Math-Net.Ru

All Russian mathematical portal

A. Kononov, Yu. Kovalenko, Approximate schedules for non-migratory parallel jobs in speed-scaled multiprocessor systems, *Sib. Èlektron. Mat. Izv.*, 2019, Volume 16, 249–257

DOI: <https://doi.org/10.33048/semi.2019.16.016>

Use of the all-Russian mathematical portal Math-Net.Ru implies that you have read and agreed to these terms of use

<http://www.mathnet.ru/eng/agreement>

Download details:

IP: 34.229.108.108

April 26, 2023, 13:56:46

СИБИРСКИЕ ЭЛЕКТРОННЫЕ
МАТЕМАТИЧЕСКИЕ ИЗВЕСТИЯ

Siberian Electronic Mathematical Reports

<http://semr.math.nsc.ru>

Том 16, стр. 249–257 (2019)

DOI 10.33048/semi.2019.16.016

УДК 519.8

MSC 90B35

APPROXIMATE SCHEDULES FOR NON-MIGRATORY
PARALLEL JOBS IN SPEED-SCALED MULTIPROCESSOR
SYSTEMS

A. KONONOV, YU. KOVALENKO

ABSTRACT. We consider a problem of scheduling rigid parallel jobs on variable speed processors so as to minimize the total energy consumption. Each job is specified by its processing volume and the required number of processors. We propose new constant factor approximation algorithms for the non-migratory cases when all jobs have a common release time and/or a common deadline.

Keywords: approximation algorithm, speed scaling, schedule, parallel job, migration.

1. INTRODUCTION

Modern processors can vary their speeds dynamically when jobs are executed. The instantaneous power required to run a job at speed s is defined by s^α , where $\alpha > 1$ is a constant. The energy consumption is the power integrated over time. We assume that the release times and deadlines are given for jobs and the aim is to calculate the speeds of the jobs and to construct a feasible schedule so as to minimize the total energy consumption. There exist various variants of the speed scaling scheduling, depending on the type of jobs and processors, and system characteristics. One of the algorithmic and complexity study of this area is devoted to revising classical scheduling problems with dynamic speed scaling (see e.g. [2, 4, 8, 10, 12, 20, 21] and others).

In this paper we study some particular cases of the speed scaling scheduling of parallel jobs, each of them requires several processors simultaneously [11]. The

KONONOV, A., KOVALENKO, YU., APPROXIMATE SCHEDULES FOR NON-MIGRATORY PARALLEL JOBS IN SPEED-SCALED MULTIPROCESSOR SYSTEMS.

© 2019 KONONOV A., KOVALENKO YU.

This work was supported by Russian Foundation for Basic Research (project N 17-07-00513).

Received June, 3, 2018, published February, 26, 2019.

motivation to consider parallel jobs consists in the fact that some jobs can not be performed asynchronously on modern computers. Such situation takes place in testing and reliable computing, parallel applications on graphics cards, computer control systems and others.

The energy effective scheduling has been widely investigated for single-processor jobs (see e.g. reviews [1, 12]). The preemptive single-processor setting is polynomially solvable. The algorithms developed in [18, 20] have running time $O(n^2)$. However, the same problem without preemptions is NP-hard [5] even in the case of tree structured jobs. Antoniadis et al. [5] proposed an $2^{4\alpha-3}$ -approximation algorithm for this case and an $2^{5\alpha-4}$ -approximation algorithm for the general non-preemptive case on one processor. Later, Bampis et al. [7] presented an algorithm that achieves $2^{\alpha-1}(1+\varepsilon)^\alpha \tilde{B}^\alpha$ -approximation for the latter problem, where $\tilde{B}_\alpha$ is the generalized Bell number.

For the migratory problem, where m parallel processors are available and all jobs are of the single-processor type, polynomial time algorithms were given in [2, 4, 8, 20]. As far as we know, the algorithm [20] has the best time complexity $O(n^3)$ among the algorithms above.

