¹ Unite and Lead: Finding Disjunctive Cliques for ² Scheduling Problems

- 3 Konstantin Sidorov¹ 🖂 🏠 💿
- 4 Delft University of Technology, The Netherlands
- 5 Imko Marijnissen¹ 🖂 🏠 💿
- 6 Delft University of Technology, The Netherlands
- 7 Emir Demirović 🖂 🏠 🕒
- $_{\rm 8}$ $\,$ Delft University of Technology, The Netherlands $\,$

⁹ — Abstract

Constraint programming solvers have seen much success in scheduling problems owing to their 10 11 efficient reasoning over constraints to solve complex problems in practice. Many algorithms have been proposed for propagating information from a *single* constraint. However, inferring and exchanging 12 information across multiple constraints can provide deeper insight into the global structure of a 13 problem. In this work, we propose to exchange information amongst constraints by inferring the 14 disjointness of tasks in scheduling problems from many constraints. We do this by (i) augmenting 15 existing propagators, such as the Cumulative and nogoods, to report when pairs of tasks are disjoint, 16 and (ii) leveraging this information by introducing the SelectiveDisjunctive propagator which 17 generates a lower bound on the earliest completion time of cliques of disjoint tasks to determine 18 19 conflicts. This allows us to aggregate disjointness information spanning multiple constraints to gain a better global overview of the problem, as well as more precise local information. We also identify a 20 problem structure where an LCG solver reasoning over Cumulative constraints separately, without 21 any reformulations, requires an exponential amount of time to prove infeasibility, which we both 22 justify theoretically and show empirically; on the other hand, our approach solves those instances 23 in polynomial time. On particular known RCPSP and RCPSP/max benchmarks, our approach 24 significantly reduces the number of conflicts required to prove optimality when resource contention 25 26 is high. Additionally, we discover new lower bounds for 16 RCPSP/max instances (closing six of them) and four RCPSP instances (closing one), as well as new upper bounds for two RCPSP/max 27 instances and four RCPSP instances. Furthermore, we empirically analyse our proposed approach 28 to determine which features are beneficial for performance, showing that *finding* cliques is one of 29 the main bottlenecks and that detecting disjointness during search can lead to improved bounds 30 on certain instances, but it generally negatively impacts learning. This work paves the way for 31 reasoning over the disjointness of tasks inferred from a variety of standard constraints to discover 32 novel information sourced from multiple constraints during search. 33

 $_{34}$ $\,$ 2012 ACM Subject Classification $\,$ Theory of computation \rightarrow Constraint and logic programming

Keywords and phrases Constraint Programming, Lazy Clause Generation, Propagation, Scheduling,
 Cumulative, Disjunctive

- ³⁷ Digital Object Identifier 10.4230/LIPIcs.CP.2025.12
- ³⁸ Supplementary Material The experimental results reported in this paper are obtained with:
- ³⁹ Software (Pumpkin and processing scripts): https://doi.org/10.5281/zenodo.15624366
- 40 Dataset (Instances and solver logs): https://doi.org/10.5281/zenodo.15624416
- ⁴¹ **Funding** Konstantin Sidorov: supported by the TU Delft AI Labs program as part of the XAIT lab.
- ⁴² *Imko Marijnissen*: supported by the NWO/OCW, as part of the Quantum Software Consortium
- ⁴³ programme (project number 024.003.037 / 3368).

31st Intern

© Konstantin Sidorov, Imko Marijnissen and Emir Demirović;

licensed under Creative Commons License CC-BY 4.0

¹ Both authors contributed equally to this research.

³¹st International Conference on Principles and Practice of Constraint Programming (CP 2025). Editor: Maria Garcia de la Banda; Article No. 12; pp. 121–12:24

Leibniz International Proceedings in Informatics LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

12:2 Unite and Lead: Finding Disjunctive Cliques for Scheduling Problems

Acknowledgements The authors would like to thank Maarten Flippo for the productive discussions
 on scheduling applications for constraint programming during the early phases of this project.

⁴⁶ **1** Introduction

⁴⁷ Scheduling is one of the important classes of problems in the field of optimisation that can ⁴⁸ be loosely described as the problem of scheduling a set of tasks to satisfy certain constraints ⁴⁹ while optimising *some* objective. Constraint programming (CP) has been successful in solving ⁵⁰ scheduling problems. The key advantage of CP is its native support for constraints that ⁵¹ capture task resource usage and efficient propagation algorithms to prune the search space, ⁵² such as time-tabling [29] or energetic reasoning [2]. Another promising technique is exploiting ⁵³ the disjointness of tasks, as identified by Gay et al. [16] to infer bounds on tasks.

A common trait of the previously mentioned works is that they reason over a *single* resource, since reasoning jointly over multiple resources significantly increases the computational cost. However, the benefits of stronger inference obtained by aggregating and exchanging information across *multiple* resource constraints may outweigh the drawbacks, an opportunity underexplored in constraint programming for scheduling.

⁵⁹ To illustrate the issues with single-resource reasoning, consider the following example:

► Example 1. We have two tasks on different resources and aim to minimise the latest finish time (makespan). Individually, neither resource constraint allows us to make any non-trivial inferences about the makespan. However, if we combine this with the information that the two tasks cannot be executed at the same time, then we can see that they need to be executed sequentially, which allows us to derive a tighter bound on the makespan of these two tasks.

To address this, we propose a principled solution for aggregating and exchanging information across constraints by reasoning over the disjointness of task pairs, based on both static incompatibilities (i.e., present in the original model) and dynamic incompatibilities (i.e., discovered during search). Our approach consists of three core components:

Variable creation We create Boolean variables representing whether two tasks are disjoint.
 Disjointness mining We extend the existing propagators for constraints defined in the

problem to infer whether two tasks should be disjoint and allow propagators to use the
Boolean variables to perform additional propagation.

Conflict detection We introduce the SelectiveDisjunctive propagator that aggregates
 this disjointness information by deriving lower bounds on the earliest completion time of
 disjoint cliques of tasks to detect conflicts.

⁷⁶ By doing so, we can detect inconsistent states for disjoint tasks by (a) aggregating information ⁷⁷ from the *whole* problem in contrast to the more orthodox focus on single constraints, (b) ⁷⁸ allowing propagators to use this shared information, and (c) computing conflicting task groups ⁷⁹ *during search*, thus using information obtained later in the search. As finding maximum ⁸⁰ cliques is an NP-hard problem [23], we propose a heuristic to find disjoint cliques of tasks ⁸¹ based on minimising the growth of the interval spanned by the clique when adding a task.

To support the usefulness of our contributions, we identify an instance structure for which state-of-the-art solvers are guaranteed to incur an exponential number of conflicts unless they reformulate the problem. While state-of-the-art LCG solvers also exhibit exponential runtime in practice, our approach can solve the same instances in polynomial time.

To evaluate our approach, we implemented it in Pumpkin [12] and ran experiments on well-known scheduling benchmarks (RCPSP and RCPSP/max). Comparing our approach to baseline Pumpkin and Google OR-Tools CP-SAT [31] (a state-of-the-art CP solver that has

won many editions of the MiniZinc Challenge since 2017) shows that our approach works 89 well when resource contention is high, with some instances exhibiting improvements of three ٩n orders of magnitude. For RCPSP/max, our approach discovers 16 new lower bounds (and 91 closes six instances) and two new upper bounds; whereas for RCPSP, our approach yields 92 four new upper bounds, four new lower bounds, and one closed instance. Our ablation study 93 shows that the main bottleneck is finding the cliques but that it can be crucial to find the 94 right cliques, whereas disjointness mining hampers performance due to poor learning on most 95 instances. 96

To summarise, we address the insufficiency of *local* information by reasoning over the 97 disjointness of tasks to infer global information. We do this by (1) introducing Boolean 98 variables representing the disjointness between two tasks and extending existing propagators 99 to infer the values of these variables and use them for additional propagation, and (2) 100 introducing the SelectiveDisjunctive propagator that heuristically finds disjoint cliques to 101 detect conflicts unattainable by current approaches due to their focus on individual constraints. 102 We identify an instance structure for which current techniques are guaranteed to incur an 103 exponential number of conflicts, while our approach only requires a polynomial number of 104 them. We experimentally evaluate our approach and observe that it is most beneficial for 105 instances with high resource contention, leading to orders-of-magnitude improvements for 106 specific benchmarks. Through an ablation study, we find that dynamically determining which 107 disjoint cliques to focus on can be beneficial and that looking for disjointness during search 108 inhibits the learning capabilities of the solver. This work paves the way for aggregating 109 information across constraints to derive *global* information about the problem. 110

The rest of the paper is structured as follows. Section 2 introduces the concepts of constraint programming relevant to our work. Section 3 reviews the previous research on disjunctive reasoning. We present and motivate our strategy for discovering and reasoning about disjunctive tasks in Section 4. We evaluate our approach on a range of scheduling problems in Section 5. Finally, we conclude and review further research directions in Section 6.

