

Non-clairvoyant Scheduling for Weighted Flow Time and Energy on Speed Bounded Processors

Sze-Hang Chan¹ Tak-Wah Lam¹ Lap-Kei Lee¹ Hing-Fung Ting¹ Peng Zhang²

¹Department of Computer Science
University of Hong Kong, Hong Kong
Email: {shchan, twlam, lklee, hfting}@cs.hku.hk

² School of Computer Science and Technology
Shandong University, China
Email: algzhang@sdu.edu.cn

Abstract

We consider the online scheduling problem of minimizing total weighted flow time plus energy on a processor that can scale its speed dynamically between 0 and some maximum speed T . In the past few years this problem has been studied extensively under the clairvoyant setting, which requires the size of a job to be known when it is released [1, 4, 5, 8, 12, 15, 16, 17]. For the non-clairvoyant setting, despite its practical importance, the progress is relatively limited. Only recently an online algorithm LAPS is known to be $O(1)$ -competitive for minimizing (unweighted) flow time plus energy in the infinite speed model (i.e., $T = \infty$) [10, 11]. This paper makes two contributions to the non-clairvoyant scheduling. First, we resolve the open problem that the unweighted result of LAPS can be extended to the more realistic model with bounded maximum speed. Second, we show that another non-clairvoyant algorithm WRR is $O(1)$ -competitive when weighted flow time is concerned.

Keywords: online algorithms, speed scaling, energy-efficient scheduling, weighted flow time

1 Introduction

Energy consumption has become an important concern for the design of modern microprocessors. Manufacturers like Intel and IBM are now producing pro-

cessors that can support *dynamic speed scaling*, which would allow operating systems to manage the power by scaling the processor speed dynamically. Running jobs slower saves more energy, yet it takes longer time. Taking speed scaling and energy usage into consideration makes job scheduling more complicated than before. The challenge arises from the conflicting objectives of optimizing some quality of service (QoS) of the schedule and minimizing the energy usage.

The theoretical study of speed scaling was initiated by Yao, Demers and Shenker [21]. They considered a model where a processor can vary its speed s between 0 and infinity dynamically, and it consumes energy at the rate s^α , where α is a constant (commonly believed to be 2 or 3 [2, 19]). Under this infinite speed model, Yao et al. studied the deadline scheduling and gave an online algorithm that is $O(1)$ -competitive for minimizing the energy for completing all jobs. Algorithms with better ratios were later obtained by Bansal, Kimbrel and Pruhs [7] and Bansal, Chan, Pruhs and Katz [6]. The best ratio now is $4^\alpha/2\sqrt{e\alpha}$ for general α and is about 6.7 when $\alpha = 3$. The infinite speed model is a convenient model to work with. Among others, it allows an online algorithm to catch up arbitrarily fast and recover any over-conservative decision on speed. However, this is not a practical model. Recently, Chan et al. [9] and Bansal et al. [4] have obtained several interesting results on speed scaling for a speed bounded processor, where the maximum speed T is a fixed constant.

Flow and energy. The study of speed scaling and energy-efficient scheduling goes beyond deadline scheduling. When scheduling jobs without deadlines, a commonly used QoS measure is the total flow time of jobs. The flow time (or simply the flow) of a job is the time elapsed since the job is released until it is completed. Note that job preemption is allowed, and a preempted job can be resumed later at the point of preemption. Assume jobs are equally important, it is natural to find a schedule that minimizes the total flow time (which is also referred to as the to-

T.W. Lam is partly supported by HKU Grant 7176104.
P. Zhang is partly supported by NSFC 60970003.

Copyright ©2010, Australian Computer Society, Inc. This paper appeared at the 16th Computing: the Australasian Theory Symposium (CATS2010), Brisbane, Australia. Conferences in Research and Practice in Information Technology (CRPIT), Vol. 109, A. Potanin and A. Viglas, Ed. Reproduction for academic, not-for profit purposes permitted provided this text is included.

tal/average response time in the literature). When jobs have varying importance or weights, it is more meaningful to minimize the total weighted flow time.

Minimizing flow time and minimizing energy usage are orthogonal objectives. To understand their tradeoff, Albers and Fujiwara [1] initiated the study of minimizing a linear combination of flow and energy. The intuition is that, from an economic viewpoint, users are willing to pay a certain (say, ρ) units of energy to reduce one unit of flow time. By changing the units of time and energy, one can further assume that $\rho = 1$ and thus would like to minimize flow time plus energy, or in general, weighted flow time plus energy.

Clairvoyant scheduling for flow plus energy.

The problem of minimizing flow plus energy has attracted a lot of attention [1, 4, 5, 8, 10, 15, 16, 17]. These works mainly focus on the clairvoyant setting which assumes that the size of a job is known when the job is released. The work of Albers and Fujiwara [1] focused on jobs of unit size. Bansal, Pruhs and Stein [8] were the first to consider jobs of arbitrary sizes. In the infinite speed model, they gave an algorithm that is $O((\frac{\alpha}{\ln \alpha})^2)$ -competitive for minimizing weighted flow plus energy. Bansal, Chan, Lam and Lee [4] later adapted the BPS algorithm to the bounded speed model. The competitive ratio remains $O((\frac{\alpha}{\ln \alpha})^2)$ when the algorithm uses a processor with maximum speed $(1 + \frac{\ln \alpha - \ln \ln \alpha}{\alpha})T$.¹ Very recently, Bansal, Chan and Pruhs [5] improved the analysis of BPS; their work implies that BPS is indeed $O(\frac{\alpha}{\ln \alpha})$ -competitive; again, the maximum speed is relaxed as before. It is worth-mentioning that the recent lower bound result on weighted flow [3] implies that without relaxing the maximum speed, no online algorithm can be constant competitive (in terms of α) for weighted flow plus energy. However, the extra speed requirement does not apply to unweighted flow.

A drawback of the BPS algorithm is that it scales the speed according to the fraction of unfinished work and thus makes the speed keep changing over time. This is practically undesirable. Recently, focusing on (unweighted) flow plus energy, Lam et al. [17] studied a different speed scaling algorithm AJC (Active Job Count), which scales the speed as a function of the number of active jobs. In other words, AJC changes the speed only at discrete times (at job arrival and completion). Technically, Lam et al. derived a potential analysis to directly account for integral flow (instead of via fractional flow) and show that AJC when coupled with the job selection algorithm SRPT (shortest remaining processing time) is indeed $O(\frac{\alpha}{\log \alpha})$ -competitive for (unweighted) flow plus energy. This result holds in both the infinite and bounded speed models. For the latter, unlike BPS, AJC does not demand relaxation of maximum speed.

¹In general, when using a processor with maximum speed $(1 + \epsilon)T$ for any $\epsilon > 0$, the competitive ratio is $\max\{(1 + 1/\epsilon), (1 + \epsilon)\}(2 + o(1))\alpha / \ln \alpha$.