The preemptive multiprocessor setting without migration is NP-hard in the strong sense as proved in [3]. Albers et al. [3] provided $(\alpha^\alpha 2^{4\alpha})$ -approximation algorithms for jobs with unit works and arbitrary deadlines, or arbitrary works and agreeable deadlines, and an $2(2 - \frac{1}{m})^\alpha$ -approximation algorithm for jobs with common release times, or common deadlines. Moreover, they showed that the problem with unit works is polynomially solvable for agreeable deadlines. Chen et al. [9] proposed a greedy algorithm with an 1.13 approximation guarantee for the case when all jobs must be executed in one time interval. In [6], for non-preemptive instances an $m^\alpha (\sqrt[n]{n})^{\alpha-1}$ -approximation algorithm has been presented, that explore the idea of transforming an optimal preemptive schedule to a non-preemptive one.

Scheduling of multiprocessor jobs has been extensively investigated in the case of regular time criteria (see e.g. the book of Drozdowski [11]), but for the criterion of energy minimization it is poorly studied. Recently, an approximation algorithm has been proposed in [15] for the speed scaling scheduling of rigid parallel jobs with migration. The algorithm returns a solution within an additive error $\varepsilon > 0$ and runs in time polynomial in m , $1/\varepsilon$ and the input size. Note that this algorithm is pseudopolynomial and it is based on solving a linear configuration program using the Ellipsoid method. In [16], we developed a strongly polynomial algorithm that achieves $(2 - \frac{1}{m})^{\alpha-1}$ -approximation ratio for the same problem. As well we showed that most of the NP-hardness proofs for scheduling problems with the maximum lateness criterion may be easily transformed to their speed scaling counterparts [15].

2. PROBLEM STATEMENT AND OUR RESULTS

We assume that a computer system consists of m parallel identical processors, which can dynamically change the speed. A set $\mathcal{J} = \{1, \dots, n\}$ of parallel jobs is given. Each job $j \in \mathcal{J}$ has a release time r_j , a deadline d_j and a processing volume (work) W_j . The number of processors simultaneously required by job j is called *job size* and denoted by $size_j$. Any subset of parallel processors of the given size can be used to execute job j . Jobs with such property are called *rigid jobs* [11]. Migration

of a job j among different subsets of $size_j$ processors is disallowed. Job preemption might or might not be allowed in the exploring of scheduling in this paper.

The speed of a job j is the rate at which its work is completed. A continuous spectrum of processor speeds is available. The power consumed when running at speed s is s^α , where α is a constant close to 3 [12]. The energy used is power integrated over time. Each of m processors may operate at variable speed, but if processors execute the same job simultaneously then all these processors run at the same speed.

A schedule is required to be feasible in the sense that each processor executes at most one job at the time and each job is processed in the required work between its release time and deadline. The problem is to find a feasible schedule that minimizes the total energy consumed on all the processors. The preemptive and nonpreemptive variants of non-migration energy-efficient scheduling of rigid jobs [7, 11] are denoted by $P|size_j, pmtn*, r_j, d_j|E$ and $P|size_j, r_j, d_j|E$, respectively.

Let us assume that all jobs have a common release time r and/or a common deadline d . We present strongly polynomial time algorithms achieving constant factor approximation guarantees for these particular cases. Our algorithms consist of two stages. At the first stage we obtain a lower bound on the minimal energy consumption and calculate intermediate execution times of jobs. Then, at the second stage, we determine final speeds of jobs and schedule them.

A lower bound on the objective function and intermediate processing times of jobs can be found in $O(n^3)$ time using the method developed in [16]. The method is based on a reduction of the speed scaling problem to the special min-cost max-flow problem [20]. Here we propose more effective approaches for the considered problem instances.

