



HAL
open science

Mechanism design for aggregating energy consumption and quality of service in speed scaling scheduling

Christoph Dürr, Łukasz Jeż, Oscar Carlos Vasquez Perez

► **To cite this version:**

Christoph Dürr, Łukasz Jeż, Oscar Carlos Vasquez Perez. Mechanism design for aggregating energy consumption and quality of service in speed scaling scheduling. WINE 2013: The 9th Conference on Web and Internet Economics, Dec 2013, Cambridge, MA, United States. pp.134-145, 10.1007/978-3-642-45046-4_12 . hal-00907832

HAL Id: hal-00907832

<https://hal.science/hal-00907832>

Submitted on 22 Nov 2013

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Mechanism design for aggregating energy consumption and quality of service in speed scaling scheduling

Christoph Dürr¹, Lukasz Jeż², and Oscar C. Vásquez³

¹ CNRS, LIP6, Université Pierre et Marie Curie, Paris, France.

² Institute of Computer Science, University of Wrocław, Poland, and DIAG, Sapienza University of Rome, Italy.

³ LIP6 and Industrial Engineering Department, University of Santiago of Chile.

Abstract. We consider a strategic game, where players submit jobs to a machine that executes all jobs in a way that minimizes energy while respecting the jobs' deadlines. The energy consumption is then charged to the players in some way. Each player wants to minimize the sum of that charge and of their job's deadline multiplied by a priority weight. Two charging schemes are studied, the *proportional cost share* which does not always admit pure Nash equilibria, and the *marginal cost share*, which does always admit pure Nash equilibria, at the price of overcharging by a constant factor.

1 Introduction

In many computing systems, minimizing energy consumption and maximizing quality of service are opposed goals. This is also the case for the speed scaling scheduling model considered in this paper. It has been introduced in [9], and triggered a lot of work on offline and online algorithms; see [1] for an overview.

The online and offline optimization problem for minimizing flow time while respecting a maximum energy consumption has been studied for the single machine setting in [14, 2, 5, 8] and for the parallel machines setting in [3]. For the variant where an aggregation of energy and flow time is considered, polynomial approximation algorithms have been presented in [7, 4, 11].

In this paper we propose to study this problem from a different perspective, namely as a strategic game. In society many ecological problems are either addressed in a centralized manner, like forcing citizens to sort household waste, or in a decentralized manner, like tax incentives to enforce ecological behavior. This paper proposes incentives for a scheduling game, in form of an energy cost charging scheme.

Consider a scheduling problem for a single processor, that can run at variable speed, such as the modern microprocessors Intel SpeedStep, AMD PowerNow! or IBM EnergyScale. Higher speed means that jobs finish earlier at the price of a higher energy consumption. Each job has some workload, representing a number of instructions to execute, and a release time before which it cannot be

scheduled. Every user submits a single job to a common processor, declaring the jobs parameters, together with a deadline, that the player chooses freely.

The processor will schedule the submitted jobs preemptively, so that all release times and deadlines are respected and the overall energy usage is minimized. The energy consumed by the schedule needs to be charged to the users. The individual goal of each user is to minimize the sum of the energy cost share and of the requested deadline weighted by the user’s priority, which represents a quality of service coefficient. This individual priority weight implies a conversion factor that allows of aggregation of deadline and energy.

In a companion paper [15] we study this game from the point of view of the game regulator, and compare different ways to organize the game which would lead to truthfulness. In this paper we focus on a particular game setting, described in the next section.

2 The model

Formally, we consider a non-cooperative game with n players and a regulator. The regulator manages the machine where the jobs are executed. Each player has a job i with a workload w_i , a release time r_i and a priority p_i , representing a quality of service coefficient. The player submits its job together with a deadline $d_i > r_i$ to the regulator. Workloads, release times and deadlines are public information known to all players, while quality of service coefficients can be private.

