

SELFISHMIGRATE: A Scalable Algorithm for Non-clairvoyantly Scheduling Heterogeneous Processors

[Extended Abstract]

Sungjin Im

EECS Department,
University of California
Merced, CA, USA 95344.

Email: sim3@ucmerced.edu

Janardhan Kulkarni

Department of Computer Science,
Duke University, 308 Research Drive,
Durham, NC, USA 27708.

Email: {kulkarni, kamesh}@cs.duke.edu

Kamesh Munagala

Department of Computer Science,
University of Pittsburgh
Pittsburgh, PA, USA 15260

Email: kirk@cs.pitt.edu

Kirk Pruhs

Abstract—We consider the classical problem of minimizing the total weighted flow-time for unrelated machines in the online *non-clairvoyant* setting. In this problem, a set of jobs J arrive over time to be scheduled on a set of M machines. Each job j has processing length p_j , weight w_j , and is processed at a rate of ℓ_{ij} when scheduled on machine i . The online scheduler knows the values of w_j and ℓ_{ij} upon arrival of the job, but is not aware of the quantity p_j . We present the *first* online algorithm that is *scalable* ($(1 + \epsilon)$ -speed $O(\frac{1}{\epsilon^2})$ -competitive for any constant $\epsilon > 0$) for the total weighted flow-time objective. No non-trivial results were known for this setting, except for the most basic case of identical machines. Our result resolves a major open problem in online scheduling theory. Moreover, we also show that no job needs more than a logarithmic number of migrations.

We further extend our result and give a scalable algorithm for the objective of minimizing total weighted flow-time plus energy cost for the case of unrelated machines. In this problem, each machine can be sped up by a factor of $f_i^{-1}(P)$ when consuming power P , where f_i is an *arbitrary* strictly convex power function. In particular, we get an $O(\gamma^2)$ -competitive algorithm when all power functions are of form s^γ . These are the *first* non-trivial non-clairvoyant results in any setting with heterogeneous machines.

The key algorithmic idea is to let jobs migrate selfishly until they converge to an equilibrium. Towards this end, we define a game where each job's utility which is closely tied to the instantaneous increase in the objective the job is responsible for, and each machine declares a policy that assigns priorities to jobs based on when they migrate to it, and the execution speeds. This has a spirit similar to coordination mechanisms that attempt to achieve near optimum welfare in the presence of selfish agents (jobs). To the best of our knowledge, this is the first work that demonstrates the usefulness of ideas from coordination mechanisms and Nash equilibria for designing and analyzing online algorithms.

I. INTRODUCTION

Many computer architects believe that architectures consisting of heterogeneous processors will be the dominant architectural design in the future: Simulation studies indicate that, for a given area and power budget, heterogeneous multiprocessors can offer an order of magnitude better performance for typical workloads [10], [30], [32], [29].

Looking at the consequences of Moore's Law even further in the future, some computer architectures are projecting that we will transition from the current era of multiprocessor scaling to an era of "dark silicon", in which switches become so dense that it is not economically feasible to cool the chip if all switches are simultaneously powered. [18]. One possible architecture in the dark silicon era would be many specialized processors, each designed for a particular type of job. The processors that are on any point of time should be those that are best suited for the current tasks.

It is recognized by the computer systems community [10] and the algorithms community that scheduling these future heterogeneous multiprocessor architectures is a major challenge. It is known that some of the standard scheduling algorithms for single processors and homogeneous processors can perform quite badly on heterogeneous processors [20]. A scalable algorithm (defined below) is known if somehow the scheduler was clairvoyant (able to know the size of a job when it arrives) [11]; however, this knowledge is generally not available in general purpose computing settings. A scalable algorithm is also known if all jobs were of equal importance [25]; however, the whole *raison d'être* for heterogeneous architectures is that there is generally heterogeneity among the jobs, most notably in their *importance/priorities*.

Therefore, a major open question in the area of online scheduling, both in theory and practice, is the design of scalable online algorithms for scheduling heterogeneous processors with arbitrary power functions, when job sizes are not known in advance (non-clairvoyant setting), and jobs have different priorities (weighted setting). The typical objective that has been studied in this context is minimizing weighted delay, or weighted flow-time. This problem generalizes most previous work [3], [11], [15], [25], [20], [25] in online single and multiple machine scheduling considered recently by the algorithms community. As we discuss below, the algorithmic techniques developed for these problems (clairvoyant setting, unweighted setting, etc) do not extend in any natural way to the most general setting with weights,

non-clairvoyance, and energy, leading to a fundamental gap in our understanding of classical scheduling algorithms. In particular, we ask:

Do different variants of multiple machine scheduling considered in literature require different algorithmic techniques and analyses, or is there one unifying technique for them all?

In this paper, we close this gap, and obtain *one unifying algorithmic technique for them all*, achieving the *first* scalable non-clairvoyant algorithms for scheduling jobs of varying importance on heterogeneous machines, even with arbitrary power functions. The interesting aspect of our work, as we discuss in Section I-B below, is that it provides an algorithmically simple and conceptually novel framework for multiple machine scheduling as a coordination game (see [14], [9]), where jobs have (virtual) utility functions for machines based on delays they contribute to, and machines announce scheduling policies, and treat migration of jobs into them as new arrivals. In hindsight, we believe this provides the correct, unifying way of viewing all recent algorithmic results [3], [15], [25], [20] in related and simpler models.

A. Our Results

We adopt the competitive analysis framework in online algorithms. We say that an online schedule is s -speed c -competitive if it is given s times faster machines and is c -competitive when compared to the optimal scheduler with no speed augmentation. The goal is to design a $O(1)$ -competitive algorithm with the smallest extra speed. In particular, a *scalable algorithm*, which is $(1 + \epsilon)$ -speed $O(1)$ -competitive for any fixed $\epsilon > 0$, is considered to be essentially the best result one can hope for in the competitive analysis framework for machine scheduling [28]. (It is known that without resource augmentation no online algorithm can have bounded competitive ratio for our problem [19], [11].)

Unrelated Machine Scheduling.

We first consider the general *unrelated* machine model where each job j can be processed at rate $\ell_{ij} \geq 0$ on each machine i ; if $\ell_{ij} = 0$, then job j cannot be processed on machine i . In a feasible schedule, each job can be scheduled by at most one machine at any point in time. Preemption is allowed without incurring any cost, and so is migration to a different machine. A job j 's flow-time $F_j := C_j - r_j$ measures the length of time between when the job arrives at time r_j and when the job is completed at C_j . When each job j has weight w_j , the total weighted flow-time is $\sum_j w_j F_j$; the unweighted case is where all jobs weights are one. The online scheduler does not know the job size p_j until the job completes. For a formal problem statement, please refer to Section II.

Our first result is the following. For comparison, we

note that no constant speed, constant competitive result was known for the weighted case, except when all machines are identical.

