_d; \quad (15a)$$

$$\ell_d(\ell_d + 1)K \geq \alpha((\ell_d - 1)(P + K - d) - 2\ell_d P_{\ell_{\text{KS}}}), \quad \text{if } \ell = \ell_d. \quad (15b)$$

Again from (8), when $\ell < \ell_d$, one has $P_\ell < P + K - d$, which together with (14) implies $2\ell_d P_{\ell_{\text{KS}}} > (\ell_d - 1)P_\ell$ and hence the right-hand side of inequality (15a) is negative. Also inequality (14) implies that the right-hand sides of inequality (15b) is non-positive. Since the left-hand sides of both inequalities in (15) are obviously non-negative, both inequalities in (15) hold indeed.

Moving from situation of $K \leq d$ to $K > d$, all the arguments above remain exactly the same except inequalities (12) and (14) are replaced, respectively, by $2P_{\ell_{\text{KS}}} \geq P$ and $2\ell_d P_{\ell_{\text{KS}}} \geq (\ell_d - 1)P$, while inequality pairs (13) and (15) are replaced, respectively, by (13a) and (15a) without the condition “if $\ell < \ell_d$ ”. \square

We are now in the position to state our main result. Recall that we have assumed that $\ell_d \geq 2$ in our consideration for KS fairness. Since $(\ell_d + 1)/(2\ell_d) \leq 2/3$ for all odd $\ell_d \geq 3$, the following theorem is immediate.

Theorem 3. *For problem 1 $|d_j^B = d|(\sum C_j^A, T_{\max}^B)$, we have $\text{PoF}_{\text{KS}} = \frac{2}{3}$.*

Proof. We only need to show that the bound of $2/3$ can be attained. As part of the problem specification, parameter $\alpha > 0$ is exogenously given and fixed, while parameter $d \geq 0$ is part of the input. We are to specify a set of instances with d as a parameter (and hence $d = 0$ as a special case) and demonstrate that these instances make the PoF arbitrarily close to $2/3$ for any $d \geq 0$. Consider the following instance \mathcal{I}_ε of the problem indexed by sufficiently small $\varepsilon > 0$. While the B -job has length K , agent A has an odd number $\ell_d \geq 3$ of jobs, with $\ell_d - 1$ jobs having length p and one job having length $p + \varepsilon$. We choose K and p such that $d \leq K$ and $K/\alpha < p$. Then the system optimal schedule σ^* sequences the jobs according to the Smith’s rule (Smith, 1956) of WSPT (weighted shortest processing time), that is, σ^* starts with the B -job and finishes with the longest A -job, with all the remaining $\ell_d - 1$ A -jobs in the middle. Therefore, we have $u^A(\sigma^*) = 0$




Figure 3: Schedules σ_{PF} and σ^* in the example on PoF tightness

and $u^B(\sigma^*) = (\ell_d p + \varepsilon + K) - K = \ell_d p + \varepsilon$ and hence $U(\sigma^*) = \alpha(\ell_d p + \varepsilon)$. It is easy to check that there exists a unique KS fair schedule, obtained with $\ell_{KS} = (\ell_d - 1)/2$. One has $u^A(\sigma_{KS}) = \frac{\ell_d + 1}{2}K$ and $u^B(\sigma_{KS}) = (\ell_d p + \varepsilon + K) - ((\frac{\ell_d + 1}{2})p + K) = \frac{\ell_d - 1}{2}p + \varepsilon$, and hence $U(\sigma_{KS}) = \frac{\ell_d + 1}{2}K + \alpha(\frac{\ell_d - 1}{2}p + \varepsilon)$. As a result, we have

$$\frac{U(\sigma^*) - U(\sigma_{KS})}{U(\sigma^*)} = \frac{\frac{\ell_d + 1}{2}(p\alpha - K)}{(\ell_d p + \varepsilon)\alpha},$$

which approaches $(\ell_d + 1)/(2\ell_d)$ (i.e., $2/3$ if we set $\ell_d = 3$) as p goes to infinity and ε to zero. \square

3.3 Proportional Fairness

Let us now consider the price of proportional fairness. It is known from Nicosia, Pacifici, and Pferschy (2017) that if a proportionally fair solution exists, then PoF_{PF} is at most $1/2$. We show that for problem $1|d_j^B = d|(\sum C_j^A, T_{\max}^B)$ this bound can actually be attained, hence the bound is tight.

