

# Multiprocessor Speed Scaling for Jobs with Arbitrary Sizes and Deadlines<sup>\*</sup>

Paul C. Bell<sup>1</sup> and Prudence W.H. Wong<sup>2</sup>

<sup>1</sup> Department of Computer Science, Loughborough University.  
P.Bell@lboro.ac.uk

<sup>2</sup> Department of Computer Science, University of Liverpool.  
pwong@liverpool.ac.uk

**Abstract.** In this paper we study energy efficient deadline scheduling on multiprocessors in which the processors consumes power at a rate of  $s^\alpha$  when running at speed  $s$ , where  $\alpha \geq 2$ . The problem is to dispatch jobs to processors and determine the speed and jobs to run for each processor so as to complete all jobs by their deadlines using the minimum energy. The problem has been well studied for the single processor case. For the multiprocessor setting, constant competitive online algorithms for special cases of unit size jobs or arbitrary size jobs with agreeable deadlines have been proposed [4]. A randomized algorithm has been proposed for jobs of arbitrary sizes and arbitrary deadlines [13]. We propose a deterministic online algorithm for the general setting and show that it is  $O(\log^\alpha P)$ -competitive, where  $P$  is the ratio of the maximum and minimum job size.

## 1 Introduction

**Energy efficient deadline scheduling.** Energy consumption has become an important concern in the design of modern processors, not only for battery-operated mobile devices with single processors but also for server farms or laptops with multi-core processors. A popular technology to reduce energy usage is *dynamic speed scaling* (see e.g., [4, 7, 8, 21]) where the processor can vary its speed dynamically. The power consumption is modelled by  $s^\alpha$  when the processor runs at speed  $s$ , where  $\alpha$  is typically 2 or 3 [11, 20]. Running a job slower saves energy, yet it takes longer to finish the job. The challenge arises from the conflicting objectives of providing good “quality of service” (QoS) and conserving energy. Deadline feasibility is a common QoS measure for job scheduling. Jobs with arbitrary sizes and deadlines arrive at unpredictable times and they are to be run on some processor. Preemption is allowed with no penalty.

The study of speed scaling was initiated by Yao et al. [21]. They studied deadline scheduling on a single processor in which jobs with arbitrary sizes and deadlines arrive online and the aim is to finish all jobs by their deadlines using the minimum amount of energy. The decision at any time is to determine which job

---

<sup>\*</sup> This work is partially supported by EPSRC Grant EP/E028276/1.

to run and at what speed. They gave an optimal offline algorithm and a simple online algorithm AVR which is  $2^{\alpha-1}\alpha^\alpha$ -competitive and they also proposed an online algorithm OA. Bansal, Kimbrel and Pruhs [8] later showed that OA is  $\alpha^\alpha$ -competitive. They also gave a  $2(\alpha/(\alpha-1))^\alpha e^\alpha$ -competitive algorithm, which is called the BKP algorithm and is better than OA when  $\alpha > 5$ . The result is further improved to  $4^\alpha/(2\sqrt{e\alpha})$ -competitive by the qOA algorithm [7].

The problem of energy efficient scheduling has also been studied for other QoS measures. The problem of minimizing flow time and energy has attracted a lot of attention [3, 5, 6, 9, 13, 18, 19]. Energy efficient scheduling has also been extended to the setting with sleep states [14, 16, 17]. The literature also contains results on other aspects of energy efficient scheduling, see [1, 2, 15].

**Energy efficient multiprocessor scheduling.** The problem of energy efficient deadline scheduling becomes NP-hard in the multiprocessor setting, even when all the jobs have the same arrival times and deadlines. In the multiprocessor setting, in addition to determining processor speeds, a job dispatching algorithm is required to assign jobs to processors. Albers et al. [4] have extended the study to the multiprocessor setting and they study the special cases of unit-size jobs or jobs with agreeable deadlines (jobs arriving earlier have earlier deadlines). If jobs have unit-size and agreeable deadlines, Round Robin (RR) is optimal. For the case of unit-sized jobs with arbitrary deadlines or arbitrary-sized jobs with agreeable deadlines, they gave an  $\alpha^\alpha 2^{4\alpha}$ -competitive algorithm. Their algorithm, called Classified Round Robin (CRR) first classifies jobs according to the density of the job (the ratio of the job size to the duration between arrival and deadline), and then schedules jobs in each class independently using RR. All jobs (of different classes) dispatched on a processor are run at a speed determined by AVR. The case for jobs of arbitrary sizes and arbitrary deadlines is left as an open question.

