

Article

A Polynomial Algorithm for Sequencing Jobs with Release and Delivery Times on Uniform Machines

Nodari Vakhania^{1,†,‡} and Frank Werner^{2,†,‡}*¹ Centro de Investigacion en Ciencias, UAEMor; nodari@uaem.mx² Faculty of Mathematics, Otto-von-Guericke University Magdeburg; frank.werner@ovgu.de

* Correspondence: frank.werner@ovgu.de; Tel.: +49-391-675-2025 (F.W.)

Abstract: The problem of sequencing n equal-length non-simultaneously released jobs with delivery times on m uniform machines to minimize the maximum job completion time is considered. To the best of our knowledge, the complexity status of this classical scheduling problem remained open up to the date. We establish its complexity status positively by showing that it can be solved in polynomial time. We adopt for the uniform machine environment the general algorithmic framework of the analysis of behavior alternatives developed earlier for the identical machine environment. The proposed algorithm has the time complexity $O(\gamma m^2 n \log n)$, where γ can be any of the magnitudes n or q_{\max} , the maximum job delivery time. In fact, n can be replaced by a smaller magnitude $\kappa < n$, which is the number of special types of jobs (it becomes known only upon the termination of the algorithm).

Keywords: scheduling; uniform machines; release time; delivery time; time complexity; algorithm

1. Introduction

In this paper, we consider a basic optimization problem of scheduling jobs with release and delivery times on uniform machines with the objective to minimize the makespan. More precisely, n jobs from the set $J = \{1, 2, \dots, n\}$ are to be processed by m parallel uniform machines (or processors) from the set $M = \{1, 2, \dots, m\}$. Job $j \in J$ is available from its release time r_j , it needs a continuous (integer) processing time p , which is the time that it needs on a slowest machine. We assume that the machines in the set M are ordered by their speeds, the fastest machines first, i.e., $s_1 \geq s_2 \geq \dots \geq s_m$ are the corresponding machine speeds, s_i being the speed of machine i . Without loss of generality, we assume that $s_m = 1$, and the processing time of job j on machine i is an integer p/s_i . Job j has one more parameter, the delivery time q_j , an integer number which represents the amount of additional time units which are necessary for the full completion of job j after it completes on the machine. So notice that the delivery of job j consumes no machine time (the delivery is accomplished by an independent agent).

Now we define a feasible schedule S as a function that assigns to each job j a starting time t_j^S and a machine i from set M , such that $t_j^S \geq r_j$ and $t_j^S \geq t_k^S + p/s_i$, for any job k scheduled earlier on the same machine (the first inequality ensures that a job cannot be started before its release time, and the second reflects the restriction that each machine can handle only one job at any time). The completion time of job j in schedule S is $c_j^S = t_j^S + p/s_i$ and the full completion time of job j in schedule S is $\mathcal{C}_j^S = c_j^S + q_j$ (the full completion time of job j takes into account the delivery time of that job, whereas the completion time of job j does not depend on its delivery time). The objective is to find an optimal schedule, i.e., a feasible schedule S that minimizes the maximal full job completion time

$$C_{\max}(S) = \max_j \mathcal{C}_j$$

commonly referred to as the makespan.

27 The studied multiprocessor optimization problem, described below, is commonly
 28 abbreviated as $Q|p_j = p, r_j, q_j|C_{\max}$ (its version with identical parallel machines is
 29 abbreviated as $P|p_j = p, r_j, q_j|C_{\max}$, the first field specifies the machine environment, the
 30 second one the job parameters, and the third one the objective function).

31 It is well-known that there is an equivalent (perhaps more traditional) formulation
 32 of the above described problem, in which instead of the delivery time q_j , every job j has
 33 its due-date d_j . The *lateness* of job j in the schedule S is $L_j^S = c_j^S - d_j$. Then the objective
 34 becomes to minimize the maximum job lateness L_{\max} , i.e., find a feasible schedule S_{OPT}
 35 in which the maximum job lateness is not more than in any other feasible schedule, i.e.,
 36 S_{OPT} is an *optimal* schedule. The equivalence is easily established by associating with
 37 each job delivery time a corresponding due-date, and vice-versa, see e.g., Bratley et al.
 38 [1]). The version of the problem with due-dates with identical and uniform machine
 39 environments are commonly abbreviated as $P|r_j, d_j|L_{\max}$ and $Q|r_j, d_j|L_{\max}$, respectively.

40 For the problem considered, each machine from a group of parallel uniform ma-
 41 chines is characterized by its own speed, independent from a particular job that can be
 42 assigned to it, unlike a machine from a group of unrelated machines whose speed is
 43 job-dependent. Because of the uniform speed characteristic, scheduling problems with
 44 uniform machines are essentially easier than scheduling problems with unrelated ma-
 45 chines, whereas scheduling problems with identical machines are easier than those with
 46 uniform ones. Here we consider the setting with equal-length jobs the complexity status
 47 of which remained open. At the same time, the version of this problem with identical
 48 machines is known to be polynomially solvable. In this paper, we settle the complexity
 49 status of the problem $Q|p_j = p, r_j, q_j|C_{\max}$ by presenting a polynomial algorithm.