Let $\alpha > 0$ be exogenously fixed for defining the system utility. Given any input parameter $d \geq 0$, consider an instance in which agent A has m (with m even) identical unit-length jobs and agent B has a single job of length K , where $K \geq d$ and $K > \alpha$. Figure 3 illustrates the fair schedule σ_{PF} and the overall optimal schedule σ^* , which is according to the Smith's rule, in WSPT order, as we mentioned in the proof of Theorem 3. We have

$$\begin{aligned} f_\infty^A &= \frac{m(m+1)}{2} + Km, & f_\infty^B &= K + m; \\ f^A(\sigma^*) &= \frac{m(m+1)}{2}, & f^B(\sigma^*) &= K; \\ f^A(\sigma_{PF}) &= \frac{m(m+1)}{2} + \frac{mK}{2}, & f^B(\sigma_{PF}) &= K + \frac{m}{2}; \end{aligned}$$

from which we obtain

$$\begin{aligned} u^A(\sigma^*) &= Km, & u^B(\sigma^*) &= 0; \\ u^A(\sigma_{PF}) &= \frac{mK}{2}, & u^B(\sigma_{PF}) &= \frac{m}{2}. \end{aligned}$$

Hence $U(\sigma^*) = Km$ and $U(\sigma_{PF}) = \frac{m(K+\alpha)}{2}$. Therefore,

$$\text{PoF}_{PF} = \frac{U(\sigma^*) - U(\sigma_{PF})}{U(\sigma^*)} = 1 - \frac{m(K+\alpha)/2}{Km} = \frac{K-\alpha}{2K},$$

which approaches to $1/2$ as K goes to infinity. This shows the following result.

Theorem 4. For problem $1|d_j^B = d|(\sum C_j^A, T_{\max}^B)$, we have $\text{PoF}_{PF} = \frac{1}{2}$.

Existence of proportionally fair schedules

Now we show that the existence of a proportionally fair schedule can be easily determined. As before, let us focus on the situation of $K \leq d$, as the situation of $d < K$ is much easier.

Rewriting (5), a proportionally fair schedule σ_{PF} exists if and only if there is an index ℓ_{PF} such that, for each ℓ , it holds

$$\frac{u^A(\ell)}{u^A(\ell_{\text{PF}})} + \frac{u^B(\ell)}{u^B(\ell_{\text{PF}})} \leq 2. \quad (16)$$

Recall from Table 1 that, if $\ell < \ell_d$ then $u^B(\ell) = P_\ell$, while $u^B(\ell_d) = P + K - d$. Hence, for all $\ell < \ell_d$, (16) becomes

$$\frac{\ell_d - \ell}{\ell_d - \ell_{\text{PF}}} + \frac{P_\ell}{P_{\ell_{\text{PF}}}} \leq 2, \quad (17)$$

while for $\ell = \ell_d$ one has

$$\frac{P + K - d}{P_{\ell_{\text{PF}}}} \leq 2. \quad (18)$$

If an ℓ_{PF} exists such that (17) is satisfied for all $\ell < \ell_d$, and such that (18) is also satisfied, then $\sigma_{\ell_{\text{PF}}}$ is proportionally fair. Let us first rule out the cases of $\ell_d = 0$ and $\ell_d = 1$ by giving the following proposition.

Lemma 4. *If $K \leq d$, a proportionally fair schedule can exist only if $\ell_d \geq 2$.*

Proof. As we have noted in the definition of proportional fairness, for a schedule to be proportionally fair, the utility of each agent should be positive. In our case, the utility $K(\ell_d - \ell_{\text{PF}})$ of agent A cannot be positive if $\ell_d = 0$. On the other hand, if $\ell_d = 1$, then setting $\ell = 0$ in (17) gives $\ell_{\text{PF}} \leq \frac{1}{2}$, which can only be satisfied if $\ell_{\text{PF}} = 0$, leading to $P_{\ell_{\text{PF}}} = 0$, the utility of agent B . \square

We therefore assume $\ell_d \geq 2$. Setting $\ell = 0$ and hence $P_\ell = 0$ in (17), we get

$$\frac{\ell_d}{\ell_d - \ell_{\text{PF}}} \leq 2. \quad (19)$$

Setting $\ell = 1$ in (17) yields

$$\frac{\ell_d}{\ell_d - \ell_{\text{PF}}} - \frac{1}{\ell_d - \ell_{\text{PF}}} + \frac{P_1}{P_{\ell_{\text{PF}}}} \leq 2. \quad (20)$$