Recently, Greiner, Nonner and Souza [13] have shown that any  $\beta$ -competitive algorithm for a single processor yields a randomized  $\beta B_\alpha$ -competitive algorithm, where  $B_\alpha$  is the  $\alpha$ th Bell number [10] and this result holds for jobs of arbitrary size and arbitrary deadlines. This means that the existing algorithms [7, 8, 21] for single processors lead to randomized online algorithms in the multiprocessor setting. Yet it is still an open question to have a competitive deterministic algorithm for the general case of jobs with arbitrary sizes and arbitrary deadlines.

**Our contribution.** In this paper we study the generalized problems in the multiprocessor setting where jobs have arbitrary sizes and arbitrary deadlines and give a deterministic online algorithm. We first show that the Classified Round Robin algorithm (CRR) [4] does not scale well when jobs have arbitrary sizes and deadlines. The competitive ratio is at least  $m^{\alpha-1}$ , where  $m$  is the number of processors. We then consider a natural extension of CRR and propose a non-migratory deterministic job dispatching algorithm, called *Dual-Classified Round Robin* (DCRR), which classifies jobs in terms of both density and sizes. We show that DCRR coupled with AVR is  $2^{4\alpha}(\log^\alpha P + \alpha^\alpha 2^{\alpha-1})$ -competitive where  $P$  is the ratio between the maximum and minimum job size. Note that

the competitive ratio is independent of  $m$  and holds even against an optimal migratory offline algorithm.

Roughly speaking, to analyze the performance of DCRR, we round the density and size of a job to the boundaries that define the classes, and show that the performance on the general set is no more than a constant factor of that on such a “nice” job set. This idea is similar to the proof in [4], which rounds only the density of the jobs. We further show that for a nice job set, the classification of DCRR means that the jobs in the same class satisfy the property of agreeable deadlines, making the analysis easier. We are then able to show that the competitive ratio of DCRR depends on the number of classes, which is related to  $\log P$ .

**Organization of the paper.** The rest of the paper is organized as follows. In Section 2, we define the problem and give some preliminary results. In Section 3, we review an existing algorithm CRR and show that it does not work well for jobs of arbitrary sizes and deadlines. In Section 4, we describe and analyze our algorithm DCRR. Finally, we conclude in Section 5.

## 2 Preliminaries

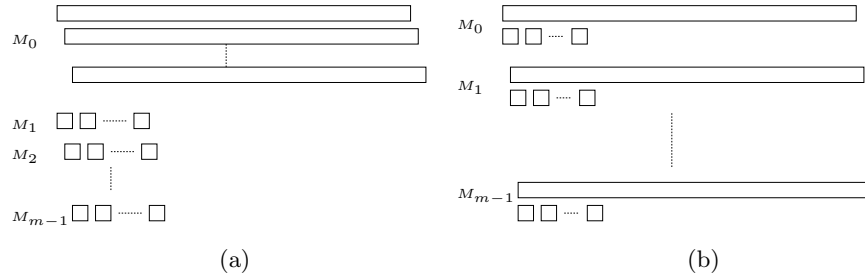
We are to schedule a set of jobs onto  $m$  processors  $M_0, M_1, \dots, M_{m-1}$ . Preemption is allowed without penalty. The speed of each processor can be varied. When running at speed  $s$ , a processor processes  $s$  units of work and consumes  $s^\alpha$  units of energy in each time unit, where  $\alpha \geq 2$ .

We denote the release time, deadline and size of a job  $j$  as  $r(j)$ ,  $d(j)$ , and  $w(j)$ , respectively. The *span* of job  $j$  is  $\text{span}(j) = d(j) - r(j)$  and the *density*  $\text{den}(j) = \frac{w(j)}{d(j) - r(j)}$ . A job  $j$  is called *active* at time  $t$  if  $r(j) \leq t \leq d(j)$ .

The problem is to dispatch the jobs to processors, and for each processor, to determine which job and at what speed to run at any time. The objective is to complete all jobs by their deadlines using the minimum energy.