50 The algorithm that we describe here essentially relies on the algorithmic framework
 51 proposed earlier for the version of the problem with identical machines in [2]. The
 52 framework is based on the analysis of some nice structural properties of specially created
 53 schedules which are analyzed in terms of the so-called behavior alternatives. The
 54 framework resulted in an $O(q_{\max}mn \log n + O(m\kappa n))$ time algorithm, where q_{\max} is the
 55 maximum job delivery time and $\kappa < n$ is a parameter which is known only after the
 56 termination of the algorithm. Each schedule is easily created by a well-known greedy
 57 algorithm commonly referred to as Largest Delivery Time heuristic (LDT-heuristic for
 58 short): iteratively, among all released jobs, it schedules one with the largest delivery
 59 time. The algorithm from [2] carries out the enumeration of LDT-schedules (ones created
 60 by the LDT-heuristic) - it is known that there is an optimal LDT-schedule. Based on the
 61 established properties, the set of LDT-schedules is reduced to a subset of polynomial
 62 size which yields a polynomial time overall performance. Although the LDT-heuristic
 63 applied to a problem instance with uniform machines does not provide the desirable
 64 properties, it can be modified to a similar method that takes into account the uniform
 65 speed characteristic. While scheduling identical machines, the minimum completion
 66 time of each next selected job is always achieved on the machine, which is the next to
 67 the machine to which the previously selected job is assigned. With uniform machines,
 68 this is not necessarily the case, for example, the next machine can be much slower than
 69 the current one. Hence, the time moment at which the job will complete on each of
 70 the machines needs to be additionally determined and then this job can be assigned
 71 to a machine on which the above minimum is reached. Here, this modified version
 72 is referred to as the LDTC-heuristic and (a schedule created by this heuristic as an
 73 LDTC-schedule). Instead of enumerating the LDT-schedules (as in [2]), the algorithm
 74 proposed here enumerates LDTC-schedules. Some properties for the identical machine
 75 environment which do not immediately hold for uniform machines are reformulated
 76 in terms of uniform machines, which allows to maintain the basic framework from [2]
 77 which, as suggested earlier, turned out to be sufficiently flexible.

78 Similarly as there exists an optimal LDT-schedule for the identical machine environ-
 79 ment, there exists an optimal LDTC-schedule for the uniform machine environment. The
 80 complete enumeration of the LDTC-schedules is avoided by the generalization of nice

81 properties of LDT-schedules to LDTC-schedules for uniform machines. These properties
 82 are obtained via the analysis of the behavior alternatives from [2] that are generalized
 83 for uniform machines. The proposed algorithm has the time complexity $O(\gamma m^2 n \log n)$,
 84 where γ can be any of the magnitudes n or q_{\max} , the maximum job delivery time. In fact,
 85 n can be replaced by a smaller magnitude κ , the number of special types of jobs; this is
 86 the same parameter κ as for the algorithm from [2] which becomes known only when
 87 the algorithm halts. The running time of the proposed algorithm is worse than that of
 88 the one from [2], basically, because of the cost of the LDTC-heuristic which is repeatedly
 89 used during the solution process.

90 The remainder of this paper is as follows. In Section 2, we give a brief literature
 91 review. Section 3 presents some necessary preliminaries. Then the basic algorithmic
 92 framework is given in Section 4. Section 5 discusses the performance analysis of the
 93 developed algorithm. Finally, Section 6 gives some concluding remarks.

94 2. Literature Review

95 If the job processing times are arbitrary, then the problem is known to be strongly
 96 NP-hard, even if there is only a single machine $1|r_j, d_j|L_{\max}$ [3], see e.g., McMahon
 97 & Florian [4] for an efficient enumerative algorithm for the problem $1|r_j, d_j|L_{\max}$. For
 98 this problem with so-called embedded jobs, where the data fulfill special conditions,
 99 Vakhania presented fast-polynomial algorithms in [5]. For the single machine case,
 100 Baptiste gave an $O(n^7)$ algorithm for the problem $1|r_j, p_j = p|\sum T_j$ [6] and also an
 101 algorithm of the same complexity for the problem $1|r_j, p_j = p|\sum w_j U_j$ [7] of minimizing
 102 the weighted number of late jobs. Chrobak et al. [8] have derived an algorithm of
 103 improved complexity $O(n^5)$ for the case of unit weights, i.e., for the problem $1|r_j, p_j =$
 104 $p|\sum U_j$. Later, Vakhania [9] gave an $O(n^2 \log n)$ algorithm for this problem. Note
 105 that for the problem $1|r_j, p_j, pmtn|\sum U_j$ with arbitrary processing times and allowed
 106 preemptions, Vakhania [10] derived an $O(n^3 \log n)$ algorithm.

107 One may consider a slight relaxation of problems $1|r_j, d_j|L_{\max}$, $P|r_j, d_j|L_{\max}$ and
 108 $Q|r_j, d_j|L_{\max}$ in which one looks for a schedule in which no job completes after its due-
 109 date. Such a feasibility setting with a single machine was considered by Garey et al. [11].
 110 They have proposed an $O(n^2 \log n)$ algorithm which has further been improved to an
 111 $O(n \log n)$ one by using a very sophisticated data structure. The authors have relied
 112 on the concept of a so-called forbidden region, an artificial gap in a schedule in which
 113 it is forbidden to start any job. Later Simons and Warmuth [12] have constructed an
 114 $O(n^2 m)$ time algorithm for the feasibility setting with the identical machine environment
 115 also using the concept of forbidden regions. (It can be mentioned that the minimization
 116 version of the problem can be solved by applying an algorithm for the feasibility problem
 117 by repeatedly increasing the due-dates of all jobs until a feasible schedule with the
 118 modified due dates is found. Using binary search makes such a reduction procedure
 119 more efficient and reduces the reduction cost to $O(\log(np/m))$.)

120 Dessousky et al. [13] considered scheduling problems on uniform machines with
 121 simultaneously released jobs (i.e., with $r_j = 0$ for every job j) and with different objective
 122 criteria. They proposed fast polynomial-time algorithms for these problems, in particular,
 123 for the criterion L_{\max} . In fact, the LDTC-heuristic is an adaptation of an optimal solution
 124 method that the authors in [13] constructed for the criterion L_{\max} .