The regulator implements some *cost sharing mechanism*, which is known to all users. This mechanism defines a cost share function b_i specifying how much player i is charged. The penalty of player i is the sum of two values: his *energy cost share* $b_i(w, r, d)$ defined by the mechanism, where $w = (w_1, \dots, w_n)$, $r = (r_1, \dots, r_n)$ and $d = (d_1, \dots, d_n)$, and his *waiting cost*, which can be either $p_i d_i$ or $p_i(d_i - r_i)$; we use the former waiting cost throughout the article but all our results apply to both. The sum of all player’s penalties, i.e., energy cost shares and waiting costs will be called the *utilitarian social cost*.

The regulator computes a minimum energy schedule for a single machine in the speed scaling model, which stipulates that at any point in time t the processor can run at arbitrary speed $s(t) \geq 0$; for a time interval I , the workload executed in I is $\int_{t \in I} s(t) dt$, while the energy consumed is $\int_{t \in I} s(t)^\alpha dt$ for some fixed physical constant $\alpha \in [2, 3]$ characteristic for a device [6]. The sum of the energy used by this optimum schedule and of all the players’ waiting costs will be called the *effective social cost*.

The minimum energy schedule can be computed in time $O(n^2 \log n)$ [10] and has (among others) the following properties [16]. The jobs in the schedule are executed by preemptive earliest deadline first order (EDF), and the speed $s(t)$ at which they are processed is piecewise linear. Preemptive EDF means that at every time point among all jobs which are already released and not yet completed, the job with the smallest deadline is executed, using job indices to break ties.

The cost sharing mechanism defines the game completely. Ideally, we would like the game and the mechanism to have the following properties.

existence of pure Nash equilibria This means that there is a strategy profile vector d such that no player can unilaterally deviate from their strategy d_i while strictly decreasing their penalty.

budget balance The mechanism is c -budget balanced, when the sum of the cost shares is no smaller than the total energy consumption and no larger than c times the energy consumption.

In the sequel we introduce and study two different cost sharing mechanisms, namely PROPORTIONAL COST SHARING where every player pays exactly the cost generated during the execution of his job, and MARGINAL COST SHARING where every player pays the increase of energy cost generated by adding this player to the game.

3 Proportional cost sharing

The proportional cost sharing is the simplest budget balanced cost sharing scheme one can think of. Every player i is charged exactly the energy consumed during the execution of his job. Unfortunately this mechanism does not behave well as we show in Theorem 1.

Fact 1 *In a single player game, the player's penalty is minimized by the deadline*

$$r_1 + w_1(\alpha - 1)^{1/\alpha} p_1^{-1/\alpha}.$$

Proof. If player 1 chooses deadline $d_1 = r_1 + x$ then the schedule is active between time r_1 and $r_1 + x$ at speed w_1/x . Therefore his penalty is

$$p_1(r_1 + x) + x^{1-\alpha} w_1^\alpha.$$

Deriving this expression in x , and using the fact that the penalty is concave in t for any $x > 0$ and $\alpha > 0$, we have that the optimal x for the player will set to zero the derivative. This implies the claimed deadline. \square

If there are at least two players however, the game does not have nice properties as we show now.

Theorem 1. *The PROPORTIONAL COST SHARING does not always admit a pure Nash equilibrium.*

The proof consists of a very simple example: there are 2 identical players with identical jobs, say $w_1 = w_2 = 1$, $r_1 = r_2 = 0$ and $p_1 = p_2 = 1$. First we determine the best response of player 1 as a function of player 2, then we conclude that there is no pure Nash equilibrium.

argument	value	applicable range
$d_1^{(1)} = (\alpha - 1)^{1/\alpha}$	$g_1(d_2) = \alpha(\alpha - 1)^{1/\alpha-1}$	$d_2 \geq 2(\alpha - 1)^{1/\alpha}$
$d_1^{(2)} = \frac{d_2}{2}$	$g_2(d_2) = d_2/2 + (d_2/2)^{1-\alpha}$	$d_2 \leq 2(\alpha - 1)^{1/\alpha}$
$d_1^{(3)} = 2 \left(\frac{\alpha-1}{2}\right)^{1/\alpha}$	$g_3(d_2) = \alpha \left(\frac{\alpha-1}{2}\right)^{1/\alpha-1}$	$\left(\frac{\alpha-1}{2}\right)^{1/\alpha} \leq d_2 \leq 2 \left(\frac{\alpha-1}{2}\right)^{1/\alpha}$
$d_1^{(4)} = d_2 + (\alpha - 1)^{1/\alpha}$	$g_4(d_2) = d_2 + \alpha (\alpha - 1)^{1/\alpha-1}$	$d_2 \leq (\alpha - 1)^{1/\alpha-1}$

Table 1. The local minimum in the range of f corresponding to f_i is a function of α and d_2 , which we denote by $d_1^{(i)}$. The value at such local minimum is again a function of α and d_2 , which we denote by $g_i(d_2)$. These are only potential minima: they exist if and only if the condition given in the last column is satisfied.