Consider any job set  $\mathcal{J}$ . For any algorithm  $A$ , we overload the symbol  $A(\mathcal{J})$  to mean both the schedule of  $A$  on  $\mathcal{J}$  and the energy required by the schedule. Let  $\text{OPT}_1$  and  $\text{OPT}_m$  denote the optimal schedule on a single processor and  $m$  processors, respectively. In [4], it has been shown that  $\text{OPT}_1(\mathcal{J})/m^{\alpha-1} \leq \text{OPT}_m(\mathcal{J})$ . We further lower bound the value  $\text{OPT}_m(\mathcal{J})$ . At any time  $t$ , the speed of AVR on a processor is the sum of the densities of all active jobs at  $t$  scheduled on this processor. It has been shown in [21] that  $\text{AVR}_1(\mathcal{J}) \leq \alpha^\alpha 2^{\alpha-1} \text{OPT}_1(\mathcal{J})$ , implying  $\text{AVR}_1(\mathcal{J}) \leq \alpha^\alpha 2^{\alpha-1} m^{\alpha-1} \text{OPT}_m(\mathcal{J})$ . Let  $\text{MIN}(\mathcal{J})$  be the minimum energy to run each job of  $\mathcal{J}$  independently of other jobs, i.e.,  $\text{MIN}(\mathcal{J}) = \sum_{j \in \mathcal{J}} (\text{den}(j))^\alpha \text{span}(j)$ . Then, we have  $\text{MIN}(\mathcal{J}) \leq \text{OPT}_m(\mathcal{J})$ . We summarize these bounds on  $\text{OPT}_m(\mathcal{J})$  in the following lemma.

**Lemma 1 ([4]).** *Consider any job set  $\mathcal{J}$ . (a)  $\text{OPT}_m(\mathcal{J}) \geq \text{OPT}_1(\mathcal{J})/m^{\alpha-1}$ . (b) (i)  $\text{MIN}(\mathcal{J}) \leq \text{OPT}_m(\mathcal{J})$ ; (ii)  $\text{AVR}_1(\mathcal{J}) \leq \alpha^\alpha 2^{\alpha-1} m^{\alpha-1} \text{OPT}_m(\mathcal{J})$ .*



**Fig. 1.** In the adversary, all jobs have density 1. The span and size of the  $m$  large jobs is  $k$  and the  $m^2 - m$  small jobs is  $\epsilon$ . (a) CRR schedules all the large jobs to processor  $M_0$  and  $m$  small jobs to each of  $M_1, \dots, M_{m-1}$ . (b) The optimal schedule dispatches one large job and  $m - 1$  small jobs to each processor.

### 3 Classified Round Robin (CRR)

In this section, we review the algorithm CRR which is  $\alpha^\alpha 2^{4\alpha}$ -competitive for the special case in which jobs are of unit size, or jobs are of arbitrary sizes but agreeable deadlines [4]. We show that CRR is no longer constant competitive when the jobs have arbitrary sizes and arbitrary deadlines.

Let  $\Delta$  be the maximum density of the jobs in  $\mathcal{J}$ . CRR classifies jobs with density  $\Delta$  into density-class-0, and jobs with density in  $[\Delta/2^k, \Delta/2^{k-1})$  into density-class- $k$ , for some positive integer  $k$ . Jobs within each class are dispatched to processors by round-robin independently. For each processor, the speed is the sum of the densities of the unfinished jobs dispatched to that processor (i.e., AVR) and the processor processes these jobs by splitting the speed equally among them.

The following theorem shows a lower bound for CRR when jobs are of arbitrary sizes and deadlines. Figure 1 shows the CRR schedule and the optimal schedule for the adversary.

**Theorem 1.** *For arbitrary size jobs with arbitrary deadlines, CRR has a competitive ratio of at least  $m^{\alpha-1}$ .*

*Proof.* Let  $\epsilon > 0$  be a small positive value and  $k > 0$  be an arbitrary large value. Given  $m$  processors, define a job set  $\mathcal{J}$  of  $m^2$  jobs such that for any  $1 \leq i \leq m^2$ , the release time of job  $j_i$  is  $i\epsilon$ . For all jobs  $j_i$  with  $i \bmod m \neq 0$ , we set the span of the job to be  $\epsilon$ . For all jobs  $j_i$  with  $i \bmod m = 0$ , we set the span of the jobs to be  $k$ . We further set the sizes of all jobs to be the same as their span, in other words, all jobs have density 1.