125 For a uniform machine environment with allowed preemptions ($pmtn$), the problem
 126 $Q|r_j, pmtn|C_{\max}$ is polynomially solvable even for arbitrary processing times [14], while
 127 a polynomial algorithm for the problem $Q|r_j, p_j = p, pmtn|\sum C_j$ with minimizing total
 128 weighted completion time in the case of equal processing times has been given in [15].
 129 The case of unrelated machines is very hard. A polynomial algorithm exists for the
 130 problem $R|r_j, pmtn|L_{\max}$ with allowed preemptions and minimizing maximum lateness,
 131 even for the case of arbitrary processing times [16]. If preemptions are forbidden,
 132 Vakhania et al. [17] gave a polynomial algorithm for the case of minimizing the makespan
 133 when only two processing times p and $2p$ are possible (i.e., for the problem $R|p \in$

134 $\{p, 2p\} | C_{max}$. Note that the case of two arbitrary processing times p and q is known
 135 to be NP-hard [18]. For the special case of identical parallel machines, there exist
 136 several works for the same setting as considered in this paper but for more complicated
 137 objective functions regarding the complexity status. In particular, the problems $P|r_j, p_j =$
 138 $p | \sum w_j C_j$ of minimizing the weighted sum of completion times [19] and $P|r_j, p_j = p | \sum T_j$
 139 of minimizing total tardiness [20] can be polynomially solved by a reduction to a linear
 140 programming problem. In [21], Vakhania presented an $O(n^3 \log n)$ algorithm for the
 141 problem $P|r_j, p_j = p | \sum U_j$ of minimizing the number of late jobs. His blossmore
 142 ('branch less, cut more') algorithm uses a solution tree, where the branching and cutting
 143 criteria are based on the analysis of behavior alternatives. Moreover, the problem
 144 $P|r_j, p_j = p | \sum f_j(C_j)$ can also be polynomially solved for the case that f_j is an arbitrary
 145 non-decreasing function such that the difference $f_i - f_j$ is monotonic for any indices
 146 i and j [22]. The authors also applied a linear programming approach. It can also be
 147 mentioned that a detailed survey on parallel machine scheduling problems with equal
 148 processing times has been given in [23].

149 3. Preliminaries

150 This section contains some useful properties, necessary terminology and concepts,
 151 some of which were introduced in [2] for identical machines.

152 **LDTC-heuristic.** We first describe the LDTC-heuristic, an adaptation of the LDT-
 153 heuristic for uniform machines. As earlier briefly noted, unlike an LDT-schedule, an
 154 LDTC-schedule is not defined by a mere permutation of the given n jobs since the
 155 machine to which the next selected job is assigned depends on the machine speed.
 156 Starting from the minimal job release time, the current scheduling time is iteratively set
 157 as the minimum release time among all yet unscheduled jobs. Iteratively, among all jobs
 158 released by the current scheduling time, the LDTC-heuristic determines one with the
 159 largest delivery time (a most *urgent* one) and schedules it on the machine on which the
 160 earliest possible completion time of this job is attained (ties can be broken by selecting
 161 the machine with the minimum index). Note that in an LDTC-schedule S , a machine will
 162 contain an idle-time interval (a *gap*) if and only if there is no unscheduled job released by
 163 the current scheduling time. The running time of the modified heuristic is the same as
 164 that of LDT-heuristic with an additional factor of m due to the machine selection at each
 165 iteration (which is not required for the uniform machine environment), which results in
 166 the time complexity $O(nm \log n)$.

167 Let σ be the LDTC-schedule obtained by the LDTC-heuristic for the initially given
 168 problem instance. As we will see in the following subsection, we may generate alter-
 169 native LDTC-schedules by iteratively modifying the originally given problem instance.
 170 The next property of an LDTC-schedule easily follows from the definition of the LDTC-
 171 heuristic (and the equality of the job processing times).

172 **Property 1.** *If in an LDTC-schedule S , job j is scheduled after job i , i.e., the ordinal number of*
 173 *job j in S is larger than that of job i , then $c_j^S \geq c_i^S$.*

174 Next, we give another easily seen important property of an LDTC-schedule S on
 175 which the proposed method essentially relies. Let A be a set of, say k jobs, all of which
 176 being released by time moment t , and let π be a permutation of k jobs, all of which being
 177 also released by time t (recall that all these jobs have equal length). Let, further, $S(A)$
 178 be a partial LDTC-schedule constructed for the jobs in the set A , and let $S(\pi)$ be a list
 179 schedule constructed for the permutation π .

180 **Property 2.** *The completion time of every machine in both schedules $S(A)$ and $S(\pi)$ is the same.*
 181 *Moreover, the i th scheduled job in the schedule $S(A)$ starts and completes at the same time as the*
 182 *i th scheduled job in the schedule $S(\pi)$.*

183 The above property also holds for a group of identical machines and is helpful
 184 for the generalization of the earlier results for identical machines from [2] to uniform
 185 machines. Roughly, ignoring the job release times, the property states that two list
 186 schedules constructed for two different permutations with the same number of jobs have
 187 the same structure. Although the starting and completion times of the jobs scheduled
 188 in the same position are the same in both schedules, the full completion times will not
 189 necessarily be the same (this obviously depends on the delivery times of these jobs).

190 **A block.** An independent part in a schedule S is commonly referred to as a *block* in
 191 the scheduling literature. We define it as a largest fragment of schedule S such that for
 192 each two successively scheduled jobs i and j , job j starts no later than job i finishes (jobs
 193 i and j can be scheduled on the same or different machines). It follows that there is a
 194 single block that starts schedule S and finishes this schedule. If these blocks coincide,
 195 then there is a single block in the schedule S , otherwise, each next block is “separated”
 196 by the previous one with gaps on each of the machines. Here a zero length gap between
 197 jobs i and j will be distinguished in case job j is scheduled at time $r_j = t_i(S)$ on the same
 198 machine as job i (it immediately succeeds job i on that machine). A block B (with at
 199 least two elements) possesses the following property that will be used later. Suppose
 200 the i th scheduled job j is removed from that block and the LDTC-heuristic is applied to
 201 the remaining jobs of the block. Then in the resultant (partial) schedule, the processing
 202 interval of the i th scheduled job overlaps with the earlier processing interval of job j in
 203 block B .

204 Some additional definitions are required to specify how the proposed algorithm
 205 creates and evaluates the feasible schedules. At any stage of the execution of the algo-
 206 rithm, independently whether a new LDTC-schedule will be generated or not, depends
 207 on specific properties of the LDTC-schedules already generated by that stage. The
 208 definitions below are helpful for the determination of these properties.

An overflow job. In an LDTC-schedule S , let o be a job realizing the maximum full
 completion time of as job, i.e.,

$$C_o(S) = C_{max}(S), \quad (1)$$

209 and let $B(S)$ be the *critical block* in the schedule S , i.e., the block containing the earliest
 210 scheduled job o satisfying equation (1). The *overflow job* $o(S)$ in the schedule S is the
 211 latest completed job on the machine in the block $B(S)$ satisfying condition (1), i.e., one
 212 with the maximum $c_o(S)$ (further ties are broken by selecting the job scheduled on the
 213 machine with the largest index).