Problems $P|size_j, pmtn*, r_j = r, d_j = d|E$ and $P|size_j, r_j = r, d_j = d|E$ are strongly NP-hard even in the case of single-processor jobs with arbitrary processing volumes as proved in [3]. Simple reductions from 2-PARTITION and 3-PARTITION imply that problem $P|size_j, pmtn*, r_j = r, d_j = d|E$ is ordinary NP-hard and problem $P|size_j, r_j = r, d_j = d|E$ is NP-hard in the strong sense even if all jobs have unit processing volumes. Using the approach from [15] (Section 3. NP-Hardness Results), we can show that problem $P|size_j, r_j = r, d_j|E$ is strongly NP-hard even in the case of two processors. The proof is almost a step by step reproduction of the NP-hardness proof for $P2|size_j, d_j|L_{\max}$ (see Theorem A1 in [17]).

The paper is organized as follows. In Section 3 we propose an $(2 - \frac{1}{m})^{\alpha-1}$ -approximation algorithm for the non-preemptive scheduling where all jobs are available in one time interval. In Section 4, $(3 - \frac{4}{m+1})^{\alpha-1}$ -approximate schedules are constructed for the preemptive and non-preemptive problems with jobs sharing a common release time (or symmetrically, a common deadline). The last section contains the concluding remarks.

3. COMMON RELEASE DATE AND DEADLINE

In this section we consider the non-preemptive case of the problem where all jobs arrive at time $r = 0$ and have a shared global deadline d .

The first stage. Now we find auxiliary durations of jobs and a lower bound on the objective function in $O(n \log n)$ time. Let m' denote the current number

of unoccupied processors and $\mathcal{J}'$ be the set of currently considered jobs. Initially $\mathcal{J}' := \mathcal{J}$ and $m' := m$.

We enumerate the jobs one by one in order of non-increasing works. If the current job i has $W_i \geq \frac{\sum_{j \in \mathcal{J}'} W_j \text{size}_j}{m'}$, then we assign duration $p_i := d$ for this job, and set $\mathcal{J}' := \mathcal{J}' \setminus \{i\}$ and $m' := m' - \text{size}_i$. After that we go to the next job. Otherwise, all jobs $l \in \mathcal{J}'$ satisfy the inequality $W_l < \frac{\sum_{j \in \mathcal{J}'} W_j \text{size}_j}{m'}$, and we assign durations $p_l := \frac{W_l m' d}{\sum_{j \in \mathcal{J}'} W_j \text{size}_j}$ for them.

The presented approach guarantees that $p_j \leq d$, $\sum_{j \in \mathcal{J}} p_j \text{size}_j = md$, and gives the lower bound on the objective function equal to $\sum_{j \in \mathcal{J}} \text{size}_j W_j^\alpha p_j^{1-\alpha}$. At the second stage we use the “non-preemptive list-scheduling” algorithm [19] to construct a feasible schedule in interval $[0, d]$.

The second stage. Whenever a subset of processors falls idle, the “non-preemptive list-scheduling” algorithm schedules a job that does not require more processors than are available, until all jobs in $\mathcal{J}$ are assigned. The time complexity of the algorithm is $O(n^2)$.

We claim that the length of the constructed schedule is at most $(2 - \frac{1}{m})d$ (see Lemma 1 below). By increasing the speed of each job in $(2 - \frac{1}{m})$ times we obtain a schedule of the length at most d . The total energy consumption is increased by a factor $(2 - \frac{1}{m})^{\alpha-1}$ in comparison with the lower bound. As a result, we have

Theorem 1. *A $(2 - \frac{1}{m})^{\alpha-1}$ -approximate schedule can be found in $O(n^2)$ time for problems $P|\text{size}_j, \text{pmtn}^*, r_j = r, d_j = d|E$ and $P|\text{size}_j, r_j = r, d_j = d|E$.*

Using the results from [14], we conclude that the approximation ratio of $(2 - \frac{1}{m})$ for the “non-preemptive list-scheduling” algorithm is tight even if all jobs have single-processor type. As a result, the energy consumption is increased in $(2 - \frac{1}{m})^{\alpha-1}$ times when we put the resulting schedule inside the interval $[0, d]$. Therefore, the approximation ratio of our algorithm is also tight.