Recently, Bansal, Chan and Pruhs [5] adapted AJC and gave a tighter analysis; they showed that the competitive ratio is 3 for minimizing (unweighted) flow plus energy (even when α is as small as 3, this result is still better than the $O(\frac{\alpha}{\log \alpha})$ bound, which is equal to 3.25). Again, no extra maximum speed is needed.

For weighted flow plus energy, it remains an open problem whether AJC (or any scaling algorithm that changes the speed discretely) can lead to a competitive guarantee.

Non-clairvoyant scheduling for flow plus energy.

All of the above results assume clairvoyance. In the non-clairvoyant setting, the size of a job is not known when the job is released; it is only known when the job is completed. This is a natural assumption from the viewpoint of operating systems. Non-clairvoyant flow time scheduling (on a fixed-speed processor) has been an interesting problem itself (e.g., [14, 18]). Chan et al. [10] initiated the study of non-clairvoyant speed scaling. Under the infinite speed model, they consider an algorithm LAPS which scales the speed as AJC and selects some most recently released jobs to share the processor. LAPS is shown to be $4\alpha^3(1 + (1 + 3/\alpha)^\alpha) = O(\alpha^3)$ -competitive for (unweighted) flow plus energy. Furthermore, they showed that no algorithm can be $O(\alpha^{1/3-\epsilon})$ -competitive for any $\epsilon > 0$, illustrating that the non-clairvoyant setting is more difficult than the clairvoyant setting. Recently, Chan et al. [11] improved the analysis of LAPS and reduce the competitive ratio to $O(\frac{\alpha^2}{\log \alpha})$. Yet, not much is known for the bounded speed model, let alone weighted flow plus energy.

Our contributions. This paper considers non-clairvoyant scheduling on a processor whose maximum speed T is a fixed constant. In the first part, we adapt the algorithm LAPS to run on such a processor and show that it is $8\alpha^2 = O(\alpha^2)$ -competitive for (unweighted) flow plus energy when the maximum speed is relaxed to $(1 + \frac{1}{\alpha-1})T$. Note that unlike the clairvoyant setting, even for unweighted jobs, extra maximum speed is necessary to achieve constant competitiveness. This inherits from the lower bound result on non-clairvoyant (fixed-speed) scheduling by Motwani, Philips and Torng [18].

In the clairvoyant setting, existing results on the bounded speed model take advantage of a local property that the online algorithm in concern accumulates at most $O(T^\alpha)$ jobs more than the optimal offline algorithm [4, 5, 17]. With this local property, it is relatively easy to adapt the analysis in the infinite speed model. In the non-clairvoyant setting, such a local property is no longer valid for algorithms like LAPS. To analyze these algorithms in the bounded speed model, we exploit a more “global” accounting of the rate of change of flow plus energy. Instead of using the above property, we integrate the maximum speed constraint into the potential analysis.

The second result of this paper concerns the more difficult general case where jobs have arbitrary weights. Under the bounded speed model, we give the first competitive algorithm called WRR (Weighted Round Robin) for weighted flow plus energy; the competitive ratio is $O(3^\alpha/\epsilon)$ when using a processor with maximum speed $(3 + \epsilon)T$, where $0 < \epsilon \leq \frac{3}{\alpha}$. Motivated by AJC, WRR uses a generalized AJC for speed scaling; i.e., the speed is a function of the total weight (instead of the number) of active jobs. Recall that all existing clairvoyant results [4, 8] on weighted flow plus energy are based on the BPS algorithm, which scales the speed continuously. Our result of WRR gives, as a by-product, the first competitive clairvoyant algorithm for weighted flow plus energy that changes the speed discretely, but the competitive ratio is worse than that of BPS if α is large.

2 Definitions and Notations

We study job scheduling on a single processor. Jobs arrive over time in an online fashion; we have no information about a job before it arrives. For any job j , we use $r(j)$ and $p(j)$ to denote its release time and work requirement (or size). In some case, each job j may have a weight $w(j)$. We consider the *non-clairvoyant* setting, in which when a job j arrives, we only know its weight $w(j)$ (if any) but not its size $p(j)$. And $p(j)$ is known only when j is completed. At any time, the processor can scale its speed between 0 and a *maximum speed* T . When running at speed s , the processor processes s units of work per unit time and consumes energy at the rate s^α , where $\alpha > 1$ is a fixed constant. Preemption is allowed; a job can be preempted and later resumed at the point of preemption without any penalty.

Flow and energy. Consider any job set I and some schedule S of I . At any time t , for any job j , we let $q(j, t)$ be the remaining work of j at t . A job j is an *active job* if it has been released but not yet completed, i.e., $r(j) \leq t$ and $q(j, t) > 0$. The flow $F(j)$ of a job j is the time elapsed since j arrives and until it is completed. The *total flow* F is equal to $\sum_{j \in I} F(j)$, which is equivalent to

$$\int_0^\infty n(t) dt$$

where $n(t)$ is the total number of active jobs at time t . The energy usage E is $\int_0^\infty (s(t))^\alpha dt$, where $s(t)$ is the processor speed at time t . The objective is to minimize the sum of total flow and energy usage, denoted by $G = F + E$.

In general, when jobs have different weights, we generalize the notion of total flow as follows. The *total weighted flow* \widehat{F} is equal to $\sum_{j \in I} w(j)F(j)$, or equivalently,

$$\int_0^\infty w(t) dt$$

where $w(t)$ is the total weight of active jobs at time t . The objective becomes minimizing total weighted flow plus energy, denoted by $\widehat{G} = \widehat{F} + E$.

3 Minimizing unweighted flow plus energy with bounded maximum speed

In this section, we consider jobs without weights and aim at minimizing total flow plus energy in the bounded speed model. We adapt the non-clairvoyant algorithm LAPS (Latest Arrival Processor Sharing) which was first given in [10] for the infinite speed model. When using a processor with maximum speed $(1 + \delta)T$, where $\delta = \frac{1}{\alpha-1}$, this algorithm is $O(\alpha^2)$ -competitive for flow plus energy. Below is the definition of LAPS, which assumes using a processor with maximum speed $(1 + \delta)T$ for some $\delta > 0$.

Algorithm LAPS. Let $0 < \beta \leq 1$ be any real. Consider any time t . The processor speed is set to $s_a(t) = \min((n_a(t))^{1/\alpha}, (1 + \delta)T)$, where $n_a(t)$ is the total number of active jobs at t . The processor processes the $\lceil \beta n_a(t) \rceil$ active jobs with the latest release time (ties are broken by job ids) by splitting the processor speed equally among these jobs.

We compare LAPS with an optimal offline algorithm OPT using a processor with maximum speed T . Our main result is the following theorem.