Lemma 1. *Given the second player's choice d_2 , the penalty of the first player as a function of his choice d_1 is given by*

$$f(d_1) = \begin{cases} f_1(d_1) = d_1 + d_1^{1-\alpha} & \text{if } d_1 \leq \frac{d_2}{2} \\ f_2(d_1) = d_1 + \left(\frac{d_2}{2}\right)^{1-\alpha} & \text{if } \frac{d_2}{2} \leq d_1 \leq d_2 \\ f_3(d_1) = d_1 + \left(\frac{d_1}{2}\right)^{1-\alpha} & \text{if } d_2 \leq d_1 \leq 2d_2 \\ f_4(d_1) = d_1 + (d_1 - d_2)^{1-\alpha} & \text{if } d_1 \geq 2d_2 \end{cases} \quad (1)$$

The local minima of $f(d_1)$ are summarized in Table 1, and the penalties corresponding to player 1 picking these minima are illustrated in Figure 1.




Fig. 1. First player's penalty (in bold) when choosing his best response as a function of second player's strategy d_2 , here for $\alpha = 3$.

Proof. Formula (1) follows by a straightforward case inspection. Then, to find all the local minima of f , we first look at the behavior of each of f_i , finding their

local minima in their respective intervals, and afterwards we inspect the border points of these intervals.

Range of f_1 : The derivative of f_1 is

$$f_1'(d_1) = 1 - (\alpha - 1)d_1^{-\alpha} ,$$

whose derivative in turn is positive for $\alpha > 1$. Therefore, f_1 has a local minimum at $d_1^{(1)}$ as specified. Since we require that this local minimum is within the range where f coincides with f_1 , the necessary and sufficient condition is $d_1^{(1)} \leq \frac{d_2}{2}$.

Range of f_2 : f_2 is an increasing function, and therefore it attains a minimum value only at the lower end of its range, $d_1^{(3)}$. However, if $d_1^{(2)}$ is to be a local minimum of f , there can be no local minimum of f in the range of f_1 (immediately to the left), so the applicable range of $d_1^{(2)}$ is the complement of that of $d_1^{(1)}$.

Range of f_3 : The derivative of f_3 is

$$f_3'(d_1) = 1 - \frac{\alpha - 1}{2}(d_1/2)^{-\alpha} , \quad (2)$$

whose derivative in turn is positive for $\alpha > 1$. Hence, f_3 has a local minimum at $d_1^{(3)}$ as specified. The existence of this local minimum requires $d_2 \leq d_1^{(3)} \leq 2d_2$, which is equivalent to $\frac{d_1^{(3)}}{2} \leq d_2 \leq d_1^{(3)}$.

Range of f_4 : The derivative of f_4 is

$$f_4'(d_1) = 1 - (\alpha - 1)(d_1 - d_2)^{-\alpha} , \quad (3)$$

whose derivative in turn is positive for $\alpha > 1$. Hence, f_4 has a local minimum at $d_1^{(4)}$ as specified. The existence of this local minimum requires $d_1^{(4)} \geq 2d_2$.