¹¹⁶ **2** Preliminaries

Constraint programming framework A Constraint Satisfaction Problem (CSP) consists of 117 a tuple $(\mathcal{X}, \mathcal{C}, \mathcal{D})$ where \mathcal{X} is the set of variables, \mathcal{C} is the set of constraints which specify 118 the relations between variables, and \mathcal{D} is the *domain* which specifies for each variable which 119 values it can take; we refer to the lowest and highest domain values as *lower* and *upper* 120 bounds and denote them LB(x) and UB(x), respectively [32]. A solution \mathcal{I} is a mapping that 121 maps each variable in \mathcal{X} to a single value in the domain of that variable in \mathcal{D} which satisfies 122 all of the constraints in \mathcal{C} . An *atomic constraint* (i.e. *atomic predicate*) is a predicate over a 123 single integer variable $x \in \mathcal{X}$ and value v signified by $[x \otimes v]$ where $\otimes \in \{=, \leq, \geq, \neq\}$. 124

¹²⁵ Constraint programming (CP) is a paradigm for solving CSPs; CP solvers enforce ¹²⁶ constraints through *propagators*, each represented with a function $f : \mathcal{D} \to \mathcal{D}'$ (where ¹²⁷ $\mathcal{D}' \subseteq \mathcal{D}$) which removes values from \mathcal{D} infeasible under the constraints in \mathcal{C} . After applying ¹²⁸ the propagators, the solver makes a *decision* which creates several subproblems by splitting ¹²⁹ the domain of a variable into two or more parts. This process of applying propagators and ¹³⁰ making decisions is performed until either a solution \mathcal{I} is found, the problem is found to be ¹³¹ unsatisfiable, or a termination criterion is met.

During the search process, CP solvers can use a technique called *nogood learning* [10]. A nogood is a partial assignment that cannot be extended to a full solution, typically rendered with an implication of the form $N = p_1 \wedge \cdots \wedge p_k \implies \bot$ with p_j being the atomic constraints

12:4 Unite and Lead: Finding Disjunctive Cliques for Scheduling Problems

encoding the partial assignment; a nogood is *unsatisfied* if no p_j is falsified. The purpose of adding nogoods is to prevent the search process from (re-)exploring parts of the search space. In this work, we use the lazy clause generation framework [14] to derive these nogoods. In this framework, it is required that the propagators explain their inferences in terms of atomic constraints, that is, they produce *explanations* of the form $e_1 \wedge \cdots \wedge e_m \implies q$.

¹⁴⁰ Cumulative constraint Cumulative is a constraint useful for many scheduling problems to model limited renewable resources, such as available work-hours. However, determining the satisfiability of a single Cumulative constraint is NP-complete [2], meaning that even reasoning efficiently over a single Cumulative requires reasoning over relaxations of the problem.

We define the Cumulative in Definition 2, the variable s_i encodes the start time of task *i*, and $(s_i + d_i)$ the finish time of *i*. Thus, in this constraint, we implicitly associate each task $i \in T$ with an interval $[s_i, s_i + d_i)$. We also introduce notation for the bounds of tasks and sets of tasks.

▶ Definition 2. Let T be a set of tasks, and let a task $i \in T$ be defined by its start variable s_i , resource usage $r_i \in \mathbb{Z}_{\geq 0}$, and duration $d_i \in \mathbb{Z}_{\geq 0}$. Finally, let $C \in \mathbb{Z}_{\geq 0}$ be the capacity of the resource. Then the Cumulative constraint is the condition that at any time point τ the cumulative resource usage of intervals $[s_i, s_i + d_i)$ covering τ does not exceed the capacity:

$$\forall \tau \in \mathbb{Z} : \sum_{i \in T : s_i \le \tau < s_i + d_i} r_i \le C.$$
(1)

Given a task $i \in T$, we denote its: earliest start time $\text{EST}_i = \text{LB}(s_i)$, earliest completion time $\text{ECT}_i = \text{LB}(s_i) + d_i$, latest start time $\text{LST}_i = \text{UB}(s_i)$, latest completion time $\text{LCT}_i = \text{UB}(s_i) + d_i$, and energy $e_i = r_i \times d_i$.

Given a set of tasks $\Omega \subseteq T$, we denote its: total duration $d_{\Omega} = \sum_{i \in \Omega} d_i$, earliest start time $\text{EST}_{\Omega} = \min_{i \in \Omega} \text{LB}(s_i)$, latest completion time $\text{LCT}_{\Omega} = \max_{i \in \Omega} \text{UB}(s_i) + d_i$, earliest completion time $\text{ECT}_{\Omega} = \max_{\Omega' \subseteq \Omega} \text{EST}_{\Omega'} + d_{\Omega'}$, and total energy $e_{\Omega} = \sum_{i \in \Omega} e_i$

¹⁶⁰ An important concept for the inference over Cumulative constraints is the *mandatory part* ¹⁶¹ (Figure 1a), which is a time interval that is covered by a task regardless of its placement:

▶ Definition 3. Given a Cumulative constraint over tasks T with capacity C, we denote the mandatory part MP_i of a task $i \in T$ as the interval MP_i := [LST_i, ECT_i). We define a height at a time point τ as total consumption of mandatory parts covering τ , that is, Height(τ) := $\sum_{i \in T : \tau \in MP_i} r_i$. We also define a reduced height at a time point τ without $\Omega \subseteq T$ as Height⁻(τ, Ω) := $\sum_{i \in T \setminus \Omega : \tau \in MP_i} r_i$. Finally, we define a profile Pr_x as a rectangle (a_x, b_x, h_x) where $\forall \tau \in [a_x, b_x]$: Height(τ) = h_x and we assume that all profiles are maximal (i.e. there are no adjacent profiles with the same height).

Disjunctive (No-Overlap) constraint In this paper, we focus on a special case of Cumulative where the resource has unit capacity C = 1 and each task has unit resource usage $r_i = 1$, known as Disjunctive. Many inference procedures that are intractable for Cumulative constraints can be efficiently executed for Disjunctive constraints. The core inference that we use is infeasibility checking (Theorem 4), as illustrated by Figure 1b.

▶ Theorem 4 (Overload checking). Let $Ω \subseteq T$ be a set of variables bound by a Disjunctive constraint such that $d_Ω > LCT_Ω - EST_Ω$. Then the Disjunctive constraint has no feasible

¹⁷⁶ solutions [37], and any feasible solution of the CSP satisfies

$$\lim_{1 \le \Omega} \left(\bigwedge_{i \in \Omega} [\![s_i \ge \mathrm{EST}_{\Omega} - \delta^-]\!] \right) \wedge \left(\bigwedge_{i \in \Omega} [\![s_i \le \mathrm{LCT}_{\Omega} - d_\omega + \delta^+]\!] \right) \implies \bot$$

$$(2)$$

for arbitrary $\delta^-, \delta^+ \in \mathbb{Z}_{\geq 0}$ satisfying $\operatorname{LCT}_{\Omega} - \operatorname{EST}_{\Omega} + \delta^- + \delta^+ < d_{\Omega}$ [38].





(a) Mandatory part illustration for a *single* task *i*.



Figure 1 Key concepts related to Cumulative and Disjunctive constraints.

3 Related Work

Much work focused on inferring information based on the Cumulative efficiently [29, 34]. 180 One of the approaches of primary interest in this work is edge-finding [29] as it is similar in 181 its conflict detection procedure (also known as *input/output consistency tests* [13, Chapter 182 4). A common trait of these approaches [22, 18] is that the earliest completion time of 183 $\Omega \subseteq T$ is used to make inferences. For example, overload checking enforces the condition 184 $ECT_{\Omega} > LCT_{\Omega} \implies \bot$, while edge finding extends this rule to $ECT_{\Omega \cup \{i\}} > LCT_{\Omega} \implies \Omega < i$, 185 updating the bounds of t such that it does not overflow the interval bounded by the tasks in 186 Ω . While edge-finding takes into account the earliest start times and latest finish times of 187 the tasks (and horizontally elastic edge-finding [18] takes into account the maximum resource 188 usage of the tasks), this type of reasoning does not take into account the disjointness of tasks 189 and its influence on the earliest completion time of a set of tasks. Furthermore, edge-finding 190 reasons over a single Cumulative, whereas we propose gathering information across several 191 Cumulative constraints to infer the earliest completion time. 192

Detecting disjoint pairs of tasks within a Cumulative constraints is also a viable reasoning 193 strategy [16], known as time-table disjunctive reasoning. This strategy can detect disjoint 194 task pairs dynamically by taking into account the time points at which other tasks are 195 guaranteed to consume *some* amount of resource. When it detects such disjointness using 196 time-tabling [29], it determines when the bounds of one of the tasks can be updated. The 197 main difference with our approach is that time-table disjunctive reasoning only reasons about 198 a single Cumulative constraint while we reason jointly over multiple constraints, and our 199 approach infers disjointness from other constraints besides the Cumulative. 200

To the best of our knowledge, the only work that addresses the reasoning over multiple Cumulative constraints is the work by Beldiceanu and Carlsson [5]. This work introduces a multi-resource Cumulatives constraint as a generalisation of the Cumulative, admitting negative resource consumption and lower/upper bounds on the cumulative resource usage.