Recall that $P_1 = p_m$, i.e., P_1 is the length of the longest job. We have

$$P_{\ell_{\text{PF}}} \leq \ell_{\text{PF}} P_1,$$

and from (19), $\ell_d - \ell_{\text{PF}} \geq \ell_{\text{PF}}$, so that

$$P_{\ell_{\text{PF}}} \leq (\ell_d - \ell_{\text{PF}}) P_1,$$

which can be rewritten as

$$-\frac{1}{\ell_d - \ell_{\text{PF}}} + \frac{P_1}{P_{\ell_{\text{PF}}}} \geq 0.$$

The latter expression implies that the left-hand side of (20) is not smaller than the left-hand side of (19). In general, let $LHS(\ell)$ denote the left hand side of (17). Recalling that $P_\ell - P_{\ell-1} = p_{m-\ell+1}$, one can write (17) for any ℓ as

$$LHS(\ell) = LHS(\ell - 1) - \frac{1}{\ell_d - \ell_{\text{PF}}} + \frac{p_{m-\ell+1}}{P_{\ell_{\text{PF}}}} \leq 2. \quad (21)$$

In order to find a proportionally fair schedule, we can simply focus on the maximum of $LHS(\ell)$. From (21) we observe that, as long as

$$\frac{p_{m-\ell+1}}{P_{\ell_{PF}}} > \frac{1}{\ell_d - \ell_{PF}},$$

$LHS(\ell)$ is an increasing function. Thereafter, it starts decreasing. Due to the fact that the jobs of the agent A are SPT-ordered, it means that $LHS(\ell)$ is indeed concave. Now, we notice that writing (17) for $\ell = \ell_{PF}$, the value of the left hand side is precisely

$$LHS(\ell_{PF}) = \frac{\ell_d - \ell_{PF}}{\ell_d - \ell_{PF}} + \frac{P_{\ell_{PF}}}{P_{\ell_{PF}}} = 2.$$

As a consequence, for σ_{PF} to be a fair schedule, it must hold that $LHS(\ell)$ reaches its maximum for $\ell = \ell_{PF}$, since otherwise there is another value $\bar{\ell}$ such that $LHS(\bar{\ell}) > LHS(\ell) = 2$, and (17) would be violated. To this purpose, since $LHS(\ell)$ is concave, it is sufficient to verify that $LHS(\ell_{PF} - 1) \leq LHS(\ell_{PF})$ and $LHS(\ell_{PF} + 1) \leq LHS(\ell_{PF})$. Recalling (19), the following result holds.

Lemma 5. *If $K \leq d$, then a proportionally fair schedule exists if and only if there exists an integer $2 \leq \ell_{PF} \leq \lfloor \ell_d/2 \rfloor$ such that (18) is satisfied and*

$$p_{m-\ell_{PF}} \leq \frac{P_{\ell_{PF}}}{\ell_d - \ell_{PF}} \leq p_{m-\ell_{PF}+1}. \quad (22)$$

If such an integer exists, then schedule $\sigma_{\ell_{PF}}$ is proportionally fair.

Moving from situation $K \leq d$ to $K > d$, inequality (18) is no longer relevant and we only need to be concerned with (17) for *all* ℓ , including $\ell = \ell_d = m$. Therefore, the following lemma follows immediately.

Lemma 6. *If $K > d$, then a proportionally fair schedule exists if and only if there exists an integer $2 \leq \ell_{PF} \leq \lfloor \ell_d/2 \rfloor$, such that (22) is satisfied. If such an integer exists, then schedule $\sigma_{\ell_{PF}}$ is proportionally fair.*

With Lemmas 5 and 6, in order to find a proportionally fair schedule if it exists, one can then try out all values of ℓ_{PF} from 2 to $\lfloor \ell_d/2 \rfloor$, and check whether (22) (and (18) if $K \leq d$) are satisfied. We observe that while the leftmost and rightmost terms of (22) decrease with ℓ_{PF} , the central term increases with ℓ_{PF} . As a consequence, the value of ℓ_{PF} satisfying (22) (if any) can be found through a binary search between 2 and $\lfloor \ell_d/2 \rfloor$. Consequently, we have established the following theorem.