Now let us consider the border points of the ranges of each f_i . Since f_2 is strictly increasing, the border point of the ranges of f_2 and f_3 is not a local minimum of f . This leaves only the border point $d_1^{(2)} = 2d_2$ of the ranges of f_3 and f_4 to consider. Clearly, $d_1^{(2)}$ is a local minimum of f if and only if $f_3'(d_1^{(2)}) \leq 0$ and $f_4'(d_1^{(2)}) \geq 0$. However, by (2), $f_3'(d_1^{(2)}) = 2 - (\alpha - 1)d_2^{-\alpha}$, and by (3), $f_4'(d_1^{(2)}) = 2 - 2(\alpha - 1)d_2^{-\alpha} < f_3'(d_1^{(2)})$, so $d_1^{(2)}$ is not a local minimum of f either. \square

Note that the range of g_1 is disjoint with the ranges of g_3 and g_4 , and with the exception of the shared border value $2(\alpha - 1)^{1/\alpha}$, also with the range of g_2 . However, the ranges of g_2 , g_3 and g_4 are not disjoint. Therefore, we now focus on their shared range, and determine which of the functions gives rise to the true local minimum (the proof is omitted due to space constraints).

Lemma 2. *The function $g_3(d_2)$ is constant, the function $g_4(d_2)$ is an increasing linear function, and the function $g_2(d_2)$ is decreasing for $d_2 < d_1^{(3)}$. Moreover, there exist two unique values*

$$d_2^\dagger = \alpha(\alpha - 1)^{1/\alpha-1}(2^{1-1/\alpha} - 1) \quad \text{such that} \quad g_4(d_2^\dagger) = g_3(d_2^\dagger) , \quad (4)$$

$$d_2^\ddagger \in \left(d_2^\dagger, d_1^{(3)} \right) \quad \text{such that} \quad g_2(d_2^\ddagger) = g_3(d_2^\ddagger) . \quad (5)$$

With Lemma 1 and Lemma 2, whose statements are summarized in Table 1 and Figure 1, we can finally determine what is the best response of the first player as a function of d_2 .

Lemma 3. *The best response for player 1 as function of d_2 is*

$$\begin{aligned} d_1^{(4)} &= d_2 + (\alpha - 1)^{1/\alpha} & \text{if} & & 0 < d_2 \leq d_2^\dagger , \\ d_1^{(3)} &= 2 \left(\frac{\alpha - 1}{2} \right)^{1/\alpha} & \text{if} & & d_2^\dagger < d_2 \leq d_2^\ddagger , \\ d_1^{(2)} &= \frac{d_2}{2} & \text{if} & & d_2^\ddagger < d_2 \leq 2(\alpha - 1)^{1/\alpha} , \\ d_1^{(1)} &= (\alpha - 1)^{1/\alpha} & \text{if} & & 2(\alpha - 1)^{1/\alpha} < d_2 . \end{aligned}$$

Proof. The proof consists in determining which of the applicable local minima of f is the global minimum for each range of d_2 . Again, the cases are depicted in Figure 1.

case (i) $0 < d_2 \leq d_2^\dagger$: In this case, we claim that the best response of player 1 is

$$d_1^{(4)} = d_2 + (\alpha - 1)^{1/\alpha} .$$

First we prove that

$$d_2^\dagger \in \left(\left(\frac{\alpha - 1}{2} \right)^{1/\alpha}, (\alpha - 1)^{1/\alpha-1} \right) .$$

The upper bound hold since

$$\begin{aligned} \alpha(\alpha - 1)^{1/\alpha-1}(2^{1-1/\alpha} - 1) &< (\alpha - 1)^{1/\alpha-1} \\ \alpha(2^{1-1/\alpha} - 1) &< 1, \end{aligned}$$

holds for $\alpha \geq 2$.

The lower bound holds since,

$$\begin{aligned} \alpha(\alpha - 1)^{1/\alpha-1}(2^{1-1/\alpha} - 1) &> \left(\frac{\alpha - 1}{2} \right)^{1/\alpha} \\ \frac{\alpha}{\alpha - 1}(2^{1-1/\alpha} - 1) &> 2^{-1/\alpha} \\ \frac{\alpha}{\alpha - 1}(2 - 2^{1/\alpha}) &> 1 \end{aligned}$$




Fig. 2. Best response of player 1 as function of d_2 , and best response of player 2 as function of d_1 . Here for $\alpha = 3$.

holds for $\alpha \geq 2$.

In fact, both inequalities are true even for $\alpha > 1$, but as we require $\alpha \geq 2$ due to Lemma 2, we settle for simpler proofs.

