

*Answer-Set Programming for Lexicographical Makespan Optimisation in Parallel Machine Scheduling**

THOMAS EITER, TOBIAS GEIBINGER, NYSRET MUSLIU and
JOHANNES OETSCH

Institute of Logic and Computation, Vienna University of Technology (TU Wien), Austria
(e-mails: eiter@kr.tuwien.ac.at, tgeibing@dbai.tuwien.ac.at, musliu@dbai.tuwien.ac.at,
oetsch@kr.tuwien.ac.at)

PETER SKOČOVSKÝ and DARIA STEPANOVA

Bosch Center for AI,
Robert Bosch Campus 1, D-71272 Renningen, Germany
(e-mails: fixed-term.peter.skocovsky@de.bosch.com, daria.stepanova@de.bosch.com)

submitted 23 March 2022; revised 20 September 2022; accepted 2 January 2023

Abstract

We deal with a challenging scheduling problem on parallel machines with sequence-dependent setup times and release dates from a real-world application of semiconductor work-shop production. There, jobs can only be processed by dedicated machines, thus few machines can determine the makespan almost regardless of how jobs are scheduled on the remaining ones. This causes problems when machines fail and jobs need to be rescheduled. Instead of optimising only the makespan, we put the individual machine spans in non-ascending order and lexicographically minimise the resulting tuples. This achieves that all machines complete as early as possible and increases the robustness of the schedule. We study the application of answer-set programming (ASP) to solve this problem. While ASP eases modelling, the combination of timing constraints and the considered objective function challenges current solving technology. The former issue is addressed by using an extension of ASP by difference logic. For the latter, we devise different algorithms that use multi-shot solving. To tackle industrial-sized instances, we study different approximations and heuristics. Our experimental results show that ASP is indeed a promising knowledge representation and reasoning (KRR) paradigm for this problem and is competitive with state-of-the-art constraint programming (CP) and Mixed-Integer Programming (MIP) solvers.

KEYWORDS: answer-set programming, parallel machine scheduling, lexicographical optimisation

1 Introduction

We consider a scheduling problem on unrelated parallel machines that arises in industrial semiconductor production at Bosch. The problem is involved due to several aspects. We

*This is an extended version of a paper (Eiter *et al.* 2021) that appeared in the Proceedings of the 18th International Conference on Principles of Knowledge Representation and Reasoning (KR'21).

have to deal with sequence-dependent setup times and job release dates on the one hand; on the other hand, setup times, release dates, and job durations are machine dependent, and jobs can only be scheduled on dedicated machines. Our goal is to maximise the *throughput*, which can be defined as the number of jobs processed per time unit. In principle, minimising the *makespan*, that is, the total length of the schedule, accomplishes this. However, we have to deal with *high machine dedication*, that is, many jobs can be processed only by few machines. When dealing with machines with high dedication, few machines determine the makespan almost regardless of how jobs are scheduled on the remaining ones. This is not ideal when jobs need to be rescheduled because of, for example, machine failure, and domain experts expressed the requirement that “*all machines should complete as early as possible*” to give the scheduler freedom for rearrangements.

Instead of optimising only the makespan, we lexicographically minimise the individual machine spans. In particular, we define the *lexicographical makespan* of a schedule as the tuple of all machine spans in non-ascending order. We prefer a schedule with smaller lexicographical makespan over one with a larger one, where we use lexicographical order for comparison. A schedule with minimal lexicographical makespan has therefore also a minimal makespan, but ties are broken using machines that complete earlier.

This idea of *lexicographical optimisation* to produce schedules that show to be robust when they need to be updated has been investigated and confirmed in the context of job scheduling on identical machines in recent work (Letsios et al. 2021). While these results further support our motivation to use this objective, the algorithms and tool chains developed there cannot be used directly as our problem is significantly more complex due to machine dedication, sequence-dependent setup times, and release dates.

We are specifically interested in using *Answer-Set Programming* (ASP) (Brewka et al. 2011; Gebser et al. 2012; Lifschitz 2019), a state-of-the-art logic-based knowledge representation and reasoning (KRR) paradigm, for our scheduling problem. ASP is interesting for two reasons: first, its expressive modelling language makes it easy to concisely model the problem including the objective function. This allows one to quickly come up with a first prototype that can be evaluated by domain experts and can serve as a blueprint for more sophisticated solutions. Second and more importantly, ASP makes it relatively easy to develop solutions which can be conveniently adapted to problem variations, a feature known as *elaboration tolerance* (McCarthy 1998). This is indeed needed as it is a goal to use similar scheduling solutions for other work centers with slightly different requirements.

While ASP makes modelling easy and provides enough flexibility for future adaptations, the combination of timing constraints and the considered objective function challenges current solving technology. The former issue is addressed by using an extension of ASP with difference logic (Janhunen et al. 2017). While this solves the issue of dealing with integer domains without blowing up the size of the grounding, it makes expressing optimisation more tricky as current technology does not support complex optimisation of integer variables. We devise different algorithms that use multi-shot solving (Gebser et al. 2019) to accomplish lexicographical optimisation for our scheduling problem.

To tackle industrial-sized instances, we study different approximations and heuristics. In particular, we consider an approximate algorithm where parts of a solution are fixed after solver calls. This allows us to find near-optimal solutions in a short time. Orthogonally to the ASP model, we formulate different *heuristic rules* using a respective ASP

extension (Gebser *et al.* 2013). These rules do not alter the solution space but guide the solver with variable assignments and improve performance.

For an experimental evaluation of our algorithms, we use random instances of various sizes that are generated based on real-life scenarios. In addition, we provide an alternative solver-independent MiniZinc model that can be used by various state-of-the-art mixed-integer programming (MIP) and constraint programming (CP) solvers. The experiments aim to explore the additional costs needed when using the lexicographical makespan for optimisation instead of the standard makespan and the trade-off between performance and solution quality. Our experimental results show that the lexicographical makespan optimisers produce schedules with small makespans and thus ensure high throughput, while at the same time accomplish our other objective of early completion times for all machines. The ASP-based algorithms scale up to instances of realistic size and demonstrate that ASP is indeed a viable KRR solving paradigm for lexicographical makespan problems.

This article is an extended version of a conference paper that has been presented at KR 2021 (Eiter *et al.* 2021). In particular, it includes proofs, further encodings and additional experimental data that were not contained in the original publication, as well as an extended discussion and other minor extensions.

The rest of this paper is organised as follows. After some background on ASP in the next section, we formally define our scheduling problem including the lexicographical makespan objective in Section 3. We then present exact ASP approaches for lexicographical makespan minimisation in Section 4 and discuss approximation approaches in Section 5. Experiments are discussed in Section 6. We review relevant literature in Section 7, before we conclude in Section 8.

2 Background on ASP

Answer-Set Programming (ASP) (Brewka *et al.* 2011; Gebser *et al.* 2012; Lifschitz 2019) provides a declarative modelling language that allows one to succinctly represent search and optimisation problems, for which solutions can be computed using dedicated ASP solvers.^{1,2}

ASP is a compact relational, in essence propositional, formalism where variables in the input language are replaced by constant symbols in a preprocessing step called *grounding*. An ASP program is a (finite) set of rules of the form

$$p_1 \mid \dots \mid p_k \text{ :- } q_1, \dots, q_m, \text{ not } r_1, \dots, \text{ not } r_n. \quad (1)$$

where all p_i , q_j , and r_l are function-free first-order atoms.³ The head are all atoms before the implication symbol :- , and the body are all the atoms and negated atoms afterwards. The intuitive meaning of the rule (1) is that if all atoms q_1, \dots, q_m can be derived, and there is no evidence for any of the atoms r_1, \dots, r_n (i.e., the rule fires) then at least one of p_1, \dots, p_k has to be derived.

A rule with an empty body (i.e., $m = n = 0$) is called a *fact*, with :- usually omitted. Facts are used to express knowledge that is unconditionally true. A rule with empty head

¹potassco.org.

²www.dlvsystem.com.

³Extensions of ASP with strong negation as well as function symbols, both uninterpreted and interpreted ones, are available, which however we disregard here.

(i.e., $k = 0$) is a *constraint*. The body of a constraint cannot be satisfied by any answer set and is used to prune away unwanted solution candidates.