Theorem 5. *If the A -jobs are already sorted in SPT order, then in $O(\log m)$ time either a proportionally fair schedule can be found, or it can be certified that it does not exist.*

It is interesting to write (22) for the special case where agent A has m identical jobs of length p . In this case, $p_{m-\ell_{PF}} = p_{m-\ell_{PF}+1} = p$ and $P_{\ell_{PF}} = \ell_{PF}p$. Hence, (22) becomes

$$\frac{1}{\ell_{PF}} \leq \frac{1}{\ell_d - \ell_{PF}} \leq \frac{1}{\ell_{PF}}$$

i.e., (22) is satisfied if and only if there exists a value ℓ_{PF} such that

$$\frac{1}{\ell_d - \ell_{PF}} = \frac{1}{\ell_{PF}},$$

i.e., $\ell_{\text{PF}} = \ell_d/2$. On the other hand, (18) becomes

$$\frac{mp + K - d}{\ell_{\text{PF}}p} \leq 2,$$

i.e.,

$$(m - \ell_d)p \leq d - K,$$

which is always true according to the definition of ℓ_d . In conclusion, since ℓ_{PF} is an integer, the following result follows.

Corollary 1. *If the jobs of agent A are all identical, a proportionally fair schedule exists if and only if ℓ_d is even, and in this case $\sigma_{\text{PF}} \equiv \sigma_{\ell_d/2}$.*

4 Problem 1||($\sum C_j^A, \sum C_j^B$)

In this section we investigate the problem in which both agents want to minimize the total completion time of their jobs.

As we have already observed, problem 1| $d_j^B = d$ |($\sum C_j^A, T_{\max}^B$) with $d = 0$ is a special case of problem 1||($\sum C_j^A, \sum C_j^B$) in which agent B has only one job. Note that, in the proof of Theorem 3, $d \geq 0$ is an input parameter. By specifically setting $d = 0$, the instance \mathcal{I}_ε in the proof of Theorem 3 demonstrates that, for problem 1||($\sum C_j^A, \sum C_j^B$), we have $\text{PoF}_{\text{KS}} \geq 2/3$. However, the precise value of PoF_{KS} is still unknown.

In problem 1| $d_j^B = d$ |($\sum C_j^A, T_{\max}^B$), the number of Pareto optimal schedules is linear in the number of A-jobs. On the contrary, in problem 1||($\sum C_j^A, \sum C_j^B$), there can be an exponential number of Pareto optimal solutions, as shown by Agnetis et al. (2004). So, *finding* a proportionally fair solution, if it exists, seems a non trivial problem.

Given a two-agent scheduling problem, an instance is said to be *symmetric* if the two agents' sets of jobs are identical in terms of job lengths.

Theorem 6. *Given a symmetric instance of problem 1||($\sum C_j^A, \sum C_j^B$), it is NP-hard to decide whether there exists a proportionally fair schedule among all Pareto optimal schedules, and if such a schedule does exist, it is NP-hard to find it.*

Proof. Consider an instance of PARTITION, i.e., a set of n integers p_1, p_2, \dots, p_n , and denote $P = \sum_{i=1}^n p_i$. We define a symmetric instance of our scheduling problem in which both agents have n jobs, of length p_1, p_2, \dots, p_n , respectively. Without loss of generality, we assume $p_1 \leq p_2 \leq \dots \leq p_n$. Consider a schedule σ having the following structure. The two jobs of length p_1 are scheduled first, followed by the two jobs of length p_2, \dots , followed by the two jobs of length p_n . We say that a schedule having such a structure is an *SPT schedule*. Note that there exist exactly 2^n SPT schedules, depending on which agent's job comes first for each pair of jobs of equal length. Note that all SPT schedules are globally optimal and hence Pareto optimal.

Given any SPT schedule σ , let $S \subseteq \{1, 2, \dots, n\}$ be the set of indices j of the A-jobs such that J_j^A precedes in σ the corresponding jobs J_j^B , for $j = 1, \dots, n$. Denote by $x(S)$ the total length of the A-jobs in S . Hence, the total length of the jobs J_j^B that precede the corresponding job J_j^A is $P - x(S)$.

It can be easily shown that the total cost for agent A and for agent B are respectively

$$P + 2 \sum_{i=1}^n (n - i)p_i + (P - x(S))$$

and

$$P + 2 \sum_{i=1}^n (n-i)p_i + x(S).$$

We say that a schedule is *balanced* if both agents have the same cost, which is possible in this case if and only if $x(S) = P/2$, that is, S is a solution to the original instance of PARTITION.