Theorem 1. *When $\delta = \frac{1}{\alpha-1}$ and $\beta = \frac{1}{2\alpha}$, LAPS is $8\alpha^2$ -competitive for (unweighted) flow plus energy, using a processor with maximum speed $(1 + \delta)T$.*

To prove Theorem 1, our analysis exploits amortization and potential functions (e.g., [8, 9]). Let $G_a(t)$ and $G_o(t)$ denote the flow plus energy incurred up to time t by LAPS and OPT, respectively. We drop the parameter t when it is clear that t is the current time. To show that LAPS is c -competitive for some constant $c \geq 1$, it suffices to define a potential function $\Phi(t)$ such that the following conditions hold: (i) $\Phi = 0$ before any job is released and after all jobs are completed; (ii) Φ is a continuous function except at some discrete times (e.g., when a job arrives, or when a job is completed by LAPS or OPT), and Φ does not increase at such times; (iii) at any other time, $\frac{dG_a(t)}{dt} + \gamma \frac{d\Phi(t)}{dt} \leq c \cdot \frac{dG_o(t)}{dt}$, where γ is a positive constant (to be set to 4α). Condition (iii) is also known as the running condition.

Potential function $\Phi(t)$. Consider any time t . Let $n_a(t)$ and $n_o(t)$ be the number of active jobs in LAPS and OPT, respectively. Let $j_1, j_2, \dots, j_{n_a(t)}$ be all the active jobs in LAPS, which are ordered by release times such that $r(j_1) \leq r(j_2) \leq \dots \leq r(j_{n_a(t)})$ (ties are broken by job ids). For any job j , let $q_a(j, t)$ and $q_o(j, t)$ be the remaining work of job j

in LAPS and OPT, respectively. For each j_i , let $x_i = \max\{q_a(j_i, t) - q_o(j_i, t), 0\}$ which is the amount of work of j_i in LAPS that is lagging behind OPT. We call a job j_i *lagging* if $x_i > 0$. The potential function $\Phi(t)$ is defined as follows.

$$\Phi(t) = \sum_{i=1}^{n_a(t)} c_i \cdot x_i$$

where $c_i = \begin{cases} i^{1-1/\alpha} & \text{if } i^{1/\alpha} \leq (1 + \delta)T; \\ i/((1 + \delta)T) & \text{otherwise.} \end{cases}$

We call c_i the *coefficient* of j_i . Note that c_i is a concave function of i and monotonically increasing.

We first check Conditions (i) and (ii). Condition (i) holds since $\Phi = 0$ before any job is released and after all jobs are completed. Now we check Condition (ii). When a job j arrives, j must be non-lagging and the coefficients of all existing jobs of LAPS remain the same, so Φ does not change. When OPT completes a job, Φ does not change. When LAPS completes a job, the coefficient of any other job either stays the same or decreases, so Φ does not increase.

It remains to check the running condition (Condition (iii)). Consider any time t when Φ does not have discrete change. Let s_a and s_o be the current speeds of LAPS and OPT, respectively. Then $\frac{dG_a}{dt} = n_a + s_a^\alpha$ and $\frac{dG_o}{dt} = n_o + s_o^\alpha$. Let ℓ be the number of lagging jobs that LAPS is processing. Note that $\ell \leq \lceil \beta n_a \rceil$. For convenience, we further define another real number $\phi \leq \beta$ such that $(\beta - \phi)n_a$ is an integer equal to $\lceil \beta n_a \rceil - \ell$. Note that ϕ can be less than zero if $\ell = 0$.

To bound the rate of change of Φ , we consider how Φ changes in an infinitesimal amount of time (from t to $t + dt$), first due to LAPS only (Lemma 2), and then due to OPT (Lemma 3). We denote the rate of change of Φ due to LAPS and OPT by $\frac{d\Phi_a}{dt}$ and $\frac{d\Phi_o}{dt}$, respectively. Note that $\frac{d\Phi}{dt} = \frac{d\Phi_a}{dt} + \frac{d\Phi_o}{dt}$.

Lemma 2. $\frac{d\Phi_a}{dt} \leq -\frac{\phi}{\beta}(1 - \beta + \frac{\phi}{2})n_a$.

Proof. LAPS is processing $\lceil \beta n_a \rceil$ jobs and ℓ of them are lagging jobs. For each of these lagging jobs j_i , its lagging size x_i is changing at the rate of $-s_a/\lceil \beta n_a \rceil$ (we only consider the change due to LAPS). For a non-lagging jobs j_i , x_i does not change.

First, consider the case that $\ell \geq 1$. To upper bound $\frac{d\Phi_a}{dt}$, the worst case is that the ℓ lagging jobs have the smallest coefficients among the $\lceil \beta n_a \rceil$ latest released jobs, i.e., the jobs are $\{j_{n_a - \lceil \beta n_a \rceil + 1}, \dots, j_{n_a - \lceil \beta n_a \rceil + \ell}\}$. On the other hand, as c_i is concave, for any integers $a < b$, we have $\sum_{i=a}^b c_i / (b - a + 1) \geq (c_a + c_b)/2$. Then

$$\begin{aligned} \frac{d\Phi_a}{dt} &\leq - \sum_{i=n_a - \lceil \beta n_a \rceil + 1}^{n_a - \lceil \beta n_a \rceil + \ell} c_i \cdot \frac{s_a}{\lceil \beta n_a \rceil} \\ &\leq -\frac{\ell}{2} (c_{n_a - \lceil \beta n_a \rceil + 1} + c_{n_a - \lceil \beta n_a \rceil + \ell}) \frac{s_a}{\lceil \beta n_a \rceil}. \end{aligned}$$

Recall that $s_a = \min(n_a^{1/\alpha}, (1 + \delta)T)$. By the definition of c_i , we have $c_i s_a \geq i$ for any $1 \leq i \leq n_a$. Furthermore, $\ell = \lceil \beta n_a \rceil - (\beta - \phi)n_a$ and hence $n_a - \lceil \beta n_a \rceil + \ell = n_a - (\beta - \phi)n_a$. Thus,

$$\begin{aligned} \frac{d\Phi_a}{dt} &\leq - \left(\frac{\ell}{2\lceil \beta n_a \rceil} \right) ((n_a - \lceil \beta n_a \rceil + 1) + (n_a - \lceil \beta n_a \rceil + \ell)) \\ &\leq \left(\frac{-\lceil \beta n_a \rceil + (\beta - \phi)n_a}{\lceil \beta n_a \rceil} \right) \left(\frac{(n_a - \beta n_a) + (n_a - (\beta - \phi)n_a)}{2} \right) \\ &\leq (-1 + \frac{(\beta - \phi)n_a}{\beta n_a}) (1 - \beta + \frac{\phi}{2}) n_a \quad (\because (\beta - \phi) \geq 0) \\ &= -\frac{\phi}{\beta} (1 - \beta + \frac{\phi}{2}) n_a. \end{aligned}$$

It remains to consider the case that $\ell = 0$. In this case, $\frac{d\Phi_a}{dt} = 0$. Since $0 < \beta \leq 1$, we have $(1 - \beta + \frac{\phi}{2})n_a = \frac{1}{2}((1 - \beta)n_a + (n_a - (\beta - \phi)n_a)) \geq \frac{1}{2}(n_a - (\beta - \phi)n_a) = \frac{1}{2}(n_a - \lceil \beta n_a \rceil + \ell) \geq 0$. Note also that $\phi n_a = \beta n_a - \lceil \beta n_a \rceil + \ell \leq 0$, i.e., $\phi \leq 0$. Then $-\frac{\phi}{\beta}(1 - \beta + \frac{\phi}{2})n_a \geq 0 = \frac{d\Phi_a}{dt}$. \square

Lemma 3. Assume that $\delta = \frac{1}{\alpha - 1}$. Then $\frac{d\Phi_o}{dt} \leq (1 - \frac{1}{\alpha})n_a + \frac{1}{\alpha}s_o^\alpha$.