The semantics of an ASP program P is given in terms of particular Herbrand models of its grounding $grd_C(P) = \bigcup_{r \in P} grd_C(r)$, where $grd_C(r)$ consists of all ground (variable-free) rules $r\theta$ that result from r by some substitution θ of the variables in r with elements from the set C of constants, which by default are the constants occurring in P . A (*Herbrand*) *interpretation* I is a subset of the set $HB_C(P)$ of all ground atoms $p(c_1, \dots, c_l)$ with a predicate p occurring in P and constants c_1, \dots, c_l from C . It is a (*Herbrand*) *model* of P , if each rule r in $grd_C(P)$ of form (1) is satisfied, that is, either (i) $\{p_1, \dots, p_k\} \cap I \neq \emptyset$ or (ii) $\{q_1, \dots, q_m\} \not\subseteq I$ or (iii) $\{r_1, \dots, r_n\} \cap I \neq \emptyset$. Then I is an *answer set* of P , if I is a \subseteq -minimal model of the reduct P^I of P by I , which is given by

$$P^I = \{p_1 \mid \dots \mid p_k \text{ :- } q_1, \dots, q_m \mid r \in P, \{r_1, \dots, r_n\} \cap I = \emptyset\};$$

intuitively, I must be a \subseteq -minimal model of all rule instances of P whose negative body is satisfied by I .

A common syntactic extension are choice rules of the form

$$i \{p_1, \dots, p_k\} j \text{ :- } q_1, \dots, q_m, \text{ not } r_1, \dots, \text{ not } r_n.$$

The meaning is that if the rule fires, then some subset S of p_1, \dots, p_k with $i \leq |S| \leq j$ has to be true as well; it is compiled to ordinary rules using hidden auxiliary predicates (Calimeri et al. 2020).

We use the hybrid system `clingo-dl` (Janhunen et al. 2017)⁴ that extends the ASP solver `clingo` by difference logic to deal with timing constraints. A difference constraint is an expression of the form $u - v \leq d$, where u and v are integer variables and d is an integer constant. In contrast to unrestricted integer constraints, systems of difference constraints are solvable in polynomial time. The latter are expressed in `clingo-dl` using *theory atoms* (Gebser et al. 2016). That job j starts after its release time, say 10, can be expressed as `&diff{ 0 - start(j) } <= -10`. Here, `0` and `start(j)` are integer variables, where `0` has a fixed value of 0; thus, `start(j)` must be at least 10.

We will also use an extension for *heuristic-driven solving* (Gebser et al. 2013) that allows to incorporate domain heuristics into ASP. These heuristics do not change the answer sets of a program but modify internal solver heuristics to bias search. The general form of a heuristic directive is

$$\#heuristic A : B. [w@p,m],$$

where A is an atom, B is a rule body, and w, p, m are terms. In particular, w is a weight, $@p$ is an optional priority, and m is a modifier like `true` or `false`. We will provide further details when we introduce specific heuristic rules later on.

Further features of the input language, like aggregation and optimisation statements, will be explained as we go.

3 Problem statement

We study the following scheduling problem. Given m machines and n jobs, every job needs to be processed by a single machine, and every machine can process at most one

⁴<https://github.com/potassco/clingo-dl>.

job at a time; preemption is not allowed. Some machines can only handle certain jobs, such that from the view of the latter, $cap(j)$ is the set of machines that can process job j .

We assume that a *release date* $r_{j,k}$ is specified for every job j and machine k as a non-negative integer. Release dates are machine dependent because transportation time for jobs to the machines depends on the transport system and their location. No job can start before its release date.

A specified amount of time may be required to change from one job to the next one. Specifically, we assume that $s_{i,j,k}$ is the time needed to set up job j directly after job i on machine k . Consequently, these times are referred to as *sequence-dependent setup times*. Every job j has a positive *duration* $d_{j,k}$ that depends on the machine k it is assigned to.

A *schedule* S for a problem instance is defined by:

1. an assignment a that maps each job j to a machine $k \in cap(j)$ capable of processing it;
2. for each machine k , a total order \preceq_k on the set J of jobs assigned to the machine via a . Relation \preceq_k determines the sequence in which the jobs in J are processed on k .

If each job can be processed by some machine, then some schedule for the problem instance exists.

Assume that j_1, \dots, j_l is the processing sequence of the jobs assigned to machine k in a given schedule. The *processing time* p_{j_i} of a job j_i is its duration plus the setup time for its predecessor (if one exists); that is, $p_{j_1} = d_{j_1,k}$ and $p_{j_i} = s_{j_{i-1},j_i,k} + d_{j_i,k}$, for $i > 1$. The *start time* st_{j_i} of job j_i is $r_{j_i,k}$ if $i = 1$, and $\max(r_{j_i,k}, st_{j_{i-1}} + p_{j_{i-1}})$ for $i > 1$. The *completion time* c_{j_i} of job j_i is $st_{j_i} + p_{j_i}$.

The *machine span* of k , $span(k)$, is the completion time of the last job j_l on k . A common optimisation criterion is to search for a schedule with a small *makespan*, which is the largest machine span of the schedule.

Versions of this problem have been extensively studied in the literature (Allahverdi 2015). The problem presented here abstracts the actual problem statement at Bosch to its most essential elements. We left out some details due to confidentiality. Other elements, like due dates or manual labor costs, have been omitted as they are not relevant for the objective function studied in this paper.

3.1 The lexicographical makespan objective

Recall that we are interested in computing schedules that maximise the throughput, but high machine dedication and rescheduling due to sudden machine failure render minimal makespan as a single objective function suboptimal.

Figure 1(a) illustrates this: assume all jobs scheduled to machine m_1 cannot be processed by any other machine. Thus, m_1 will always determine the makespan and the remaining jobs can be put almost arbitrarily on the remaining machines. This can lead to unnecessary workload on these machines. A more severe problem is when machines suddenly fail and their jobs need to be rescheduled. To cope with events like machine failure, the domain experts formulated the requirement that “all machines should



Fig. 1. Different schedules involving three machines and six jobs.

complete as early as possible” with the intention to give the scheduler maximal freedom in rearranging jobs with minimal decrease in throughput.

We next define the *lexicographical makespan* for lexicographical optimisation of machine spans to obtain robust schedules (Letsios et al. 2021).

Definition 1

Given a schedule S involving m machines, the *lexicographical makespan*, or *lex-makespan* for short, of S is the tuple $ms(S) = (c_1, \dots, c_m)$ of all the machine spans of S in non-ascending order.

In this definition, c_1 is a maximal machine span and hence corresponds to the makespan.

For schedules S and S' involving m machines each, S has a smaller lex-makespan than S' if $ms(S)$ is smaller than $ms(S')$ under lexicographical order, that is, on the least index i where $ms(S) = (c_1, \dots, c_m)$ and $ms(S') = (c'_1, \dots, c'_m)$ disagree, we have $c_i < c'_i$. For a set \mathcal{S} of schedules, $S \in \mathcal{S}$ is then optimal if $ms(S)$ is minimal over all schedules in \mathcal{S} .

Consider Figure 1 for illustration. We would prefer schedule (b) over (c) under the lex-makespan objective. For both schedules, the lex-makespan is given by the machine spans of m_1 , m_2 , and m_3 in that order. Both schedules have the same makespan, but schedule (b) has a smaller machine span for m_2 . If machine m_1 fails and most of the jobs can only be rescheduled to machine m_2 , schedule (b) would indeed be advantageous. It happens also earlier for schedule (b) that machines m_2 and m_3 complete all their jobs and are therefore free if new jobs need to be scheduled.

To describe the dynamics of a schedule, we define, for a time point t and schedule S , $M(S, t)$ as the number of machines that complete at or before t . We then obtain:

Proposition 1

Let S and S' be two schedules for some problem instance. Then, $ms(S) < ms(S')$ iff there is a time point t such that $M(S, t) > M(S', t)$, and, for every $t' > t$, $M(S, t') \geq M(S', t')$.

Proof

Let $s = (c_1, \dots, c_m)$ be the lexical makespan of S and $s' = (d_1, \dots, d_m)$ the one of S' . Let $(*)$ stand for the right-hand side of the proposition: there is a time point t such that $M(S, t) > M(S', t)$ and for any $t' > t$, $M(S, t') \geq M(S', t')$.