Lemma 1. *Given m processors, an interval $[0, d]$, and a set of jobs $\mathcal{J}$ with processing times $p_j \leq d$ and sizes size_j , where $\sum_{j \in \mathcal{J}} p_j \text{size}_j \leq md$. The length of the schedule S constructed by the “non-preemptive list-scheduling” algorithm is at most $(2 - \frac{1}{m})d$.*

Proof. Let T denote the length of schedule S . If at least $\frac{m+1}{2}$ processors are used at any time instance in S , we have

$$d \geq \frac{1}{m} \sum_{j \in \mathcal{J}} p_j \text{size}_j \geq \frac{m+1}{2m} T \geq \frac{T}{2 - 1/m}.$$

Otherwise, assume that I is the last time interval of schedule S with $m_I < \frac{m+1}{2}$ processors being used during I . By the construction of S there is a job j that is performed during the whole interval I . Let C_j be the completion time of j in S . It is easy to see that at every point in time during interval $[0, C_j - p_j)$ schedule S uses at least $m - \text{size}_j + 1 \geq m - m_I + 1$ processors (otherwise job j should be started earlier). Moreover, at least m_I processors are utilized in interval $[C_j - p_j, C_j)$, therefore, each job executed in interval $[C_j, T)$ requires no less than $m - m_I + 1$ processors. Thus, the total load of all processors $\sum_{j \in \mathcal{J}} p_j \text{size}_j$ is at least $m_I p_j + (T - p_j)(m - m_I + 1) \leq dm$. If $p_j \geq \frac{T}{2 - \frac{1}{m - m_I + 1}}$, then $\frac{T}{2 - \frac{1}{m}} \leq \frac{T}{2 - \frac{1}{m - m_I + 1}} \leq p_j \leq d$.

Otherwise as $1 \leq m_I \leq \frac{m}{2} < \frac{m+1}{2}$ we have

$$\begin{aligned} d &\geq \frac{1}{m} ((m - m_I + 1)T - p_j(m - 2m_I + 1)) \geq \frac{T}{m} \left(m - m_I + 1 - \frac{m - 2m_I + 1}{2 - \frac{1}{m - m_I + 1}} \right) \\ &= \frac{T(m - m_I + 1)}{(2m - 2m_I + 1)} \geq \frac{mT}{2m - 1} = \frac{T}{2 - \frac{1}{m}}. \end{aligned}$$

□

4. COMMON RELEASE DATE OR DEADLINE

In this section we study the problem without migration where all jobs are released at time $r = 0$ but have individual deadlines.

4.1. Preemptive Problem. Here we consider the case when the preemption of jobs is allowed.

The first stage. It can easily be checked that the optimal energy consumption of rigid jobs with sizes $size_j$ and works W_j on m processors is at least $\frac{1}{m^{\alpha-1}}$ times that of single-processor jobs with works $size_j W_j$ on one processor (the result is proved similarly to Lemma 1 in [3]). So, if we find an optimal solution of the latter problem, and decrease the speeds of jobs in m times, then a lower bound on the energy consumption of rigid jobs is obtained. However such approach may lead to the execution time of a job i greater than d_i .

In [21], Yao et al. showed that the preemptive problem on a single processor is solvable in polynomial time. They proposed an efficient algorithm called *YSD* that repeatedly identifies time intervals of highest density. The density of an interval I is the total work released and to be completed in I divided by the length of I . The algorithm repeatedly schedules jobs in highest density intervals and takes care of reduced subproblems.