Theorem I.1. [Section II.] *For any $\epsilon > 0$, there is a $(1 + \epsilon)$ -speed $O(1/\epsilon^2)$ -competitive non-clairvoyant algorithm for the problem of minimizing the total weighted flow-time on unrelated machines. Furthermore, each job migrates at most $O((\log W + \log n)/\epsilon)$ times, where W denotes the ratio of the maximum job weight to the minimum.*

It is well known [20] that in the non-clairvoyant setting, jobs need to be migrated to obtain $O(1)$ competitive ratio. Our algorithm migrates each job relatively small number of times. Reducing the number of migrations is not only theoretically interesting but also highly desirable in practice [12], [24].

Power Functions.

We further extend our result to unrelated machines with arbitrary power functions. Each machine i is associated with an arbitrary power function $f_i : [0, \infty) \rightarrow [0, \infty)$ which is strictly convex and has $f_i(0) = 0$. When machine i uses power $P > 0$, it runs $f_i^{-1}(P)$ times faster than its original speed – it processes job j at a rate of $\ell_{ij} f_i^{-1}(P)$. This model is widely studied [1], [7], [21], [11], [20], [15] and we consider the standard objective of minimizing the total weighted flow-time plus the total energy consumption [2].

Theorem I.2. [Section III.] *For any $\epsilon > 0$, there is a $(1 + \epsilon)$ -speed $O(1/\epsilon^2)$ -competitive non-clairvoyant algorithm for the problem of minimizing the total weighted flow-time plus total energy consumption on unrelated machines. This result holds even when each machine i has an arbitrary strictly-convex power function $f_i : [0, \infty) \rightarrow [0, \infty)$ with $f_i(0) = 0$.*

The same guarantee on the number of migrations by a job stated in Theorem I.1 holds for this setting. The theorem implies a $O(\gamma^2)$ -competitive algorithm (without resource augmentation) when each machine i has a power function $f(s) = s^\gamma$ for some $\gamma > 1$, perhaps most important power functions in practice. Our result also implies a scalable algorithm in the model where at each time instant a processor i can either run at speed s_i consuming a power P_i , or be shutdown and consume no energy.

We note that no $O(1)$ -speed $O(1)$ -competitive non-clairvoyant algorithm was known prior to our work even in the related machine setting for *any* nontrivial classes of power functions.

B. Technical Contributions: Selfish Migration and Nash Equilibrium

Our main technical contribution is a new conceptually simple game-theoretic framework for multiple machine scheduling that unifies, simplifies and generalizes previous work, both in terms of algorithm design as well as analysis

using dual fitting. Before presenting this framework, we present some difficulties an online scheduler has to overcome in the non-clairvoyant settings we consider.

An online scheduler for multiple machines consists of two scheduling components: A single-machine scheduling policy on each machine, and the global machine assignment rule which assigns jobs to machines. In the context of clairvoyant scheduling [11], [3], the authors in [3] show via a dual fitting analysis that the following algorithm is scalable: Each machine runs a scalable single-machine scheduling policy such as Highest Density First (HDF); this is coupled with a simple greedy dispatch rule that assigns arriving jobs to the machine on which they cause the least increase in flow-time to previously assigned jobs. This simple yet elegant algorithm has been very influential. In particular, the greedy dispatch rule has become standard, and been used in various scheduling settings [26], [22], [3], [34], [15]. The analysis proceeds by setting dual variables corresponding to a job to the marginal increase in total delay due to the arrival of this job, and showing that this setting is not only feasible but also extracts a constant fraction of the weighted flow-time in the objective. *The immediate-dispatch greedy rule is necessary for analysis in all the aforementioned work*, since they require the algorithm to measure each job’s effect on the system at the moment it arrives.

In a non-clairvoyant context, there are main two hurdles that arise. First, to use the greedy dispatch rule, it is crucial to measure how much a new job affects the overall delay, that is, how much the job increases the objective. To measure the increase, the scheduler must know the job size, which is not allowed in the non-clairvoyant setting. Secondly, as mentioned before, jobs must migrate to get a $O(1)$ -competitive algorithm even with any $O(1)$ -speed, and this makes it more difficult to measure how much each job is responsible for the overall delay. This difficulty appears in the two analysis tools for online scheduling, potential function [27] and dual fitting method [3], [23], [15]. Due to these difficulties, there have been very few results for non-clairvoyant scheduling on heterogeneous machines [21], [20], [25]. Further, there has been *no* work in *any* heterogeneous machines setting for the weighted flow-time objective.

1) SELFISHMIGRATE Framework

We demonstrate a simple framework SELFISHMIGRATE that addresses the above two issues in one shot. Our algorithm can be best viewed in a game theoretic setting where jobs are selfish agents, and machines declare their scheduling policies in advance.

Machine Behavior.

Each machine maintains a *virtual queue* on the current set of jobs assigned to it; newly arriving jobs are appended to the tail of this queue. In a significant departure from previous work [3], [15], [21], [20], [25], each machine treats

a migration of a job to it as an arrival, and a migration out of it as a departure. This means a job migrating to a machine is placed at the tail of the virtual queue.

Each machine runs a scheduling policy that is a modification of weighted round robin (WRR) that smoothly assigns larger speed to jobs in the tail of the queue, taking weights into account. This is a smooth extension of the algorithm Latest Arrival Processor Sharing (LAPS or WLAPS) [17], [16]. We note that the entire analysis also goes through with WRR, albeit with $(2+\epsilon)$ -speed augmentation. The nice aspect of our smooth policies (unlike WLAPS) is that we can approximate the instantaneous delay introduced by this job to jobs ahead of it in its virtual queue, *even without knowing job sizes*. This will be critical for our analysis.

Job Behavior.

Each job j has a virtual utility function, which roughly corresponds to the inverse of the instantaneous weighted delay introduced by j to jobs ahead of it in its virtual queue, and their contribution to j ’s weighted delay. Using these virtual utilities, jobs perform sequential best response (SBR) dynamics, migrating to machines (and get placed in the tail of their virtual queue) if doing so leads to larger virtual utility. Therefore, at each time instant, the job migration achieves a Nash equilibrium of the SBR dynamics on the virtual utilities. We show that our definition of the virtual utilities implies they never decrease due to migrations, arrivals, or departures, so that at any time instant the Nash equilibrium exists and is computable. (We note that at each time step, we simulate SBR dynamics and migrate each job directly to the machine that is predicted by the Nash equilibrium.)

When a job migrates to a machine, the virtual utility starts off being the same as the real speed the job receives. As time goes by, the virtual queue ahead of this job shrinks, and that behind it increases. This lowers the real speed the job receives, but its virtual utility, which measures the inverse of the impact to jobs ahead in the queue and vice versa, does not decrease. Our key contribution is to define the coordination game on the virtual utilities, rather than on the actual speed improvement jobs receive on migration. A game on the latter quantities (utility equals actual speed) need not even admit to a Nash equilibrium.

Given the above framework, our analysis proceeds by setting the dual variable for a job to the increase in overall weighted delay it causes on jobs ahead of it in its virtual queue. Our key insight is to show that Nash dynamics on virtual utilities directly corresponds to our setting of dual variables being feasible for the dual constraints, implying the desired competitive ratio. This overall approach is quite general, even extending to energy constraints, and requires two key properties from the virtual utility:

- The virtual utility should correspond roughly to the

inverse of the instantaneous delay induced by a job on jobs ahead of it in its virtual queue.