12:6 Unite and Lead: Finding Disjunctive Cliques for Scheduling Problems

However, their reasoning makes limited use of the additional information provided by 205 reasoning over multiple Cumulative constraints, nor does it take into account disjointness. 206 Besides reasoning about Cumulative constraints, we consider other works that reason 207 about disjunctive tasks. One area is branch-and-bound algorithms for scheduling problems. 208 A common example is the resource-constrained scheduling problem (RCPSP), in which one 209 of the lower bounds (LB4) on makespan is the sum of durations in a group of pairwise 210 disjunctive tasks; with similar bounds being proposed for activity groups with at most two or 211 at most k of them running in parallel [24]. Thus, our approach generalises that strategy while 212 also working in the constraint programming context, where the tasks can be (a) arbitrarily 213 bounded by the solver and (b) involved in other constraints. This type of reasoning can also 214 be embedded into the search tree node representation, for example, as a scheduling scheme 215 [7], which accounts for the pairs of tasks that are either disjunctive or running in parallel. 216

We thus note that there are several methods that attempt to find relations between variables and to exploit these relations to infer information about the bounds. However, the relations discovered by the majority of state-of-the-art techniques are local to individual resource constraints or do not reason over the disjointness of tasks; contrary to that, our work aims to discover this disjointness *dynamically during search* from multiple constraints and use information *across* propagators to infer the earliest completion time of a set of tasks.

4 Our Contribution: Finding Disjunctive Cliques

This section will introduce our approach for aggregating information across constraints. We propose the following workflow:

- 1. Given a set of tasks T (bound by Cumulative constraints), introduce new variables δ_{ij}
- for all $i, j \in T$ and reify them as $\delta_{ij} \to [s_i, s_i + d_i) \cap [s_j, s_j + d_j) = \emptyset$. (Subsection 4.2) 228 **2.** Introduce the reification variables δ into other constraints (such as the Cumulative and 229 difference logic) and extend their propagation logic to derive implications of the form 230 [variable bounds] $\Longrightarrow \delta_{ij}$, and propagate variable bounds based on δ . (Subsection 4.4) 231 **3.** Introduce our novel SelectiveDisjunctive propagator which derives lower bounds on

the earliest competion time of sets of pairwise disjoint tasks given (a) pairs of tasks $\{i, j\}$ such that $\delta_{ij} = 1$, and (b) the bounds on start times to discover conflicts. (Subsection 4.3) Our approach allows us to gain a *global* view across multiple constraints rather than relying on single constraints to independently make local inferences which causes limited information exchange between propagators. Thus, our approach allows them to communicate more elaborate facts between propagators and a CP solver, such as the disjointness of a pair of tasks, as opposed to existing methods focusing on a *local* view.

²³⁹ 4.1 Motivating Example: 3*n* Problem

Let us first look at a case where reasoning over a single resource is insufficient: Suppose that we are scheduling six tasks, each with unit duration. At every time unit, we are given a fixed amount of resources of types I, II, and III. The amount of resources consumed by each task is illustrated in Figure 2a. *How much time has to pass before all six tasks are completed?*

This is possible to do in six units of time by sequentially scheduling the tasks in any order. What cannot be inferred by reasoning over single constraints (such as time-tabling or energetic reasoning) is that this schedule is the fastest possible (in terms of the latest completion time). To see why, observe that the two yellow tasks cannot be run in parallel due to the shortage of resource I, and the same is true for both the two green tasks and the two blue tasks by similar reasoning. On the other hand, scheduling a yellow and a

250 green task in parallel is not possible, as it would overflow the resource II; the same holds

for a green and a blue task, as well as for a blue and a yellow task. By this point, we have considered *all* pairs of tasks and concluded for each of them that they cannot be run in

parallel, making it impossible to improve upon the sequential schedule.



(a) Conflict graph of the motivating example. Edge colour indicates resource, solid edges connect tasks jointly consuming 140%, and dotted edges connect tasks jointly consuming 110%. All tasks of the same colour have the same resource usage.

(b) Number of conflicts generated by CP-SAT solving 3nproblem instances with d = 1.

Figure 2 Motivating example: scheduling problem with tasks that are pairwise disjoint due to *various* resource constraints and how CP-SAT performs on these instances.

253

The previous reasoning relies only on the partition of tasks into three groups such that choosing one task from every two groups implies disjointness; this property holds if we consider groups of size n. We generalise the reasoning above not only in terms of group size n but also by duration d and capacity parameters p, q, M:

▶ Definition 5. Given a group size $n \in \mathbb{N}_{\geq 1}$, duration $d \in \mathbb{N}_{\geq 1}$, and parameters $p, q, M \in \mathbb{N}_{\geq 1}$ with 2p < M < p + q < 2q, the 3n problem instance is the constraint satisfaction problem defined over variables X, Y, Z with |X| = |Y| = |Z| = n and domains $\mathcal{D}(\cdot) \equiv [1, (3n - 1)d]$:

 $\begin{array}{ll} & \texttt{cumulative}(X+Y;[d]^{2n};[p]^n+[q]^n;M) & \texttt{cumulative}(Y+Z;[d]^{2n};[p]^n+[q]^n;M) \\ & \texttt{cumulative}(Z+X;[d]^{2n};[p]^n+[q]^n;M) \end{array}$

As we established, any 3n instance is unsatisfiable, but to establish this without search, a solver needs information about all resources. However, each resource individually only implies that any solution occupies some time segment of duration $2n \times d$, and the remaining (n-1)d time units have to be resolved by search. This suggests that any valid sequence of inferences that operates on Cumulative constraints separately has exponential length in terms of n.

We prove this, and the 3n problem thus serves a role similar to pigeonhole principle formulas [19], but for the Cumulative models instead of propositional formulas. We show that any unsatisfiability proof of a 3n problem in the C-RES proof system [21] has length exponential in n, which shows that any LCG solver requires an exponential number of steps to prove unsatisfiability, as long as it introduces neither new constraints nor new variables:

▶ **Theorem 6** (3*n* problem intractability). Let \mathcal{P} be a C-RES proof of a CSP instance with all clauses valid for any single constraint in Definition 5 for d = 1. Then $|\mathcal{P}| = \Omega(1.05^n)$.

Proof Sketch. First, we show that \mathcal{P} can be encoded, with extra assumptions, as a pigeonhole problem proof over 3n variables (Appendix A). Next, we adapt the argument by Beame and Pitassi [4] to verify the exponential lower bound (Appendix B).

12:8 Unite and Lead: Finding Disjunctive Cliques for Scheduling Problems

We evaluate this insight empirically by running 3n instances for $n \leq 10$ and $d \leq 5$ with five 279 solvers: two LCG solvers (Google OR-Tools CP-SAT [31] and Chuffed [9]), a non-learning 280 CP solver (Gecode [17]), and two branch-and-cut solvers (CBC [15] and HiGHS [20]). We 281 ran each solver with a 12-hour time limit, 1 CPU, and 4000 MB of RAM. Both the LCG and 282 non-learning solvers exhibit exponential behaviour, as shown on Figure 2b for CP-SAT, and 283 only one of the runs $n \ge 6$ proved unsatisfiability. These results suggest that LCG solvers 284 indeed resort to exhaustive enumeration on the 3n problem. Branch-and-cut solvers have 285 exhibited better performance in solving those instances, demonstrating another similarity 286 with pigeonhole principle formulas. HiGHS solved all instances with $d \leq 3$ in seven minutes. 287 CBC has performed worse, as it has failed to solve instances with $n \ge 7$, however, it has 288 successfully solved d = 1 for $n \leq 6$ in less than two seconds. 289

290 4.2 SelectiveDisjunctive Constraint

Now that we have motivated the necessity for using *global* information, we describe how we use information concerning the disjointness of tasks across *multiple* constraints by introducing the SelectiveDisjunctive constraint. The SelectiveDisjunctive constraint (i) generalises the disjunctive constraint in a way that allows other propagators to indicate disjoint task pairs, and (ii) restricts the variables indicating this disjointness to be true when a pair of intervals in the assignment is disjoint. We formalise this intuition in Definition 7.

▶ Definition 7. Let T be a set of tasks, and Δ be the $n \times n$ matrix of Boolean variables indexed by T. Then a SelectiveDisjunctive (T, Δ) constraint is a condition that for any two tasks $i, j \in T, i \neq j$ either $\delta_{ij} = 0$ or $[s_i, s_i + d_i) \cap [s_j, s_j + d_j] = \emptyset$.

Introducing these variables and SelectiveDisjunctive does not prune any valid solutions, but it allows performing two new operations: (a) fixing the variables from Δ and (b) using these Δ variables in other propagations. This allows us to encode facts—such as 'a pair of tasks cannot overlap given a partial assignment of other tasks'—that are not possible to encode concisely using atomic constraints over the original model variables. More specifically, we can now observe that Definition 7 embeds the standard Disjunctive constraint (or even exponentially many of them); to formalise this, we first need the following notion:

Definition 8. Given a search state with domains \mathcal{D}' and a SelectiveDisjunctive constraint, a conflict graph $\mathcal{G}|_{\mathcal{D}'}$ is a graph with a vertex for each task $i \in T$ and an edge for tasks $\{i, j\}$ such that $\mathcal{D}'(\delta_{ij}) = \{1\}$ (ommitting \mathcal{D} when it is clear from context).