We claim that an SPT schedule σ is proportionally fair if and only if it is balanced. Moreover, if such a schedule exists, it is the unique proportionally fair schedule. Let σ^* be a balanced SPT schedule. Since the instance is symmetric, we let $f_\infty = f_\infty^A = f_\infty^B$ and $f = f^A(\sigma^*) = f^B(\sigma^*)$. Also, for any other schedule σ , we can assume with no loss of generality that $f^A(\sigma) = f - \delta$ and $f^B(\sigma) = f + \Delta$, where $\Delta \geq \delta$ (since otherwise σ would have a smaller total cost than σ^* , which is impossible as σ^* is globally optimal). Condition (5) can be therefore written as

$$\frac{f_\infty - (f - \delta)}{f_\infty - f} + \frac{f_\infty - (f + \Delta)}{f_\infty - f} \leq 2,$$

which clearly holds due to $\Delta \geq \delta$. So, if an SPT balanced schedule exists, it is proportionally fair.

Now we show that no non-balanced schedule σ' can be proportionally fair. In fact, suppose that σ' is proportionally fair and one has $f^A(\sigma') = f - \delta$ and $f^B(\sigma') = f + \Delta$. Due to the symmetry of the instance (and to the symmetry axiom), the schedule σ'' obtained by exchanging the positions of the two homologous jobs in each pair should also be proportionally fair, for which one would have $f^A(\sigma'') = f + \Delta$ and $f^B(\sigma'') = f - \delta$. However, in Nicosia, Pacifici, and Pferschy (2017) it is proved that for any symmetric instance, if a proportionally fair solution exists, then it is unique, so neither σ' nor σ'' can be proportionally fair.

In conclusion, in our instance a proportionally fair schedule exists if and only if the original instance of PARTITION is a yes-instance. \square

We remark that in problem 1|| $(\sum C_j^A, \sum C_j^B)$, deciding whether a proportionally fair solution exists is not known to be in NP. This is why Theorem 6 only establishes NP-hardness and not NP-completeness.

With arguments very similar to those in the proof of Theorem 6, the following result can also be established.

Theorem 7. *Given a symmetric instance of problem 1|| $(\sum C_j^A, \sum C_j^B)$, it is NP-hard to find a KS fair schedule among all Pareto optimal schedules.* \square

5 Discussion and conclusions

We have defined a utility function in terms of scheduling cost and, with this utility, studied some fairness issues. However, there can be different utilities. Fehr and Schmidt (1999) and Cui, Raju, and Zhang (2007) consider utilities that incorporate fairness concerns of individual agents. Below we briefly explore the possibility of having another utility that can be used to address fairness as a *global* concern.

5.1 A different utility function

Suppose we have the following as an alternative definition of utility,

$$u^i(\sigma) = \frac{1}{f_i(\sigma)}, \text{ for } i = A, B. \quad (23)$$

This definition satisfies the fundamental properties that a utility function is positive and non-increasing in the cost. However, we argue that such a choice of utility function is not appropriate.

Let us consider proportional fairness. Under definition (23), proportional fairness definition (5) can be rewritten as

$$\frac{f^A(\sigma_{\text{PF}}) - f^A(\sigma)}{f^A(\sigma)} + \frac{f^B(\sigma_{\text{PF}}) - f^B(\sigma)}{f^B(\sigma)} \leq 0 \quad \text{for all } \sigma. \quad (24)$$

In order to investigate the consequences of (24), let us consider the following example.

Example 2. Consider an instance of problem 1|| $(\sum C_j^A, C_{\max}^B)$, in which agent A has m identical unit-time jobs, while agent B has a single job, of length βm , for some rational $\beta > 0$. Clearly, there are exactly $m + 1$ Pareto optimal schedules, completely characterized by the number of unit-time jobs scheduled after the job of agent B . As usual, σ_ℓ denotes the schedule in which ℓ jobs follow the job of agent B , and hence $m - \ell$ precede it. We have

$$f^A(\sigma_\ell) = \frac{m(m+1)}{2} + \beta \ell m, \quad \text{and} \quad f^B(\sigma_\ell) = \beta m + m - \ell.$$