Proof. To upper bound $\frac{d\Phi_o}{dt}$, the worst case is that OPT is processing the job j_{n_a} with the largest coefficient c_{n_a} . Thus, $\frac{d\Phi_o}{dt} \leq c_{n_a} s_o$.

If $n_a^{1/\alpha} \leq (1 + \delta)T$, we have $c_{n_a} = n_a^{1-1/\alpha}$ and hence $\frac{d\Phi_o}{dt} \leq n_a^{1-1/\alpha} s_o$. We apply the Young's Inequality [20], which is stated in Lemma 5 below, by setting $p = 1/(1 - \frac{1}{\alpha})$, $q = \alpha$, $x = n_a^{1-1/\alpha}$ and $y = s_o$. Then $\frac{d\Phi_o}{dt} \leq n_a^{1-1/\alpha} s_o \leq (1 - \frac{1}{\alpha})n_a + \frac{1}{\alpha}s_o^\alpha$.

If $n_a^{1/\alpha} > (1 + \delta)T$, we have $c_{n_a} = \frac{n_a}{(1 + \delta)T}$. Recall that $s_o \leq T$ and $\delta = \frac{1}{\alpha - 1}$. We conclude that $\frac{d\Phi_o}{dt} \leq \left(\frac{n_a}{(1 + \delta)T} \right) s_o \leq \left(\frac{n_a}{1 + \delta} \right) = (1 - \frac{1}{\alpha})n_a$. The lemma thus follows. \square

We are now ready to prove the running condition, which together with Conditions (i) and (ii), implies Theorem 1.

Lemma 4. Assume that $\delta = \frac{1}{\alpha - 1}$, $\beta = \frac{1}{2\alpha}$, and $\gamma = 4\alpha$. At any time when Φ does not have discrete change, $\frac{dG_a}{dt} + \gamma \frac{d\Phi}{dt} \leq 8\alpha^2 \cdot \frac{dG_o}{dt}$.

Proof. LAPS is processing $\lceil \beta n_a \rceil - \ell$ non-lagging jobs, which are also active jobs in OPT. Thus, $n_o \geq \lceil \beta n_a \rceil - \ell = (\beta - \phi)n_a$. Note that $\frac{dG_a}{dt} = n_a + s_a^\alpha \leq 2n_a$, and $\frac{dG_o}{dt} = n_o + s_o^\alpha \geq (\beta - \phi)n_a + s_o^\alpha$. By Lemmas 2 and 3, $\frac{dG_a}{dt} + \gamma \frac{d\Phi}{dt} - 8\alpha^2 \cdot \frac{dG_o}{dt} \leq 2n_a + \gamma \cdot (1 - \frac{1}{\alpha})n_a + \frac{\gamma}{\alpha}s_o^\alpha - \gamma \frac{\phi}{\beta}(1 - \beta + \frac{\phi}{2})n_a - 8\alpha^2(\beta - \phi)n_a - 8\alpha^2 s_o^\alpha$. We set $\beta = \frac{1}{2\alpha}$ and $\gamma = 4\alpha$. Then the above inequality

becomes

$$\begin{aligned}
 & \frac{dG_a}{dt} + \gamma \frac{d\Phi}{dt} - 8\alpha^2 \cdot \frac{dG_a}{dt} \\
 & \leq (4 - 8\alpha^2)s_o^\alpha + n_a(2 + (4\alpha - 4) - \\
 & \quad \phi(8\alpha^2 - 4\alpha + 4\alpha^2\phi) - (4\alpha - 8\alpha^2\phi)) \\
 & \leq n_a(-2 - 4\alpha^2\phi^2 + 4\alpha\phi) \\
 & = n_a(-1 - (2\alpha\phi - 1)^2) \\
 & \leq 0.
 \end{aligned}$$

□

Below is the formal statement of Young's Inequality, which is used in the proof of Lemma 3.

Lemma 5 (Young's Inequality [20]). *For positive reals p, q, x, y where $\frac{1}{p} + \frac{1}{q} = 1$, $xy \leq \frac{1}{p}x^p + \frac{1}{q}y^q$.*

4 Minimizing weighted flow plus energy with bounded maximum speed

In this section, we consider jobs with arbitrary weights and give a non-clairvoyant algorithm WRR (Weighted Round Robin) that is $O(3^\alpha/\epsilon)$ -competitive for weighted flow plus energy, when using a processor with maximum speed $(3 + \epsilon)T$, where $0 < \epsilon \leq \frac{3}{\alpha}$. Note that when $\epsilon = \frac{3}{\alpha}$, the competitive ratio becomes $O(\alpha 3^\alpha)$. The algorithm WRR scales its speed based on the total weight of active jobs and shares the processor among the active jobs according to the ratio of their weights. Below is the formal definition of WRR, which assumes using a processor with maximum speed $(3 + \epsilon)T$ for any $\epsilon > 0$.

Algorithm WRR. Consider any time t . The processor speed is set to $s_a(t) = (3 + \epsilon) \cdot \min((w_a(t))^{1/\alpha}, T)$, where $w_a(t)$ is the total weight of active jobs at t . The processor processes all active jobs such that each active job j receives processor speed equal to $s(t) \cdot (w(j)/w_a(t))$.

We compare WRR against an optimal offline algorithm OPT that uses a processor with maximum speed T . Our main result is the following theorem. As a remark, when $\epsilon = \frac{3}{\alpha}$, the competitive ratio is $(6\alpha + 4)(1 + 3^\alpha(1 + \frac{1}{\alpha})^\alpha) = O(\alpha 3^\alpha)$.