We can make two key observations based on Definition 8: (1) a conflict graph can be 310 seen as an *exponentially large* collection of Disjunctive constraints, and (2) this is also a 311 dynamically updating collection of Disjunctive constraints, provided that the other parts of 312 a constraint solving engine can propagate the domains of reification variables. On top of this, 313 it suggests that any reasoning for the conventional Disjunctive constraint can be deployed 314 for any clique of the conflict graph. In Subsection 4.3, we discuss how to find these cliques in 315 the conflict graph and how to apply the overload checking rule for the Disjunctive in this 316 context. 317

4.3 Overload checking for SelectiveDisjunctive

³¹⁹ To clarify the intuition behind our inference strategy, consider the following example:

Example 9. Let i, j, k be tasks with possible start times $s_i \in [0, 2], s_j \in [0, 2], s_k \in [2, 3]$ and durations $d_i = d_j = 2, d_k = 3$ bound by a SelectiveDisjunctive constraint such that

³²² $\delta_{ij} = \delta_{jk} = 1$ in the current search state. We can observe that any feasible assignment ³²³ also satisfies Disjunctive($[s_i, s_j, s_k], [d_i, d_j, d_k]$). However, this constraint is infeasible by ³²⁴ overload checking for $\{s_i, s_j, s_k\}$: these three tasks jointly cover seven time units, but their ³²⁵ EST is 0 and LCT is 6, leaving only six time units among these three tasks.

³²⁶ ► Note 10. The example above purposefully does not specify why the Δ-variables are true: ³²⁷ the same reasoning can be carried out as soon as those intervals are established to be disjoint. ³²⁸ We discuss the specific strategies for deriving those facts in Subsection 4.4.

We formalise this intuition in Proposition 11 by reformulating Theorem 4 in the context of the SelectiveDisjunctive constraint.

³³¹ ► Proposition 11 (Selective overload checking). Let T be the set of tasks bound by a ³³² SelectiveDisjunctive constraint, and suppose that $\Omega \subseteq T$ is a clique in the conflict ³³³ graph G induced by the current variable domains. Then this constraint is infeasible if ³³⁴ $d_{\Omega} > LCT_{\Omega} - EST_{\Omega}$, and any feasible solution of the CSP satisfies

$$(\bigwedge_{i,j\in\Omega,\ i\neq j}\delta_{ij})\wedge\left(\bigwedge_{\omega\in\Omega}[\![s_{\omega}\geq\mathrm{EST}_{\Omega}-\delta^{-}]\!]\right)\wedge\left(\bigwedge_{\omega\in\Omega}[\![s_{\omega}\leq\mathrm{LCT}_{\Omega}-d_{\omega}+\delta^{+}]\!]\right)\implies \bot (3)$$

 $\text{ for arbitrary } \delta^-, \delta^+ \in \mathbb{Z}_{\geq 0} \text{ satisfying } \operatorname{LCT}_{\Omega} - \operatorname{EST}_{\Omega} + \delta^- + \delta^+ < d_{\Omega}.$

Unlike the Disjunctive overload checking, propagation of the SelectiveDisjunctive 337 requires an additional step before overload checking can be performed: finding cliques. 338 Finding a clique leading to a conflict, even if it exists, is NP-complete² and we thus propose a 330 heuristic in Algorithm 1 that dynamically explores some of the cliques present in the conflict 340 graph to balance the trade-off between finding "good" cliques and the time spent finding 341 them. The heuristic algorithm considers each node as a root for a clique \mathcal{C} and then adds 342 tasks $j \in T \setminus C$ which minimise $\operatorname{LCT}_{C \cup \{j\}} - \operatorname{EST}_{C \cup \{j\}}$ (breaking ties in favour of tasks with 343 a longer duration) while retaining the clique property. The time complexity of Algorithm 1 344 is $\mathcal{O}(|T|^4)$. 345

346 4.4 Disjointness mining

We described how to deploy the SelectiveDisjunctive propagator when some of the reification variables have been assigned, but we have not yet described how to infer the domain of these reification variables. In this section, we describe strategies for *disjointness mining*, that is, discovering, possibly during the search, pairs of tasks $i, j \in T$ that are inferred to be disjoint.

Domain disjointness One direct source of disjointness is the current bounds of variables. The rule can be formally stated as $\forall i, j \in T$: $[\text{EST}_i, \text{LCT}_i) \cap [\text{EST}_j, \text{LCT}_j) = \emptyset \implies [\![\delta_{ij} = 1]\!]$. We perform this check during the selective overload checking (Subsection 4.3). Given a propagation of δ_{ij} via this detection, we explain the propagation according to Equation 4.



$$\begin{cases} [[s_i \leq \text{EST}_j - d_i]] \land [[s_j \geq \text{EST}_j]] & \text{if } \text{LCT}_i \leq \text{EST}_j \\ [[s_j \leq \text{EST}_i - d_j]] \land [[s_i \geq \text{EST}_i]] & \text{otherwise} \end{cases} \implies [[\delta_{ij} = 1]]$$

$$\tag{4}$$

² This can be shown by reduction from the maximum clique problem

Data: A set of tasks T with intervals $[\text{EST}_i, \text{LCT}_i)$ and durations d_i for all $i \in T$. **Data:** A conflict graph $\mathcal{G} = (T, E)$. **Result:** A set $\Omega \subseteq T$ satisfying $d_{\Omega} > \text{LCT}_{\Omega} - \text{EST}_{\Omega}$ or \perp if none was found. if $\exists e \in E : d_e > \operatorname{LCT}_e - \operatorname{EST}_e$ then return e; /* Check for binary conflicts */ for $t_0 \in T$ do $\Omega \leftarrow \{t_0\} ;$ /* Ω is a clique containing t_0 */ /* [S,T) is an interval covered by Ω */ $S, T \leftarrow \text{EST}_{t_0}, \text{LCT}_{t_0}$; while $d_{\Omega} \leq LCT_{\Omega} - EST_{\Omega}$ do // Collect tasks that are disjoint with Ω $\Omega^+ \leftarrow \{j \in T \setminus \Omega : [S,T) \cap [\text{EST}_i, \text{LCT}_i) \neq \emptyset \land \forall i \in \Omega, \{i,j\} \in E\};$ if $\Omega^+ = \emptyset$ then break; /* Cannot extend Ω into a clique, try another root */ $t^{+} \leftarrow$ an element of $\arg\min_{j\in\Omega^{+}} \left(\max(T, \operatorname{LCT}_{j}) - \min(S, \operatorname{EST}_{j}), -d_{j} \right);$ $\Omega \leftarrow \Omega \cup \{t^+\}, \ S \leftarrow \min(S, \mathrm{EST}_t^+), \ T \leftarrow \max(T, \mathrm{LCT}_t^+); \ /* \text{ Extend } \Omega */$ if $d_{\Omega} > \text{LCT}_{\Omega} - \text{EST}_{\Omega}$ then return Ω ; return \perp ;

Algorithm 1 Heuristic conflict discovery procedure.

371

³⁵⁷ **Difference logic** Scheduling problems often include tasks dependent on the execution of ³⁵⁸ other tasks. These dependencies can be provided in the form of precedence constraints, which ³⁵⁹ provide static information about the disjointness. For example, to encode that j starts δ ³⁶⁰ time units after i, we can use the difference logic constraint $s_i + \delta \leq s_j$.

These constraints encode *stronger* conditions than disjointness; for example, precedence can be seen as disjointness with ordering. Thus, we can deduce the disjointness variables from the difference constraint for appropriate parameter values (this occurs at the root level and requires no explanation). We formalise this in the following proposition:

Proposition 12. A pair of intervals $[s_i, s_i + d_i), [s_j, s_j + d_j)$ is disjoint (i.e., $\delta_{ij} = 1$ in any feasible solution) if any of the following constraints are implied by the CSP: (a) $s_j - s_i \ge \delta$ for some constant $\delta \ge d_i$, or (b) $s_i - s_j \ge \delta$ for some constant $\delta \ge d_j$.

³⁶⁸ **Cumulative** There are several **Cumulative** propagation techniques that reason about when ³⁶⁹ the combination of scheduling a task *i* and a set of tasks $\Omega \subseteq T \setminus \{i\}$ leads to resource ³⁷⁰ overflows to infer relations between tasks. In this work, we use resource profiles to determine when a pair of tasks is disjoint. An example of this reasoning can be seen in Example 13.



Figure 3 Example where profiles (marked in grey) cause disjointness between all tasks i, j, k.

▶ **Example 13.** In addition to the assumptions of Example 9, suppose that the problem has three Cumulative constraints as visualised in Figure 3. We can derive from Resource

1 that *i* and *j* are disjoint ($\delta_{ij} = 1$); if they are not, then they overlap at some point $\tau^* \in [\max(\text{EST}_i, \text{EST}_j) = 0, \min(\text{LCT}_i, \text{LCT}_j) = 4]$. However, this would cause an overflow of the capacity, implying disjointness. Similar reasoning derives that $\delta_{jk} = 1$ and $\delta_{ik} = 1$.