We propose a modification of the algorithm *YSD* [21] to obtaining a lower bound on the minimal energy consumption for the considered problem. At the same time, for each rigid job $i \in \mathcal{J}$ we find an intermediate duration p_i such that

$$(1) \quad p_i \leq d_i,$$

$$(2) \quad \sum_{j \in \mathcal{J}_i} p_j size_j \leq m d_i, \quad \mathcal{J}_i = \{j \in \mathcal{J} : d_j \leq d_i\}.$$

Now we construct a special schedule for jobs of works $size_j W_j$ on one processor, which will assure conditions (1) and (2) for the corresponding rigid jobs. Let $J(t, t') := \{j \in \mathcal{J} : t \leq r_j < d_j \leq t'\}$ be the jobs which need to be processed in some time interval $[t, t')$. Initially we have intervals of the form $[0, d_j)$, $j \in \mathcal{J}$, as all jobs release at time $r = 0$.

Modified YSD Algorithm

Step 1. Repeat steps 1.1 and 1.2 until $\mathcal{J} = \emptyset$:

Step 1.1. Let $[t, t')$ be the interval with maximum density, i.e., that maximizes

$$\frac{\sum_{j \in J(t, t')} \text{size}_j W_j}{t' - t}.$$

Step 1.2. If the inequality

$$(3) \quad \frac{W_i}{d_i} \leq \frac{\sum_{j \in J(t, t')} \text{size}_j W_j}{(t' - t)m}$$

holds for all $i \in J(t, t')$, then process these jobs in interval $[t, t')$ with speed equal to the maximum density, i.e., set processing time $p_i^1 := \frac{\text{size}_i W_i (t' - t)}{\sum_{j \in J(t, t')} \text{size}_j W_j}$ for each job $i \in J(t, t')$. Then remove the jobs $J(t, t')$ from $\mathcal{J}$, and adjust the remaining jobs as if the time interval $[t, t')$ does not exist, i.e., set $r_j := t'$ for each job $j \in \mathcal{J}$. Endpoints and densities of intervals are updated for Step 1.1.

Otherwise enumerate all jobs from $J(t, t')$ for which inequality (3) is violated. The current job i is assigned in interval $[d_i - \frac{d_i \text{size}_i}{m}, d_i)$, i.e., $p_i^1 := \frac{d_i \text{size}_i}{m}$. Then remove job i from $\mathcal{J}$ and delete interval $[d_i - p_i^1, d_i)$ from the further consideration, i.e., set $d_j := d_j - p_i^1$ for $d_j \geq d_i$ and $d_j := d_i - p_i^1$ for $d_i - p_i^1 \leq d_j < d_i$.

Step 2. Return the resulting durations p_i^1 of jobs $j \in \mathcal{J}$.

At least one job is deleted at each call of Step 1.2. Removing a job requires $O(n)$ additional operations to update information about the remaining jobs and intervals. Therefore, the running time of the algorithm is $O(n^2)$.

The intermediate processing times p_i of rigid jobs $i \in \mathcal{J}$ are calculated as $\frac{p_i^1 m}{\text{size}_i}$, and give the lower bound $\sum_{j \in \mathcal{J}} \text{size}_j W_j^\alpha p_j^{1-\alpha}$ on the energy consumption. The conditions (1) and (2) hold for the computed processing times.

At the second stage we use the “preemptive earliest deadline list-scheduling” algorithm to construct a feasible schedule of problem $P|\text{size}_j, \text{pmt}n^*, r_j = 0, d_j|E$.

The second stage. The “preemptive earliest deadline list-scheduling” algorithm schedules jobs in order of non-decreasing deadlines as follows. If $\text{size}_i > \frac{m}{2}$, then job i is assigned at the end of the current schedule. Otherwise we start job i at the earliest time instant when size_i processors are idle and process it during p_i time, ignoring intervals of jobs with $\text{size}_j > \frac{m}{2}$. The time complexity of the algorithm is $O(n^2)$.