(\Rightarrow) Assume that $s < s'$. Consider the least $i \in \{1, \dots, m\}$ with $c_i < d_i$. Observe that $M(S, c_i) \geq m - i + 1$ while $M(S', c_i) = m - i$. Thus, there is a time point $t = c_i$ such that $M(S, t) > M(S', t)$. Since $c_j = d_j$ for any $j < i$, it follows that, for any time point $t' > t$, $M(S, t') \geq M(S', t')$. Hence, $(*)$ holds.

(\Leftarrow) Towards a contradiction, assume that both $s \geq s'$ and $(*)$ hold. We distinguish between the two cases $s = s'$ and $s' < s$. If $s = s'$ then there cannot be any time point t with $M(S, t) > M(S', t)$, a contradiction to $(*)$. If $s' < s$, the left-to-right side of the proposition implies that there is a time point t such that $M(S', t) > M(S, t)$, and, for any $t' > t$, $M(S', t') \geq M(S, t')$. Again, a contradiction to $(*)$. \square

For problems involving many machines, hierarchically minimising all the machine spans can be excessive if the overall makespan is dominated by few machines only. However, comparing lex-makespans allows for a rather natural parametrisation, namely an integer l that defines the number of components to consider in the comparison.

Definition 2

Given schedules S and S' involving m machines each and an integer l , $1 \leq l \leq m$, we say $ms(S) = (c_1, \dots, c_m)$ is smaller than $ms(S') = (c'_1, \dots, c'_m)$ under *parametrised lexicographical order*, in symbols $ms(S) \leq_l ms(S')$, if under lexicographical order $(c_1, \dots, c_l) \leq (c'_1, \dots, c'_l)$.

Note that for a schedule with m machines, we obtain the makespan objective if $l = 1$ and the full lex-makespan if $l = m$.

4 An exact ASP model with difference logic

A problem instance is described by ASP facts using some fixed predicate names. We illustrate this by an example with one machine m_1 and two jobs j_1, j_2 . The machine is capable of processing all jobs and all release dates are 0. The setup time is 4 when changing from job j_1 to j_2 and 2 vice versa. Both jobs have duration 5. The according facts are

```
machine(m1).
cap(m1, j1). cap(m1, j2).
job(j1). duration(j1, m1, 5). release(j1, m1, 0).
job(j2). duration(j2, m1, 5). release(j2, m1, 0).
setup(j1, j2, m1, 4). setup(j2, j1, m1, 2).
```

Any problem instance can be described using this format.

Next, we present the ASP encoding for computing minimal schedules. The entire program is given in Figure 2.

The encoding consists of three parts: Lines 1–10 qualitatively model feasible sequences of jobs on machines, while the quantitative model for completion times is realised in Lines 12–22 with difference logic; we avoid by this doing integer arithmetic in the Boolean ASP constraints, which would blow up the size of the grounding. Finally, the optimisation is accomplished by Lines 24–26. Intuitively, we guess an upper bound for each machine span, which is then minimised. The directive in Line 26 assigns each span a priority equal to their value, thus ensuring that the highest span is the most important followed by the second highest and so forth.

The first line of Figure 2 expresses that each job is assigned to a machine capable of processing it. The notation $\mathbf{asg}(J, M) : \mathbf{cap}(M, J)$ means that in the grounding step for each value j of the global variable J (as it occurs in the body), $\mathbf{asg}(J, M)$ is replaced by all atoms $\mathbf{asg}(j, m)$ for which $\mathbf{cap}(j, m)$ can be derived.

```

1 1 {asg(J,M): cap(M,J)} 1 :- job(J).
2
3 before(J1,J2,M) | before(J2,J1,M) :- asg(J1,M), asg(J2,M), J1 < J2.
4
5 1 {first(J,M): asg(J,M)} 1 :- asg(_,M).
6 1 {last(J,M): asg(J,M)} 1 :- asg(_,M).
7 1 {next(J1,J2,M): before(J1,J2,M)} 1 :- asg(J2,M), not first(J2,M).
8 1 {next(J2,J1,M): before(J2,J1,M)} 1 :- asg(J2,M), not last(J2,M).
9 :- first(J1,M), before(J2,J1,M).
10 :- last(J1,M), before(J1,J2,M).
11
12 &diff{0 - c(J1)} <= -(T+D+S) :- asg(J1,M), next(J3,J1,M),
13                                     setup(J3,J1,M,S), duration(J1,M,D),
14                                     release(J1,M,T).
15 &diff{c(J2) - c(J1)} <= -(P+S) :- before(J2,J1,M), next(J3,J1,M),
16                                     setup(J3,J1,M,S),
17                                     duration(J1,M,P).
18 &diff{0 - c(J1)} <= -(T+D) :- asg(J1,M), duration(J1,M,D),
19                                     release(J1,M,T).
20 &diff{c(J) - cmax} <= 0 :- job(J).
21
22 &diff{c(J2) - c(J1)} <= -P :- before(J2,J1,M), duration(J1,M,P).
23
24 1 {span(M,T): int(T)} 1 :- machine(M).
25 &diff{c(J) - 0} <= S :- asg(J,M), span(M,S).
26 #minimize{T@T,M: span(M,T)}.

```

Fig. 2. ASP encoding with difference logic for lex-makespan optimisation.

We further require that the jobs assigned to a machine are totally ordered. That is, for any two distinct such jobs j_1 and j_2 , either $j_1 \prec j_2$ or $j_2 \prec j_1$ holds. This is achieved by the rule in Line 3. In Lines 5–6, the predicates `first/2` and `last/2`, representing the first and last job on each machine, respectively, are defined. Constraints in Lines 9–10 ensure that this selection is compatible with the order given by `before/3`. Furthermore, each job except the last (resp. first) has a unique successor (resp. predecessor); this is captured by `next/3` in Lines 7–8.

We use difference logic to express that jobs are put on the machines in the order defined by `next/3`. The rules in Lines 12–19 closely follow respective definitions from Section 3. Line 20 defines `cmax` as an upper bound of any completion time. In any answer-set, `cmax` will be the actual makespan since the solver will always instantiate integer variables with the smallest value possible. The redundant rule in Line 22 helps the solver to further prune the search space.

For optimisation, we “guess” a span for each machine in Line 24. Here `int/1` is assumed to provide a bounded range of integers. In Line 25, we enforce that machines complete not later than the guessed spans. The actual objective function is defined by the last line of Figure 2 notably concise: any machine contributes its span c to a cost function at priority level c . The cost function accumulates contributing values and the solver minimises answer sets by lexicographically comparing cost tuples ordered by priority.

The formal correctness of the encoding is assured by the following proposition.

Proposition 2

For every problem instance I , the schedules of I with minimal lex-makespan are in one-to-one correspondence with the optimal answer sets of the program P_I consisting of the rules in Figure 2 augmented with the fact representation of I .

Proof (Sketch)

Given a (correct) schedule S for a problem instance I , we can define an interpretation M_I that is an answer set of the program P_I . To this end, we include in M_I all atoms `asg(j,k)` such that job j is assigned to machine k in S ; all atoms `before(j1,j2,k)` such that jobs $j1$ and $j2$ are assigned in S to the same machine k and $j1$ starts before $j2$; all atoms `next(j1,j2,k)` such that jobs $j1$ and $j2$ are assigned in S to the same machine k and $j1$ starts immediately before $j2$; the atoms `first(j,k)` and `last(j,k)` where j is in S the first, respectively, last job on machine k ; and the atom `span(k,t)` for each machine k , where t is $span(k)$ according to S . Furthermore, each integer variable $c(j)$ is set to the completion time of job j in S ; `cmx` equals the maximal completion time. It can then be shown that M_I is an answer set of P , according to the semantics of `clingo-dl`. In turn, it can be shown that every answer set M of the program P_I encodes a correct schedule S for the problem instance I , which is given by the atoms `asg(j,k)`, `before(j1,j2,k)` in M .

Thus, the correct schedules for I correspond to the answer sets of P_I . By the statement `#minimize{TOT,M: span(M,T)}`, those answer sets M are optimal where the vector $v_M = (span(k_1), span(k_2), \dots, span(k_m))$ of all machine spans in the encoded schedule in decreasing order, that is, $span(k_j) \geq span(k_{j-1})$ for all $m \geq j > 1$, is lexicographical minimal; indeed, if vector v_M is smaller than vector $v_{M'}$, then the value of the objective function for M is smaller than the one for M' . Consequently, the optimal answer sets M of P_I correspond one-to one to the schedules for I with minimal lex-makespan. \square

4.1 Direct multi-shot optimisation

The performance bottleneck for the ASP approach from the previous section is grounding. In particular, the definition of the machine spans must be grounded over the entire relevant integer range. We can define machine spans in difference logic as bounds on completion times similar to the makespan:

```
&diff{ 0 - span(M)} <= 0 :- machine(M).
&diff{c(J) - span(M)} <= 0 :- asg(J,M).
```

However, multi-objective optimisation for integer variables with priorities is unfortunately not supported in the current version (1.1.1) of `clingo-dl`. Schedules with minimal lex-makespan are still computable using multiple solver calls and incrementally adding constraints.