In this example, just as in the motivating example (Subsection 4.1), neither of the three constraints is infeasible; in fact, each can only imply one disjointness relation. However, the *intersection* of those constraints is correctly declared infeasible by (i) Cumulative mining and (ii) SelectiveDisjunctive overload checking. Thus, combining the mining strategies with SelectiveDisjunctive inferences can generate conflicts that are out of reach for the conventional techniques for the Cumulative constraint.

We adapt the time-table disjunctive reasoning [16] by going over all pairs of tasks $i, j \in T$, and calculating their overlap o_{ij} . Then we check $\forall \tau \in o_{ij}$ if there exists a profile such that scheduling the tasks together with that profile would lead to an overflow. A simple implementation of this approach has time complexity $\mathcal{O}(|T|^2|TT|)$, where TT is the set of existing profiles (i.e, the time-table). This reasoning is formally stated in the following proposition:

▶ Proposition 14. Consider a Cumulative constraint on tasks T with capacity C. Let the overlap between two tasks $i, j \in T$ be given by $o_{ij} = [\max(\text{EST}_i, \text{EST}_j), \min(\text{LCT}_i, \text{LCT}_j))$. Then δ_{ij} is true for any feasible solution as long as there is no point $\tau \in o_{ij}$ such that Height⁻($\tau, \{i, j\}$) plus the resource consumption of i and j fits within the capacity:

³⁹³
$$\forall \tau \in o_{ij} : r_i + r_j + \text{Height}^-(\tau, \{i, j\}) > C \implies \llbracket \delta_{ij} = 1 \rrbracket$$

³⁹⁴ Moreover, given a propagation of δ_{ij} by a set of profiles $\mathcal{P} = \{(a_0, b_0, h_0), ..., (a_m, b_m, h_m)\},$ ³⁹⁵ $a_0 < \cdots < a_m$, we explain the propagation according to Equation 5, assuming without loss ³⁹⁶ of generality that $\text{EST}_i \leq \text{EST}_j$. For \mathcal{E}_{pr} , we use the big-step explanation [33].

$$\begin{cases} \llbracket s_j \ge a_0 \rrbracket \land \llbracket s_j \le b_m - d_j + 1 \rrbracket & \text{if } \operatorname{LCT}_i \ge \operatorname{LCT}_j \\ \llbracket s_i \le b_m - d_i + 1 \rrbracket \land \llbracket s_j \ge a_0 \rrbracket & \text{otherwise} \end{cases} \land \left(\bigwedge_{Pr \in \mathcal{P}} \mathcal{E}_{pr}(Pr) \right) \implies \llbracket \delta_{ij} = 1 \rrbracket$$
(5)

The first part of the explanation in Equation 5 is based on the observation that either one task subsumes the entire interval of the other (in which case the explanation depends only on the bounds of s_j), or there is a partial overlap where one task starts at the same time or before the other and ends at the same time or before the other (in which case the overlapping part is only defined by the upper bound on s_i and the lower bound on s_j).

Additionally, we adapt the propagation by Gay et al. [16] and propose a rule using δ in Equation 6 (the rule for updating s_i is symmetric). The rule in Equation 6 and the rule proposed by Gay et al. [16] do not subsume one another since Equation 6 can propagate based on information gained from other constraints, while the rule by Gay et al. can propagate when two tasks are not *fully* disjoint.

$$[\delta_{ij} = 1] \wedge \text{ECT}_i > \text{EST}_j \wedge \text{LST}_i < \text{ECT}_j \implies [s_j \ge \text{ECT}_i]$$

$$(6)$$

⁴⁰⁹ **Nogood** Last, we describe how to derive disjointness from nogoods: given an unsatisfied ⁴¹⁰ nogood N, tasks $i, j \in T$ are disjoint if falsifying each of the unassigned atomic constraints ⁴¹¹ implies δ_{ij} . An example can be seen in Example 15 and it is formalised in Proposition 16.

▶ Example 15. Given two tasks $i, j \in T$ such that $s_i \in [0, 9], s_j \in [7, 20]$ with durations $d_{i13} = 2, d_j = 5$, consider a nogood $N = \{p_i = [s_i \ge 5], p_j = [s_j \le 11]], ...\}$. Suppose that in the current search state, only the first two atomic constraints are unassigned, while the rest are satisfied. Thus, in any feasible solution either $\neg p_i = [s_i \le 4]$ or $\neg p_j = [s_j \ge 12]$ holds. But in either case, we can derive δ_{ij} via domain disjointness: If $\neg p_i$ is true, then LCT_i = 6 \leq EST_j = 7, implying δ_{ij} .

If $\neg p_j$ is true, then LCT_j = 11 \leq EST_i = 12, also implying δ_{ij} .

⁴¹⁹ ► **Proposition 16.** Let \mathcal{D}' be a set of domains, and suppose that a nogood N has no falsified ⁴²⁰ atomic constraints in this domain. Then a pair of tasks *i*, *j* is disjoint if for every atomic ⁴²¹ constraint $p \in N$ not satisfied with respect to \mathcal{D}' , δ_{ij} is implied by domain disjointness.

⁴²² Our procedure is a simplified version of Proposition 16 based on the two-watcher scheme [28, ⁴²³ 27]. Whenever a watcher is updated, we perform a scan to determine whether the conditions ⁴²⁴ of Proposition 16 hold while reasoning over domain disjointness; this misses out on some ⁴²⁵ propagations, but the time required is significantly reduced. The time complexity of this ⁴²⁶ approach is $\mathcal{O}(|N|)$ for each nogood N. We explain this propagation due to an unsatisfied ⁴²⁷ nogood N as $\mathcal{E}_{sn}(s_i, s_j, p_i, p_j) \land \left(\bigwedge_{p^+ \in N^+} p^+\right) \implies [\![\delta_{ij} = 1]\!]$, where $N^+ \subseteq N$ is the set of ⁴²⁸ satisfied atomic constraints, and $p_x = [\![s_x \otimes v]\!]$ is one of the two undecided atomic constraints. ⁴²⁹ See Appendix C for the definition of \mathcal{E}_{sn} .

430 **5** Experiments

We implemented our approach as part of Pumpkin [12]. First, we present an ablation study 431 in Subsection 5.1, which shows that (a) dynamic disjointness mining hinders the learning of 432 the solver but that it can be beneficial on specific instances, (b) an alternative efficient sorting 433 heuristic has a significant impact on the search but also leading to massive slowdowns on 434 some instances, and (c) enabling the propagation in Equation 6 does not consistently change 435 the performance. Next, we run a comparison on RCPSP and RCPSP/max benchmarks in 436 Subsection 5.2 with baseline Pumpkin and CP-SAT, which shows that our approach can 437 accelerate search by orders of magnitude on certain instances. Additionally, we compare our 438 results with bounds reported in the literature and show that our approach discovers new 439 lower bounds on 16 RCPSP/max instances and four RCPSP instances, as well as new upper 440 bounds on two RCPSP/max instances and four RCPSP instances. 441

Experiment Setup The implementation of our approach with the infrastructure for running
the experiments is available in the supplementary material. We ran our experiments on *DelftBlue* [11], with each run of an instance being allocated a single core of an Intel Xeon
E5-6248R 24C 3.0GHz processor and 4000 MB of RAM with a time limit of one hour.

For our evaluation, we use two models that minimise the latest completion time of all tasks (makespan). One is **RCPSP/max**, a collection of **Cumulative** constraints and precedence relations encoded by the constraints $s_i + \gamma_{ij} \leq s_j$ with *arbitrary* constants $\gamma_{ij} \in \mathbb{Z}$; we use the test suites distributed by PSPLIB [25], namely, C, D, UBO, and SM suites. We also use **RCPSP**, a special case of RCPSP/max where all precedence relations have the form $s_i + d_i \leq s_j$; we use the data files provided in the MiniZinc benchmarks [36], which correspond to the AT [30], BL [3], Pack [8], Pack-d and KSD15-D [26], and PSPLIB [25] suites.

Given a minimisation objective \mathcal{O} , each of the solvers is run with one of the two search directions. In the **primal search**, the solver generates a series of problems with an extra assumption $[\mathcal{O} \leq o_n]$ for decreasing values of o_n . In the **dual search**, the solver generates a series of problems with an extra assumption $[\mathcal{O} \leq o_n]$ for increasing values of o_n .

We discard all the instances solved by the Pumpkin baseline within five seconds in either search direction, leaving us with 736 RCPSP instances and 349 RCPSP/max instances. The models we used and the data files are available in the supplementary material. We evaluate the rate of progress of a solver towards the best-known bound with the following metrics. First, if M(t) is the lowest makespan discovered at time $t \in [0, T]$, and M^* is the

lowest discovered makespan for this problem, then the **primal integral** is $\int_0^T \frac{M(t)-M^*}{M(t)}$ and measures how fast the solver progresses towards good solutions [6]. Conversely, if B(t) is the highest lower bound discovered at time $t \in [0, T]$, and B^* is the highest discovered lower bound for this problem, the **dual integral** is defined as $\int_0^T \frac{B^* - B(t)}{B^*}$.