Algorithm CRR classifies all  $m^2$  jobs into the same class  $C_0$  since they have the same density and dispatches jobs according to round robin by their release time. Thus the first processor receives the  $m$  jobs of large span  $k$  and large size

$k$ . The energy used by the first processor is therefore  $km^\alpha$  as  $\epsilon \rightarrow 0$  and the energy of the remaining processors approaches 0.

On the other hand, we can dispatch one large span job and  $m - 1$  small span jobs to each processor. As  $\epsilon$  tends to 0, the energy used by each processor is  $k$  and the total energy of the schedule is  $km$ . Therefore, the competitive ratio of CRR is at least  $m^{\alpha-1}$ .  $\square$

We note that even if we classify jobs according to their sizes, such a classification plus round robin still does not perform well. We give a similar adversary with  $m^2$  jobs of the same size,  $m$  of them having a small span (thus large density) and the rest with very large span. One job of small span arrives followed by  $m - 1$  large span jobs and this repeats for  $m$  times. Then CRR assigns all the small span jobs to the same processor, dominating the energy used by the algorithm. The optimal offline algorithm can dispatch one small span job to each processor, distributing the energy used much better and thus the same lower bound can be obtained.

## 4 Dual-Classified Round Robin (DCRR)

### 4.1 The Algorithm

We now describe our algorithm DCRR (Dual-Classified Round Robin). In addition to classifying jobs into density classes, DCRR also classifies jobs according to sizes. Let  $\Gamma$  be the maximum job size of a job set  $\mathcal{J}$ . Jobs with size in  $(\Gamma/2^{h+1}, \Gamma/2^h]$  are classified into size-class- $h$ , for some integer  $h \geq 0$  (note the difference from the definition of density-classes). We then define the set  $C_{k,h}$  to be the set of jobs in density-class- $k$  and size-class- $h$ . For simplicity, we assume that  $\Delta$  and  $\Gamma$  are known in advance.<sup>3</sup> With the definition of  $C_{k,h}$ , DCRR dispatches jobs in the same  $C_{k,h}$  in a round robin manner, independent of other classes. Then all jobs (of different classes) dispatched to the same processor are run using a speed determined by AVR (see Algorithm 1.1).

### 4.2 Framework of the Analysis and Nice Job Sets.

To analyze the performance of DCRR, we transform job set  $\mathcal{J}$  to a *nice* job set  $\mathcal{J}^*$  (to be defined) and show that such a transformation only increases the energy usage modestly. Furthermore, we show that for a nice job set  $\mathcal{J}^*$ , we can bound  $\text{DCRR}(\mathcal{J}^*)$  by  $\text{OPT}_m(\mathcal{J}^*)$  and in turn by  $\text{OPT}_m(\mathcal{J})$ . Then we can establish the competitive ratio of DCRR.

A job set  $\mathcal{J}^*$  is said to be a *nice job set* if every job  $j^*$  in  $\mathcal{J}^*$  satisfies the following properties.

<sup>3</sup> If  $\Delta$  and  $\Gamma$  are not known in advance, the class definition could be modified slightly. Specifically, the first job which arrives will define the initial density and size classes  $\Delta'$  and  $\Gamma'$ . New jobs may have larger sizes or density than these  $\Delta'$  and  $\Gamma'$  and thus we may have classes with a negative index, but the analysis can be seen to still hold and increasing the competitive ratio by at most a factor of 2, see [4] for further details.

**Algorithm 1.1.** Algorithm DCRR

Let  $\Delta$  and  $\Gamma$  be (respectively) the maximum density and maximum size of all jobs.

**Classification:** A job is classified into  $C_{k,h}$  if its density is in  $[\Delta/2^k, \Delta/2^{k-1})$  and its size is in  $(\Gamma/2^{h+1}, \Gamma/2^h]$ .

**Job dispatching:** Jobs of the same class  $C_{k,h}$  are dispatched (upon their arrival) to the  $m$  processors using a round-robin strategy, i.e., the  $i$ -th job in  $C_{k,h}$  is dispatched to processor- $(i \bmod m)$ , and different classes are handled independently.

**Speed running:** The speed of each processor is determined by AVR on the jobs dispatched to that processor and the speed is split equally among these jobs (note that this gives a feasible schedule).