- SBR dynamics should monotonically improve virtual utility, leading to a Nash equilibrium that corresponds exactly to satisfying the dual constraints.

Our main (and perhaps quite surprising) contribution is to show the existence of such a virtual utility function for WRR and its scalable modifications, when coupled with the right notion of virtual queues. In hindsight, we believe this framework is the right way to generalize the greedy dispatch rules and dual fitting analysis from previous works [3], [25], and we hope it finds more applications in complex scheduling settings.

2) Comparison with Previous Techniques

As mentioned above, the algorithmic idea based on coordination games is very different from the previous greedy dispatch rules [11], [3], [15] for clairvoyant scheduling, and also from previous work on non-clairvoyant scheduling [21], [20], [25]. We contrast our algorithm with these, highlighting the necessity of new techniques.

It is instructive to compare this framework with the scalable non-clairvoyant algorithm for *unweighted* flow-time [25]. This algorithm has the seeds of several ideas we develop here – it introduces virtual queues, a smooth variant of LAPS [17], [16] for single-machine scheduling, as well as migration based on the delay a job contributes to jobs ahead in its queue. However, this algorithm always assigns priorities to jobs in order of original arrival times, and migrates jobs to machines preserving this ordering. In essence, this algorithm mimics the clairvoyant algorithms [11], [3] that account for delay a job contributes to those that arrived ahead in time. This idea of arrival time ordering is specific to unweighted jobs, and does not extend to the weighted case or to energy constraints. In contrast, we let each machine treat migrations as new arrivals, leading us to view job migration through a game-theoretic lens. This leads to a more natural framework via instantaneous Nash equilibria, with a simple dual fitting analysis. The simplicity makes the framework extend easily to energy constraints. The resulting accounting using the delay a job j induces to those ahead of it in its virtual queue is novel – in contrast to previous work [3], [15], [21], [20], [25], the virtual queue changes with time and could possibly include jobs whose original arrival time $r_{j'}$ is later than that of job j .

To further illustrate the technical difficulty of the weighted case, let us consider Round-Robin (RR) and its variants. The work of [21] consider unweighted jobs, and gives $(2 + \epsilon)$ -speed $O(1)$ -competitive algorithm for the total (unweighted) flow-time on related machines. The work of [20] improves this to a scalable algorithm, but for simpler illustration, we focus on RR. In the RR used in [21], each of n fastest machines is distributed to all n jobs uniformly. It is not difficult to see that this fractional schedule can be

decomposed into a convex combination of feasible actual schedules. Hence, RR allows us to view multiple machines as a more powerful single machine, and this is the crux of the analysis in [21]. In contrast, the work of [20] argues that the weighted case is not as simple: in fact, they show that natural extensions of weighted round robin for related machines fail to be competitive.

C. Other Related Work

For a survey on online scheduling, please see [33]. As alluded to above, for weighted flow-time on unrelated machines, the best clairvoyant result is a $(1 + \epsilon)$ -speed $O(1/\epsilon)$ -competitive algorithm [3], [15]. For the version with power functions, the best corresponding clairvoyant result is a $(1 + \epsilon)$ -speed $O(1/\epsilon)$ -competitive algorithm [15]. In the most basic setting of multiple machines where machines are all identical, the work of [13] gives the first analysis of scalable algorithms for weighted flow-time. Intuitively, machine assignment rule should have a spirit of load balancing. Indeed, the work of [13] shows two machine assignment rules can be combined with various single machine scheduling policies to yield a scalable algorithm. One is random assignment rule, and the other is based on volume of each job class.

For the problem of non-clairvoyantly scheduling a single machine, the WLAPS (Weighted Latest Arrival Processor Sharing) algorithm [17], [5], [16] is scalable for the total weighted flow even when jobs have arbitrary speedup curves. Other non-clairvoyant scalable algorithms for the unweighted case include Shorted Elapsed Time First (SETF) and Multi-level Feedback [28]. The work of [6] extends SETF to its weighted version. While Shortest Remaining Processing Time (SRPT) is optimal for the total flow time objective on a single machine with clairvoyance, even slightly more general settings (non-clairvoyant or weighted or multiple machine settings) do not admit a $O(1)$ -competitive algorithm without resource augmentation [8], [4], [31].

II. UNRELATED MACHINE SCHEDULING

In this problem, there are m unrelated machines. Job j is processed at rate $\ell_{ij} \in [0, \infty)$ on each machine i . Each job has processing length p_j and weight w_j . The online algorithm is allowed to preempt and migrate jobs at any time with no penalty. The important constraint is that at any instantaneous time, a job can be processed only on a single machine. Job j is released at time r_j . In the non-clairvoyant online scheduling model we consider, the scheduler knows the values of ℓ_{ij} and w_j when the job arrives, but is not aware of the processing length p_j .¹ Without loss of generality we assume that weights w_j are integers.

Fix some scheduling policy \mathcal{P} . At each time instant t , each active job j with $r_j \leq t$ is assigned to some machine i . Let

¹It is easy to show that if ℓ_{ij} values are not known, then no online algorithm can have a bounded competitive ratio even with any constant speed augmentation.

$J_i(t)$ denote the set of jobs assigned to machine i . Machine i splits its processing power among the jobs in $J_i(t)$. Let $\nu_j(t)$ denote the processing power assigned to job $j \in J_i(t)$. We enforce that $\sum_{j \in J_i(t)} \nu_j(t) \leq 1$ for all i, t . Then, $j \in J_i(t)$ executes at rate $q_j(t) = \ell_{ij} \nu_j(t)$. The completion time C_j is defined as the earliest time t_j such that

$$C_j = \operatorname{argmin}_{t_j} \left(\int_{t=r_j}^{t_j} q_j(t) dt \geq p_j \right)$$

At this time, the job finishes processing and departs from the system. The objective is to find a scheduling policy that minimizes the sum of weighted flow-times $\sum_j w_j F_j$, where $F_j = C_j - r_j$ is the flow-time of job j .

In the *speed augmentation* setting, we assume the online algorithm can process job j at rate $(1 + \epsilon)\ell_{ij}$ on machine i , where $\epsilon > 0$. We will compare the resulting flow-time against an offline optimum that knows p_j and r_j at time 0, but is not allowed the extra speed. Our main result is a *scalable* algorithm that, for any $\epsilon > 0$, is $O(1/\epsilon^2)$ competitive with speed augmentation of $(1 + \epsilon)$.

A. The SELFISHMIGRATE Algorithm

Our algorithm can be best viewed as a coordination mechanism between the machines and the jobs. Each machine declares a *single machine policy* that it uses to prioritize and assign rates to arriving jobs. Given these policies, jobs migrate to machines that give them the most instantaneous utility (in a certain virtual sense). We will now define the single machine scheduling policy, and the utility function for jobs.