214 **A kernel.** Next, we define an important component in an LDTC-schedule S defined
 215 as its segment containing the overflow job $o = o(S)$ such that the jobs scheduled before
 216 job o in the block $B(S)$ have a delivery time not smaller than q_o (we will write o instead
 217 $o(S)$ when this causes no confusion). This segment of the schedule S is called its *kernel*
 218 and is denoted by $K(S)$.

219 Intuitively, on the one hand, the kernel $K(S)$ is a critical part in a schedule S , and
 220 on the other hand, it is relatively easy to arrange the kernel jobs optimally. In fact, we
 221 will explore different LDTC-schedules identifying the kernel in each of these schedules.
 222 We will also relate this kernel to the kernels of the earlier generated LDTC-schedules.
 223 We need to introduce a few more definitions.

224 **An emerging job.** Suppose that a job j of kernel $K(S)$ is *pushed* by a non-kernel job i
 225 scheduled earlier on the same machine, that is, the LDTC-heuristic would schedule job j
 226 earlier if job i was forced to be scheduled after job j . Assume further that job e scheduled
 227 before job o in the block $B(S)$ is such that $q_e < q_o$. Then job e is called a *regular emerging*
 228 *job* in the schedule S , and the latest scheduled (regular) emerging job (the one closest to
 229 job o) is called the *delaying* emerging job.

230 The following optimality condition can be established already in the initial LDTC-
 231 schedule σ .

232 **Lemma 1.** *If the initial LDTC-schedule σ contains a kernel K such that no job of that kernel is*
 233 *pushed by a non-kernel (emerging) job, then this schedule is optimal.*

234 **Proof.** By the condition, if a job j from kernel K does not start at its release time, then
 235 it is pushed by another job i from this kernel with $q_i \geq q_j$ (the latter inequality follows
 236 from the definition of a kernel). The lemma is now easily established by an interchange
 237 argument and by the fact that the makespan of the schedule σ is realized by a job from
 238 the kernel K . \square

239 3.1. Constructing alternative LDTC-schedules

240 Due to Lemma 1, from here on, it is assumed that the condition in this lemma is
 241 not satisfied, i.e., there exists an emerging job e in the schedule S (note that $e \in B(S)$ as
 242 otherwise job e may not push a job of the kernel $K(S)$). Since job e is pushing a job of
 243 kernel $K(S)$, the removal of this job may potentially decrease the starting and hence the
 244 full completion time of the overflow job $o(S)$. At the same time, note again that by the
 245 definition of a block, the omission of a job not from the block $B(S)$ may not affect the
 246 starting time of any job from the block $B(S)$. That is why we restrict here our attention
 247 to the jobs of the block $B(S)$. (Later we will introduce an alternative notion of a passive
 248 emerging job and then will use “emerging job” for either a regular or passive emerging
 249 job; until then we use “emerging job” for a “regular emerging job”.)

250 Clearly, no emerging job can actually be removed as the resultant schedule would
 251 be infeasible. Instead, to restart the jobs in the kernel $K(S)$ earlier, an emerging job e is
 252 *applied* to this kernel, i.e., it is forced to be rescheduled after all jobs of the kernel $K(S)$
 253 whereas any job, scheduled after the kernel $K(S)$ in the schedule S is maintained to
 254 be scheduled after that kernel. The application of job e is accomplished in two steps:
 255 first, the original release time of job e and that of the jobs, scheduled after kernel $K(S)$
 256 in schedule S , is increased to the release time of any job in the kernel $K(S)$. Then the
 257 resultant schedule denoted by S_e (the so-called *complementary schedule* or a *C-schedule*)
 258 is obtained by the LDTC-heuristic which is merely applied to the modified problem
 259 instance). (Such a schedule generation technique for a single-machine setting was
 260 suggested by McMahon and Florian [4].)

261 By Lemma 1, the kernel $K(S)$, a fragment of the LDTC-schedule S considered as an
 262 independent LDTC-schedule is optimal if it possesses no emerging job. Otherwise, the
 263 jobs of the kernel $K(S)$ are pushed by the corresponding emerging jobs. Some of these
 264 emerging jobs can be scheduled after the kernel $K(S)$ in an optimal complete schedule
 265 S_{OPT} . Such a rescheduling is achieved by the creation of the corresponding C-schedules
 266 (as we will see in Lemma 2, it will suffice to consider only C-schedules, i.e., S_{OPT} is a
 267 C-schedule).

268 The application of an emerging job has two “opposite” effects. On the positive
 269 side, since the number of jobs scheduled before the kernel $K(S)$ in the schedule S_e is
 270 one less than that in the schedule S , the overflow job $o(S)$ in the schedule S_e will be
 271 completed earlier than it was completed in the schedule S ; likewise, the completion time
 272 of the latest scheduled job of the kernel $K(S)$ in the schedule S_e will be less than the
 273 completion time of job $o(S)$ in the schedule S . Hence, the application of an emerging
 274 job gives a potential to improve schedule S . On the negative side, it creates a new gap
 275 within the former execution interval of job e or at a later time moment before kernel $K(S)$
 276 (see Lemma 1 in [2] for a proof for the case of identical machines, the uniform machine
 277 case can be proved similarly). Such a gap may enforce a right-shift (delay) of the jobs
 278 included after job e in the schedule S_e . So roughly, the C-schedule S_e favors the kernel
 279 $K(S)$ but creates a potential conflict for later scheduled jobs.

280 4. The basic algorithmic framework

281 In this section, we give the basic skeleton of the proposed algorithm and prove
 282 its correctness. The schedule S_{OPT} is characterized by a proper processing order of the

283 emerging jobs scheduled in between the kernels. Starting with the initial LDTC-schedule
 284 σ , an emerging job in the current LDTC-schedule is applied and a new C-schedule is
 285 created; in this schedule the kernel is again determined. The same operation is iteratively
 286 repeated as long as the established optimality conditions are not satisfied. As we will
 287 show later, it will suffice to enumerate all C-schedules to find an optimal solution to the
 288 problem.