To this end, we present Algorithm 1 for lex-makespan minimisation by using multiple solver calls.

Algorithm 1 is a standard way for multi-objective minimisation by doing a (highest priority first) hierarchical descent. Due to symmetries, showing a lack of solutions is usually more costly for this problem than finding one; this makes alternative strategies with fewer expected solver calls like binary search or exponentially increasing search steps less attractive.

We use `clingo-dl` and the encoding from Figure 2 without the optimisation statement in Lines 21–23 to implement $solve(M)$ in Algorithm 1. A handy feature is that `clingo-dl` supports *multi-shot solving* (Gebser *et al.* 2019) where parts of the solver state are kept

Algorithm 1: Lex-Makespan Optimisation**Input:** model M involving m machines and parameter l with $1 \leq l \leq m$ **Output:** schedule R for M with parametrised lex-makespan (c_1, \dots, c_l) $solve(M) \dots$ returns a solution for M or \emptyset if none is found within fixed resource limits $bound(i \circ b), \circ \in \{<, \leq\} \dots$ constraints enforcing that $c_i \circ b$ for the lex-makespan (c_1, \dots, c_m)

```

1  $(c_1, \dots, c_l) \leftarrow (0, \dots, 0)$ 
2  $R \leftarrow solve(M)$ 
3 for  $i \leftarrow 1$  to  $l$  do
4    $S \leftarrow R$ 
5   while  $S \neq \emptyset$  do
6      $R \leftarrow S$ 
7      $c_i \leftarrow$   $i$ th element of the lex-makespan of  $S$ 
8      $S \leftarrow solve(M \cup bound(i < c_i))$ 
9    $M \leftarrow M \cup bound(i \leq c_i)$ 
10 return  $R$  with lex-makespan  $(c_1, \dots, c_l)$ 

```

throughout multiple runs, thereby saving computational resources. Notably, $solve(M)$ in Algorithm 1 does not limit us to use ASP solver. We can in principle use any exact method that is capable of producing solutions for a model M in the input language of the respective system.

The constraints for $bound(i \leq b)$ are quite easy to express in ASP: that the i th component of the lex-makespan is smaller than or equal to b is equivalent to enforcing that at least $m - i + 1$ machines have a span of at most b . We can encode the latter by non-deterministically selecting $m - i + 1$ machines and enforcing that they complete not later than b :

```

(m-i+1) {sel(M): machine(M)}.
&diff{span(M) - 0} <= b :- sel(M).

```

While Algorithm 1 is guaranteed to return a schedule with minimal lex-makespan when resources for $solve(M)$ are not limited, we will in practice restrict the time spent for search in $solve(M)$ by a suitable time limit.

4.2 Domain-specific heuristics

We use two domain heuristics to improve performance by guiding search more directly to promising areas of the search space. Both heuristic directives use the modifier **true**: Whenever an atom needs to be assigned a truth value, the solver will pick the one with the highest weight among the ones with the highest priority and assigns it to true at first.

```

1  1 {asg(J,M) : cap(M,J)} 1 :- job(J).
2  1 {first(J1,M) : asg(J1,M)} 1 :- #count{J2 : asg(J2,M)} > 0.
3  1 {last(J1,M) : asg(J1,M)} 1 :- #count{J2 : asg(J2,M)} > 0.
4  1 {next(J1,J2,M) : asg(J1,M)} 1 :- asg(J2,M), not first(J2,M).
5  1 {next(J1,J2,M) : asg(J2,M)} 1 :- asg(J1,M), not last(J2,M).
6  reach(J,M) :- first(J,M).
7  reach(J2,M) :- reach(J1,M), next(J1,J2,M).
8  :- asg(J,M), not reach(J,M).
9
10 process(J,T) :- first(J,M), duration(J,M,T).
11 process(J2,D+S) :- next(J1,J2,M), duration(J2,M,D),
12     setup(J1,J2,M,S).
13 start(J,T) :- int(T), first(J,M), release(J,M,T).
14 start(J2,T) :- int(T), next(J1,J2,M),
15     T = #max{R : release(J2,M,R) ; C : compl(J1,C)}.
16 compl(J,S+P) :- int(T), start(J,S), process(J,P).
17 span(M,T) :- int(T), last(J,M), compl(J,T).
18
19 #minimize{ T@T,M : span(M,T) }.

```

Fig. 3. Plain ASP encoding for lexical makespan minimisation.

Recall that job durations depend on the machines. The first heuristic expresses the idea to assign jobs to machines if their duration is low on that machine.

```

#heuristic asg(J,M): duration(J,M,D),
    maxDuration(J,F), W=F-D. [W@2,true]
maxDuration(J,M) :- job(J), M = #max{D: duration(J,_,D)}.

```

Here, `maxDuration/2` defines the longest duration of a given job over all machines. Then, the heuristic directive gives a high weight to an atom `asg(j,m)` if the duration of j is low relative to its maximal duration. The priority level of this rule is 2; this means that the solver will try to assign jobs to machines before deciding on other atoms.

The second heuristic directive affects how jobs are put on machines. We want to avoid large setup times and follow an analogous strategy as for the first heuristic:

```

#heuristic next(J,K,M): setup(J,K,M,S), maxSetup(K,M,T),
    cap(J,M), cap(K,M), W=T-S. [W@1,true]
maxSetup(J,M,S) :- job(J), machine(M),
    S = #max{T: setup(_,J,M,T)}.

```

Atom `next(j,k,m)` gets a high weight if putting job j before k results in a relatively small setup time. We also need to make sure that machine m is actually capable of processing both jobs. The order of the jobs has a lower priority than the machine assignment.

4.3 Plain clingo model

The ASP encoding in Figure 2 can also be formulated without difference logic which makes it suitable as input for the ASP solver `clingo`. This ASP program for computing minimal schedules is given in Figure 3. There, completion times are computed directly with arithmetic aggregates. While this approach works for small instances, it does not scale for larger integer domains.

Algorithm 2: Lex-Makespan Approximation**Input:** model M involving m machines and parameter l with $1 \leq l \leq m$ **Output:** schedule S for M with parametrised lex-makespan (c_1, \dots, c_l) $opt(M) \dots$ returns the best solution for M found within fixed resource limits

```

1  $(c_1, \dots, c_l) \leftarrow (0, \dots, 0)$ 
2  $S \leftarrow \emptyset$ 
3 for  $i \leftarrow 1$  to  $l$  do
4    $S \leftarrow opt(M)$ 
5    $c_i \leftarrow$  makespan of  $S$ 
6   remove some machine  $k$  that completes at  $c_i$  and all jobs assigned to  $k$  from  $M$ 
7 return  $S$ 

```

5 ASP-based approximation

While the exact methods from the previous section have the advantage that we can run a solver until we find a guaranteed optimal solution, this works only for very small problem instances. Finding good solutions within a time limit is in practice more important than showing optimality. This is what the ASP-based approximation method we discuss next are designed to accomplish.

There is a simple way to turn the exact encoding from Figure 2 into an approximation that scales better to larger instances. It has been introduced for `clingo-dl` optimisation in the context of train scheduling (Abels et al. 2019), and we apply it for our machine scheduling application. Recall that we use `int/1` to define a range $[0, 1, \dots, n]$ of integers from which potential bounds for individual machine spans are taken from. We can instead consider integers from $[0, 1 \cdot g, \dots, i \cdot g]$ where g is the granularity of the approximation, and i is chosen such that $n \leq i \cdot g$; i can be significantly smaller than n , depending on g , and thus reduce the search space and the size of the grounding. A larger g makes the approximation more coarse, a smaller g makes it closer to the exact encoding. We will compare the exact encoding and this approximation in Section 6.

5.1 Multi-shot approximation

We next present a variant of Algorithm 1 for approximation, where we assume that we have an exact optimiser $opt(\cdot)$ which is good at finding schedules with small makespans. This optimiser can then be employed to compute schedules with small lex-makespans. The specifics are presented as Algorithm 2.