- The density  $\text{den}(j^*) = \Delta/2^k$ , for some positive integer  $k$ .
- The size  $w(j^*) = \Gamma/2^h$ , for some positive integer  $h$ .

Given a job set  $\mathcal{J}$ , we transform each job  $j \in \mathcal{J}$  into a job  $j^*$  as follows. Suppose  $j$  is in class  $C_{k,h}$ .

- We set the release time of  $j^*$  to be the same as  $j$ , i.e.,  $r(j^*) = r(j)$ .
- We round up the size of  $j$  to the maximum in the class  $C_{k,h}$ , i.e.,  $w(j^*) = \Gamma/2^h$ . Then, we have  $w(j) \leq w(j^*) \leq 2w(j)$ .
- We round down the density of  $j$  to the minimum in the class  $C_{k,h}$ , i.e.,  $\text{den}(j^*) = \Delta/2^k$ . Then, we have  $\text{den}(j)/2 \leq \text{den}(j^*) \leq \text{den}(j)$ .
- Effectively, we set the deadline  $d(j^*) = r(j^*) + (\frac{\Gamma}{2^h} \cdot \frac{2^k}{\Delta})$ .

In other words, job densities only decrease and sizes only increase. The following lemma relates the optimal schedule for  $\mathcal{J}$  and  $\mathcal{J}^*$ , as well as the DCRR schedule for  $\mathcal{J}$  and  $\mathcal{J}^*$ . The implication of the lemma is that we can focus on analyzing the performance of DCRR on nice job set  $\mathcal{J}^*$ .

**Lemma 2.** *For any job set  $\mathcal{J}$  and its corresponding nice job set  $\mathcal{J}^*$ , we have (a)  $2^\alpha \text{OPT}_m(\mathcal{J}) \geq \text{OPT}_m(\mathcal{J}^*)$ ; (b)  $\text{DCRR}(\mathcal{J}) \leq 2^\alpha \text{DCRR}(\mathcal{J}^*)$ .*

*Proof.* (a) We construct from  $\text{OPT}_m(\mathcal{J})$  a feasible schedule  $S$  for  $\mathcal{J}^*$ , and show that this increases the energy slightly. The dispatching of  $S$  follows the dispatching of  $\text{OPT}_m(\mathcal{J})$ . For any processor, at any time  $t$ ,  $S$  runs at double the speed that  $\text{OPT}_m(\mathcal{J})$  does.  $S$  is feasible for  $\mathcal{J}^*$  because  $w(j^*) \leq 2w(j)$  and  $\text{span}(j) \leq \text{span}(j^*)$ , the latter implies that whenever  $j$  is run, it is within the span of  $j^*$ . Because of the double speed,  $S = 2^\alpha \text{OPT}_m(\mathcal{J})$ . As  $S$  is a feasible schedule for  $\mathcal{J}^*$ , we have  $S \geq \text{OPT}_m(\mathcal{J}^*)$ . Then the statement follows.

(b) First we notice that a job  $j$  and its corresponding  $j^*$  belong to the same class. The release time of  $j^*$  is also kept the same as  $j$ . Therefore,  $j^*$  will be dispatched to the same processor as  $j$ . In the schedule of  $\text{DCRR}(\mathcal{J}^*)$ , at any time when the job  $j^*$  is active, it contributes  $\text{den}(j^*)$  to the speed of that processor. If we consider a schedule  $S'$  that runs double the speed at any time and on any processor as  $\text{AVR}(\mathcal{J}^*)$  does, the job  $j^*$  contributes  $2 \times \text{den}(j^*)$  to the speed. As for energy usage,  $S' = 2^\alpha \text{DCRR}(\mathcal{J}^*)$ . On the other hand, in  $\text{DCRR}(\mathcal{J})$ , at

any time that  $j$  is active, it contributes  $\text{den}(j)$  to the speed of that processor. Since  $2 \text{den}(j^*) \geq \text{den}(j)$  and  $\text{span}(j^*) > \text{span}(j)$ , for any processor, the schedule  $S'$  runs at least the same speed as  $\text{DCRR}(\mathcal{J})$ , and probably higher. Therefore,  $S' \geq \text{DCRR}(\mathcal{J})$ , and the statement follows.  $\square$