466 5.1 Ablation Study

To evaluate the influence of individual components on the performance of the whole approach, we ran the disjunctive approach, both in primal and dual search directions, with *all* sixteen combinations of enabling and disabling the following features:

470 **Cumulative mining** Enable the mining for **Cumulative** as described in Subsection 4.4.

⁴⁷¹ **Nogood mining** Enable the mining for nogoods as described in Subsection 4.4.

⁴⁷² Sorting Restrict Ω^+ in Algorithm 1 to tasks j with $LCT_j \leq \min_{i \in \Omega} LCT_i$.

⁴⁷³ **Propagation** Execute the propagation shown in Equation 6 on discovered disjunctive pairs.

Figure 4a reports the relative gains from enabling each feature. For each feature, we consider *all* combinations of the remaining three features and evaluate the ratio between

the aforementioned integrals when adding the feature to this configuration and *without* the feature.





(a) Distribution of ratio of integrals of enabling each evaluated feature and *not* enabling the feature. The vertical line in the middle corresponds to no observed difference, a ratio larger than 1 corresponds to an improvement by using the feature, and values smaller than 1 correspond to performance degradation.

(b) Change in LBD split by change in # conflicts with cumulative mining compared to baseline; a negative value indicates a decrease in the LBD when using mining.

Figure 4 Impact of enabling the listed features on the solver performance.

One of the major takeaways is that cumulative mining has a predominantly negative 478 impact on search performance. As we understand, this happens because nogood learning 479 derives statements overly specific to the current search state and thus leads to less general 480 nogoods. To support this conclusion, we compare the average literal block distance (LBD) [1] 481 of nogoods produced³ with and without cumulative mining in Figure 4b, as this metric 482 indicates learned nogood quality. In 47% of all pairs of runs, mining leads to fewer conflicts, 483 while having a relatively small impact on LBD; however, in the other half of run pairs, mining 484 not only increases the number of conflicts but also exhibits a much larger increase in the 485 LBD. 486

³ If two runs reached different solutions, we compare the number of conflicts to the best common bound.

12:14 Unite and Lead: Finding Disjunctive Cliques for Scheduling Problems

On the positive side, we have discovered that cumulative mining *can* be a beneficial 487 strategy for instances where many *triples* of tasks—but few *pairs* of them—are mutually 488 disjoint. In that case, fixing the position of a few tasks can uncover a variety of disjoint 489 pairs, aiding the conflict discovery. In particular, this is the case for several instances from 490 the Pack-d collection; we ran the mining configuration of our approach on those instances 491 and observed that the non-mining implementation discovers few conflicts (and thus makes 492 little progress), whereas the mining implementation discovers several bounds tighter than 493 reported by all the other approaches (see Appendix D). 494

Another takeaway is that the sorting heuristic has a significant influence on the search. Enabling it reduces integrals for most instances, suggesting that the bottleneck of our approach is the clique search, and most of the iterations of Algorithm 1 do not yield conflicts. However, there is a group of instances for which the *dual* integral degrades by three orders of magnitude or more; for this reason, we only use the sorting heuristic when the disjunctive approach runs in the primal search direction.

Last, enabling propagation does not exhibit a consistent change in any direction; the primal search can substantially benefit from propagation, but fails to do so consistently, and the dual search is left virtually uninfluenced by the propagation. A similar conclusion holds for using nogood mining in dual search, albeit the performance of primal search becomes consistently worse with nogood mining enabled.

As a result, we use the following configurations: enable sorting and disable propagation and mining for *primal search*, and disable all additional features for *dual search*.

508 5.2 Scheduling Problem Evaluation

⁵⁰⁹ In order to establish whether our approach is beneficial, we compare the disjunctive approach ⁵¹⁰ with baseline Pumpkin and CP-SAT v9.11 using the aforementioned integrals.

We start by evaluating the influence of adding disjointness reasoning to the baseline Pumpkin. As seen from Figure 5: (1) the disjunctive approach proves much more successful when using dual search as opposed to primal search, and (2) the disjunctive approach commonly translates to *much* larger improvements in cases where it does not hinder the search, while it can also degrade the performance on a number of instances. Figure 5b suggests



(a) Primal integrals.

(b) Dual integrals.

Figure 5 Comparison of integrals of baseline and disjunctive approaches. Both axes are logarithmic; lines are evenly spaced and correspond to a tenfold relative change between the integrals.

515

⁵¹⁶ that the disjunctive approach can accelerate the search by at least *three orders of magnitude*,

which indicates that our approach can improve state-of-the-art bounds. Indeed, our approach
discovers a variety of lower bounds that are both better than the bounds previously reported
and the bounds discovered by all other evaluated approaches. Appendix D introduces a
precise criterion for reporting the bounds and the complete list of the novel bounds.

Remarkably, our approach is able to close five of the previously open RCPSP/max J30 521 instances in under a second. For each of those instances, there is a disjoint clique Ω at the 522 root level with d_{Ω} equal to the optimal bound. In all five cases, these cliques correspond 523 to various resource constraints, making them similar to the aforementioned 3n instances; 524 fittingly, CP-SAT has been able to certify only one of them after ten minutes of search, 525 whereas the other instances time out with all other approaches. To clarify this point, we 526 have reproduced and justified such a clique for one of those instances in Appendix E, thereby 527 closing one of the previously open instances without any search. 528

Both plots in Figure 5 suggest that instance features have a major influence on the change in performance. To establish which instance properties aid our approach, we introduce a measure of similarity of a Cumulative constraint to a Disjunctive constraint.

▶ Definition 17. For a Cumulative constraint over a resource with capacity C, and tasks T, we call the quantity $\text{RC} = \frac{1}{|T|} \sum_{i \in T} \frac{r_i}{C}$ the constrainedness of that Cumulative. Given a CSP, we say that its resource constrainedness RC is the maximum constrainedness among its cumulative constraints. A larger value indicates more resource contention.

Figure 6 demonstrates the impact of the resource constrainedness: there is exactly *one* instance with an RC under 15% which benefits from using the disjunctive approach, whereas instances with a "scarce" resource (i.e., $\text{RC} \ge 40\%$) account for the largest improvements and have the most variability. The performance on the instances between those two values mostly degrades, with a small number of instances seeing minor improvements.



Figure 6 Distribution of dual integral ratio of baseline and disjunctive approach for various values of resource constrainedness. Points to the right of the vertical line correspond to improvements by the disjunctive approach, and points to the left correspond to deterioration.

540

Last, we evaluate our approach against CP-SAT. Figure 7a shows that the baseline consistently performs better on PSPLIB, with worse performance on pack(d) instances. On the other hand, Figure 7b indicates that the disjunctive approach has a much smaller edge against CP-SAT, yet it can improve certain instances by many orders of magnitude.

545 **6** Conclusions

We have presented a novel way to aggregate information about the disjointness of tasks in scheduling problems across multiple constraints. We do this by creating new variables representing the disjointness between pairs of tasks, mining for this disjointness during search, and heuristically finding disjoint cliques that lead to conflicts. We show that the "standard" approaches require exponentially many conflicts to prove optimality on crafted instances,

12:16 Unite and Lead: Finding Disjunctive Cliques for Scheduling Problems



Figure 7 Comparison of Pumpkin and CP-SAT dual integrals. All axes are logarithmic; lines on the plots are evenly spaced and correspond to a tenfold relative change.

while our approach can show their infeasibility instantaneously. We also show that our 551 approach provides improvements in the order of magnitudes for benchmark instances with 552 high resource contention, indicating that reasoning across multiple constraints during search 553 is of vital importance for inferring the *global* structure of instances. 554

One direction for future work would be determining what is a good clique. Another 555 direction would be to look into incorporating more Disjunctive reasoning over the cliques. 556 Additionally, future work could focus on mining *jointness* rather than *disjointness*. 557

558		References
559	1	Gilles Audemard and Laurent Simon. Refining restarts strategies for SAT and UNSAT.
560		In Proceedings of the 18th international conference on principles and practice of constraint
561		programming, pages 118–126, Berlin, Heidelberg, 2012. Springer-Verlag.
562	2	Philippe Baptiste, Claude Le Pape, and Wim Nuijten. Satisfiability tests and time-bound
563		adjustments for cumulative scheduling problems. Annals of operations research, 92:305-333,
564		1999. doi:10.1023/a:1018995000688.
565	3	Philippe Baptiste and Claude Le Pape. Constraint propagation and decomposition techniques
566		for highly disjunctive and highly cumulative project scheduling problems. Constraints: an
567		international journal, $5(1/2)$:119–139, January 2000. doi:10.1023/a:1009822502231.
568	4	Paul Beame and Toniann Pitassi. Simplified and improved resolution lower bounds. In
569		Proceedings of the 37th annual symposium on foundations of computer science, FOCS '96,
570		page 274, USA, 1996. IEEE Computer Society.
571	5	Nicolas Beldiceanu and Mats Carlsson. A new multi-resource cumulatives constraint with
572		negative heights. In Proceedings of the 8th international conference on principles and practice
573		of constraint programming, CP '02, pages 63–79, Berlin, Heidelberg, 2002. Springer-Verlag.
574	6	Timo Berthold. Measuring the impact of primal heuristics. Operations research letters,
575		41(6):611-614, 1 November 2013. doi:10.1016/j.orl.2013.08.007.
576	7	Peter Brucker, Sigrid Knust, Arno Schoo, and Olaf Thiele. A branch and bound algorithm for
577		the resource-constrained project scheduling problem. European journal of operational research,
578	_	107(2):272–288, 1 June 1998. doi:10.1016/s0377-2217(97)00335-4.
579	8	Jacques Carlier and Emmanuel Néron. On linear lower bounds for the resource constrained pro-
580		ject scheduling problem. European journal of operational research, 149(2):314–324, 1 September
581	-	2003. doi:10.1016/s0377-2217(02)00763-4.
582	9	Geoffrey Chu, Peter J Stuckey, Andreas Schutt, Thorsten Ehlers, Graeme Gange, and Kathryn
583		Francis. Chuffed: The chuffed CP solver. https://github.com/chuffed/chuffed.