Algorithm 2 uses $opt(\cdot)$ to recompute and improve parts of a solution by fixing the jobs on the machine with the highest span after each solver call. Similar to Algorithm 1, we require multiple solver calls but with a profound difference: the number of solver calls to the makespan optimiser is bounded by the number of machines, and after each solver call, the problem instance is significantly simplified and thus easier to solve.

As `clingo-dl` allows to directly minimise a single integer variable, we can implement the makespan optimiser $opt(\cdot)$ directly using our difference logic encoding and minimise `cmax`. However, we opted for using multi-shot solving again to reuse heuristic values and learned clauses from previous solver runs.

Table 1. *Machines (m) and jobs (n) per instance class*

	M3	M5	M10	M15	M20
<i>m</i>	3	5	10	15	20
<i>n</i>	5–50	10–50	50–200	100–200	150–200

6 Experimental evaluation

We now provide an experimental evaluation of the approaches for lex-makespan optimisation from above. Notably, any approach that produces schedules with small makespan will also produce small lex-makespans. Our primary goal is to investigate the difference in schedule quality when spending all resources for makespan optimisation versus distributing them for lex-makespan optimisation to different priorities. As we cannot disclose real instances from the semi-conductor production application, we use random instances of realistic size and structure instead. In addition to experiments with ASP-based solvers, to compare with other solving paradigms, we provide a solver-independent model in MiniZinc (Nethercote *et al.* 2007). This model can be solved by many CP and MIP solvers. However, for a direct comparison, we also model our problem directly in `cplex` and `coptimizer`.

6.1 Problem instances

We generated 500 benchmark instances of different sizes randomly. The random instance generator is designed however to reflect relevant properties of the real instances. The generator is based on previous work in the literature (Vallada and Ruiz 2011b), but also produces instances with high machine dedication and amends older benchmarks, which were designed for different objective, with random release dates. The 500 instances can be grouped into five classes, shown in Table 1, of 100 instances each. The instance generator and all the encoding and algorithms are online available.⁵

The machine capabilities were assigned uniformly at random for half of the instances in every class: for each job, a random number of machines were assigned as capable. For the other half, we assigned the capabilities such that 80% of the jobs can only be performed by 20% of the machines. We refer to the latter setting as high-dedication and the former as low-dedication.

For each job j and any machine k , the duration $d_{j,k}$, setup time $s_{j,i,k}$ for any other job i , and release date $r_{j,k}$ were drawn uniformly at random from $[10, 500]$, $[0, 100]$, and $[0, r_{max}]$, respectively, where

$$r_{max} = \frac{1}{m} \sum_{1 \leq j \leq n} \frac{1}{|cap(j)|} \left(\sum_{k \in cap(j)} d_{j,k} + \sum_{1 \leq j' \leq n, k \in cap(j')} s_{j',j,k} \right).$$

6.2 A solver-independent MiniZinc model

As an alternative to ASP, we implemented a solver-independent model for schedule optimisation in the well-known high-level modelling language MiniZinc for constraint

⁵<http://www.kr.tuwien.ac.at/research/projects/bai/tplp22.zip>.

$$\text{alldifferent_except_0}(p_{1 \leq i \leq n}) \tag{2}$$

$$\sum_{i=1}^n (p_i = 0) \leq \text{nvalue}(a_{1 \leq i \leq m}) \tag{3}$$

$$a_i \in \text{cap}(i) \tag{4} \quad 1 \leq i \leq n$$

$$p_i \neq 0 \rightarrow a_i = a_{p_i} \tag{5} \quad 1 \leq i \leq n$$

$$p_i \neq i \tag{6} \quad 1 \leq i \leq n$$

$$c_i \geq (\text{max}(c_j, r_{i,a_i}) + s_{j,i,a_i} + d_{i,a_i}) \cdot (p_i = j) \tag{7} \quad 1 \leq i, j \leq n, i \neq j$$

$$c_i \geq r_{i,a_i} + d_{i,a_i} \tag{8} \quad 1 \leq i \leq n$$

$$s_k = \text{max}_{1 \leq i \leq n} (c_i \cdot (a_i = k)) \tag{9} \quad 1 \leq k \leq m$$

$$\text{alldifferent}(l_{1 \leq i \leq m}) \tag{10}$$

$$l_i > l_j \rightarrow s_i \geq s_j \tag{11} \quad 1 \leq i, j \leq m$$

Fig. 4. Solver-independent MiniZinc model for lex-makespan optimisation.

satisfaction and optimisation problems. MiniZinc models, after being compiled into FlatZinc, can be used by a wide range of solvers. We provide a direct model that represents our best attempt at using MiniZinc. It thus serves only as a first baseline and more advanced models may improve performance. Our model of the problem statement from Section 3 and the objective function are as follows.

For each job $i \in \{1, \dots, n\}$, we use the following decision variables: $a_i \in \{1, \dots, m\}$ for its assigned machine, $p_i \in \{0, \dots, n\}$ representing its predecessor or 0 if it has none, and $c_i \in \{0, \dots, h\}$ denoting its completion time where h is the scheduling horizon. For each machine $j \in \{1, \dots, m\}$, we use $s_j \in \{1, \dots, h\}$ denoting its span, and its level in the ordering of spans $l_j \in \{1, \dots, m\}$.

The constraints of the MiniZinc model are given in Figure 4. The global Constraint (2) ensures that no two jobs have the same predecessor, while (3) enforces that the number of first jobs is less or equal to the number of assigned machines. The latter is determined through a global `nvalue` constraint returning the number of different values in $a_{1 \leq i \leq m}$. Constraint (4) ensures that every job is assigned a capable machine, and Constraint (5) ensures that a job’s predecessor is on the same machine. Constraint (6) expresses that no job is its own predecessor, while (7) ensures that each job starts after its predecessor and the corresponding setup time. Finally, (8) enforces that every first job starts after its release date.

Defining an objective function for lex-makespan minimisation is more intricate. For this, we use Constraints (9–11). Here (9) defines the span for each machine to be the latest completion time of any job scheduled on it, while (10–11) ensure that each machine is assigned a different level and the levels order the machines with respect to their spans.

The objective function for minimising the lex-makespan can then be expressed as

$$\text{min} \sum_{i=1}^m h^{l_i-1} \cdot s_i.$$

Intuitively, the levels represent the priorities for optimisation. By assigning the span of machine i the weight h^{l_i-1} , it is more important than all spans of machines on lower levels.

$$\sum_{1 \leq i \leq m} a_{j,i} = 1 \quad 1 \leq j \leq n \quad (12)$$

$$\sum_{i \in \text{cap}(j)} a_{j,i} = 1 \quad 1 \leq j \leq n \quad (13)$$

$$\sum_{1 \leq j \leq n} p_{j,0,i} \leq 1 \quad 1 \leq i \leq m \quad (14)$$

$$\sum_{0 \leq k \leq n} p_{j,k,i} = a_{j,i} \quad 1 \leq j \leq n, 1 \leq i \leq m \quad (15)$$

$$\sum_{1 \leq k \leq n} p_{k,j,i} = a_{j,i} \quad 1 \leq j \leq n, 1 \leq i \leq m \quad (16)$$

$$c_j + h \cdot (1 - p_{j,k,i}) \geq c_k + d_{j,i} + s_{k,j,i} \quad 1 \leq j, k \leq n, 1 \leq i \leq m \quad (17)$$

$$c_j \geq (r_{j,i} + d_{j,i}) \cdot a_{j,i} \quad 1 \leq j \leq n, 1 \leq i \leq m \quad (18)$$

$$c_j + h \cdot (1 - p_{j,k,i}) \geq r_{j,i} + d_{j,i} + s_{k,j,i} \quad 1 \leq j, k \leq n, 1 \leq i \leq m \quad (19)$$

$$s_i^M + h \cdot (1 - a_{j,i}) \geq c_j \quad 1 \leq j \leq n, 1 \leq i \leq m \quad (20)$$

$$(1 - l_{i,k}) \cdot h + s_k^L \geq s_i^M \quad 1 \leq k \leq l, 1 \leq i \leq m \quad (21)$$

$$s_i^L \geq s_j^L \quad 1 \leq i < j \leq l \quad (22)$$

$$\sum_{1 \leq k \leq l} l_{i,k} = 1 \quad 1 \leq i \leq m \quad (23)$$

$$\sum_{1 \leq i \leq m} l_{i,k} = 1 \quad 1 \leq k \leq l \quad (24)$$

Fig. 5. MIP model used for lex-makespan optimisation in `cplex`.