Theorem 6. *Using a processor with maximum speed $(3 + \epsilon)T$, where $0 < \epsilon \leq \frac{3}{\alpha}$, WRR is c -competitive for weighted flow plus energy, where $c = (\frac{18}{\epsilon} + 4)(1 + (3 + \epsilon)^\alpha) \leq (\frac{18}{\epsilon} + 4)(1 + 3^\alpha\epsilon) = O(3^\alpha/\epsilon)$.*

The rest of this section is devoted to proving Theorem 6. Let $\widehat{G}_a(t)$ and $\widehat{G}_o(t)$ be the weighted flow plus energy incurred up to time t by WRR and OPT, respectively. We drop the parameter t when it is clear that t is the current time. Similar to Section 3, to

prove that WRR is c -competitive, we derive a potential function $\Phi(t)$ that satisfies the following conditions: (i) $\Phi = 0$ before any job is released and after all jobs are completed; (ii) Φ is a continuous function except at some discrete times where Φ does not increase; (iii) *Running condition*: at any other time, $\frac{d\widehat{G}_a(t)}{dt} + \gamma \frac{d\Phi(t)}{dt} \leq c \cdot \frac{d\widehat{G}_o(t)}{dt}$, where γ is a positive constant (to be set to $(2 - \frac{1}{\alpha})(1 + (3 + \epsilon)^\alpha)$).

Potential function $\Phi(t)$. Consider any time t . For any job j , let $q_a(j, t)$ and $q_o(j, t)$ be the remaining work of j at t in WRR and OPT, respectively. An active job j in WRR is *lagging* if WRR has processed less on j than OPT at time t . Let $L = \{j_1, j_2, \dots, j_\ell\}$ be the set of lagging jobs in WRR, ordered in descending order of their lagging size x_i (recall that $x_i = q_a(j_i, t) - q_o(j_i, t)$). In other words, $x_1 \geq x_2 \geq \dots \geq x_\ell$. We define the potential function $\Phi(t)$ as follows.

$$\Phi(t) = \sum_{i=1}^{\ell} c_i \cdot x_i$$

where $c_i = \begin{cases} \left(\sum_{k=1}^i w(j_k)\right)^{1-1/\alpha} & \text{if } \sum_{k=1}^i w(j_k) \leq T^\alpha; \\ \frac{2\alpha}{2\alpha-1} \left(\sum_{k=1}^i \frac{w(j_k)}{T}\right) & \text{otherwise.} \end{cases}$

We call c_i the *coefficient* of j_i . Note that c_i is monotonically increasing with i , but c_i is not a concave function.

We first check Conditions (i) and (ii). Condition (i) holds since $\Phi = 0$ before any job is released and after all jobs are completed. Now, we show that Condition (ii) holds. When a job j_i joins L or leaves L (e.g., WRR completes j_i), x_i must be zero and j_i must be at the end of L , so the coefficients of all other jobs do not change and Φ does not change. When two jobs j_i and j_j change their order in L , it must be the case that $x_i = x_j$ and thus the value of Φ does not change.

It remains to check the running condition (Condition (iii)). Consider any time t when Φ does not have discrete change. Let w_a and w_o be the total weight of active jobs at t in WRR and OPT, respectively. Let $w_\ell = \sum_{i=1}^{\ell} w(j_i)$ be the total weight of jobs in L . Note that $w_\ell \leq w_a$. Furthermore, let s_a and s_o be the current speeds of WRR and OPT, respectively. As stated in Section 2, $\frac{d\widehat{G}_a}{dt} = w_a + s_a^\alpha$ and $\frac{d\widehat{G}_o}{dt} = w_o + s_o^\alpha$. We will divide the analysis into cases depending on whether w_ℓ is small or big.

To bound the rate of change of Φ , we consider how Φ changes first due to OPT only (Lemma 7) and then due to WRR (Lemma 8). We denote the rate of change of Φ due to OPT and WRR by $\frac{d\Phi_o}{dt}$ and $\frac{d\Phi_a}{dt}$, respectively. Note that $\frac{d\Phi}{dt} = \frac{d\Phi_o}{dt} + \frac{d\Phi_a}{dt}$.

Lemma 7. *If $w_\ell \leq T^\alpha$, $\frac{d\Phi_o}{dt} \leq \frac{1}{\alpha}s_o^\alpha + (\frac{\alpha-1}{\alpha})w_\ell$; if $w_\ell > T^\alpha$, $\frac{d\Phi_o}{dt} \leq (\frac{2\alpha-1}{2\alpha-1})w_\ell$.*

Proof. To upper bound $\frac{d\Phi_o}{dt}$, observe that the worst case is when OPT is processing the job j_ℓ with the largest coefficient c_ℓ . Then x_i is increasing at the rate of s_o (we only consider the change due to OPT) and hence $\frac{d\Phi_o}{dt} \leq c_\ell s_o$. When $w_\ell \leq T^\alpha$, $c_\ell = w_\ell^{1-1/\alpha}$ and thus $\frac{d\Phi_o}{dt} \leq w_\ell^{1-1/\alpha} s_o$. We apply the Young's Inequality (Lemma 5 in Section 3) with $p = \alpha$, $q = \alpha/(\alpha - 1)$, $x = s_o$ and $y = w_\ell^{1-1/\alpha}$. Then we have

$$\frac{d\Phi_o}{dt} \leq \frac{1}{\alpha} s_o^\alpha + \left(\frac{\alpha-1}{\alpha}\right) w_\ell.$$

When $w_\ell > T^\alpha$, $c_\ell = \left(\frac{2\alpha}{2\alpha-1}\right)\left(\frac{w_\ell}{T}\right)$. Since $s_o \leq T$, $\frac{d\Phi_o}{dt} \leq c_\ell s_o \leq c_\ell T \leq \left(\frac{2\alpha}{2\alpha-1}\right) w_\ell$. \square

Lemma 8. If $w_\ell \leq T^\alpha$, $\frac{d\Phi_a}{dt} \leq -\left(\frac{\alpha}{2\alpha-1}\right) w_\ell^{2-1/\alpha} \left(\frac{s_a}{w_a}\right)$; if $w_\ell > T^\alpha$, $\frac{d\Phi_a}{dt} \leq -\left(\frac{\alpha}{2\alpha-1}\right) (3+\epsilon) \left(\frac{w_\ell^2}{w_a}\right)$.

Proof. To upper bound $\frac{d\Phi_a}{dt}$, note that each job $j_i \in L$ is being processed at the rate of $s_a \cdot \frac{w(j_i)}{w_a}$ (we only consider the change due to WRR), and thus x_i is changing at the rate of $-s_a \cdot \frac{w(j_i)}{w_a}$. To ease discussion, let $y_i = \sum_{k=1}^i w(j_k)$. Note that $y_0 = 0$, $y_\ell = w_\ell$, and for any $1 \leq i \leq \ell$, $y_i - y_{i-1} = w(j_i)$.