These bounds imply that in case (i) player 1 chooses the minimum among the 3 local minima $d_1^{(2)}$, $d_1^{(3)}$, and $d_1^{(4)}$, where the middle one is only an option for $(\frac{\alpha-1}{2})^{1/\alpha} \leq d_2 \leq d_2^\dagger$. It follows from Lemma 2 that the last option always dominates: by (5), for every $(\frac{\alpha-1}{2})^{1/\alpha} \leq d_2 < d_2^\dagger$, we have $g_3(d_2) < g_2(d_2)$, and by (4), for every $(\frac{\alpha-1}{2})^{1/\alpha} \leq d_2 \leq d_2^\dagger$, we have $g_4(d_2) < g_3(d_2)$. This concludes the analysis for case (i).

case (ii) $d_2^\dagger < d_2 \leq d_2^\ddagger$: In this case, we claim that the best response of player 1 is

$$d_1^{(3)} = 2 \left(\frac{\alpha - 1}{2} \right)^{1/\alpha} .$$

First we observe that by Lemma 2 (5),

$$d_2^\dagger < d_1^{(4)} ,$$

which rules out $d_1^{(1)}$ as a choice for player 1, leaving only $d_1^{(2)}$, $d_1^{(3)}$, and $d_1^{(4)}$. Again, Lemma 2 implies that $d_1^{(4)}$ dominates other choices: by (5), we have $g_3(d_2) < g_2(d_2)$ for all $(\frac{\alpha-1}{2})^{1/\alpha} \leq d_2 < d_2^\dagger$, and by (4), we have $g_3(d_2) < g_4(d_2)$ for all $d_2 > d_2^\dagger$.

Note that for $\alpha = 2$, the range of this case is empty.

case (iii) $d_2^\dagger < d_2 \leq 2(\alpha - 1)^{1/\alpha}$: For this range, only $d_1^{(2)}$ and $d_1^{(3)}$ are viable choices for player 1, and Lemma 2 (5) implies that $d_1^{(2)}$ dominates. Therefore

first player's best response is

$$d_1^{(2)} = \frac{d_2}{2} .$$

case (iv) $2(\alpha - 1)^{1/\alpha} < d_2$: For this range, the only viable choice for player 1 is

$$d_1^{(1)} = (\alpha - 1)^{1/\alpha} ,$$

which is therefore his best response.

This concludes the proof of the lemma. □

By the symmetry of the players, the second player's best response is in fact an identical function of d_1 as the one stated in Lemma 3. By straightforward inspection it follows that there is no fix point (d_1, d_2) to this game, which implies the following theorem, see figure 2 for illustration.

4 Marginal cost sharing

In this section we propose a different cost sharing scheme, that improves on the previous one in the sense that it admits pure Nash equilibria, however for the price of overcharging by at most a constant factor.

Before we give the formal definition we need to introduce some notations. Let $\text{OPT}(d)$ be the energy minimizing schedule for the given instance, and $\text{OPT}(d_{-i})$ be the energy minimizing schedule for the instance where job i is removed. We denote by $E(S)$ the energy cost of schedule S .

In the marginal cost sharing scheme, player i pays the penalty function

$$p_i d_i + E(\text{OPT}(d)) - E(\text{OPT}(d_{-i})).$$

This scheme defines an exact potential game by construction [12]. Formally, let n be the number of players, $D = \{d | \forall j : d_j > r_j\}$ be the set of action profiles (deadlines) over the action sets D_i of each player.

Let us denote the effective social cost corresponding to a strategy profile d by $\Phi(d)$. Then

$$\Phi(d) = \sum_{i=1}^n p_i d_i + E(\text{OPT}(d)).$$

Clearly, if a player i changes its strategy d_i and his penalty decreases by some amount Δ , then the effective social cost decreases by the same amount Δ , because $E(\text{OPT}(d_{-i}))$ remains unchanged.

4.1 Existence of Equilibria

While the best response function is not continuous in the strategy profile, precluding the use of Brouwer's fixed-point theorem, existence of pure Nash equilibria can nevertheless be easily established.

To this end, note that the global minimum of the effective social cost, if it exists, is a pure Nash equilibrium. Its existence follows from (1) compactness of a non-empty sub-space of strategies with bounded social cost and (2) continuity of Φ .