### 4.3 Analysis of DCRR

With Lemma 2, the analysis of DCRR on a general job set  $\mathcal{J}$  can be done via the analysis of DCRR on  $\mathcal{J}^*$ . Recall that  $\text{MIN}(\mathcal{J})$  is the minimum energy to run each job of  $\mathcal{J}$  independently of other jobs, i.e.,  $\text{MIN}(\mathcal{J}) = \sum_{j \in \mathcal{J}} (\text{den}(j))^\alpha \text{span}(j)$ . First, we show in Lemma 3 a property about how DCRR dispatches jobs in a class to the  $m$  processors. Then, in Lemma 4, we relate the sum of energy usage of AVR on jobs DCRR dispatched to each machine with  $\text{MIN}(\mathcal{J})$  and  $\text{AVR}_1(\mathcal{J})$ . Finally, together with Lemma 2, we can then conclude in Theorem 2 the competitive ratio of DCRR.

The following is a modification to a lemma from [4]. Since all spans within a class are identical, they have agreeable deadlines and the same proof follows as is shown in [4].

**Lemma 3.** *For any time  $t$ , DCRR assigns to each processor at most  $\lceil C_{k,h}(t)/m \rceil$  jobs from  $\mathcal{J}^*$ , where  $C_{k,h}(t)$  is the set of jobs from  $C_{k,h}$  active at time  $t$ .*

Let  $\mathcal{J}_i^*$  be the subset of  $\mathcal{J}^*$  that is dispatched to processor  $i$  by DCRR. Then  $\text{DCRR}(\mathcal{J}^*) = \sum_{1 \leq i \leq m} \text{AVR}_1(\mathcal{J}_i^*)$ . We now relate  $\sum_{1 \leq i \leq m} \text{AVR}_1(\mathcal{J}_i^*)$  with  $\text{MIN}(\mathcal{J}^*)$  and  $\text{AVR}_1(\mathcal{J}^*)$ .

**Lemma 4.** *For any nice job set  $\mathcal{J}^*$ , the following inequality holds*

$$\sum_{1 \leq i \leq m} \text{AVR}_1(\mathcal{J}_i^*) \leq 2^{2\alpha} ((\log^\alpha P^*) \text{MIN}(\mathcal{J}^*) + \text{AVR}_1(\mathcal{J}^*)/m^{\alpha-1})$$

where  $P^* = \frac{\max \{w(j) | j \in \mathcal{J}^*\}}{\min \{w(j) | j \in \mathcal{J}^*\}}$ .

*Proof.* We adapt the proof of CRR in [4]. Let  $C_{k,h,i}(t)$  for  $1 \leq i \leq m$  be the set of jobs from class  $C_{k,h}$  assigned to processor  $i$  active at time  $t$  dispatched by DCRR. Let  $s_i(t)$  denote the speed of the average rate AVR algorithm on processor  $i$  at time  $t$ . Since the speed of AVR is the sum of densities of all active jobs at each time point, we see that:

$$s_i(t) = \sum_{k \geq 0} \sum_{h \geq 0} |C_{k,h,i}(t)| \frac{\Delta}{2^k}. \quad (1)$$

Running jobs according to the *Earliest Deadline First* policy yields a feasible schedule. Let  $s(t)$  denote the speed of AVR for the whole job set  $\mathcal{J}^*$ . Then  $s(t) = \sum_{k \geq 0} \sum_{h \geq 0} |C_{k,h}(t)| \Delta / 2^k$ .

Fix a time  $t \geq 0$  and a processor  $1 \leq i \leq m$ . Let  $K_1$  be the set of job class indices  $(k, h)$  such that  $|C_{k,h,i}(t)| = 1$  and  $K_2$  be the indices  $(k, h)$  such

that  $|C_{k,h,i}(t)| \geq 2$ . Define  $k_1 = \min\{k | (k, h) \in K_1\}$  for some  $h \geq 0$  and  $P^* = \frac{\max\{w(j) | j \in \mathcal{J}^*\}}{\min\{w(j) | j \in \mathcal{J}^*\}}$ . Using Equation (1) and Lemma 3, we see that