First, consider $w_\ell \leq T^\alpha$. In this case, for each job $j_i \in L$, $c_i = y_i^{1-1/\alpha}$.

$$\begin{aligned} \frac{d\Phi_a}{dt} &= \sum_{i=1}^{\ell} y_i^{1-1/\alpha} \cdot \left(-s_a \cdot \frac{w(j_i)}{w_a}\right) \\ &= -\frac{s_a}{w_a} \sum_{i=1}^{\ell} y_i^{1-1/\alpha} \cdot (y_i - y_{i-1}) \\ &\quad (\text{since } y_i - y_{i-1} = w(j_i)) \\ &\leq -\frac{s_a}{w_a} \sum_{i=1}^{\ell} \int_{y_{i-1}}^{y_i} x^{1-1/\alpha} dx \\ &\quad (\text{since } x^{1-1/\alpha} \text{ is monotonically increasing}) \\ &\leq -\frac{s_a}{w_a} \int_0^{y_\ell} x^{1-1/\alpha} dx \\ &= -\frac{s_a}{w_a} \left(\frac{y_\ell^{2-1/\alpha}}{2-1/\alpha}\right) \\ &= -\left(\frac{\alpha}{2\alpha-1}\right) w_\ell^{2-1/\alpha} \left(\frac{s_a}{w_a}\right) \end{aligned}$$

Now, consider $w_\ell > T^\alpha$. In this case, $w_a \geq w_\ell > T^\alpha$ and hence $s_a = (3+\epsilon) \cdot \min(w_a^{1/\alpha}, T) = (3+\epsilon)T$. Note that $y_\ell = w_\ell > T^\alpha$. We let $g < \ell$ be the largest integer such that $y_g \leq T^\alpha$. Then

$$\begin{aligned} \frac{d\Phi_a}{dt} &= \sum_{i=1}^{\ell} c_i \cdot \left(-s_a \cdot \frac{w(j_i)}{w_a}\right) \\ &= -\left(\sum_{i=1}^g w(j_i) y_i^{1-1/\alpha} + \sum_{i=g+1}^{\ell} \frac{2\alpha}{2\alpha-1} \frac{w(j_i) y_i}{T}\right) \\ &\quad \cdot \left(\frac{s_a}{w_a}\right) \\ &\leq -\left(\int_0^{y_g} x^{1-1/\alpha} dx + \frac{2\alpha}{(2\alpha-1)T} \int_{y_g}^{y_\ell} x dx\right) \cdot \left(\frac{s_a}{w_a}\right) \\ &= -\left(\left(\frac{\alpha}{2\alpha-1}\right) y_g^{2-1/\alpha} + \frac{\alpha}{(2\alpha-1)T} (y_\ell^2 - y_g^2)\right) \cdot \left(\frac{s_a}{w_a}\right) \\ &\leq -\left(\frac{\alpha}{(2\alpha-1)T} (y_g^2 + y_\ell^2 - y_g^2)\right) \cdot \left(\frac{(3+\epsilon)T}{w_a}\right) \\ &\quad (\text{since } y_g^{1/\alpha} \leq T \text{ and } s_a = (3+\epsilon)T) \\ &= -\left(\frac{\alpha}{2\alpha-1}\right) (3+\epsilon) \left(\frac{y_\ell^2}{w_a}\right) \\ &= -\left(\frac{\alpha}{2\alpha-1}\right) (3+\epsilon) \left(\frac{w_\ell^2}{w_a}\right). \end{aligned}$$

\square

We are ready to show the following running condition, which together with Conditions (i) and (ii), implies Theorem 6.

Lemma 9. Assume that $\gamma = (2 - \frac{1}{\alpha})(1 + (3 + \epsilon)^\alpha)$. At any time when Φ does not have discrete change, $\frac{d\hat{G}_a}{dt} + \gamma \frac{d\Phi}{dt} \leq c \cdot \frac{d\hat{G}_o}{dt}$, where $c = (\frac{18}{\epsilon} + 4)(1 + (3 + \epsilon)^\alpha)$.

Proof. The analysis is divided into three cases depending on whether $w_a > T^\alpha$ and whether $w_\ell > T^\alpha$. In each case, we further divide the analysis depending on whether $w_\ell > (1 - \beta)w_a$, where $\beta = \epsilon/(6 + 2\epsilon)$. It is useful to note that $(3 + \epsilon)(1 - \beta)^2 \geq 3$.

Case 1: $w_a \leq T^\alpha$. In this case, $s_a = (3 + \epsilon) \cdot \min(w_a^{1/\alpha}, T) = (3 + \epsilon)w_a^{1/\alpha}$. Since $w_\ell \leq w_a$, we also have $w_\ell \leq T^\alpha$.

If $w_\ell > (1 - \beta)w_a$, then by Lemmas 7 and 8,

$$\begin{aligned} \frac{d\hat{G}_a}{dt} + \gamma \frac{d\Phi}{dt} &\leq (w_a + s_a^\alpha) + \frac{\gamma}{\alpha} s_o^\alpha + \gamma \cdot \left(\frac{\alpha-1}{\alpha}\right) w_\ell \\ &\quad - \gamma \cdot \left(\frac{\alpha}{2\alpha-1}\right) w_\ell^{2-1/\alpha} (3+\epsilon) w_a^{1/\alpha-1} \\ &\leq (1 + (3+\epsilon)^\alpha) w_a + \frac{\gamma}{\alpha} s_o^\alpha + \gamma \cdot \left(\frac{\alpha-1}{\alpha}\right) w_a \\ &\quad - \gamma \cdot \left(\frac{\alpha}{2\alpha-1}\right) ((1-\beta)w_a)^{2-1/\alpha} (3+\epsilon) w_a^{1/\alpha-1} \\ &\leq \frac{\gamma}{\alpha} s_o^\alpha + ((1 + (3+\epsilon)^\alpha) + \gamma \cdot \left(\frac{\alpha-1}{\alpha}\right)) \\ &\quad - \gamma \cdot \left(\frac{\alpha}{2\alpha-1}\right) (1-\beta)^{2-1/\alpha} (3+\epsilon) w_a. \end{aligned}$$

Choosing $\gamma = (2 - \frac{1}{\alpha})(1 + (3 + \epsilon)^\alpha)$, we have $(1 + (3 + \epsilon)^\alpha) = \gamma \cdot \left(\frac{\alpha}{2\alpha-1}\right)$. By the definition of β and the fact that $0 < \beta < 1$, we have $(1 - \beta)^{2-1/\alpha} (3 +$

$\epsilon) > (1 - \beta)^2(3 + \epsilon) \geq 3$. Furthermore, observe that $\frac{2\alpha}{2\alpha-1} > 1 > \frac{\alpha-1}{\alpha}$. Therefore,

$$\begin{aligned} & 1 + (3 + \epsilon)^\alpha + \gamma \left(\frac{\alpha-1}{\alpha} \right) - \gamma \left(\frac{\alpha}{2\alpha-1} \right) (1 - \beta)^{2-\frac{1}{\alpha}} (3 + \epsilon) \\ & \leq \gamma \left(\frac{\alpha}{2\alpha-1} \right) (1 + 2 - 3) \\ & = 0 \end{aligned}$$

Since $\frac{\gamma}{\alpha} = \frac{1}{\alpha} \cdot (2 - \frac{1}{\alpha})(1 + (3 + \epsilon)^\alpha) \leq 2(1 + (3 + \epsilon)^\alpha)$ and $c = (\frac{18}{\epsilon} + 4)(1 + (3 + \epsilon)^\alpha)$, it follows that $c \geq \frac{\gamma}{\alpha}$ and thus $\frac{d\hat{G}_a}{dt} + \gamma \frac{d\Phi}{dt} \leq \frac{\gamma}{\alpha} s_o^\alpha \leq c \cdot \frac{d\hat{G}_o}{dt}$.