289 We associate a complete feasible C-schedule with each node in a *solution tree* T ,
 290 the initial LDTC-schedule being associated with the root. Aiming to avoid a brutal
 291 enumeration of all C-schedules, we carry out a deeper study of the structure of the
 292 problem and some additional useful properties of LDTC-schedules. In fact, our solution
 293 tree T consists of a single chain of C-schedules. We will refer to a node of the tree as a *stage*
 294 (since each node represents a particular stage in the algorithm with the corresponding
 295 LDTC-schedule). We let $T_h = (S^0, \dots, S^h)$ be the sequence of C-schedules generated
 296 by stage h . Thus, S^0 is the initial LDTC-schedule, and the schedule S^h of stage $h > 0$,
 297 the immediate successor of schedule S^{h-1} , is obtained by one of the extension rules as
 298 described below.

299 In the schedule S^h , the overflow job $o(S^h)$, the delaying emerging job l and the
 300 kernel $K(S^h)$ are determined. Using the *normal extension rule*, we let $S^{h+1} := S_l^h$, where
 301 l is the delaying emerging job in the schedule S^h . Alternatively, the schedule S^{h+1} is
 302 constructed from the schedule S^h by the *emergency extension rule* as described in the
 303 following subsection.

304 4.1. Types of emerging jobs and the extended behavior alternatives

305 **A marched emerging job.** An emerging job may be in different possible states. It
 306 is useful to distinguish these states and treat them accordingly. Suppose that e is an
 307 emerging job in the schedule S^g and it is applied by stage h , $h > g$ (in a predecessor-
 308 schedule of schedule S^h). Then job e is called *marched* in the schedule S^h if $e \in B(S^h)$.
 309 Intuitively, the existence of a marched job in the schedule S^h indicates an “interference”
 310 of the kernel $K(S^h)$ with an earlier arranged part of the schedule preceding that kernel.

311 **A stuck emerging job.** Suppose that job e is marched in the schedule S^h and
 312 $E(S^h) = \emptyset$, where $B(S^h)$ is a non-primary block. Then job e is called *stuck* in the schedule
 313 S^h if either it is scheduled before job $o(S^h)$ or $e = o(S^h)$ (observe that any job stuck in
 314 the schedule S^h belongs to the kernel $K(S^h)$).

315 **Block evolution in the solution tree T .** Although $E(S^h) = \emptyset$, since $B(S)$ is a non-
 316 primary block, a “potential” regular emerging job might be “hidden” in some block
 317 preceding block $B(S^h)$, in the schedule S^h . In general, a block in the schedule S^h can be a
 318 part of a larger block from the schedule S^g for some $g < h$. Recall that the application of
 319 an emerging job e in a C-schedule S yields the raise of a new gap in the C-schedule S_e .
 320 As it can be straightforwardly seen, this may cause a separation or the *splitting* of the
 321 critical block $B(S)$ into two (or even more) new blocks. Likewise, since job e may push
 322 the following jobs in the schedule S_e , because of the forced right-shift of these jobs, two
 323 or more blocks may *merge* forming a bigger block consisting of the jobs from the former
 324 blocks.

325 We will refer to blocks from two different C-schedules as *congruent* if both of them
 326 are formed by the same set of jobs. A block in a C-schedule, which is congruent to a
 327 block from the initial LDTC-schedule, will be referred to as a *primary* block. Observe that
 328 a non-primary block may arise because of either block splitting or/and block merging
 329 and that all blocks in the initial LDTC-schedule are primary.

330 If the block B is arisen as a result of an application of an emerging job e , then this
 331 block is said to be a non-primary block *of* job e , and the latter job is said to be the *splitting*
 332 job of B . Note that, since the application of an emerging job not necessarily causes a block
 333 split, a non-primary block of job e may contain some other already applied emerging jobs
 334 (recall that before each application of an emerging job, the current release time of that

335 job is increased accordingly; e.g., if job e is applied ι times, its release time is modified ι
 336 times).

337 We will refer to the blocks arisen after the splitting of one particular block as the
 338 *direct descendants* of this block (the latter block is the *direct predecessor* of the former ones).
 339 If $B \in S$ is the direct predecessor of the blocks B_1, \dots, B_k , then $B \subseteq B_1 \cup \dots \cup B_k$, but not
 340 vice versa (because of a possible block merging).

341 A *descendant* of a block is its direct descendant recurrently, any descendant of a
 342 descendant of a block is also a descendant of the latter block. If a block is a descendant
 343 of another block, then the latter is a *predecessor* of the former block. Two or more blocks
 344 are said to be *relative* if they have at least one common predecessor block.

345 **A passive emerging job.** A *passive emerging job* e in the schedule S^h is a ("hidden")
 346 regular emerging job from the schedule S^g , i.e., job e belongs to a block from schedule
 347 S^g , relative to block $B(S^h)$ (preceding this block), such that $q_e < q_{o(S^h)}$.

348 **Extended behavior alternatives.** Now we define two extended behavior alterna-
 349 tives which, together with the five basic behavior alternatives were introduced earlier in
 350 [2] (Section 2.3). Suppose that there exists no regular emerging job in the C-schedule S
 351 (i.e., the block $B(S)$ starts with the kernel $K(S)$), and there is no stuck job in this schedule.
 352 Then an *exhaustive instance of alternative (a)* (abbreviated EIA(a)) in the schedule S is said
 353 to occur. If now there exists a stuck job in the schedule S , then an *extended instance of*
 354 *alternative (b)* with this job in the schedule S (abbreviated EIA(b)) is said to occur (there
 355 may exist more than one job stuck in the schedule S^h).

356 The first above behavior alternative immediately yields an optimal solution, and
 357 the second one indicates that some rearrangement of the already applied emerging
 358 jobs might be required. The first and the second behavior alternatives, respectively, are
 359 treated in the following lemma and in the next subsection, respectively.