Let ℓ_{PF} denote the value of ℓ corresponding to a proportionally fair solution. Writing (24) for this example, one has, after some algebra, that for $\sigma_{\ell_{\text{PF}}}$ to be a fair schedule, it must hold for each $\ell \neq \ell_{\text{PF}}$ that

$$\frac{\beta m(\ell_{\text{PF}} - \ell)}{m(m+1)/2 + \beta \ell m} + \frac{\ell - \ell_{\text{PF}}}{\beta m + m - \ell} \leq 0. \quad (25)$$

Note that $\ell - \ell_{\text{PF}}$ is a non-zero common factor of the two terms on the left-hand side of inequality (25), which can be both positive and negative if $\ell_{\text{PF}} \notin \{0, m\}$ (i.e., $0 < \ell_{\text{PF}} < m$). In such a case, by taking $\ell = m$ and $\ell = 0$, respectively, we obtain from (25) the following two inequalities:

$$m\beta^2 - m\beta - \frac{m+1}{2} \geq 0, \quad \text{and} \quad m\beta^2 + m\beta - \frac{m+1}{2} \leq 0,$$

which imply that

$$\beta \geq \frac{1}{2} + \frac{\sqrt{3m^2 + 2m}}{2m}, \quad \text{and} \quad \beta \leq -\frac{1}{2} + \frac{\sqrt{3m^2 + 2m}}{2m}.$$

Clearly the above two inequalities cannot be satisfied at the same time. Therefore, we conclude that no schedule σ_ℓ can be proportionally fair unless $\ell_{\text{PF}} \in \{0, m\}$.

The conclusion of the above example is a bit disappointing. Following the classical definition of proportionally fair schedule, we find that either a fair schedule does not exist, or it is one for the two extreme schedules¹, which to any intuitive notion of fairness appears indeed as the most unfair! This suggests that (23) is not a very sensible choice of utility function, at least as long as our scheduling setting is concerned. This can perhaps be explained as follows. Clearly $u^i(\sigma)$ becomes increasingly insensitive to increases in cost, while on the contrary it is highly sensitive when the cost is small. Such cost-proneness is not in accordance with the classical view of scheduling objectives, in which the cost is typically linear in certain figures (completion time, tardiness, number of tardy jobs, etc.), hence more in line with a risk-neutrality attitude.

¹Recall that, in this case, the two extreme schedules correspond to the two scenarios $\ell = 0$ and $\ell = m$, where one agent obtains his best solution while the other gets his worst.

Problem	Fair solution	Existence	Computation	PoF
$1 d_j^B = d (\sum C_j^A, T_{\max}^B)$	Kalai-Smorodinsky	always	$O(\log m)^\dagger$	$\frac{2}{3}$
	Proportionally fair	$O(\log m)^\dagger$	$O(\log m)^\dagger$	$\frac{1}{2}$
$1 (\sum C_j^A, \sum C_j^B)$	Kalai-Smorodinsky	always	NP-hard	$\geq \frac{2}{3}$
	Proportionally fair	NP-hard	NP-hard	$\frac{1}{2}$

[†] A-Jobs are assumed to be already in SPT order.

Table 2: Summary of results

5.2 Future research

In this paper we have carried out a preliminary investigation of the concept of price of fairness in single-machine scheduling problems. The main findings are summarized in Table 2. The third and fourth columns respectively display the complexity of establishing whether a fair schedule exists and, if so, finding a fair schedule. We have that the price of proportional fairness is smaller than the price of Kalai-Smorodinsky fairness, but the former solution does not always exist.

Different from previous studies, our work adopts a generalized definition of system utility, expressed as a weighted sum of the agent utilities. We prove that PoF_{KS} and PoF_{PF} are independent of particular choice of the weights in the system utility. This also shows, as a byproduct (when α tends to 0 or $+\infty$), that PoF also bounds the relative loss in one-agent utility of a fair solution with respect to the best solution for that agent.

We have restricted ourselves to problems in which one of the agents wishes to minimize the sum of the completion times of his jobs. The only open problem in this respect is determining whether PoF_{KS} in problem $1 || (\sum C_j^A, \sum C_j^B)$ is strictly larger than $2/3$. We expect that further research would determine the complexity of finding fair solutions and the values of the price of fairness in different scenarios, such as $1 || (C_{\max}^A, \sum w_j^B C_j^B)$ or $1 || (T_{\max}^A, T_{\max}^B)$ when jobs have individual due dates.