If $w_\ell \leq (1 - \beta)w_a$, we simply adapt the bound of $\frac{d\Phi_a}{dt}$ in Lemma 8 as $\frac{d\Phi_a}{dt} \leq 0$. Since any active job in WRR that is not lagging must also be an active job in OPT, $w_o \geq w_a - w_\ell \geq w_a - (1 - \beta)w_a \geq \beta w_a$. Choosing $\gamma = (2 - \frac{1}{\alpha})(1 + (3 + \epsilon)^\alpha)$ and recalling that $\beta = \epsilon/(6 + 2\epsilon)$, by Lemma 7,

$$\begin{aligned} \frac{d\hat{G}_a}{dt} + \gamma \frac{d\Phi}{dt} & \leq (1 + (3 + \epsilon)^\alpha)w_a + \frac{\gamma}{\alpha} s_o^\alpha + \gamma \cdot \left(\frac{\alpha-1}{\alpha} \right) w_\ell \\ & \leq \frac{\gamma}{\alpha} s_o^\alpha + ((1 + (3 + \epsilon)^\alpha) + \gamma) w_a \\ & \leq \frac{\gamma}{\alpha} s_o^\alpha + ((1 + (3 + \epsilon)^\alpha) \\ & \quad + (2 - \frac{1}{\alpha})(1 + (3 + \epsilon)^\alpha))w_a \\ & \leq \frac{\gamma}{\alpha} s_o^\alpha + 3(1 + (3 + \epsilon)^\alpha) \frac{w_o}{\beta} \\ & = \frac{\gamma}{\alpha} s_o^\alpha + (\frac{18}{\epsilon} + 4)(1 + (3 + \epsilon)^\alpha)w_o \\ & \leq c \cdot \frac{d\hat{G}_o}{dt} . \end{aligned}$$

Case 2: $w_a > T^\alpha$ and $w_\ell \leq T^\alpha$. In this case, $s_a = (3 + \epsilon) \cdot \min(w_a^{1/\alpha}, T) = (3 + \epsilon)T$. Since $w_\ell \leq T^\alpha$ and w_o is at least the total weight of non-lagging jobs in WRR (because if a job is non-lagging in WRR, it must be active in OPT), we have $w_a \leq w_o + w_\ell \leq w_o + T^\alpha$.

If $w_\ell > (1 - \beta)w_a$, then by Lemmas 7 and 8,

$$\begin{aligned} \frac{d\hat{G}_a}{dt} + \gamma \frac{d\Phi}{dt} & \leq (w_a + (3 + \epsilon)^\alpha T^\alpha) + \frac{\gamma}{\alpha} s_o^\alpha + \gamma \cdot \left(\frac{\alpha-1}{\alpha} \right) w_\ell \\ & \quad - \gamma \cdot \left(\frac{\alpha}{2\alpha-1} \right) w_\ell^{2-1/\alpha} \frac{(3+\epsilon)T}{w_a} \\ & \leq ((w_o + T^\alpha) + (3 + \epsilon)^\alpha T^\alpha) \\ & \quad + \frac{\gamma}{\alpha} s_o^\alpha + \gamma \cdot \left(\frac{\alpha-1}{\alpha} \right) T^\alpha \\ & \quad - \gamma \cdot \left(\frac{\alpha}{2\alpha-1} \right) ((1 - \beta)w_a)^{2-1/\alpha} \frac{(3+\epsilon)T}{w_a} \\ & \leq w_o + \frac{\gamma}{\alpha} s_o^\alpha + ((1 + (3 + \epsilon)^\alpha) + \gamma \left(\frac{\alpha-1}{\alpha} \right)) \\ & \quad - \gamma \left(\frac{\alpha}{2\alpha-1} \right) (1 - \beta)^{2-1/\alpha} (3 + \epsilon) T^\alpha \\ & \quad (\text{since } w_a > T^\alpha). \end{aligned}$$

Choosing $\gamma = (2 - \frac{1}{\alpha})(1 + (3 + \epsilon)^\alpha)$ and then using the same argument as in Case 1, we can argue that the coefficient of T^α is non-positive. Therefore, $\frac{d\hat{G}_a}{dt} + \gamma \frac{d\Phi}{dt} \leq w_o + \frac{\gamma}{\alpha} s_o^\alpha \leq c \cdot \frac{d\hat{G}_o}{dt}$.

If $w_\ell \leq (1 - \beta)w_a$, we simply use the bound $\frac{d\Phi_a}{dt} \leq 0$. Note that $w_a > T^\alpha$ and $w_a \geq w_\ell$. By Lemma 7, we have

$$\begin{aligned} \frac{d\hat{G}_a}{dt} + \gamma \frac{d\Phi}{dt} & \leq (w_a + (3 + \epsilon)^\alpha T^\alpha) + \frac{\gamma}{\alpha} s_o^\alpha + \gamma \cdot \left(\frac{\alpha-1}{\alpha} \right) w_\ell \\ & \leq \frac{\gamma}{\alpha} s_o^\alpha + ((1 + (3 + \epsilon)^\alpha) + \gamma) w_a . \end{aligned}$$

Choosing $\gamma = (2 - \frac{1}{\alpha})(1 + (3 + \epsilon)^\alpha)$ and then using the same argument as in Case 1, we can show that $\frac{d\hat{G}_a}{dt} + \gamma \frac{d\Phi}{dt} \leq c \cdot \frac{d\hat{G}_o}{dt}$.

Case 3: $w_a > T^\alpha$ and $w_\ell > T^\alpha$. In this case, $s_a = (3 + \epsilon)T \leq (3 + \epsilon)w_a^{1/\alpha}$.