Note that restricting the model to Constraints (2-8) and using the objective $\min \max_{1 \leq i \leq n}(c_i)$ expresses the scheduling problem with the standard makespan objective.

CPLEX and CP optimizer models

We also encoded our problem in the modelling languages of `cplex` and `coptimizer`. As for the MiniZinc model, the encodings represent our best attempts at using respective modelling languages. The MIP model for `cplex` uses the following variables.

- $a_{j,i} \in \{0, 1\}$ indicating that job j is assigned to machine i ,
- $p_{j,k,i} \in \{0, 1\}$ representing that k is the predecessor of j (k can be zero) on machine i ,
- c_j the completion time of job j ,
- s_i^M the span of machine i ,
- s_l^L the span of level l , and
- $l_{i,l} \in \{0, 1\}$ indicating that machine i has level l .

The last variable is obviously only needed for the lex-makespan objective.

The model with all constraints is given in Figure 5. Constraints (12) and (13) ensure that each job is assigned exactly one capable machine. Constraint (14) enforces that each machine has at most one first job. Constraint (15) states that each job is either the first job on its assigned machine or has exactly one predecessor. The next Constraint (16)

$$\begin{aligned}
 p_j = 0 \vee p_k = 0 \vee p_j \neq p_k & & 1 \leq j < k \leq n & (25) \\
 p_j \neq j & & 1 \leq j \leq n & (26) \\
 a_j \in \text{cap}(j) & & 1 \leq j \leq n & (27) \\
 \sum_{1 \leq i \leq m} [\text{count}((a_j)_{1 \leq j \leq n}, i) > 0] \geq \text{count}((p_j)_{1 \leq j \leq n}, 0) & & & (28) \\
 p_j \neq k \vee a_j = a_k & & 1 \leq j, k \leq n, j \neq k & (29) \\
 c_j \geq (\max(r_{j,a_j}, c_k) + d_{j,a_j} + s_{k,j,a_j}) \cdot p_j = k & & 1 \leq j, k \leq n, j \neq k & (30) \\
 c_j \geq r_{j,a_j} + d_{j,a_j} & & 1 \leq j \leq n & (31) \\
 s_i^M + (h \cdot [a_j \neq i]) \geq c_j & & 1 \leq i \leq m, 1 \leq j \leq n & (32) \\
 s_i^L \geq s_j^L & & 1 \leq i < j \leq l & (33) \\
 l_i \neq l_j & & 1 \leq i < j \leq n & (34) \\
 [l_i \neq j] \cdot h + s_j^L \geq s_i^M & & 1 \leq i \leq m, 1 \leq j \leq l & (35)
 \end{aligned}$$

Fig. 6. Native `coptimizer` Model for lex-makespan optimisation.

enforces that each job has at most one successor on its machine. Constraint (17) ensures that each job starts after its predecessor. Job starting after their release is enforced through (18) and (19). Machine spans are ensured to be bigger than all completion times on the machine by Constraint (20). Constraints (21–24) make sure that the level spans correspond to the machine spans and are descending.

Optimising the lex-makespan objective can now be achieved by using the lexicographic minimisation functionality of `cplex` to minimise the level spans $s_i^L (1 \leq i \leq l)$. We can also use this model to optimise makespan by simply dropping Constraints (21–24) and minimising the maximum machine span.

The decision variables of the `coptimizer` model are of a slightly different form.

- a_j indicates the assigned machine of job j that is, $1 \leq a_j \leq m$,
- p_j is the predecessor of j or zero if there is none,
- c_j the completion time of job j ,
- s_i^M the span of machine i ,
- s_i^L the span of level l , and
- l_i indicating the level of machine i .

The constraints for `coptimizer` are given in Figure 6⁶. Constraint (25) ensures that predecessors are all different or zero. The next Constraint (26) enforces that no job is its own predecessor. Constraint (12) makes sure that each job is assigned a capable machine. The number of first jobs is bounded by the number of assigned machines through Constraint (14). Constraint (29) ensures that the predecessor of a job is assigned the same machine. Jobs starting after their predecessor and their release is enforced through (30) and (31). Machine spans are ensured to be bigger than all completion times on the machine by Constraint (32). Lastly, Constraints (33–35) make sure that the level spans correspond to the machine spans and are descending.

⁶For a proposition P , $[P] = 1$ if P is true and $[P] = 0$ otherwise.

Similarly to `cplex`, optimising lex-makespan can be accomplished by using the lexicographic minimisation functionality of `coptimizer` and the model can also be easily modified to minimise makespan.

6.3 Systems

We use `clingo-dl` version (1.1.1) for `solve(·)` and `opt(·)` in Algorithms 1 and 2, respectively. For both algorithms, the time limit for optimising any level of the lex-makespan can be set to a geometric sequence with ratio 0.5. Thus, half of the total time limit is spent on optimising the highest priority level, a quarter on the next level, etc.

We use a FlatZinc linearisation of the MiniZinc model to compare with four MIP/CP solvers (`cplex` 12.10⁷, `coptimizer` 20.1⁷, `gcode` 6.3.0⁸, , and `or-tools` 7.8⁹) against the hybrid ASP approach. All solvers could be used for makespan minimisation. Regarding the lex-makespan, `gcode` could not produce any solutions due to numerical issues and both `cplex` and `coptimizer` wrongly reported optimality for some solutions. This is a technical issue that is probably due to the translation of MiniZinc model to FlatZinc or too high values for the lex-makespan objective function. Due to this, we also encoded the problem and both objectives in the native modelling languages of `cplex` and `coptimizer`.¹⁰ Those direct approaches had no problems with wrongly reported optimal solutions and their results are included below. In difference, `or-tools` had no issues with the MiniZinc model and was thus run using this model. We generally used all solvers in single-threaded mode and with default settings.

6.4 Experimental results and discussion

All experiments were conducted on a cluster with 13 nodes, where each node has two Intel Xeon CPUs E5-2650 v4 (max. 2.90 GHz, 12 physical cores, no hyperthreading) and 256 GB RAM. For each run, we set a memory limit of 20 GB and all solvers only used one solving thread. The application at the production site requires schedule computation within 300 s that is, 5 min, which we adopted as time limit. However, in order to offer additional insights, we also performed experiments with run times of up to 15 min.

Figures 7 and 8 give an overview of the performance of solvers on all instances. We compare the approaches for makespan and lex-makespan optimisation . For each solver, we report the number of instances where some solution was found (*feasible*), a minimal solution among all approaches was found (*best*), and a globally minimal solution was found (*optimal*).

The makespan comparisons provide an important baseline as every approach that aims at improving lex-makespan necessarily involves makespan minimisation, and any good lex-makespan optimisers also needs to produce small makespans. With the practice-oriented timeout of 5 min, the `clingo-dl` approach finds solutions for most of the instances and shows very good performance for the number of best and optimal solutions compared to `cplex`, `coptimizer`, `or-tools` and `gcode`. At least when using our MiniZinc model,

⁷<https://www.ibm.com/analytics/cplex-optimizer>.

⁸<https://www.gcode.org>.

⁹<https://developers.google.com/optimisation>.

¹⁰The encodings are available online at <http://www.kr.tuwien.ac.at/research/projects/bai/tplp22.zip>.




Fig. 7. Different systems on all instances for makespan (top) and lex-makespan (bottom) with a time limit of 5 min.

`or-tools` and `gecode` have difficulties to find solutions for a large proportion of the instances. The performance of `cplex` and `cpoptimizer` is better, but it is still behind `clingo-dl`. It should be noted that we did not investigate further improvements for those solvers. In the setting with 15 min run time, `cpoptimizer` finds slightly more best solutions than `clingo-dl` which still solved the most instances to optimality. For the other solvers, the longer time out does not seem to make much difference.