1) Single Machine Policy: Weighted Ranked Processor Sharing, WRPS (k)

This policy is parametrized by an integer k (that we will later set to $1/\epsilon$) and $\eta > 1$ that captures the speed augmentation (and that we set later to $1 + 3\epsilon$). Fix some machine i and time instant t . Recall that $J_i(t)$ denotes the set of jobs assigned to this machine at time t . Let $W(i, t)$ denote their total weight, *i.e.*, $W(i, t) = \sum_{j \in J_i(t)} w_j$. The machine maintains a *virtual queue* on these jobs.

We now introduce some notation based on these virtual queues. Let $\sigma(j, t)$ denote the machine to which job j is assigned at time t . Therefore, $i = \sigma(j, t)$ if and only if $j \in J_i(t)$. Let $\pi_j(t)$ denote the rank of j in the virtual queue of $i = \sigma(j, t)$, where the head of the virtual queue has rank 1 and the tail of the queue has rank $|J_i(t)|$. Let $\mathcal{J}_j(t)$ denote the set of jobs ahead of job $j \in J_i(t)$ in the virtual queue of machine i . In other words

$$\mathcal{J}_j(t) = \{j' \mid \sigma(j', t) = \sigma(j, t) \text{ and } \pi_{j'}(t) < \pi_j(t)\}$$

Let $\mathcal{W}_j(t) = \sum_{j' \in \mathcal{J}_j(t)} w_{j'}$ denote the total weight of jobs ahead of job j in its virtual queue.

Rate Assignment.

At time instant t , the total processing rate of the machine i is divided among the jobs in $J_i(t)$ as follows. Job $j \in J_i(t)$ is assigned processing power $\nu_j(t)$ as follows:

$$\nu_j(t) := \eta \cdot \frac{(\mathcal{W}_j(t) + w_j)^{k+1} - \mathcal{W}_j^{k+1}}{W(i, t)^{k+1}} \quad (1)$$

The rate at which job $j \in J_i(t)$ is processed at time t is therefore $\ell_{ij} \nu_j(t)$. Note that $\sum_{j \in J_i(t)} \nu_j(t) = \eta$ at all time instants t and for all machines i . Note that if $k = 0$, this is exactly weighted round robin. As k becomes larger, this gives higher rate to jobs in the tail of the queue, taking the weights w_j into account. This ensures that small jobs arriving later do not contribute too much to the flow-time, hence reducing the speed augmentation. One important property of WRPS (k) is that if a new job is added to the tail of the virtual queue, then all the old jobs are slowed down by the same factor. This is one of the important characteristics of weighted round robin which ensures that for a pair of jobs weighted delay induced by each other is exactly same. WRPS(k) preserves this property to a factor of $O(k)$, and this will be crucial to our analysis.

Our policy is weighted version of the single machine scheduling policy introduced in [25]. We note that using the natural setting of $k = 0$ (weighted round robin) gives a competitive algorithm with speedup $(2 + \epsilon)$, and this is tight even for a single machine. We use a larger value of k to reduce the amount of speed augmentation needed. (We believe that WRPS(k) gives a black-box reduction from any $(2 + \epsilon)$ -speed algorithm using WRR into a scalable algorithm.)

Arrival Policy.

The behavior of the policy is the same when a job j either arrives to the system and chooses machine i , or migrates from some other machine to machine i at time t . In either case, the job j is placed at the *tail* of the virtual queue. In other words, if $J_i(t^-)$ is the set of jobs just before the arrival of job j , then we set $\sigma(j, t) = i$ and $\pi_j(t) = |J_i(t^-)| + 1$. Therefore, the virtual queue sorts the jobs in order in which they arrive *onto this machine*. Since a job could also arrive due to a migration, this is *not the same* as ordering on the r_j – a job with smaller r_j that migrates at a latter point onto machine i will occupy a relatively later position in its virtual queue.

Departure Policy.

If job j departs from machine i either due to completion or due to migrating to a different machine, the job simply deletes itself from i 's virtual queue, keeping the ordering of other jobs the same. In other words, for all jobs $j' \in J_i(t)$ with $\pi_{j'}(t) > \pi_j(t)$, the value $\pi_{j'}(t)$ decreases by 1.

2) Virtual Utility of Jobs and Selfish Migration

The virtual queues define a *virtual utility* of job as follows. Let $j \in J_i(t)$ at time t . Then its virtual utility is defined as:

$$\phi(j, t) = \frac{\ell_{ij}}{W_j(t) + w_j}$$

We interpret this utility as follows: The inverse of this quantity will be roughly in proportion to the marginal increase in instantaneous weighted delay that job j induces on jobs $\mathcal{J}_j(t)$ that are ahead of it in its virtual queue, and their contribution to the weighted delay of j . We will establish this in the Delay Lemmas below. This marginal increase is exactly what we need in order to define dual variables in our proof, and in some sense, the virtual utility is defined keeping this in mind. This utility function is similar to the utility function used for the unweighted case in [25].

At every time instant, job $j \in J_i(t)$ behaves as follows: If it were to migrate to machine $d \neq i$, it would be placed at the tail of d 's queue and would obtain virtual utility $\frac{\ell_{dj}}{W(d, t) + w_j}$. If this quantity is larger than $\phi(j, t)$, then job j migrates to machine d . This leads to the corresponding changes to the virtual queues of machine i (job j is deleted), machine d (job j is added to the tail), and the virtual utility $\phi(j, t)$ of job j (which is set to $\frac{\ell_{dj}}{W(d, t) + w_j}$). At every time instant t , this defines a game on the jobs, and starting with the configuration at the previous step, the jobs simulate sequential best response dynamics, where they sequentially migrate to improve virtual utility, till the system reaches a Nash equilibrium. In this configuration, each job is on a machine that is locally optimal for $\phi(j, t)$.

Note that if a job departs from a machine, the virtual utilities of other jobs on that machine either stay the same or increase. Further, if a job migrates to a machine, it is placed on the tail of the virtual queue, so that the virtual utilities of other jobs on the machine remain the same. This shows that sequential best response dynamics guarantees that the virtual utilities of *all* jobs are monotonically non-decreasing with time, converging to a Nash equilibrium. (Note that jobs don't actually need to execute best response dynamics since they can directly migrate to the machines corresponding to the resulting Nash equilibrium.)

When a new job arrives to the system, it simply chooses the machine i which maximizes its virtual utility, $\frac{\ell_{ij}}{W(i, t) + w_j}$, where $W(i, t)$ is the weight of jobs assigned to i just before the arrival of job j . This completes the description of the algorithm.

The following lemma is an easy consequence of the description of the algorithm.