360 **Lemma 2.** *A C-schedule S^h is optimal if an EIA(a) in it occurs.*

361 **Proof.** By the condition, there exists no stuck job in the schedule S^h . This implies that
 362 none of the jobs of the kernel $K(S^h)$ can be scheduled at some earlier time moment
 363 without causing a forced delay of a more urgent job from this kernel, and the lemma can
 364 easily be proved by an interchange argument. \square

365 4.2. Emergency extension rule

366 Throughout this subsection, we assume that there arises an EIA(b) in the schedule
 367 S^h and that there exists a passive emerging job in this schedule. Let l be the latest
 368 scheduled passive emerging job in S^h . By the definition of job l , there is a schedule S^g ,
 369 a predecessor of schedule S^h in the solution tree T , and a job $e \in B(S^h)$, stuck in the
 370 schedule S^h , such that $e \in B(S^g)$ and $l \in B(S^g)$. Let e be the latest applied job stuck
 371 in the schedule S^h . Although jobs l and e belong to different blocks in the schedule S^h ,
 372 the corresponding blocks can be merged by reverting the application(s) of job e . This
 373 can clearly be accomplished by restoring the corresponding earlier release time of job
 374 e (recall that the release time of an emerging job is increased each time it is applied).
 375 Once these blocks are merged, job l becomes a regular emerging job and hence, it can be
 376 applied.

377 Let r be the release time of an emerging job j before it is applied to a kernel K .
 378 Then job j is said to be *revised* (for K) if it is placed back before K , i.e., its release time is
 379 reassigned the value r and the LT-heuristic is applied.

380 In more detail, let B be a block relative to $B(S^h)$ in the schedule S^h containing
 381 job l . Not only B and $B(S^h)$ are different blocks, there might be a chain B_1, \dots, B_{k-1} of
 382 succeeding (relevant) blocks between the blocks B and $B(S^h)$ in the schedule S^h . Let
 383 $B_0 = B$ and $B_k = B(S^h)$. First, the blocks B_{k-1} and B_k are merged by reverting the
 384 application of the corresponding emerging job. Then the resulting block is similarly
 385 merged with block B_{k-2} , and so on. In general, to merge the block B' with its successive
 386 (relative) block B'' , the corresponding release time of one of the currently applied jobs

387 scheduled in block B'' (scheduled between the jobs of B' and B'' before its application) is
 388 restored, i.e., this job is *revised*.

389 Note that by the definition, the revision of the splitting job of blocks B' and B''
 390 will merge these two blocks. Similarly, the revision of any other applied emerging job,
 391 scheduled in B'' , will have the same effect. Among all such jobs with the largest delivery
 392 time, the latest scheduled one in block B'' will be referred to as the *active splitting* job for
 393 the blocks B' and B'' .

394 The blocks B_{k-1} and B_k are merged by the revision of the active splitting job of
 395 these two blocks which is scheduled in block B_k . Similarly, the active splitting job of
 396 block B_{k-2} is revised to merge block B_{k-2} with the earlier obtained block, and so on, this
 397 process continues until all blocks from the chain $B_0, B_1, \dots, B_{k-1}, B_k$ are merged.

398 We denote the resultant merged block by $\mathcal{B}(l)$ (this block, ending with the jobs from
 399 block $B(S^h)$ can, in general, be non-primary), and we will refer to the above described
 400 procedure as the *chain of revisions* for the passive emerging job l . Observe that, although
 401 the outcome of this procedure somewhat resembles the traditional backtracking, it is
 402 still different as it keeps untouched the “intermediate” applications that could have been
 403 earlier carried out between the reverted applications.

404 Let $Rev_l(S^h)$ be the C-schedule, obtained from schedule S^h by the chain of revisions
 405 for job l . Observe that job l changes its status from a passive to a regular emerging job
 406 in this schedule, and hence the emergency extension rule applies job l in the schedule
 407 $Rev_l(S^h)$ to kernel $K(Rev_l(S^h))$, setting $S^{h+1} := (Rev_l(S^h))_l$ (we will also use the shorter
 408 notation $S^h_{+l} = (Rev_l(S^h))_l$ for the resultant schedule).

409 4.3. The description of the algorithm and its correctness

410 We give the following Algorithm 1 and prove its correctness.

411 **Algorithm 1:**

412 **Step 1:** Set $h = 0$ and $S^h := \sigma$.

413 **Step 2:** If the condition of Lemma 1 holds, then return the schedule σ and stop.

414 **Step 3:** Set $h = h + 1$.

415 **Step 4:** { *stopping rule* } If in the schedule S^h , either (i) there exists neither a regular nor a
 416 passive emerging job or there occurs an EIA(a), then return a schedule from the tree T
 417 with the minimum makespan and stop;

418 **Step 5:** { *normal extension rule* } : if in the schedule S^h , there occurs no EIA(b), then
 419 $S^{h+1} := S^h_l$, where l is the regular delaying emerging job else

420 { *emergency extension rule* } if in the schedule S^h there occurs an EIA(b), then $S^{h+1} := S^h_{+l}$,
 421 where l is the passive delaying emerging job.

422 **Step 6:** goto Step 3.

423 Now we can prove the following theorem.

424 **Theorem 1.** For some stage h , the C-schedule S^h is an optimal schedule S_{OPT} .

425 **Proof.** First, we show that an optimal schedule is a C-schedule. Suppose that the
 426 condition in Lemma 1 is not satisfied for the schedule $S^0 = \sigma$, and that the C-schedule
 427 S^h , $h > 0$, is not optimal. Then, due to Property 2, in any feasible schedule S with a
 428 better makespan, the number of jobs scheduled before the kernel $K(S^h)$ in block $B(S^h)$
 429 must be one less than in the latter schedule (otherwise, a job from the kernel $K(S^h)$ will
 430 not have a smaller full completion time in the schedule S' than job $o(S^h)$ in the schedule
 431 S^h). Hence, some job $e \in B(S^h)$ included before the kernel $K(S^h)$ in the schedule S^h
 432 must be scheduled after that kernel in the schedule S . Moreover, e is to be either a
 433 regular or passive emerging job. Indeed, if e is not an emerging job, then $q_e \geq q_o$,

434 $o = o(S^h)$ or/and $e \notin B(S^h)$. In the latter case, by the definition of a block, no feasible
 435 rearrangement involving jobs of a non-critical block may decrease the full completion
 436 time of the overflow job o . Otherwise, $t_e^{(S^h)e} \geq t_o^{S^h}$ and due to inequality $q_e \geq q_o$ and
 437 Proposition 2, we obtain $|(S^h)_e| \geq |S^h|$, which implies that the makespan of any feasible
 438 schedule in which job e is scheduled after the kernel $K(S^h)$ cannot be less than that of
 439 schedule S^h . Then the correctness of the stopping rule clearly follows from Lemma 2.