For the lex-makespan comparisons, we use `or-tools`, `cplex`, `cpoptimizer` as well as `clingo-dl`, the `clingo-dl` approximation described in Section 5 with granularity $g = 10$, and Algorithms 1 and 2. While `or-tools`, `cplex` and `clingo-dl` struggle now to find feasible solutions in the 5 min setting, our multi-shot approaches shine in comparison with `cpoptimizer` trailing closely behind. The `clingo-dl` approximation does find more feasible solutions than normal `clingo-dl`. However, it finds the least number of best solutions when compared to the others. While the approximation does indeed improve performance of `clingo-dl` on bigger instances, it performs worse on small instances and the results for the bigger instances are dwarfed by the other approaches. If Algorithm 1 is used without geometric timeouts, it reports optimal solutions for quite a number of instances and finds the most best solutions. This, when comparing with Algorithms 1 and 2 with geometric timeouts, is not surprising as more time, for many instances all




Fig. 8. Different systems on all instances for makespan (top) and lex-makespan (bottom) with a time limit of 15 min.

the time, is spent on minimising the first component of the lex-makespan. With the 15 min run time as shown in Figure 8, the results are generally similar. However, here **cpoptimizer** manages to overtake Algorithm 1 on the number of best solutions found even though the latter still proves more instances to be optimal. Furthermore, with the longer run time, **cplex** does find significantly more feasible solutions. Overall, our ASP-based solution is still largely competitive with the commercial solver **cpoptimizer**.

Figures 7 and 8 do not tell us much about the quality differences of the schedules produced by the different algorithms. As our informal objective is that machines complete as early as possible, we show in Figures 9 and 10 graphs for the ratio of machines that completed as a function of schedule time, that is, $f(t) = M(S, t)/m$. This allows us to compare different approaches on the same instance classes, where the x-axis is schedule time in seconds and the y-axis is the average of $M(S, t)/m$ over the instances. The curves for different algorithms can be interpreted as follows: the earlier a curve reaches 1 (all machines completed), the smaller is the average makespan of the instances. The shape of the curve reveals details about the quality of the schedule prior to this point. For our informal objective, a steep incline of this curve is desired – the earlier it gets ahead and stays ahead, the better.




Fig. 9. Machine completion rate over time for makespan and lex-makespan (5 min run time).

We again start by discussing the results for the 5 min time limit. We only consider Algorithms 1 and 2 with geometric timeouts against plain makespan optimisation with `clingo-dl` in Figure 9. We show instance classes of increasing size from left to right and compare instances of type “low dedication” in the upper row and “high dedication” in




Fig. 10. Machine completion rate over time for makespan and lex-makespan (15 min run time).

the lower row. All approaches produce small makespans and thus schedules with a high throughput. This is worth emphasising since only half of the time limit is used here for makespan optimisation by Algorithms 1 and 2. However, when comparing the shape of the curves, the lex-makespan optimisers show their strengths for meeting our informal

objective; the difference to makespan is subtle for Algorithm 1 but more pronounced for Algorithm 2. While Algorithm 2 finds fewer minimal solutions than Algorithm 1 according to Figure 7, it tends to get ahead the earliest in terms of completed machines when considering the execution of the schedules, especially for high dedication instances. We can quantify this by looking at the average ratio of machines finished at each point in time. On average, Algorithm 1 improves this metric by 6.3% whereas Algorithm 2 shows an improvement of 9.53%.

The graphs for the longer run time of 15 min shown in Figure 10 are largely similar. However, the improvements of Algorithm 2 are slightly more pronounced, especially for the instances with 15 machines.

In summary, `clingo-dl` performs best among all considered approaches for plain makespan optimisation. For lex-makespan, Alg. 1 performs best when the time limit is 5 min, while `cpoptimizer` finds more often better solutions with the longer time limit of 15 min.

7 Related work

7.1 Parallel machine scheduling

Many variants of parallel machine scheduling problem, for example (Allahverdi et al. 2008; Allahverdi 2015), have been studied extensively in the literature. Previous publications have considered eligibility of machines, for example (Afzalirad and Rezaeian 2016; Perez-Gonzalez et al. 2019; Bektur and Saraç 2019), machine-dependent processing time, for example (Vallada and Ruiz 2011a; Avalos-Rosales et al. 2013; Allahverdi 2015), and sequence dependent setup times, for example (Vallada and Ruiz 2011a; Perez-Gonzalez et al. 2019; Fanjul-Peyro et al. 2019; Gedik et al. 2018).

7.2 Lexicographical makespan

The idea to use lexicographical makespan optimisation to obtain robust schedules for identical parallel machines comes from Letsios et al. (2021) but has not been used, to the best of our knowledge, when setup times are present. The general idea of optimising not only the element that causes the highest costs but also the second one and so on to obtain robustness, fairness, or balancedness is studied under the notion of min-max optimisation for a various combinatorial problems (Burkard and Rendl 1991; Ogryczak and Śliwiński 2006).

There are several related objective functions that can be used to achieve similar effects as minimising the lex-makespan. Load balancing can be used to obtain balanced resource utilisation by equalising the workload on machines (Rajakumar et al. 2004; Yildirim et al. 2007; Sabuncu and Simsek 2020). One measure for this is to minimise the relative percentage of imbalance in workload (Rajakumar et al. 2004), which is determined based on the difference between a machine span and the makespan. Other approaches try to minimise the difference between the largest and the smallest machine span (Ouazene et al. 2014). Note that the machines m_2 and m_3 in Figure 1(a) are balanced under this notion, but this does not ensure that the machines finish as early as they could.

7.3 ASP applications

Sabuncu and Simsek (2020) provide a novel formulation of a machine scheduling problem in ASP. Their approach bears some similarity to ours, but their objective is to balance the workload of the given machines. In general, load balancing can lead to schedules where, for example, longer than necessary setup times are used to artificially prolong machine spans for reducing imbalances. Another idea is to minimise workload instead of balancing it. While this achieves short processing times, it does not ensure that all machines finish as early as possible either. Similar to the workload, minimising the sum of machine spans does not prevent that jobs are scheduled in an unbalanced way; Figure 1(b) and (c) serve as an example of two schedules with the same total machine span. Minimising a non-linear sum of machine spans like their squares comes close to our informal objective but is different from the lex-makespan as it does not guarantee a minimal makespan (Walter 2017). Another way to obtain compact schedules is it to minimise the total completion times (Weng *et al.* 2001). This, however, can pull short jobs to the front of the schedule, which can adversely interfere with avoiding large setup times.

7.4 ASP extensions

Extending ASP with difference logic is just one way to blend integer constraints and ASP and there are several other approaches (Lierler 2014; Gebser *et al.* 2009). We evaluated ASP with full integer constraints with `clingcon`, but performance on our problem was poor. The fast propagation enabled by the lower computational complexity of difference logic seems a big advantage here. The `clingo-dl` system has indeed been used for the related problem of job-shop scheduling and makespan optimisation (Janhunen *et al.* 2017; El-Kholany and Gebser 2020). Train scheduling for the Swiss Federal Railways is another application of `clingo-dl` that involves routing, scheduling, and complex optimisation (Abels *et al.* 2019). However, we are solving a different problem with a more complex objective function and also compare to other solving paradigms.

8 Conclusion

We studied the application of hybrid ASP with difference logic to solve a challenging parallel machine scheduling problem with setup times in industrial semi-conductor production at Bosch. As objective function, we used the lex-makespan which generalises the canonical makespan to a tuple of machine spans and aims at accomplishing short completion times for all machines. Semi-conductor production involves not only one but several connected work centers that solve similar problems. Having a flexible ASP solution for one that can easily be adapted to others is highly desirable. To make ASP perform up to par, we appropriated advanced techniques like difference constraints, multi-shot solving, domain heuristics, and approximations for our application.

For the experimental evaluation, we considered random instances of realistic size and structure. We further implemented a solver-independent MiniZinc model as well as direct encodings that we used for comparisons with MIP and CP solvers. The results show that the objective of short completion times for machines is well achieved. Performance is improved by using approximations without significant deterioration of the schedules

produced. It is encouraging that the ASP approaches turn out to be competitive with commercial MIP and CP solvers which are, at least to some extent, engineered for industrial scheduling problems.

We plan to study meta-heuristics for lex-makespan optimisation in combination with ASP and methods to combine the lex-makespan with other common objective functions.

Acknowledgements

We would like to thank the reviewers for their comments and Michel Janus, Andrej Gisbrecht, and Sebastian Bayer for very helpful discussions on the scheduling problems at Bosch, as well as Roland Kaminski, Max Ostrowski, Torsten Schaub, and Philipp Wanko for their help, valuable suggestions, and feedback on constraint ASP. Johannes Oetsch was supported by funding from the Bosch Center for Artificial Intelligence.