Lemma II.1. *For all jobs j , $\phi(j, t)$ is non-decreasing over the interval $t \in [r_j, C_j]$.*

B. Analysis of SELFISHMIGRATE

We first write a linear programming relaxation of the problem $\text{LP}_{\text{primal}}$ described below which was first given by [3], [19]. It has a variable x_{ijt} for each machine $i \in [m]$, each job $j \in [n]$ and each unit time slot $t \geq r_j$. If the machine i processes the job j during the whole time slot t , then this variable is set to 1. The first constraint says that every job has to be completely processed. The second constraint says that a machine cannot process more than one unit of jobs during any time slot. Note that the LP allows a job to be processed simultaneously across different machines.

$$\begin{aligned} \text{Min} \quad & \sum_j \sum_i \sum_{t \geq r_j} \left(\frac{\ell_{ij}(t - r_j)}{p_j} + 1 \right) \cdot w_j \cdot x_{ijt} \\ & \sum_i \sum_{t \geq r_j} \frac{\ell_{ij} \cdot x_{ijt}}{p_j} \geq 1 \quad \forall j \\ & \sum_{j: t \geq r_j} x_{ijt} \leq 1 \quad \forall i, t \\ & x_{ijt} \geq 0 \quad \forall i, j, t : t \geq r_j \end{aligned} \quad (\text{LP}_{\text{primal}})$$

It is easy to show that the above LP lower bounds the optimal flow-time of a feasible schedule within factor 2. We use the dual fitting framework to analyze SELFISHMIGRATE. We write the dual of $\text{LP}_{\text{primal}}$ as,

$$\begin{aligned} \text{Max} \quad & \sum_j \alpha_j - \sum_i \sum_t \beta_{it} \\ \frac{\ell_{ij} \cdot \alpha_j}{p_j} - \beta_{it} & \leq \frac{w_j \ell_{ij}(t - r_j)}{p_j} + w_j \quad \forall i, j, t : t \geq r_j \\ \alpha_j & \geq 0 \quad \forall j \\ \beta_{it} & \geq 0 \quad \forall i, t \end{aligned} \quad (\text{LP}_{\text{dual}})$$

We will show that there is a feasible solution to LP_{dual} that has objective $O(\epsilon^2)$ times the total weighted flow-time of SELFISHMIGRATE, provided we augment the speed by $\eta = (1 + 3\epsilon)$. From now on, we will assume that each processor in SELFISHMIGRATE has η extra speed when processing jobs.

1) Instantaneous Delay and Setting Dual Variables

Recall that each machine runs WRPS (k) with $k = 1/\epsilon$, and we assume without loss of generality that $1/\epsilon$ is an integer. We define the *instantaneous weighted delay* induced by job j on jobs ahead of j in its virtual queue (the set $\mathcal{J}_j(t)$) as follows:

$$\delta_j(t) = \frac{1}{\eta} \left(\sum_{j' \in \mathcal{J}_j(t)} (w_{j'} \cdot \nu_j(t) + w_j \cdot \nu_{j'}(t)) + w_j \cdot \nu_j(t) \right)$$

This quantity sums the instantaneous weighted delay that $j' \in \mathcal{J}_j(t)$ induces on j and vice versa, plus the delay seen by j due to itself. Note that $\delta_j(t)$ is equal to $\frac{1}{\eta} \left((W_j(t) + w_j) \cdot \nu_j(t) + w_j \cdot \sum_{j' \in \mathcal{J}_j(t)} \nu_{j'}(t) \right)$.

Define,

$$\Delta_j = \int_{t=r_j}^{C_j} \delta_j(t) dt$$

as the cumulative weighted delay induced by j on jobs ahead of it in its virtual queue and vice versa. Note that the set $\mathcal{J}_j(t)$ changes with t and can include jobs that are released after job j . It is an easy exercise to check that $\sum_j w_j F_j = \sum_j \Delta_j$. Our way of accounting for weighted delay is therefore a significant departure from previous work that either keeps $\mathcal{J}_j(t)$ the same for all t (clairvoyant algorithms), or preserves orderings based on arrival time.

We now perform the dual fitting. We set the variables of the LP_{dual} as follows. We set β_{it} proportional to the total weight of jobs alive on machine i at time t , i.e., $\beta_{it} = \frac{1}{k+3} W(i, t)$. We set $\alpha_j = \frac{1}{k+2} \Delta_j$, i.e., proportional to the cumulative weighted delay induced by j on jobs ahead of it in its virtual queue.

We first bound the dual objective as follows (noting $k = 1/\epsilon$ and $\eta = 1 + 3\epsilon$):

$$\begin{aligned} & \sum_j \alpha_j - \sum_{i,t} \beta_{it} = \sum_j \frac{\Delta_j}{k+2} - \sum_{i,t} \frac{W(i,t)}{k+3} \\ &= \epsilon \left(\sum_j \frac{\Delta_j}{1+2\epsilon} - \sum_{i,t} \frac{W(i,t)}{1+3\epsilon} \right) \\ &= \epsilon \cdot \sum_j w_j F_j \cdot \left(\frac{1}{1+2\epsilon} - \frac{1}{1+3\epsilon} \right) \\ &= O(\epsilon^2) \sum_j w_j F_j \end{aligned} \quad (2)$$

In the rest of the analysis, we show that this setting of dual variables satisfies the dual constraints.

2) Delay Lemmas

The dual constraints need us to argue about the weighted delay induced by j till any point t . For this purpose, we define for any $t^* \in [r_j, C_j]$ the following:

$$\Delta_j^1(t^*) = \int_{t=r_j}^{t^*} \delta_j(t) dt \quad \text{and} \quad \Delta_j^2(t^*) = \int_{t=t^*}^{C_j} \delta_j(t) dt$$

The following propositions have elementary proofs which have been omitted.

Proposition II.2. Consider any integer $k \geq 0$, and $0 \leq w \leq 1$, then $(1-w)^k \geq 1-kw$.

Proposition II.3. Consider any integer $k \geq 0$, and $w, w' \geq 0$, then $(w+w')^k \geq w^k + kw^{k-1}w'$.

The first Delay Lemma bounds the quantity $\Delta_j^1(t^*)$ as follows:

Lemma II.4 (First Delay Lemma). For any time instant $t^* \in [r_j, C_j]$ and for any job j ,

$$\Delta_j^1(t^*) \leq (k+2) \cdot w_j \cdot (t^* - r_j)$$

Proof:

$$\begin{aligned} & \Delta_j^1(t^*) \\ &= \frac{1}{\eta} \int_{t=r_j}^{t^*} \left(\nu_j(t) \cdot (\mathcal{W}_j(t) + w_j) + w_j \cdot \left(\sum_{j' \in \mathcal{J}_j(t)} \nu_{j'}(t) \right) \right) dt \\ &\leq \frac{1}{\eta} \int_{t=r_j}^{t^*} \left(\eta \cdot \frac{(\mathcal{W}_j(t) + w_j)^{k+1} - \mathcal{W}_j(t)^{k+1}}{W(\sigma(j,t), t)^{k+1}} \cdot (\mathcal{W}_j(t) + w_j) + w_j \cdot \eta \right) dt \\ &\quad \text{[Definition of } \nu] \\ &\leq \int_{t=r_j}^{t^*} \left(\frac{(\mathcal{W}_j(t) + w_j)^{k+1} - \mathcal{W}_j(t)^{k+1}}{(\mathcal{W}_j(t) + w_j)^{k+1}} \cdot (\mathcal{W}_j(t) + w_j) + w_j \right) dt \\ &\quad \text{[Since } W(\sigma(j,t), t) \geq \mathcal{W}_j(t) + w_j] \\ &= \int_{t=r_j}^{t^*} \left(\left(1 - \left(1 - \frac{w_j}{\mathcal{W}_j(t) + w_j} \right)^{k+1} \right) \cdot (\mathcal{W}_j(t) + w_j) + w_j \right) dt \\ &\leq \int_{t=r_j}^{t^*} \left(\frac{w_j \cdot (k+1)}{(\mathcal{W}_j(t) + w_j)} \cdot (\mathcal{W}_j(t) + w_j) + w_j \right) dt \\ &\quad \text{[Proposition II.2]} \\ &= (k+2) \cdot w_j \cdot (t^* - r_j) \end{aligned}$$