440 It remains to show that the search in the space of the C-schedules is correctly
 441 organized. There are two extension rules. The normal extension rule is used at stage h if
 442 there exists a regular emerging job in the schedule S^h . In this case, the delaying emerging
 443 job l is applied, i.e., $S^{h+1} := S_l^h$. Consider an alternative feasible schedule S_e^h , where
 444 e is another emerging job (above we have shown that only emerging jobs need to be
 445 considered). It is easy to see that the left-shift of the kernel jobs in the schedule S_e^h cannot
 446 be more than that in the schedule S_l^h , and the forced right-shift for the jobs scheduled
 447 after job e in the schedule S_e^h cannot be less than that of the jobs scheduled after job l
 448 in the schedule S_l^h (recall that $p_e = p_l$). Hence, the schedule S_e^h is dominated by the
 449 schedule S_l^h unless job l gets stuck at a later stage $h' > h$. In the latter case, the emergency
 450 extension rule revises first job l . In the resultant C-schedule, the passive delaying job
 451 converts into a regular delaying emerging job. Then the emergency extension rule
 452 applies this (converted) regular regular delaying emerging job. We complete the proof
 453 by repeatedly applying the above reasoning for the normal extension rule. \square

454 5. Performance analysis

455 It is not difficult to see that the direct application of Algorithm 1 of the previous
 456 section may yield the generation of some redundant C-schedules: the jobs from the
 457 same kernel K including the overflow job o may be forced to be right-shifted after they
 458 are already "arranged" (i.e., the corresponding emerging job(s) are already applied to
 459 that kernel) due to the arrangement accomplished for the kernel K' preceding kernel
 460 K . As a result (because of the application of an emerging job for the kernel K'), one or
 461 more redundant C-schedules in which a job from the kernel K repeatedly becomes an
 462 overflow job might be created. Such an unnecessary rearrangement of the portion of a
 463 C-schedule between the kernels K' and K is avoided by restricting the number of jobs
 464 that are allowed to be scheduled in that portion. This number becomes well-defined
 465 after the first disturbance of this portion caused by the application of an emerging job
 466 for the kernel K' . This issue was studied in detail for the case of identical machines in
 467 [2] (see Section 4.1). It can be readily verified that the basic estimations for the case of
 468 identical machines similarly hold for uniform machines. In particular, the number of
 469 the enumerated C-schedules remains the same for uniform machines. A complete time
 470 complexity analysis requires a number of additional concepts and definitions from [2]
 471 and would basically repeat the arguments for identical machines.

472 Recall that we use a different schedule generation mechanism for identical and
 473 uniform machines: the LDT-heuristic applied for the schedule generation in the identical
 474 machine environment is replaced by the LDTC-heuristic for the uniform machine envi-
 475 ronment. LDT-schedules possess a number of nice properties used in the algorithm from
 476 [2]. LDTC-schedules also possess such necessary useful properties (Properties 1 and 2)
 477 that allowed us to use the basic framework from [2]. While generating an LDT-schedule
 478 for identical machines, every next job is scheduled on the next available machine (the
 479 next to the last machine m being machine 1) and the starting and completion time of each
 480 next scheduled job is not smaller than that of all previously scheduled ones. In some
 481 sense, the generalization of these properties are Properties 1 and 2, which still assure
 482 that the structural pattern of the generated schedules is kept, and it does not depend
 483 on which particular jobs are being scheduled in a particular time interval (note that this
 484 would not be the case for an unrelated machine environment). This allowed us to use the
 485 basic framework from [2] for the uniform machine environment (for instance, intuitively,

486 while restricting the number of the scheduled jobs between two successive kernels, no
487 matter which particular jobs are being scheduled between these kernels).

488 Another “redundancy issue” occurs when a series of emerging jobs are successively
489 applied to the same kernel K without reaching the desired result, i.e., the applied emerg-
490 ing jobs become new overflow jobs in the corresponding C-schedules (see Section 4.2 in
491 [2]). In Lemma 7 from the latter reference, it is shown that this yields an additional factor
492 of p in the running time of the algorithm. This result also holds for the uniform machine
493 environment. The magnitude p remains valid for the uniform machine environment
494 as the difference between the completion times of two successively scheduled jobs on
495 the same machine cannot be more than p (recall that p is the processing time of any job
496 on the slowest machine m). The desired result follows since the delivery time of each
497 emerging job next applied to kernel K is strictly less than that of the previously applied
498 one (see the proof of Lemma 7 in [2]).

499 The above results yield the same bound $O(\gamma m)$ on the number of the enumerated
500 C-schedules as in [2]), where γ can be either n or q_{\max} (see Lemma 8 from Section 4.3 and
501 Theorem 2 from Section 6.1, in [2]). In fact, γ is the total number of the applied emerging
502 jobs, a magnitude, that can be essentially smaller than n . This yields the overall cost
503 $O(\gamma m^2 n \log n)$ due to the cost $O(mn \log n)$ of the LDTC-heuristic (instead of $O(n \log n)$
504 for the LDT-heuristic). A further refinement of the overall time complexity accomplished
505 in [2] is not possible for the uniform machine environment. In the algorithm from [2],
506 while generating every next C-schedule, instead of applying the LDT-heuristic to the
507 whole set of jobs, it is only applied to the jobs from a small part of the current C-schedule,
508 the so-called critical segment (a part containing its kernel), and the remaining jobs are
509 scheduled in linear time just by right-shifting the jobs following the critical segment by
510 the required amount of time units (conserving their current processing order). This is not
511 possible for uniform machines as such an obtained schedule will not necessarily remain
512 an LDTC-schedule, i.e., a linear time rescheduling will not provide the desired structure.