For (2), note that $\sum_i p_i d_i$ is clearly continuous in d , and hence $\Phi(d)$ is continuous if $E(\text{OPT}(d))$ is. The continuity of the latter is clear once considers all possible relations of the deadlines chosen by the players.

For (1), let d' be any (feasible) strategy profile such that $d_i > r_i$ for each player i . The subspace of strategy profiles d such that $\Phi(d) \leq \Phi(d')$ is clearly closed, and bounded due to the $p_i d_i$ terms. Thus it is a compact subspace that contains the global minimum of Φ .

4.2 Convergence can take forever

In this game the strategy set is infinite. Moreover, the convergence time can be infinite as we demonstrate below in Theorem 2. Notice that this also proves that in general there are no dominant strategies in this game.

Theorem 2. *For the game with the marginal cost sharing mechanism, the convergence time to reach a pure Nash equilibrium can be unbounded.*

Proof. The proof is by exhibiting again the same small example, with 2 players, release times 0, unit weights, unit penalty factors, and $\alpha > 2$.

For this game there are two pure Nash equilibria, the first one is

$$d_1 = \left(\frac{\alpha - 1}{2}\right)^{1/\alpha}, \quad d_2 = d_1 + (\alpha - 1)^{1/\alpha},$$

while the second one is symmetric for players 1 and 2.

In the remainder of the proof, we assume that player 1 chooses a deadline which is close to the pure Nash equilibrium above. By analyzing the best responses of the players, we conclude that after a best response of player 2, and then of player 1 again, he chooses a deadline which is even closer to the pure Nash equilibrium above but different from it, leading to an infinite convergence sequence of best responses. The proofs of the following two lemmas are omitted.

Now suppose $d_1 = \delta \left(\frac{\alpha - 1}{2}\right)^{1/\alpha}$ for some $1 < \delta < 2^{1/\alpha}$. What is the best response for player 2?

Lemma 4. *Given the first player's choice d_1 , the penalty of the second player as a function of his choice d_2 is given by*

$$h(d_2, d_1) = \begin{cases} h_1(d_2, d_1) = d_2 + d_2^{1-\alpha} + (d_1 - d_2)^{1-\alpha} - d_1^{1-\alpha} & \text{if } d_2 \leq \frac{d_1}{2} \\ h_2(d_2, d_1) = d_2 + (2^\alpha - 1)d_1^{1-\alpha} & \text{if } \frac{d_1}{2} \leq d_2 \leq d_1 \\ h_3(d_2, d_1) = d_2 + 2^\alpha d_2^{1-\alpha} - d_1^{1-\alpha} & \text{if } d_1 \leq d_2 \leq 2d_1 \\ h_4(d_2, d_1) = d_2 + (d_2 - d_1)^{1-\alpha} & \text{if } d_2 \geq 2d_1, \end{cases}$$

and the best response for player 2 as function of d_1 is

$$d_1 + (\alpha - 1)^{1/\alpha} = (\alpha - 1)^{1/\alpha}(1 + 2^{-1/\alpha}\delta) \quad (6)$$

From now on we assume that player 2 chooses $d_2 = d_1 + (\alpha - 1)^{1/\alpha} = (\alpha - 1)^{1/\alpha}(1 + 2^{-1/\alpha}\delta)$. What is the best response for player 1?

Lemma 5. *Given the second player's choice d_2 , the penalty of the first player as a function of his choice d_1 is given by $h(d_1, d_2)$ and the best response for player 1 is*

$$d_1 = \delta' \left(\frac{\alpha - 1}{2} \right)^{1/\alpha},$$

for some $\delta' \in (1, \delta)$.

This concludes the proof of Theorem 2. □

4.3 Bounding total charge

In this section we bound the total cost share for the MARGINAL COST SHARING SCHEME, by showing that it is at least $E(\text{OPT}(d))$ and at most α times this value. In fact we show a stronger claim for individual cost shares.