Let $p_j(t^*) = \int_{t=t^*}^{C_j} \ell_{\sigma(j,t)} \cdot \nu_j(t) dt$ denote the residual size of job j at time t^* . The second Delay Lemma states that total marginal increase in the algorithm's cost due to job j till its completion is upper bounded by the marginal increase in the algorithm's cost if the job j stays on machine $\sigma(j, t^*)$ till its completion. However, as noted before, marginal increase in the cost of the algorithm on a single machine is inversely proportional to the job's virtual speed. The proof of the second Delay Lemma hinges crucially on the fact that a job selfishly migrates to a new machine only if its virtual utility increases. In fact, the statement of this lemma implies the correctness of our setting of virtual utility. ■

Lemma II.5 (Second Delay Lemma). For any time instant $t^* \in [r_j, C_j]$ and for any job j , let $i^* = \sigma(j, t^*)$ denote the machine to which job j is assigned at time t^* . Then:

$$\Delta_j^2(t^*) \leq \frac{1}{\eta} \cdot \frac{k+2}{k+1} \cdot \frac{p_j(t^*)}{\phi(j, t^*)} \leq \frac{1}{\eta} \cdot \frac{k+2}{k+1} \cdot p_j \cdot \frac{\mathcal{W}_j(t^*) + w_j}{l_{i^*j}}$$

Proof:

$$\begin{aligned} & \Delta_j^2(t^*) \\ &= \frac{1}{\eta} \int_{t=t^*}^{C_j} \left(\nu_j(t) \cdot (\mathcal{W}_j(t) + w_j) + w_j \cdot \left(\sum_{j' \in \mathcal{J}_j(t)} \nu_{j'}(t) \right) \right) dt \\ &= \frac{1}{\eta} \int_{t=t^*}^{C_j} \left(\nu_j(t) \cdot (\mathcal{W}_j(t) + w_j) + \eta \cdot w_j \cdot \frac{\mathcal{W}_j(t)^{k+1}}{W(\sigma(j,t), t)^{k+1}} \right) dt \\ &= \frac{1}{\eta} \int_{t=t^*}^{C_j} \nu_j(t) \cdot \left(\mathcal{W}_j(t) + w_j + w_j \cdot \frac{\mathcal{W}_j(t)^{k+1}}{(\mathcal{W}_j(t) + w_j)^{k+1} - \mathcal{W}_j(t)^{k+1}} \right) dt \\ &\leq \frac{1}{\eta} \int_{t=t^*}^{C_j} \nu_j(t) \cdot \left(\mathcal{W}_j(t) + w_j + \frac{\mathcal{W}_j(t) + w_j}{k+1} \right) dt \\ &\quad \text{[Proposition II.3]} \end{aligned}$$

$$\begin{aligned}
&= \frac{1}{\eta} \cdot \frac{k+2}{k+1} \cdot \int_{t=t^*}^{C_j} \ell_{\sigma(j,t)j} \cdot \nu_j(t) \cdot \frac{1}{\phi(j,t)} dt \\
&\leq \frac{1}{\eta} \cdot \frac{k+2}{k+1} \cdot \frac{1}{\phi(j,t^*)} \cdot \int_{t=t^*}^{C_j} \ell_{\sigma(j,t)j} \cdot \nu_j(t) dt \\
&\quad \text{[Lemma II.1]} \\
&= \frac{1}{\eta} \cdot \frac{k+2}{k+1} \cdot \frac{p_j(t^*)}{\phi(j,t^*)} \leq \frac{1}{\eta} \cdot \frac{k+2}{k+1} \cdot p_j \cdot \frac{\mathcal{W}_j(t^*) + w_j}{l_{i^*j}}
\end{aligned}$$

■

Note that the previous two lemmas imply the following by summation:

Lemma II.6. *For any time instant $t \in [r_j, C_j]$ and job j that is assigned to machine $i^* = \sigma(j, t)$, we have:*

$$\begin{aligned}
\Delta_j &= \Delta_j^1(t) + \Delta_j^2(t) \\
&\leq (k+2) \cdot w_j \cdot (t - r_j) + \frac{1}{\eta} \cdot \frac{k+2}{k+1} \cdot p_j \cdot \frac{\mathcal{W}_j(t) + w_j}{l_{i^*j}}
\end{aligned}$$

3) Checking the Feasibility of Constraints

Now it remains to prove that constraints of LP_{dual} are satisfied. To see this, fix job j and time instant t . We consider two cases.

Case 1: Machine $i = \sigma(j, t)$. Then

$$\begin{aligned}
&\alpha_j - \frac{p_j}{\ell_{ij}} \beta_{it} \\
&= \frac{\Delta_j}{k+2} - \frac{p_j}{\ell_{ij}} \cdot \frac{W(i, t)}{k+3} \\
&\leq w_j \cdot (t - r_j) + \frac{p_j}{\eta \cdot (k+1)} \cdot \frac{\mathcal{W}_j(t) + w_j}{\ell_{ij}} \\
&\quad - \frac{p_j}{\ell_{ij}} \cdot \frac{W(i, t)}{k+3} \quad \text{[Lemma II.6]} \\
&\leq w_j \cdot (t - r_j) \quad \text{[since } \eta = 1 + 3\epsilon, k = 1/\epsilon]
\end{aligned}$$

Case 2: Machine $i \neq \sigma(j, t)$. Then

$$\begin{aligned}
&\alpha_j - \frac{p_j}{\ell_{ij}} \beta_{it} \\
&= \frac{\Delta_j}{k+2} - \frac{p_j}{\ell_{ij}} \cdot \frac{W(i, t)}{k+3} \\
&\leq w_j \cdot (t - r_j) + \frac{1}{\eta} \cdot \frac{p_j}{k+1} \cdot \frac{\mathcal{W}_j(t) + w_j}{\ell_{\sigma(j,t)j}} - \frac{p_j}{\ell_{ij}} \cdot \frac{W(i, t)}{k+3} \\
&\quad \text{[Lemma II.6]} \\
&\leq w_j \cdot (t - r_j) \\
&\quad + \left(\frac{p_j}{\eta \cdot (k+1)} \cdot \frac{W(i, t) + w_j}{\ell_{ij}} - \frac{p_j}{k+3} \cdot \frac{W(i, t)}{\ell_{ij}} \right) \\
&\leq w_j \cdot (t - r_j) + \frac{w_j p_j}{\ell_{ij}} \quad \text{[since } \eta = 1 + 3\epsilon, k = 1/\epsilon]
\end{aligned}$$

The penultimate inequality follows since the machine $\sigma(j, t)$ maximizes the virtual utility of job j at time t . Therefore, the dual constraints are satisfied for all time instants t and all jobs j , and we derive that SELFISHMIGRATE is $(1 + \epsilon)$ -speed augmentation, $O(1/\epsilon^2)$ -competitive against $\text{LP}_{\text{primal}}$, completing the proof of the first part of Theorem I.1.