$$\begin{aligned}
s_i(t) &= \sum_{(k,h) \in K_1} \frac{\Delta}{2^k} + \sum_{(k,h) \in K_2} |C_{k,h,i}(t)| \frac{\Delta}{2^k} \\
&\leq (\log P^*) \frac{\Delta}{2^{k_1-1}} + \sum_{(k,h) \in K_2} \left\lceil \frac{|C_{k,h}(t)|}{m} \right\rceil \frac{\Delta}{2^k} && \text{(by Lemma 3)} \\
&\leq (\log P^*) \frac{\Delta}{2^{k_1-1}} + \sum_{(k,h) \in K_2} \frac{2|C_{k,h}(t)|}{m} \frac{\Delta}{2^k} \\
&\leq 4 \cdot \max \left\{ (\log P^*) \frac{\Delta}{2^{k_1}}, \frac{s(t)}{m} \right\} && (2)
\end{aligned}$$

We shall integrate  $s_i(t)^\alpha$  first over all  $t$  when the first term of Equation (2) is dominating to give an upper bound on required energy of:

$$(4 \log(P^*))^\alpha \sum_{k \geq 0} \sum_{h \geq 0} |C_{k,h} \cap \mathcal{J}_i^*| \left(\frac{\Delta}{2^k}\right)^\alpha \left(2^{k-h} \frac{\Gamma}{\Delta}\right)$$

Integrating  $s_i(t)^\alpha$  when the second term of Equation (2) is dominating gives  $\left(\frac{\Delta}{m}\right)^\alpha \text{AVR}_1(\mathcal{J}^*)$ . Summing over  $1 \leq i \leq m$  shows that  $\sum_{i=1}^m \text{AVR}_1(\mathcal{J}_i^*) \leq 4^\alpha ((\log^\alpha P^*) \text{MIN}(\mathcal{J}^*) + m^{1-\alpha} \text{AVR}_1(\mathcal{J}^*))$  as required.  $\square$

Together with Lemma 2, we can conclude the competitive ratio of DCRR in the following theorem.

**Theorem 2.** *For an arbitrary job set  $\mathcal{J}$ , the competitive ratio of algorithm DCRR is at most  $2^{4\alpha}(\log^\alpha P + \alpha^\alpha 2^{\alpha-1})$ , where  $P$  is the ratio between the maximum and minimum job size.*

*Proof.* By Lemma 4, we know that:

$$\sum_{1 \leq i \leq m} \text{AVR}_1(\mathcal{J}_i^*) \leq 2^{2\alpha} ((\log^\alpha P^*) \text{MIN}(\mathcal{J}^*) + \text{AVR}_1(\mathcal{J}^*)/m^{\alpha-1}).$$

Since  $\text{MIN}(\mathcal{J}^*) \leq \text{OPT}_m(\mathcal{J}^*)$  and  $\text{AVR}_1(\mathcal{J}^*) \leq \alpha^\alpha 2^{\alpha-1} m^{\alpha-1} \text{OPT}_m(\mathcal{J}^*)$  by Lemma 1 (b) (ii), we therefore conclude that

$$\text{DCRR}(\mathcal{J}^*) \leq \sum_{1 \leq j \leq m} \text{AVR}_1(\mathcal{J}_j^*) \leq 2^{2\alpha} \text{OPT}_m(\mathcal{J}^*) ((\log^\alpha P^*) + \alpha^\alpha 2^{\alpha-1}).$$

By Lemma 2 (a) and (b) above,  $\text{DCRR}(\mathcal{J}) \leq 2^\alpha \text{DCRR}(\mathcal{J}^*)$  and  $\text{OPT}_m(\mathcal{J}^*) \leq 2^\alpha \text{OPT}_m(\mathcal{J})$ . Then, we have

$$\text{DCRR}(\mathcal{J}) \leq 2^{4\alpha} \text{OPT}_m(\mathcal{J}) ((\log^\alpha P^*) + \alpha^\alpha 2^{\alpha-1}).$$

Note that from the proof of Lemma 4,  $\log P^*$  is essentially the number of size classes used by DCRR which does not change under  $\mathcal{J}$  or  $\mathcal{J}^*$ , therefore  $\log P$  and  $\log P^*$  can be taken to be equal and the theorem holds.  $\square$



## 5 Conclusion

We extend the study of energy efficient deadline scheduling on multiprocessor to jobs with arbitrary sizes and deadlines. We analyze the performance of the deterministic algorithm DCRR. In the proof of Theorem 2, the  $\log P$  factor comes in the case  $K_1$ , yet we note that this bound is rather loose and we believe that this can be improved. On the other hand, one may consider how DCRR can be coupled with OA instead of AVR to improve the results. Another open question is to consider speed bounded processors [12], in which case, not all the jobs can be completed by their deadlines. The concern becomes to maximize the throughput (number of jobs completed by their deadlines) and to minimize the energy used to achieve this throughput. The problem has been considered in the single processor setting [5, 12]. It would be interesting to derive algorithms that are competitive both in throughput and energy in the multiprocessor setting.