We claim that the completion time C_j of each job j in the constructed schedule is at most $\left(3 - \frac{4}{m+1}\right) d_j$ (see Lemma 2 below). Hence an increasing of the speeds in $\left(3 - \frac{4}{m+1}\right)$ times yields a feasible schedule. The total energy consumption is increased by a factor $\left(3 - \frac{4}{m+1}\right)^{\alpha-1}$.

Obviously, by interchanging release times and deadlines, the presented algorithm can also handle to the case of jobs with individual release times but a common deadline. As a result, we have

Theorem 2. A $\left(3 - \frac{4}{m+1}\right)^{\alpha-1}$ -approximate schedule can be found in $O(n^2)$ time for problems $P|\text{size}_j, \text{pmt}n^*, r_j = r, d_j|E$ and $P|\text{size}_j, \text{pmt}n^*, r_j, d_j = d|E$.

Lemma 2. Given m processors and a set of jobs $\mathcal{J}$ with deadlines d_i , processing times $p_i \leq d_i$ and sizes size_i , where $\sum_{j \in \mathcal{J}_i} p_j \text{size}_j \leq m d_i$ for each $i \in \mathcal{J}$ with

$\mathcal{J}_i = \{j \in \mathcal{J} : d_j \leq d_i\}$. The completion time C_i is at most $\left(3 - \frac{4}{m+1}\right) d_i$ for each job $i \in \mathcal{J}$ in the schedule S constructed by the “preemptive earliest deadline list-scheduling” algorithm.

Proof. We consider an arbitrary deadline d_i , where job i has the maximal completion time C_i in schedule S among all jobs with deadline equal to d_i .

Note that $C_j \leq C_i$ for all $j \in \mathcal{J}_i$. Let S_i denote the part of schedule S , which contains only jobs from $\mathcal{J}_i$ and occupies interval $[0, C_i)$. We will show that $C_i \leq \left(3 - \frac{4}{m+1}\right) d_i$.

If at least $\frac{m+1}{2}$ processors are used at any time instance in subschedule S_i , we have

$$d_i \geq \frac{1}{m} \sum_{j \in \mathcal{J}_i} p_j \text{size}_j \geq \frac{m+1}{2m} C_i \geq \frac{C_i}{2 - \frac{1}{m}}.$$

Otherwise assume that l is the last job in subschedule S_i , that requires $\text{size}_l \leq \frac{m}{2}$ processors. It is easy to see that all time-slots in intervals $[0, C_l - p_l)$ and $[C_l, C_i)$ use at least $\frac{m+1}{2}$ processors, and at least size_l processors are utilized in interval $[C_l - p_l, C_l)$. Therefore, the total load of all processors in subschedule S_i is at least

$$\frac{m+1}{2} (C_i - p_l) + \text{size}_l p_l \leq \sum_{j \in \mathcal{J}_i} p_j \text{size}_j \leq d_i m.$$

If $p_l \geq \frac{C_i}{3 - \frac{4}{m+1}}$, then $C_i \leq \left(3 - \frac{4}{m+1}\right) d_l \leq \left(3 - \frac{4}{m+1}\right) d_i$.

Otherwise as $\text{size}_l \geq 1$ we have

$$\begin{aligned} m d_i &\geq \frac{m+1}{2} C_i - p_l \left(\frac{m+1}{2} - \text{size}_l \right) \geq \frac{m+1}{2} C_i - p_l \frac{m-1}{2} \\ &\geq C_i \left(\frac{m+1}{2} - \frac{m-1}{2 \left(3 - \frac{4}{m+1}\right)} \right) = C_i \frac{(m+1)(3m-1) - (m-1)(m+1)}{2(3m-1)} \\ &= C_i \frac{m(m+1)}{3m-1} = \frac{m C_i}{3 - \frac{4}{m+1}} \end{aligned}$$

□

4.2. Non-Preemptive Problem. In this subsection we consider the case when preemptions are disallowed and each job requires no more than $\frac{m}{2}$ processors.