Polynomial Time Algorithm and Minimizing Reassignments. A careful observation of the analysis reveals that to satisfy dual constraints, each job need not be on the machine which gives the highest virtual utility. We can change the policy SELFISHMIGRATE so that a job migrates to a different machine only if its virtual utility increases by a factor of at least $(1 + \epsilon)$. Note that this does not change the monotonicity properties of the virtual utility of a job, hence the entire analysis follows (with the speed augmentation η increased by a factor of $1 + \epsilon$). This also implies that for any job j , the total number of migrations is at most $(\log_{(1+\epsilon)} W + \log_{(1+\epsilon)} n)$, where W is the ratio of the maximum weight of all jobs to the minimum weight. Omitting the simple details, we complete the proof of Theorem I.1.

III. WEIGHTED FLOW-TIME AND ENERGY FOR ARBITRARY POWER FUNCTIONS

In this section we present a simple extension to SELFISHMIGRATE to get a scalable algorithm for minimizing the sum of weighted flow-time and energy for arbitrary power functions. The problem formulation is the same as in Section II, with an added feature. Each machine i can be run at a variable rate $S(i, t)$ by paying an energy cost of $f_i(S(i, t))$, where f_i is a machine dependent, convex increasing function (also called as power function). The rate $S(i, t)$ can be partitioned among the jobs $J_i(t)$, so that $\sum_{j \in J_i(t)} \nu_j(t) \leq S(i, t)$. As before, job $j \in J_i(t)$ runs at speed $q_j(t) = \nu_j(t) \times \ell_{ij}$.

As in Section II, we define the completion time C_j of job j to satisfy $\int_{t=r_j}^{C_j} q_j(t) dt = p_j$. As before, preemption and migration of jobs are allowed without any penalty, but each job must be assigned to a single machine at every instant of time. The scheduler is not aware of the processing lengths p_j . Our objective is to minimize sum of weighted flow-time and energy consumed:

$$\text{Objective} = \sum_j w_j F_j + \sum_i \int_t f_i(S(i, t)) dt$$

In a resource augmentation analysis, we assume that the online algorithm gets $(1 + \epsilon)$ more speed, for any $\epsilon > 0$, for *consuming the same energy*. Alternatively, the offline benchmark has to pay a cost of $f_i((1 + \epsilon)s)$ if it runs machine i at a rate of s . Speed augmentation is required to achieve meaningful competitive ratios for the case of arbitrary power functions. To elaborate on this point, consider a function $f(s)$ that takes an infinitesimal value in the interval $0 \leq s \leq 1$ and sharply increases when $s > 1$. For such a power function, any competitive online scheduler has to be optimal at each instant of time unless we give it more resources. A scalable algorithm in the speed augmentation setting implies algorithms with small competitive ratios when the energy cost function can be approximated by

polynomials. In particular, the result translates to an $O(\gamma^2)$ -competitive algorithm (without any resource augmentation) when the power function is $f_i(s) = s^\gamma$.

Let g_i be the inverse of power function f_i . Note that g is an increasing concave function. Before we describe our algorithm, we make the following simple observation regarding concave functions.

Proposition III.1. *For any increasing concave function g with $g(0) = 0$, $\frac{g(w)}{w}$ is decreasing in w .*

A. The SELFISHMIGRATE-ENERGY Algorithm

Our algorithm SELFISHMIGRATE-ENERGY, is very similar to the algorithm SELFISHMIGRATE, and we only outline the differences with Section II. The most important difference is the policy that sets the speeds of the machines.

Speed Scaling Policy:

We set the speed of machine i at time t , denoted by $S(i, t)$, such that the total energy cost is equal to the total weight of jobs at time t .

$$\begin{aligned} f_i(S(i, t)) &= W(i, t) \quad \text{or equivalently,} \\ S(i, t) &= g_i(W(i, t)) \end{aligned} \quad (3)$$

This is same as the speed scaling policy used in [7], [3].

Single Machine Policy.

The remaining components of the algorithm remain similar to SELFISHMIGRATE algorithm. We briefly mention the differences.

Each machine runs WRPS (k) where $k = \frac{1}{\epsilon}$. In this policy, the notions of virtual queues, rank of a job, and the arrival and departure policies (with associated notation) – remain the same. In particular, a job that arrives or migrates to a machine are placed at the tail of the virtual queue and assigned the highest rank on the machine. At time instant t , the total processing rate $S(i, t) = g_i(W(i, t))$ of the machine i is divided among the jobs in $J_i(t)$ as follows.

$$\nu_j(t) := g_i(W(i, t)) \cdot \frac{(\mathcal{W}_j(t) + w_j)^{k+1} - \mathcal{W}_j^{k+1}}{W(i, t)^{k+1}} \quad (4)$$

As before, this implies job $j \in J_i(t)$ is processed at rate $\ell_{ij} \nu_j(t)$.

Virtual Utility of Jobs and Selfish Migration.

Consider a job $j \in J_i(t)$ at time t . Its virtual utility is defined as:

$$\phi(j, t) = g_i(\mathcal{W}_j(t) + w_j) \cdot \frac{\ell_{ij}}{\mathcal{W}_j(t) + w_j}$$

Using this virtual utility, the jobs perform sequential best response dynamics, migrating to other machines if it improves its virtual utility. As before, this leads to a Nash equilibrium

every step.² If a job moves out of a machine, the weights $\mathcal{W}_{j'}(t)$ of other jobs on the machine either stay the same or decrease. Using Proposition III.1, this implies the virtual utility of other jobs on the machine either remains the same or increases. Therefore, similar to Lemma II.1, we easily get the monotonicity of the virtual utilities of jobs.

Lemma III.2. *For all jobs j , $\phi(j, t)$ is non-decreasing over the interval $t \in [r_j, C_j]$.*

Due to space constraints we omit the analysis of Theorem I.2. The analysis uses a convex programming relaxation due to [3], [15], but the overall flow of the analysis remains very similar. In particular, the monotonicity property of $\phi(j, t)$ again plays a crucial role in the dual fitting. We note that Theorem I.2 immediately implies the following corollary.

Corollary III.3. *There is $O(\gamma^2)$ -competitive non-clairvoyant algorithm for minimizing weighted flow-time plus energy, when each machine follows a polynomial power function $f(s) = s^\gamma$.*

We will include the complete analysis in the full version of this paper.

Acknowledgments

The authors thank Naveen Garg, Benjamin Moseley, Ravishankar Krishnaswamy, Anupam Gupta, Minghong Lin, Michael Nugent, and Neal Barcelo for initial discussions about this line of research over the last few years.

Im was supported by NSF award CCF-1409130. Kulkarni was supported by NSF awards CCF-1008065 and IIS-0964560. Munagala was supported by NSF awards CCF-0745761, CCF-1348696, and CCF-1408784. Pruhs was supported in part by NSF grants CCF-1115575, CNS-1253218, CCF-1421508, and an IBM Faculty Award.