Theorem 3. *For every player i , its marginal costshare is at least its proportional costshare and at most α times the proportional costshare.*

Proof. Fix a player i , and denote by S_{-i} the schedule obtained from $\text{OPT}(d)$ when all executions of i are replaced by idle times. Clearly we have the following inequalities.

$$E(\text{OPT}(d_{-i})) \leq E(S_{-i}) \leq E(\text{OPT}(d))$$

Then the marginal cost share of player i can be lower bounded by

$$E(\text{OPT}(d)) - E(\text{OPT}(d_{-i})) \geq E(\text{OPT}(d)) - E(S_{-i}).$$

According to [16] the schedule OPT can be obtained by the following iterative procedure. Let S be the support of a partial schedule. For every interval $[t, t']$ we define its domain $I_{t,t'} := [t, t'] \setminus S$, the set of included jobs $J_{t,t'} := \{j : [r_j, d_j] \subseteq [t, t']\}$, and the density $\sigma_{t,t'} := \sum_{j \in J_{t,t'}} w_j / |I_{t,t'}|$. The procedure starts with $S = \emptyset$, and while not all jobs are scheduled, selects an interval $[t, t']$ with

maximal density, and schedules all jobs from $J_{t,t'}$ in earliest deadline order in $I_{t,t'}$ at speed $\sigma_{t,t'}$ adding $I_{t,t'}$ to S .

For the upper bound, let $t_1 < t_2 < \dots < t_\ell$ be the sequence of all release times and deadlines for some $\ell \leq 2n$. Clearly both schedules S run at uniform speed in every interval $[t_{k-1}, t_k)$. For every $1 \leq k \leq n$ let s_k be the speed of S in $[t_{k-1}, t_k)$, and s'_k the speed of S' in the same interval.

From the algorithm above it follows that every job is scheduled at constant speed, so let s_a be the speed at which job i is scheduled in $\text{OPT}(d)$. It also follows that if $s_k > s_a$, then $s'_k = s_k$, and if $s_k \leq s_a$, then $s'_k \leq s_k$.

We establish the following upper bound.

$$\begin{aligned}
E(\text{OPT}(d)) - E(\text{OPT}(d_{-i})) &= \sum_{k=1}^{\ell} s_k^\alpha (t_k - t_{k-1}) - s_k'^\alpha (t_k - t_{k-1}) \\
&= \sum (t_k - t_{k-1}) (s_k^\alpha - (s_k - (s_k - s'_k))^\alpha) \\
&= \sum (t_k - t_{k-1}) s_k^\alpha \left(1 - \left(1 - \frac{s_k - s'_k}{s_k} \right)^\alpha \right) \\
&\leq \sum (t_k - t_{k-1}) s_k^\alpha \left(1 - \left(1 - \alpha \frac{s_k - s'_k}{s_k} \right) \right) \\
&= \sum (t_k - t_{k-1}) \alpha s_k^{\alpha-1} (s_k - s'_k) \\
&\leq \alpha s_a^{\alpha-1} \sum (t_k - t_{k-1}) (s_k - s'_k) \\
&= \alpha s_a^{\alpha-1} w_i \\
&= \alpha (E(\text{OPT}(d)) - E(S_{-i})).
\end{aligned}$$

The first inequality uses the generalized Bernoulli inequality, and the last one the fact that for all k with $s_k \neq s'_k$ we have $s_k \leq s_a$.

The theorem follows from the fact that $s_a^{\alpha-1} w_i$ is precisely the proportional cost share of job i in $\text{OPT}(d)$. \square

A tight example is given by n jobs, each with workload $1/n$, release time 0 and deadline 1. Clearly the optimal energy consumption is 1 for this instance. The marginal cost share for each player is $1 - (1 - 1/n)^\alpha$. Finally we observe that the total marginal cost share tends to α , i.e.

$$\lim_{n \rightarrow +\infty} n - n(1 - 1/n)^\alpha = \alpha.$$

5 A note on cross-monotonicity

We conclude this paper by a short note on *cross-monotonicity*. This is a property of cost sharing games, stating that whenever new players enter the game, the cost share of any fixed player does not increase. This property is useful for stability in the game, and is the key to the Moulin carving algorithm [13], which selects a set of players to be served for specific games.