References

- ABELS, D., JORDI, J., OSTROWSKI, M., SCHAUB, T., TOLETTI, A. AND WANKO, P. 2019. Train scheduling with hybrid ASP. In *Proceedings of the 15th International Conference on Logic Programming and Nonmonotonic Reasoning (LPNMR 2019)*, vol. 11481. Lecture Notes in Computer Science. Springer, 3–17.
- AFZALIRAD, M. AND REZAEIAN, J. 2016. Resource-constrained unrelated parallel machine scheduling problem with sequence dependent setup times, precedence constraints and machine eligibility restrictions. *Computers and Industrial Engineering* 98, 40–52, .
- ALLAVERDI, A. 2015. The third comprehensive survey on scheduling problems with setup times/costs. *European Journal of Operational Research* 246, 2, 345–378.
- ALLAVERDI, A., NG, C. T., CHENG, T. C. E. AND KOVALYOV, M. Y. 2008. A survey of scheduling problems with setup times or costs. *European Journal of Operational Research* 187, 3, 985–1032.
- AVALOS-ROSALES, O., ALVAREZ, A. M. AND ÁNGEL-BELLO, F. 2013. A reformulation for the problem of scheduling unrelated parallel machines with sequence and machine dependent setup times. In *Proceedings 23rd International Conference on Automated Planning and Scheduling (ICAPS 2013)*, 278–282.
- BEKTUR, G. AND SARAÇ, T. 2019. A mathematical model and heuristic algorithms for an unrelated parallel machine scheduling problem with sequence-dependent setup times, machine eligibility restrictions and a common server. *Computers and Operations Research* 103, 46–63.
- BREWKA, G., EITER, T. AND TRUSZCZYŃSKI, M. 2011. Answer set programming at a glance. *Communications of the ACM* 54, 12, 92–103, 2011.
- BURKARD, R. E. AND RENDL, F. 1991. Lexicographic bottleneck problems. *Operations Research Letters* 10, 5, 303–308, 1991.
- CALIMERI, F., FABER, W., GEBSER, M., IANNI, G., KAMINSKI, R., KRENNWALLNER, T., LEONE, N., MARATEA, M., RICCA, F. AND SCHAUB, T. 2020. Asp-core-2 input language format. *Theory and Practice of Logic Programming*, 20, 2, 294–309.
- EITER, T., GEIBINGER, T., MUSLIU, N., OETSCH, J., SKOČOVSKÝ, P. AND STEPANOVA, D. 2021. Answer-set programming for lexicographical makespan optimisation in parallel machine scheduling. In *Proceedings of the International Conference on Principles of Knowledge Representation and Reasoning*, vol. 18, 280–290.
- EL-KHOLANY, M. AND GEBSER, M. 2020. Job shop scheduling with multi-shot ASP. In *Proceedings of the 4th Workshop on Trends and Applications of Answer Set Programming (TAASP 2020)*.

- FANJUL-PEYRO, L., RUIZ, R. AND PEREA, F. 2019. Reformulations and an exact algorithm for unrelated parallel machine scheduling problems with setup times. *Computers & Operations Research* 101, 173–182.
- GEBSEER, M., OSTROWSKI, M. AND SCHAUB, T. 2009. Constraint answer set solving. In *Proceedings of the 25th International Conference on Logic Programming (ICLP 2009)*, vol. 5649. Lecture Notes in Computer Science. Springer, 235–249.
- GEBSEER, M., KAMINSKI, R., KAUFMANN, B. AND SCHAUB, T. 2012. Answer set solving in practice. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 6, 3, 1–238.
- GEBSEER, M., KAUFMANN, B., ROMERO, J., OTERO, R., SCHAUB, T. AND WANKO, P. 2013. Domain-specific heuristics in answer set programming. In *Proceedings of the 27th AAAI Conference on Artificial Intelligence (AAAI 2013)*. AAAI Press, 350–356.
- GEBSEER, M., KAMINSKI, R., KAUFMANN, B., OSTROWSKI, M., SCHAUB, T. AND WANKO, P. 2016. Theory solving made easy with Clingo 5. In *Tech. Comm. 32nd International Conference on Logic Programming (ICLP 2016)*, vol. 52. *OASICs*. Schloss Dagstuhl-Leibniz-Zentrum f. Informatik, 2:1–2:15.
- GEBSEER, M., KAMINSKI, R., KAUFMANN, B. AND SCHAUB, T. 2019. Multi-shot ASP solving with clingo. *Theory and Practice of Logic Programming* 19, 1, 27–82.
- GEDIK, R., KALATHIA, D., EGILMEZ, G. AND KIRAC, E. 2018. A constraint programming approach for solving unrelated parallel machine scheduling problem. *Computers & Industrial Engineering* 121, 139–149.
- JANHUNEN, T., KAMINSKI, R., OSTROWSKI, M., SCHELLHORN, S., WANKO, P. AND SCHAUB, T. 2017. Clingo goes linear constraints over reals and integers. *Theory and Practice of Logic Programming* 17, 5–6, 872–888.
- LETSIOS, D., MISTRY, M. AND MISENER, R. 2021. Exact lexicographic scheduling and approximate rescheduling. *European Journal of Operational Research* 290, 2, 469–478.
- LIERLER, Y. 2014. Relating constraint answer set programming languages and algorithms. *Artificial Intelligence* 207, 1–22.
- LIFSCHITZ, V. 2019. *Answer Set Programming*. Springer.
- MCCARTHY, J. 1998. Elaboration tolerance. In *Common Sense*, vol. 98.
- NETHERCOTE, N., STUCKEY, P. J., BECKET, R., BRAND, S., DUCK, G. J. AND TACK, G. 2007. MiniZinc: Towards a standard CP modelling language. In *International Conference on Principles and Practice of Constraint Programming*, vol. 4741. Lecture Notes in Computer Science. Springer, 529–543.
- OGRYCZAK, W. AND ŚLIWIŃSKI, T. 2006. On direct methods for lexicographic min-max optimization. In *Proceedings of the 6th International Conference on Computational Science and Its Applications (ICCSA 2006)*, vol. 3982. Lecture Notes in Computer Science. Springer, 802–811.
- OUAZENE, Y., YALAOU, F., CHEHADE, H. AND YALAOU, A. 2014. Workload balancing in identical parallel machine scheduling using a mathematical programming method. *International Journal of Computational Intelligence Systems* 7, Suppl. 1, 58–67.
- PEREZ-GONZALEZ, P., FERNANDEZ-VIAGAS, V., GARCÍA, M. Z. AND FRAMIÑAN, J. M. 2019. Constructive heuristics for the unrelated parallel machines scheduling problem with machine eligibility and setup times. *Computers and Industrial Engineering* 131, 131–145.
- RAJAKUMAR, S., ARUNACHALAM, V. AND SELLADURAI, V. 2004. Workflow balancing strategies in parallel machine scheduling. *The International Journal of Advanced Manufacturing Technology* 23, 5–6, 366–374.
- SABUNCU, O. AND SIMSEK, M. C. 2020. Solving assembly line workload smoothing problem via answer set programming. In *Proceedings of the 13th Workshop on Answer Set Programming and Other Computing Paradigms (ASPOCP 2020)*, vol. 2678. CEUR Workshop Proceedings. CEUR-WS.org.

- VALLADA, E. AND RUIZ, R. 2011a. A genetic algorithm for the unrelated parallel machine scheduling problem with sequence dependent setup times. *European Journal of Operational Research* 211, 3, 612–622.
- VALLADA, E. AND RUIZ, R. 2011b. A genetic algorithm for the unrelated parallel machine scheduling problem with sequence dependent setup times. *European Journal of Operational Research* 211, 3, 612–622.
- WALTER, R. 2017. A note on minimizing the sum of squares of machine completion times on two identical parallel machines. *Central European Journal of Operations Research* 25, 1, 139–144.
- WENG, M. X., LU, J. AND REN, H. 2001. Unrelated parallel machine scheduling with setup consideration and a total weighted completion time objective. *International Journal of Production Economics* 70, 3, 215–226.
- YILDIRIM, M. B., DUMAN, E., KRISHNAN, K. K. AND SENNIAPPAN, K. 2007. Parallel machine scheduling with load balancing and sequence dependent setups. *International Journal of Operations Research* 4, 1, 42–49.