A lower bound on the objective function and intermediate processing times of jobs are calculated as in Subsection 4.1. But at the second stage we use the “non-preemptive earliest deadline list-scheduling” algorithm to construct a schedule. This algorithm assigns jobs in the schedule as soon as possible in order of non-decreasing deadlines, and its time complexity is also $O(n^2)$.

The property of Lemma 2 holds as well in the presence of the “non-preemptive earliest deadline list-scheduling” algorithm for problem $P|\text{size}_j, r_j = 0, d_j|E$ with $\text{size}_j \leq \frac{m}{2}$. So, an increasing of the job speeds in $\left(3 - \frac{4}{m+1}\right)$ times in the constructed schedule leads to a feasible solution. Thus the following theorem holds.

Theorem 3. $A \left(3 - \frac{4}{m+1}\right)^{\alpha-1}$ -approximate schedule can be found in $O(n^2)$ time for problems $P|size_j, r_j = r, d_j|E$ and $P|size_j, r_j, d_j = d|E$ with $size_j \leq \frac{m}{2}$.

5. CONCLUSION

We have studied the energy minimization under a global release time or a global deadline constraint. Strongly polynomial time approximation algorithms are developed for the rigid jobs with no migration. Our algorithms have constant factor approximation guarantees.

Further research might address the approaches to the problems with more complex structure, where processors are heterogeneous and jobs have alternative execution modes with different characteristics.

REFERENCES

- [1] S. Albers, *Energy-efficient algorithms*, Communications of the ACM, **53**:5 (2010), 86–96. DOI: 10.1145/1735223.1735245 MR2364973 .
- [2] S. Albers, A. Antoniadis, G. Greiner, *On multi-processor speed scaling with migration*, Journal of Computer and System Sciences, **81** (2015), 1194–1209. DOI: 10.1016/j.jcss.2015.03.001 MR3354788
- [3] S. Albers, F. Müller, S. Schmelzer, *Speed scaling on parallel processors*, Algorithmica, **68**:2 (2007), 404–425. DOI: 10.1007/s00453-012-9678-7
- [4] E. Angel, E. Bampis, F. Kacem, D. Letsios, *Speed scaling on parallel processors with migration*, 18th International European Conference on Parallel and Distributed Computing, Lecture Notes in Computer Science, **7484** (2012), 128–140. DOI: 10.1007/978-3-642-32820-6_15; Zbl 1312.68029
- [5] A. Antoniadis, C.C. Huang, *Non-preemptive speed scaling*, J. Sched., **16**: 4 (2013), 385–394. DOI: 10.1007/s10951-013-0312-6; MR3079660
- [6] E. Bampis, A. Kononov, D. Letsios, G. Lucarelli, I. Nemparis, *From preemptive to non-preemptive speed-scaling scheduling*, Discrete Applied Mathematics, **181** (2015), 11–20. DOI: 10.1016/j.dam.2014.10.007; MR3284507
- [7] E. Bampis, A. Kononov, D. Letsios, G. Lucarelli, M. Sviridenko, *Energy efficient scheduling and routing via randomized rounding*, J. Sched., **21**: 1 (2018), 35–51. DOI: 10.1007/s10951-016-0500-2; MR3763430
- [8] B.D. Bingham, M.R. Greenstreet, *Energy optimal scheduling on multiprocessors with migration*, International Symposium on Parallel and Distributed Processing with Applications (ISPA’08), 2018, IEEE, 153–161. DOI: 10.1109/ISPA.2008.128
- [9] J. Chen, H. Hsu, K. Chuang, C. Yang, A. Pang, T. Kuo, *Multiprocessor energy-efficient scheduling with task migration considerations*, 16th Euromicro Conference on Real-Time Systems (ECRTS2004), 2004, IEEE, 101–108. DOI: 10.1109/EMRTS.2004.1311011
- [10] V. Cohen-Addad, Z. Li, C. Mathieu, I. Milis, *Energy-efficient algorithms for non-preemptive speed-scaling*, International Workshop on Approximation and Online Algorithms (WAOA 2014), Lecture Notes in Computer Science, **8952** (2015), 107–118. Zbl 06512659; DOI: 10.1007/978-3-319-18263-6_10
- [11] M. Drozdowski, *Scheduling for Parallel Processing*, London: Springer-Verlag, 2009. Zbl 1187.68090
- [12] M.E.T Gerards, J.L. Hurink, P.K.F. Hölzenspies, *A survey of offline algorithms for energy minimization under deadline constraints*, J. Sched., **19** (2016), 3–19. DOI: 10.1007/s10951-015-0463-8; Zbl 1341.90046
- [13] G. Greiner, T. Nonner, A. Souza, *The bell is ringing in speed-scaled multiprocessor scheduling*, Theory of Computing Systems, **54**:1 (2014), 24–44. DOI: 10.1007/s00224-013-9477-9; Zbl 1310.68048
- [14] B. Johannes, *Scheduling parallel jobs to minimize the makespan*, J. Sched., **9** (2006), 433–452. DOI: 10.1007/s10951-006-8497-6; Zbl 1154.90464