In the game that we consider, the minimum energy of an optimal schedule for a set S of jobs contrasts with many studied games, where serving more players becomes more cost effective, because the used equipment is better used.

Consider a very simple example of two identical players, submitting their respective jobs with the same deadline 1. Suppose the workload of each job is w , then the minimum energy necessary to schedule one job is w^α , while the cost to serve both jobs is $(2w)^\alpha$, meaning that the cost share increase whenever a second player enters the game. Therefore the marginal cost sharing scheme is not cross-monotonic.

6 Acknowledgements

We would like to thank anonymous referees for remarks and suggestions on an earlier version of this paper.

Christoph Dürr and Oscar C. Vásquez were partially supported by grant ANR-11-BS02-0015. Lukasz Jeż was partially supported by MNiSW grant N N206 368839, 2010-2013, EU ERC project 259515 PAAI, and FNP Start scholarship.

References

1. S. Albers. Energy-efficient algorithms. *Communications of the ACM*, 53(5):86–96, 2010.
2. Susanne Albers and Hiroshi Fujiwara. Energy-efficient algorithms for flow time minimization. *ACM Transactions on Algorithms (TALG)*, 3(4):49, 2007.
3. Eric Angel, Evripidis Bampis, and Fadi Kacem. Energy aware scheduling for unrelated parallel machines. In *Green Computing and Communications (GreenCom), 2012 IEEE International Conference on*, pages 533–540. IEEE, 2012.
4. Nikhil Bansal, Ho-Leung Chan, Dmitriy Katz, and Kirk Pruhs. Improved bounds for speed scaling in devices obeying the cube-root rule. *Theory OF Computing*, 8:209–229, 2012.
5. Nikhil Bansal, Tracy Kimbrel, and Kirk Pruhs. Speed scaling to manage energy and temperature. *Journal of the ACM (JACM)*, 54(1):3, 2007.
6. David M Brooks, Pradip Bose, Stanley E Schuster, Hans Jacobson, Prabhakar N Kudva, Alper Buyuktosunoglu, J Wellman, Victor Zyuban, Manish Gupta, and Peter W Cook. Power-aware microarchitecture: Design and modeling challenges for next-generation microprocessors. *Micro, IEEE*, 20(6):26–44, 2000.
7. Rodrigo A. Carrasco, Garud Iyengar, and Clifford Stein. Energy aware scheduling for weighted completion time and weighted tardiness. Technical report, arxiv.org, 2011.
8. Sze-Hang Chan, Tak-Wah Lam, and Lap-Kei Lee. Non-clairvoyant speed scaling for weighted flow time. In *Algorithms-ESA 2010*, pages 23–35. Springer, 2010.
9. S. Irani and K.R. Pruhs. Algorithmic problems in power management. *ACM SIGACT News*, 36(2):63–76, 2005.
10. M.g Li, Andrew C. Yao, and Frances F. Yao. Discrete and continuous min-energy schedules for variable voltage processor. In *Proceedings of the National Academy of Sciences of the United States of America*, volume 103 of *PNAS'06*, pages 3983–3987. National Academy of Sciences, 2006.

11. Nicole Megow and José Verschae. Dual techniques for scheduling on a machine with varying speed. In *Proc. of the 40th International Colloquium on Automata, Languages and Programming (ICALP)*, 2013.
12. D. Monderer and L.S. Shapley. Potential games. *Games and Economic Behavior*, 14:124–143, 1996.
13. H. Moulin and S. Shenker. Strategyproof sharing of submodular costs: budget balance versus efficiency. *Economic Theory*, 18(3):511–533, 2001.
14. Kirk Pruhs, Patchrawat Uthaisombut, and Gerhard Woeginger. Getting the best response for your erg. *ACM Transactions on Algorithms (TALG)*, 4(3):38, 2008.
15. O.C. Vasquez. Energy in computing systems with speed scaling: optimization and mechanisms design. Technical report, arxiv.org, 2012.
16. F. Yao, A. Demers, and S. Shenker. A scheduling model for reduced cpu energy. In *Proceedings of the 36th Annual Symposium on Foundations of Computer Science, FOCS '95*, pages 374–382, Washington, DC, USA, 1995. IEEE Computer Society.