- Rina Dechter. Learning while searching in constraint-satisfaction-problems. In *Proceedings* of the fifth AAAI national conference on artificial intelligence, AAAI'86, pages 178–183,
 Philadelphia, Pennsylvania, 1986. AAAI Press.
 Delft High Parformance Computing Control Delft Plue 2024
- ⁵⁸⁷ 11 Delft High Performance Computing Centre. DelftBlue, 2024.
- Emir Demirović, Maarten Flippo, Imko Marijnissen, Konstantin Sidorov, and Jeff Smits. Pump kin: a lazy clause generation constraint solver in Rust. https://github.com/ConSol-Lab/
 Pumpkin, 2024.
- ⁵⁹¹ 13 Ulrich Dorndorf. Consistency tests. In *Project scheduling with time windows*, chapter 4, pages
 ⁵⁹² 31-65. Physica-Verlag HD, Heidelberg, 2002. doi:10.1007/978-3-642-57506-8_4.
- Thibaut Feydy and Peter J Stuckey. Lazy clause generation reengineered. In *Proceedings of* the 15th international conference on principles and practice of constraint programming, CP'09,
 pages 352–366, Berlin, Heidelberg, 2009. Springer-Verlag.
- John Forrest, Ted Ralphs, Stefan Vigerske, Haroldo Gambini Santos, John Forrest, Lou
 Hafer, Bjarni Kristjansson, jpfasano, EdwinStraver, Jan-Willem, Miles Lubin, rlougee, a andre, jpgoncal, Samuel Brito, h-i-gassmann, Cristina, Matthew Saltzman, tosttost, Bruno
 Pitrus, Fumiaki Matsushima, Patrick Vossler, Ron @ Swgy, and to-st. coin-or/Cbc: Release
 releases/2.10.12, 2024. doi:10.5281/ZENOD0.13347261.
- Steven Gay, Renaud Hartert, and Pierre Schaus. Time-table disjunctive reasoning for the
 cumulative constraint. In *Integration of AI and OR techniques in constraint programming*, Lec ture notes in computer science, pages 157–172, Cham, 2015. Springer International Publishing.
 doi:10.1007/978-3-319-18008-3_11.
- Gecode Team. Gecode: Generic constraint development environment. http://www.gecode.org,
 2006.
- Vincent Gingras and Claude-Guy Quimper. Generalizing the edge-finder rule for the cumulative constraint. In *Proceedings of the twenty-fifth international joint conference on artificial intelligence*, IJCAI'16, pages 3103–3109. AAAI Press, 9 July 2016. doi:10.5555/3061053.
 3061056.
- Armin Haken. The intractability of resolution. *Theoretical computer science*, 39:297–308, 1985.
 doi:10.1016/0304-3975(85)90144-6.
- Q Huangfu and J A J Hall. Parallelizing the dual revised simplex method. Mathematical
 programming computation, 10(1):119-142, March 2018. doi:10.1007/s12532-017-0130-5.
- Joey Hwang and David G Mitchell. 2-way vs. d-way branching for CSP. In Principles and
 practice of constraint programming CP 2005, Lecture notes in computer science, pages 343–357,
 Berlin, Heidelberg, 1 October 2005. Springer Berlin Heidelberg. doi:10.1007/11564751_27.
- Roger Kameugne, Laure Pauline Fotso, Joseph Scott, and Youcheu Ngo-Kateu. A quad ratic edge-finding filtering algorithm for cumulative resource constraints. Constraints: an
 international journal, 19(3):243–269, July 2014. doi:10.1007/s10601-013-9157-z.
- Richard M Karp. Reducibility among combinatorial problems. In Complexity of computer computations, pages 85–103, Boston, MA, 1972. Springer US. doi:10.1007/978-1-4684-2001-2\
 _9.
- 62424Robert Klein and Armin Scholl. Computing lower bounds by destructive improvement: An625application to resource-constrained project scheduling. European journal of operational research,626112(2):322-346, January 1999. doi:10.1016/s0377-2217(97)00442-6.
- Rainer Kolisch and Arno Sprecher. PSPLIB a project scheduling problem library. European
 journal of operational research, 96(1):205–216, January 1997. doi:10.1016/s0377-2217(96)
 00170-1.
- Oumar Koné, Christian Artigues, Pierre Lopez, and Marcel Mongeau. Event-based MILP
 models for resource-constrained project scheduling problems. Computers & operations research,
 38(1):3–13, 1 January 2011. doi:10.1016/j.cor.2009.12.011.
- Joao Marques-Silva, Ines Lynce, and Sharad Malik. Conflict-driven clause learning SAT solvers.
 In *Handbook of satisfiability*, Frontiers in artificial intelligence and applications, chapter 4. IOS
 Press, 2 February 2021. doi:10.3233/faia200987.

12:18 Unite and Lead: Finding Disjunctive Cliques for Scheduling Problems

- Matthew W Moskewicz, Conor F Madigan, Ying Zhao, Lintao Zhang, and Sharad Malik.
 Chaff: engineering an efficient SAT solver. In *Proceedings of the 38th design automation conference*. ACM, 2002. doi:10.1109/dac.2001.935565.
- Wim Nuijten. Time and resource constrained scheduling: a constraint satisfaction approach.
 PhD thesis, Technische Universiteit Eindhoven, Eindhoven, 1994. doi:10.6100/IR431902.
- Ramón Alvarez-Valdés Olaguíbel and José Manuel Tamarit Goerlich. Heuristic algorithms
 for resource-constrained project scheduling: A review and an empirical analysis. Advances in
 project scheduling, pages 113–134, 1989.
- 644 31 Laurent Perron and Frédéric Didier. CP-SAT, 7 May 2024.
- Francesca Rossi, Peter van Beek, and Toby Walsh, editors. Handbook of constraint programming.
 Foundations of artificial intelligence. Elsevier Science, London, England, 18 August 2006.
 doi:10.1016/s1574-6526(06)x8001-x.
- Andreas Schutt, Thibaut Feydy, Peter J Stuckey, and Mark G Wallace. Explaining the
 cumulative propagator. *Constraints: an international journal*, 16(3):250–282, July 2011.
 doi:10.1007/s10601-010-9103-2.
- Andreas Schutt, Thibaut Feydy, Peter J Stuckey, and Mark G Wallace. Solving RCPSP/max
 by lazy clause generation. *Journal of scheduling*, 16(3):273–289, June 2013. doi:10.1007/
 s10951-012-0285-x.
- Beter J Stuckey. RCPSP. https://people.eng.unimelb.edu.au/pstuckey/rcpsp/. Accessed:
 2025-3-29.
- Peter J Stuckey, Thibaut Feydy, Andreas Schutt, Guido Tack, and Julien Fischer. The
 MiniZinc Challenge 2008–2013. AI magazine, 35(2):55–60, 1 June 2014. doi:10.1609/aimag.
 v35i2.2539.
- ⁶⁵⁹ **37** Petr Vilím. $o(n \log n)$ filtering algorithms for unary resource constraint. In Jean-Charles Régin ⁶⁶⁰ and Michel Rueher, editors, *Proceedings of CP-AI-OR 2004*, volume 3011 of *Lecture Notes in* ⁶⁶¹ *Computer Science*, pages 335–347, Nice, France, April 2004. Springer-Verlag.
- Petr Vilím. Computing explanations for the unary resource constraint. In Roman Barták and
 Michela Milano, editors, Integration of AI and OR techniques in constraint programming for
 combinatorial optimization problems, second international conference, CPAIOR 2005, Prague,
 Czech Republic, May 30 June 1, 2005, volume 3524 of Lecture Notes in Computer Science,
 pages 396-409. Springer, 2005. doi:10.1007/11493853_29.
- Petr Vilím. Timetable edge finding filtering algorithm for discrete cumulative resources. In
 Proceedings of the 8th international conference on integration of Ai and OR techniques in
 constraint programming for combinatorial optimization problems, CPAIOR'11, pages 230–245,
 Berlin, Heidelberg, 2011. Springer-Verlag.
- 40 Petr Vilím, Philippe Laborie, and Paul Shaw. Failure-directed search for constraint-based
 scheduling. In CPAIOR '15: proceedings of the 12th international conference on integration
 of AI and OR techniques in constraint programming for combinatorial optimization problems,
 Barcelon, Spain, 2015. Springer-Verlag.