- [15] A. Kononov, Y. Kovalenko, *On speed scaling scheduling of parallel jobs with preemption*, International Conference on Discrete Optimization and Operations Research (DOOR-2016), Lecture Notes in Computer Science, **9869** (2016), 309–321. DOI: 10.1007/978-3-319-44914-2_25; Zbl 1380.68077
- [16] A. Kononov, Y. Kovalenko, *An approximation algorithm for preemptive speed scaling scheduling of parallel jobs with migration*, International Conference on Learning and Intelligent Optimization (LION-11), Lecture Notes in Computer Science, **10556** (2017), 351–357. DOI: 10.1007/978-3-319-69404-7_30
- [17] C-Y. Lee, X. Cai, *Scheduling one and two-processor tasks on two parallel processors*, IIE Transactions, **31**: 5 (1999), 445–455. DOI: 10.1023/A:1007501324572
- [18] M. Li, F. Yao, H. Yuan, *An $O(n^2)$ algorithm for computing optimal continuous voltage schedules*, International Conference on Theory and Applications of Models of Computation (TAMC 2017), Lecture Notes in Computer Science, **10185** (2017), 389–400. DOI: 10.1007/978-3-319-55911-7_28
- [19] E. Naroska, U. Schwiegelshohn, *On an on-line scheduling problem for parallel jobs*, Information Processing Letters, **81**: 6 (2002), 297–304. DOI: 10.1016/S0020-0190(01)00241-1; Zbl 1013.68027
- [20] A. Shioura, N. Shakhlevich, V. Strusevich, *Energy saving computational models with speed scaling via submodular optimization*, Third International Conference on Green Computing, Technology and Innovation (ICGCTI2015), 2015, Serdang, 7–18.
- [21] F. Yao, A. Demers, S. Shenker, *A scheduling model for reduced CPU energy*, 36th Annual Symposium on Foundation of Computer Science (FOCS 1995), 1995, IEEE, 374–382. DOI: 10.1109/SFCS.1995.492493; Zbl 0938.68533

ALEXANDER KONONOV
 SOBOLEV INSTITUTE OF MATHEMATICS,
 4, PR. KOPTYUGA AVE.,
 NOVOSIBIRSK, 630090, RUSSIA
E-mail address: `alvenko@math.nsc.ru`

YULIA KOVALENKO
 SOBOLEV INSTITUTE OF MATHEMATICS,
 4, PR. KOPTYUGA AVE.,
 NOVOSIBIRSK, 630090, RUSSIA
E-mail address: `julia.kovalenko.ya@yandex.ru`