513 6. Concluding remarks

514 The complexity status of a long-standing open multiprocessor scheduling prob-
515 lem was established. Constructively, a polynomial-time solution was proposed. The
516 employed here approach of the analysis of behavior alternatives from [2] turned out
517 to be sufficiently general so that it was possible to extend it also to a uniform machine
518 environment. At the same time, as we have mentioned earlier, this approach cannot be
519 extended to the unrelated machine environment, mainly because the structural pattern
520 of the generated schedules will depend on which particular jobs are scheduled in a
521 particular time interval on each machine from a group of unrelated machines. However,
522 the approach might be extensible for shop scheduling problems. It is a challenging
523 question whether it can also be extended to the case when there are two allowable job
524 processing times (this turned out to be possible for the single-machine environment
525 see [24]) and for a much more general setting with mutually divisible job processing
526 times (this turned out to be also possible for the single-machine environment, a maximal
527 polynomially solvable special case of problem $1|r_j, d_j|L_{\max}$ dealt with recently in [25])
528 for identical and uniform machine environments.

529 **Author Contributions:** conceptualization, N.V.; investigation, N.V. and F.W.; writing–original
530 draft preparation, N.V. and F.W.

531 **Funding:** The first author was supported by DAAD grant for Research Stays for University
532 Academics and Scientists (2020) 57507438.

533 **Conflicts of Interest:** “The authors declare no conflict of interest.”

References

1. Bratley, P.; Florian, M.; Robillard, P. On sequencing with earliest start times and due–dates with application to computing bounds for $(n/m/G/F_{\max})$ problem. *Naval Res. Logist. Quart.*, **1973**, *20*, 57–67.

2. Vakhania, N. A better algorithm for sequencing with release and delivery times on identical processors, *Journal of Algorithms* **2003**, *48*, 273–293.
3. Garey, M.R.; Johnson, D.S. *Computers and Intractability: A Guide to the Theory of NP-completeness*. Freeman, San Francisco (1979).
4. McMahon, G; Florian, M. On scheduling with ready times and due dates to minimize maximum lateness, *Operations Research* **1975**, *23*, 475–482.
5. Vakhania, N. Fast solution of single-machine scheduling problem with embedded jobs, *Theoretical Computer Science* **2019**, *782*, 91–106.
6. Baptiste, P. Scheduling equal-length jobs on identical parallel machines, *Discrete Applied Mathematics* **2000**, *103*, 21–32.
7. Baptiste, P. Polynomial time algorithms for minimizing the weighted number of late jobs on a single machine with equal processing times. *Journal of Scheduling*, **1999**, *2* (6), 245–252.
8. Chrobak, M.; Dürr, C.; Jawor, W.; Kowalik, L.; Kurowski, M. A note on scheduling equal-length jobs to maximize throughput. *Journal of Scheduling*, **2006**, *9*, 71–73.
9. Vakhania, N. A study of single-machine scheduling problem to maximize throughput, *Journal of Scheduling* **2013**, *16*, No. 4, 395–403.
10. Vakhania, N. Scheduling jobs with release times preemptively on a single machine to minimize the number of late jobs, *Operations Research Letters* **2009**, *37*, 405–410.
11. Garey, M.R.; Johnson, D.S.; Simons, B.B.; Tarjan, R.E.: Scheduling unit-time tasks with arbitrary release times and deadlines, *SIAM J. Comput.* **1981**, *10*, 256–269.
12. Simons, B.; Warmuth, M.: A fast algorithm for multiprocessor scheduling of unit-length jobs, *SIAM J. Comput.* **1989**, *18*, 690–710.
13. Dessouky, M.; Lageweg, B.J.; Lenstra, J.K.; van de Velde, S.L.: Scheduling identical jobs on uniform parallel machines, *Statistica Neerlandica* **1990**, *44*, 115–123.
14. Labetoulle, J.; Lawler, E.L.; Lenstra, J.K.; Rinnooy Kan, A.H.G.: Preemptive scheduling of uniform machines subject to release dates, in Pulleyblank, H.R. (ed.) *Progress in Combinatorial Optimization*, New York, Academic Press, **1984** 245–261.
15. Kravchenko, S.; Werner, F. Preemptive scheduling on uniform machines to minimize mean flow time, *Computers and Operations Research* **2009**, *36*, 2816–2821.
16. Lawler, E.L.; Labetoulle, J.: On preemptive scheduling of unrelated parallel processors by linear programming, *Journal of the Association of Computing Machinery* **1978**, *25* (4), 612–619.
17. Vakhania, N.; Hernandez, J.; Werner, F. Scheduling unrelated machines with two types of jobs. *Int. J. Prod. Res.* **2014**, *52* (13), 3793–3801.
18. Lenstra, J.K.; Shmoys, D.B.; Tardos, E. Approximation algorithms for scheduling unrelated machines, *Mathematical Programming*, *46*, 259 - 271 (1990).
19. Brucker, P.; Kravchenko, S. Scheduling jobs with release times on parallel machines to minimize total tardiness, OSM Reihe P, **2005**, Heft 257, Universität Osnabrück, Fachbereich Mathematik/Informatik.
20. Brucker, P.; Kravchenko, S. Scheduling jobs with equal processing times and time windows on identical parallel machines, *Journal of Scheduling* **2008**, *11* (4), 229–237.
21. Vakhania, N. Branch less, cut more and minimize the number of late equal-length jobs on identical machines, *Theoretical Computer Science* **2012**, *465*, 49–60.
22. Kravchenko, S.; Werner, F. On a parallel machine scheduling problem with equal processing times, *Discrete Applied Mathematics*, **2009**, *157*, 848–852.
23. Kravchenko, S.; Werner, F.: Parallel machine problems with equal processing times: A survey, *Journal of Scheduling* **2011**, *14* (5), 435–444.
24. Vakhania, N. "Single-Machine Scheduling with Release Times and Tails". *Annals of Operations Research*, *129*, p.253-271 2004.
25. Vakhania, N. Dynamic Restructuring Framework for Scheduling with Release Times and Due-Dates.
Journal version: *Mathematics* 2019, *7*(11), 1104; <https://doi.org/10.3390/math7111104>