References

- Agnetis, A., Billaut, J.-C., Gawiejnowicz, S., Pacciarelli, D., and A. Soukhal, *Multiagent Scheduling Models and Algorithms*, Springer, 2014.
- Agnetis A., P.B. Mirchandani, D. Pacciarelli, and A. Pacifici (2004). Scheduling problems with two competing agents, *Operations Research*, 52(2), 229–242.
- Agnetis A., G. de Pascale, and M. Pranzo (2009). Computing the Nash solution for scheduling bargaining problems, *International J. Operational Research*, 6(1), 54–69.
- Albino, V., Carbonara, N. and I. Giannoccaro, Industrial districts as complex adaptive systems, in Karlsson, C. et al (eds.), *Industrial Clusters and Inter-firm Networks*, Elgar Publishers, 58–83, 2005.
- Anshelevich E., A. Dasgupta, J. Kleinberg, E. Tardös, T. Wexler, and T. Roughgarden (2004). The Price of Stability for Network Design with Fair Cost Allocation, *The 45th Annual IEEE Symposium on Foundations of Computer Science*, 59–73.

- Bertsimas D., V. Farias, and N. Trichakis (2011). The price of fairness, *Operations Research*, 59(1), 17–31.
- Bohm P., and B. Larsen (1994). Fairness in a tradeable-permit treaty for carbon emissions reductions in Europe and the former Soviet Union, *Environmental and Resource Economics*, 4, 219–239.
- Brams, S.J., and A.D. Taylor (1996). *Fair division: From cake-cutting to dispute resolution*. Cambridge University Press.
- Caragiannis I., C. Kaklamanis, P. Kanellopoulos, and M. Kyropoulou (2012). The efficiency of fair division, *Theory of Computing Systems*, 50(4), 589–610.
- Chun Y. (1988). The equal-loss principle for bargaining problems, *Economics Letters*, 26(2), 103–106.
- Cui T.H., J.S. Raju, and Z.J. Zhang (2007). Fairness and channel coordination, *Management Science*, 53(8), 1303–1314.
- Ertogral K. and D. Wu (2000). Auction-theoretic coordination of production planning in the supply chain, *IIE Transactions*, 32(10), 931–940.
- Forsythe R., J.L. Horowitz, N.E. Savin, and Martin Sefton (1994). Fairness in Simple Bargaining Experiments, *Games and Economic Behavior*, 6(3), 347–369.
- Fehr E. and K.M. Schmidt (1999). A theory of fairness, competition, and cooperation, *The Quarterly Journal of Economics*, August, 817–868.
- Kalai E. and M. Smorodinsky (1975). Other solutions to Nash bargaining problem, *Econometrica*, 43, 513–518.
- Karsu Ö. and A. Morton (2015). Inequity averse optimization in operational research, *European Journal of Operational Research*, 245(2), 343–359.
- Kelly F.P., A.K. Maulloo, and D.K.H. Tan (1998). Rate control in communication networks: shadow prices, proportional fairness and stability, *Journal of the Operational Research Society*, 49, 237–252.
- Lang, F. and A. Fink, Collaborative machine scheduling: Challenges of individually optimizing behavior, *Concurrency and computation practice and experience*, 27(11), 2869–2888, 2015.
- Mariotti M. (1998). Nash bargaining theory when the number of alternatives can be finite, *Social Choice Welfare*, 15, 413–421.
- Naldi M., G. Nicosia, A. Pacifici, and U. Pferschy (2016). Maximin Fairness in Project Budget Allocation, *Electronic Notes in Discrete Mathematics*, 55, 65–68.
- Nash J. (1950). The bargaining problem. *Econometrica*, 18(2), 155–162.
- Nicosia G., A. Pacifici, and U. Pferschy (2017). Price of fairness for allocating a bounded resource, *European Journal of Operational Research*, 257, 933–943.
- Smith, W.E. (1956). Various optimizers for single-stage production, *Naval Research Logistics Quarterly* 3, 59–66.

- Soomer, M.J. and J. Franx (2008). Scheduling aircraft landings using airlines' preferences, *European Journal of Operational Research*, 190(1), 277–291.
- Tang, L., Zhao, X., Liu, J. and J. Y.-T. Leung (2017). Competitive two-agent scheduling with deteriorating jobs on a single parallel-batching machine, *European Journal of Operational Research*, 263(2), 401–411.
- Wellman, M.P., Walsh, W.E., Wurman, P.R. and J.K. MacKie-Mason, Auction Protocols for Decentralized Scheduling, *Games and Economic Behavior*, 35, 271–303, 2001.