If $w_\ell > (1 - \beta)w_a$, then by Lemmas 7 and 8,

$$\begin{aligned} \frac{d\hat{G}_a}{dt} + \gamma \frac{d\Phi}{dt} & \leq (1 + (3 + \epsilon)^\alpha)w_a + \gamma \cdot \left(\frac{2\alpha}{2\alpha-1} \right) w_\ell \\ & \quad - \gamma \cdot \left(\frac{\alpha}{2\alpha-1} \right) (3 + \epsilon) \frac{w_\ell^2}{w_a} \\ & \leq (1 + (3 + \epsilon)^\alpha)w_a + \gamma \cdot \left(\frac{2\alpha}{2\alpha-1} \right) w_a \\ & \quad - \gamma \cdot \left(\frac{\alpha}{2\alpha-1} \right) (3 + \epsilon) \frac{((1-\beta)w_a)^2}{w_a} \end{aligned}$$

Choosing $\gamma = (2 - \frac{1}{\alpha})(1 + (3 + \epsilon)^\alpha)$, we have $(1 + (3 + \epsilon)^\alpha) = \gamma \cdot \left(\frac{2\alpha}{2\alpha-1} \right)$. By the definition of β , $(1 - \beta)^2(3 + \epsilon) \geq 3$. Therefore, $\frac{d\hat{G}_a}{dt} + \gamma \frac{d\Phi}{dt} \leq \left(\gamma \cdot \left(\frac{\alpha}{2\alpha-1} \right) + \gamma \cdot \left(\frac{2\alpha}{2\alpha-1} \right) - \gamma \cdot \left(\frac{3\alpha}{2\alpha-1} \right) \right) w_a = 0 \leq c \cdot \frac{d\hat{G}_o}{dt}$.

If $w_\ell \leq (1 - \beta)w_a$, we simply use the bound $\frac{d\Phi_a}{dt} \leq 0$. Recall that in this case, $w_o \geq \beta w_a$. Choosing $\gamma = (2 - \frac{1}{\alpha})(1 + (3 + \epsilon)^\alpha)$ and recalling that $\beta = \epsilon/(6 + 2\epsilon)$, by Lemma 7,

$$\begin{aligned} \frac{d\hat{G}_a}{dt} + \gamma \frac{d\Phi}{dt} & \leq (1 + (3 + \epsilon)^\alpha)w_a + \gamma \cdot \left(\frac{2\alpha}{2\alpha-1} \right) w_\ell \\ & \leq 3(1 + (3 + \epsilon)^\alpha)w_a \\ & \leq 3(1 + (3 + \epsilon)^\alpha) \frac{w_o}{\beta} \\ & = (\frac{18}{\epsilon} + 4)(1 + (3 + \epsilon)^\alpha)w_o \\ & \leq c \cdot \frac{d\hat{G}_o}{dt} . \end{aligned}$$

In conclusion, the running condition is satisfied in all the three cases. \square

5 Conclusion

In this paper we have given two non-clairvoyant scheduling algorithms for minimizing flow plus energy. The first algorithm (LAPS) is $8\alpha^2$ -competitive for (unweighted) flow plus energy, when using a processor with maximum speed $\frac{\alpha}{\alpha-1}T$. The second algorithm (WRR) is $O(3^\alpha/\epsilon)$ -competitive for weighted flow plus energy, when using a processor with maximum speed $(3 + \epsilon)T$, where $0 < \epsilon \leq \frac{3}{\alpha}$. We believe that LAPS can be generalized to minimize weighted flow plus energy, and the competitive ratio would remain $O(\alpha^2)$.

References

- [1] S. Albers and H. Fujiwara. Energy-efficient algorithms for flow time minimization. *ACM Transactions on Algorithms*, 3(4):49, 2007.
- [2] D. M. Brooks, P. Bose, S. E. Schuster, H. Jacobson, P. N. Kudva, A. Buyuktosunoglu, J. D.

- Wellman, V. Zyuban, M. Gupta, and P. W. Cook. Power-aware microarchitecture: Design and modeling challenges for next-generation microprocessors. *IEEE Micro*, 20(6):26–44, 2000.
- [3] N. Bansal and H. L. Chan. Weighted flow time does not admit $O(1)$ -competitive algorithms. In *Proc. SODA*, pages 1238–1244, 2009.
- [4] N. Bansal, H. L. Chan, T. W. Lam, and L. K. Lee. Scheduling for speed bounded processors. In *Proc. ICALP*, pages 409–420, 2008.
- [5] N. Bansal, H. L. Chan, and K. Pruhs. Speed scaling with an arbitrary power function. In *Proc. SODA*, pages 693–701, 2009.
- [6] N. Bansal, H. L. Chan, K. Pruhs, and D. Katz. Improved bounds for speed scaling in devices obeying the cube-root rule. In *Proc. ICALP*, pages 144–155, 2009.
- [7] N. Bansal, T. Kimbrel, and K. Pruhs. Speed scaling to manage energy and temperature. *Journal of the ACM*, 54(1):3, 2007.
- [8] N. Bansal, K. Pruhs, and C. Stein. Speed scaling for weighted flow time. In *Proc. SODA*, pages 805–813, 2007.
- [9] H. L. Chan, W. T. Chan, T. W. Lam, L. K. Lee, K. S. Mak, and P. W. H. Wong. Energy efficient online deadline scheduling. In *Proc. SODA*, pages 795–804, 2007.
- [10] H. L. Chan, J. Edmonds, T. W. Lam, L. K. Lee, A. Marchetti-Spaccamela, and K. Pruhs. Non-clairvoyant speed scaling for flow and energy. In *Proc. STACS*, pages 255–264, 2009.
- [11] H. L. Chan, J. Edmonds, and K. Pruhs. Speed scaling of processes with arbitrary speedup curves on a multiprocessor. In *Proc. SPAA*, pages 1–10, 2009.
- [12] G. Greiner, T. Nonner, and A. Souza. The bell is ringing in speed-scaled multiprocessor scheduling. In *Proc. SPAA*, pages 11–18, 2009.
- [13] S. Irani and K. Pruhs. Algorithmic problems in power management. *ACM SIGACT News*, 32(2):63, 2005.
- [14] B. Kalyanasundaram and K. Pruhs. Minimizing flow time nonclairvoyantly. *Journal of the ACM*, 50(4):551–567, 2003.
- [15] T. W. Lam, L. K. Lee, H. F. Ting, I. K. K. To, and P. W. H. Wong. Sleep with guilt and work faster to minimize flow plus energy. In *Proc. ICALP*, pages 665–676, 2009.
- [16] T. W. Lam, L. K. Lee, I. K. K. To, and P. W. H. Wong. Competitive non-migratory scheduling for flow time and energy. In *Proc. SPAA*, pages 256–264, 2008.
- [17] T. W. Lam, L. K. Lee, I. K. K. To, and P. W. H. Wong. Speed scaling functions for flow time scheduling based on active job count. In *Proc. ESA*, pages 647–659, 2008.
- [18] R. Motwani, S. Phillips, and E. Torng. Nonclairvoyant scheduling. *Theoretical Computer Science*, 130(1):17–47, 1994.
- [19] T. Mudge. Power: A first-class architectural design constraint. *IEEE Computer*, 34(4):52–58, 2001.
- [20] J. M. Steele. *The Cauchy-Schwarz master class: An introduction to the art of mathematical inequalities*, page 136. Cambridge University Press, 2004.
- [21] F. Yao, A. Demers, and S. Shenker. A scheduling model for reduced CPU energy. In *Proc. FOCS*, pages 374–382, 1995.