## References

1. S. Albers. Algorithms for energy saving. In S. Albers, H. Alt, and S. Näher, editors, *Efficient Algorithms*, volume 5760 of *Lecture Notes in Computer Science*, pages 173–186. Springer, 2009.
2. S. Albers. Energy-efficient algorithms. *Communication ACM*, 53(5):86–96, 2010.
3. S. Albers and H. Fujiwara. Energy-efficient algorithms for flow time minimization. *ACM Transactions on Algorithms*, 3(4):49, 2007.
4. S. Albers, F. Muller, and S. Schmelzer. Speed scaling on parallel processors. In *Proceedings of ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pages 289–298, 2007.
5. N. Bansal, H. L. Chan, T. W. Lam, and L. K. Lee. Scheduling for speed bounded processors. In *Proceedings of International Colloquium on Automata, Languages and Programming (ICALP)*, pages 409–420, 2008.
6. N. Bansal, H. L. Chan, and K. Pruhs. Speed scaling with an arbitrary power function. In *Proceedings of ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 693–701, 2009.
7. N. Bansal, H. L. Chan, K. Pruhs, and D. Rogozhnikov-Katz. Improved bounds for speed scaling in devices obeying the cube-root rule. In *Proceedings of International Colloquium on Automata, Languages and Programming (ICALP)*, to appear, 2009.
8. N. Bansal, T. Kimbrel, and K. Pruhs. Speed scaling to manage energy and temperature. *Journal of the ACM*, 54(1):3, 2007.
9. N. Bansal, K. Pruhs, and C. Stein. Speed scaling for weighted flow time. In *Proceedings of ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 805–813, 2007.
10. H. W. Becker and J. Riordan. The arithmetic of Bell and Stirling numbers. *American Journal of Mathematics*, 70:385–394, 1948.
11. 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.

12. H. L. Chan, W. T. Chan, T. W. Lam, L. K. Lee, K. S. Mak, and P. W. H. Wong. Optimizing throughput and energy in online deadline scheduling. *ACM Transactions on Algorithms*, 6(1):10, 2009. Preliminary version appeared in Proceedings of Symposium on Discrete Algorithms SODA, pages 795–804, 2007.
13. G. Greiner, T. Nonner, and A. Souza. The bell is ringing in speed-scaled multiprocessor scheduling. In *Proceedings of ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pages 11–18, 2009.
14. X. Han, T.-W. Lam, L.-K. Lee, I. K.-K. To, and P. W.-H. Wong. Deadline scheduling and power management for speed bounded processors. *Theoretical Computer Science*, 411(40–42):3587–3600, 2010.
15. S. Irani and K. Pruhs. Algorithmic problems in power management. *ACM SIGACT News*, 32(2):63–76, 2005.
16. S. Irani, S. Shukla, and R. K. Gupta. Algorithms for power savings. *ACM Transactions on Algorithms*, 3(4):41, 2007.
17. 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 *Proceedings of International Colloquium on Automata, Languages and Programming (ICALP)*, pages 665–676, 2009.
18. T. W. Lam, L. K. Lee, I. K. K. To, and P. W. H. Wong. Improved multiprocessor scheduling for flow time and energy. *Journal of Scheduling*. To appear <http://dx.doi.org/10.1016/j.tcs.2010.05.035>. Preliminary version appeared in Proceedings of ACM Symposium on Parallelism in Algorithms and Architectures, pages 256–264, 2008.
19. 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 *Proceedings of European Symposium on Algorithms (ESA)*, pages 647–659, 2008.
20. T. Mudge. Power: A first-class architectural design constraint. *Computer*, 34(4):52–58, 2001.
21. F. Yao, A. Demers, and S. Shenker. A scheduling model for reduced CPU energy. In *Proceedings of IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 374–382, 1995.