REFERENCES

- [1] Susanne Albers. Energy-efficient algorithms. *Commun. ACM*, 53(5):86–96, 2010.
- [2] Susanne Albers and Hiroshi Fujiwara. Energy-efficient algorithms for flow time minimization. *ACM Transactions on Algorithms*, 3(4), 2007.
- [3] S. Anand, Naveen Garg, and Amit Kumar. Resource augmentation for weighted flow-time explained by dual fitting. In *SODA*, pages 1228–1241, 2012.
- [4] Nikhil Bansal and Ho-Leung Chan. Weighted flow time does not admit $o(1)$ -competitive algorithms. In *SODA*, pages 1238–1244, 2009.
- [5] Nikhil Bansal, Ravishankar Krishnaswamy, and Viswanath Nagarajan. Better scalable algorithms for broadcast scheduling. In *ICALP (I)*, pages 324–335, 2010.

²Note that unlike Section II where migrations can only happen on departure of jobs from the system, migrations can now also happen on arrival of jobs into the system.

- [6] Nikhil Bansal and Kirk Pruhs. Server scheduling in the weighted ℓ_p norm. In *LATIN*, pages 434–443, 2004.
- [7] Nikhil Bansal, Kirk Pruhs, and Clifford Stein. Speed scaling for weighted flow time. *SIAM J. Comput.*, 39(4):1294–1308, 2009.
- [8] Luca Becchetti and Stefano Leonardi. Nonclairvoyant scheduling to minimize the total flow time on single and parallel machines. *J. ACM*, 51(4):517–539, 2004.
- [9] Sayan Bhattacharya, Sungjin Im, Janardhan Kulkarni, and Kamesh Munagala. Coordination mechanisms from (almost) all scheduling policies. In *ITCS*, pages 121–134, 2014.
- [10] Fred A. Bower, Daniel J. Sorin, and Landon P. Cox. The impact of dynamically heterogeneous multicore processors on thread scheduling. *IEEE Micro*, 28(3):17–25, May 2008.
- [11] J. S. Chadha, N. Garg, A. Kumar, and V. N. Muralidhara. A competitive algorithm for minimizing weighted flow time on unrelated machines with speed augmentation. In *STOC*, 2009.
- [12] Sze-Hang Chan, Tak-Wah Lam, Lap-Kei Lee, and Jianqiao Zhu. Nonclairvoyant sleep management and flow-time scheduling on multiple processors. In *Proceedings of the Twenty-fifth Annual ACM Symposium on Parallelism in Algorithms and Architectures*, SPAA '13, pages 261–270, New York, NY, USA, 2013. ACM.
- [13] Chandra Chekuri, Ashish Goel, Sanjeev Khanna, and Amit Kumar. Multi-processor scheduling to minimize flow time with epsilon resource augmentation. In *STOC*, pages 363–372, 2004.
- [14] G. Christodoulou, E. Koutsoupias, and A. Nanavati. Coordination mechanisms. *Theor. Comput. Sci.*, 410(36):3327–3336, August 2009.
- [15] Nikhil R. Devanur and Zhiyi Huang. Primal dual gives almost optimal energy efficient online algorithms. In *SODA*, pages 1123–1140, 2014.
- [16] Jeff Edmonds, Sungjin Im, and Benjamin Moseley. Online scalable scheduling for the ℓ_k -norms of flow time without conservation of work. In *ACM-SIAM Symposium on Discrete Algorithms*, 2011.
- [17] Jeff Edmonds and Kirk Pruhs. Scalably scheduling processes with arbitrary speedup curves. *ACM Transactions on Algorithms*, 8(3):28, 2012.
- [18] Hadi Esmaeilzadeh, Emily R. Blem, Renée St. Amant, Karthikeyan Sankaralingam, and Doug Burger. Dark silicon and the end of multicore scaling. In *ISCA*, pages 365–376, 2011.
- [19] Naveen Garg and Amit Kumar. Minimizing average flow-time : Upper and lower bounds. In *FOCS*, pages 603–613, 2007.
- [20] Anupam Gupta, Sungjin Im, Ravishankar Krishnaswamy, Benjamin Moseley, and Kirk Pruhs. Scheduling heterogeneous processors isn't as easy as you think. In *SODA*, pages 1242–1253, 2012.
- [21] Anupam Gupta, Ravishankar Krishnaswamy, and Kirk Pruhs. Nonclairvoyantly scheduling power-heterogeneous processors. In *Green Computing Conference*, pages 165–173, 2010.
- [22] Anupam Gupta, Ravishankar Krishnaswamy, and Kirk Pruhs. Scalably scheduling power-heterogeneous processors. In *ICALP (I)*, pages 312–323, 2010.
- [23] Anupam Gupta, Ravishankar Krishnaswamy, and Kirk Pruhs. Online primal-dual for non-linear optimization with applications to speed scaling. In *WAOA*, pages 173–186, 2012.
- [24] Anupam Gupta, Amit Kumar, and Cliff Stein. Maintaining assignments online: Matching, scheduling, and flows. In *SODA*, pages 468–479, 2014.
- [25] Sungjin Im, Janardhan Kulkarni, and Kamesh Munagala. Competitive algorithms from competitive equilibria: Non-clairvoyant scheduling under polyhedral constraints. In *STOC*, 2014 (to appear).
- [26] Sungjin Im and Benjamin Moseley. Online scalable algorithm for minimizing ℓ_k -norms of weighted flow time on unrelated machines. In *SODA*, pages 95–108, 2011.
- [27] Sungjin Im, Benjamin Moseley, and Kirk Pruhs. A tutorial on amortized local competitiveness in online scheduling. *SIGACT News*, 42(2):83–97, 2011.
- [28] Bala Kalyanasundaram and Kirk Pruhs. Speed is as powerful as clairvoyance. *Journal of the ACM*, 47(4):617–643, 2000.
- [29] David Koufaty, Dheeraj Reddy, and Scott Hahn. Bias scheduling in heterogeneous multi-core architectures. In *Proceedings of the 5th European Conference on Computer Systems*, EuroSys '10, pages 125–138, New York, NY, USA, 2010. ACM.
- [30] Rakesh Kumar, Dean M. Tullsen, and Norman P. Jouppi. Core architecture optimization for heterogeneous chip multiprocessors. In *Proceedings of the 15th International Conference on Parallel Architectures and Compilation Techniques*, PACT '06, pages 23–32, New York, NY, USA, 2006. ACM.
- [31] Stefano Leonardi and Danny Raz. Approximating total flow time on parallel machines. *J. Comput. Syst. Sci.*, 73(6):875–891, 2007.
- [32] R. Merritt. CPU designers debate multi-core future. *EE Times*, Feb. 2010.
- [33] Kirk Pruhs, Jiri Sgall, and Eric Torng. *Handbook of Scheduling: Algorithms, Models, and Performance Analysis*, chapter Online Scheduling. 2004.
- [34] Nguyen Kim Thang. Lagrangian duality in online scheduling with resource augmentation and speed scaling. In *ESA*, pages 755–766, 2013.