Α Conversion of 3n problem proofs into pigeonhole principle 675 pseudo-proofs 676

For convenience, we reproduce the definition of the C-RES proof system subject to the 677 notation differences between our paper and the work of Hwang and Mitchell [21]. 678

Definition 18. Given a set of variables \mathcal{X} with domains \mathcal{D} , a **clause** is a constraint of the 679 $form \left(\bigvee_{j=1}^{N} \llbracket x_{j}^{+} = v_{j}^{+} \rrbracket\right) \vee \left(\bigvee_{j=1}^{M} \neg \llbracket x_{j}^{-} = v_{j}^{-} \rrbracket\right) \ with \ x_{j}^{\pm} \in \mathcal{X}, v_{j}^{\pm} \in \mathcal{D}(x).$ 680

▶ **Definition 19.** Given a CSP $(\mathcal{X}, \mathcal{D}, \mathcal{C})$, a sequence of clauses \mathcal{P} is an (unsatisfiability) 681 **proof** of this CSP if the last clause is empty, and every clause $\omega \in \mathcal{P}$ is produced by one of 682 the following rules: 683

Resolution ω can be represented as $\omega' \vee \omega''$ for some clauses ω', ω'' such that $\omega' \vee [x = v]$ 684 and $\omega'' \vee \neg [x = v]$ are clauses encountered earlier in the proof. We denote the result of 685

this operation as $\omega' \diamond \omega''$. 686

687

Domain clause $\omega = \bigvee_{v \in \mathcal{D}(x)} [\![x = v]\!]$ for some variable x. **Unique value clause** $\omega = \neg [\![x = u]\!] \lor \neg [\![x = v]\!]$ for some variable x and two distinct values 688 $u, v \in \mathcal{D}(x).$ 689

Constraint clause ω is **implied** by some constraint $c \in C$, that is, any solution satisfying c 690 also satisfies ω . 691

Each step of the proof is either a "consistency" clause encoding that any variable is 692 assigned to exactly one value, a constraint clause (which in our case corresponds to a 693 propagation), or a resolution of two clauses (which typically corresponds to branching or 694 clause learning). The original C-RES definition by Hwang and Mitchell also restricts \mathcal{C} to 695 only contain nogoods, and thus the constraint clauses are the negations of nogoods in \mathcal{C} . 696 Again, that difference does not impact the proof definition, because any constraint can be 697 replaced by a set of all nogoods implied by it, which impacts the size of \mathcal{C} but not the proof 698 length. 699

In this appendix, we show that the proofs of 3n instances can be encoded as the proofs 700 of the *pigeonhole formula* [19] with 3n variables, with clauses encoding the pigeonhole 701 subformulas for up to 2n variables. We use the following definition of a pigeonhole problem 702 in this text: 703

 \blacktriangleright Definition 20. A pigeonhole formula PHP_m is the unsatisfiable propositional formula 704

on $(m-1) \times m$ variables $p_{i,j}, 1 \leq i \leq m-1, 1 \leq j \leq m$ having the following clauses: 705

Pigeon clauses Any pigeon is placed in a hole: $p_{1,i} \vee p_{2,i} \vee \cdots \vee p_{m-1,i} \forall 1 \leq i \leq m-1$. 706 **Hole clauses** No hole contains two pigeons: $\bar{p}_{i,j} \vee \bar{p}_{i,k} \forall 1 \leq i \leq m-1, 1 \leq j < k \leq m$. 707

To simplify the description of the encoding, we switch to a more lenient notion of a proof 708 proposed by Beame and Pitassi [4]: 709

▶ Definition 21. Given a formula PHP_m , an assignment α is called critical if some (m-1)710 pigeons are assigned to (m-1) distinct holes, that is, $p_{1,\pi_1}, \ldots, p_{m-1,\pi_{m-1}}$ are satisfied by 711 α for some pairwise different $\pi_i \in [1,m], 1 \leq j \leq m-1$; if additionally the only assigned 712 pigeon has index k, then α is called k-critical. 713

A pair of clauses ω', ω'' is **congruent** if ω' and ω'' are equal for any critical assignment. 714 If any critical assignment satisfying clauses ω' and ω'' also satisfies another clause ρ , then 715 this clause is a critical implication by ω' and ω'' . 716

12:20 Unite and Lead: Finding Disjunctive Cliques for Scheduling Problems

One of the immediate simplifications resulting from the notion of congruency is that we
 can assume without loss of generality that all clauses defined on pigeonhole variables are
 defined without negations:

Proposition 22 (Beame and Pitassi [4], Section 3). Given a formula PHP_m, any clause ω is congruent to a **positive clause** ω⁺ obtained by replacing all literals $\bar{p}_{i,j} \in \omega$ with a negative polarity by the conjunction $\bar{p}_{i,j} \mapsto \bigvee_{k \in [1,m], k \neq i} p_{i,k}$ of variables over all pigeons except the *i*-th one.

The core result of this appendix is the reduction from 3n proofs to pigeonhole proofs:

▶ Lemma 23. Any proof \mathcal{P} of a 3n problem instance can be encoded into a sequence of clauses \mathcal{P}' with $|\mathcal{P}'| \leq |\mathcal{P}|$ such that the final clause is empty and any clause is either a pigeon clause⁴ of PHP_{3n}, a critical implication of previous clauses, or a subproblem clause of the form

$$\sigma(K,V) := \bigvee_{k \in [1,3n-1] \setminus K, v \in V} p_{k,v}, |V| > |K|$$

$$(7)$$

730 with $1 < |V| \le 2n$.

We start by replacing the original 3n problem with a simpler formulation that does not invalidate unsatisfiability proofs of 3n problems.

Definition 24. Given a set of variables X, the Alldifferent(X) constraint is true for the solutions in which any two distinct variables $x, y \in X$ are assigned to different values.

▶ Lemma 25. Any proof of a 3n problem instance is also a proof of a CSP on the same variables and constraints Alldifferent(G) with $G \in \{X \cup Y, Y \cup Z, Z \cup X\}$.

Proof. Let \mathcal{P} be a proof of the 3n problem instance, and ω is a constraint clause in \mathcal{P} with respect to the original 3n instance. We show that this is also a constraint clause with respect to the CSP in the lemma statement. This is sufficient to show that \mathcal{P} is also a valid proof of the new problem, because the resolution steps only depend on the previous clauses. As for the domain and unique value clauses, they only depend on the underlying variable definitions, which are the same between two problem instances, and are also valid for the new problem instance.

Without loss of generality, suppose that a clause ω is implied by the first constraint in Definition 5. Consider an arbitrary solution \mathcal{I} that falsifies ω and thus also falsifies the Cumulative constraint in question. Any overflow of such a constraint can be described as either scheduling two tasks from X in parallel, scheduling one task from X and one task from Y in parallel, or scheduling three tasks from Y in parallel. In each case, this assignment violates the Alldifferent $(X \cup Y)$. Thus, Alldifferent $(X \cup Y)$ is false whenever ω is false, or, in other words, ω is true whenever Alldifferent $(X \cup Y)$ is true.

⁷⁵¹ We now complete the reduction to pigeonhole proofs as follows:

Proof of Lemma 23. Introduce a one-to-one mapping between atomic predicates in \mathcal{P} and pigeonhole variables with $[x_j = v] \leftrightarrow p_{v,j}$ for some ordering of variables $\mathcal{X} = \{x_1, \ldots, x_{3n}\}$; we further refer to this as the canonical mapping. We show that the proof \mathcal{P}' produced by

⁴ Hole clauses are tautological with respect to critical assignments.

(i) replacing atomic constraints in C-RES clauses with canonical mapping and (ii) discarding tautologies satisfies the lemma conditions. More specifically, we show that each clause $\omega_P \in \mathcal{P}'$ produced from a C-RES clause $\omega_C \in \mathcal{P}$ either falls into one of the three stated categories or is tautologically true. Additionally, we enforce the following invariant: if a non-tautological clause ω_C is falsified by an assignment \mathcal{I} of (3n-1) variables to (3n-1)different values, then the critical assignment produced from it by canonical mapping falsifies ω_P .

We start with the "consistency" clauses. If $\omega_C = \llbracket x_k = 1 \rrbracket \lor \cdots \lor \llbracket x_k = 3n-1 \rrbracket$ is a domain clause, then $\omega_P = p_{1,k} \lor \cdots \lor p_{3n-1,k}$ is a valid clause in \mathcal{P}' as a pigeon clause of PHP_{3n}. Additionally, if \mathcal{I} falsifies ω_P , that means that x_k is unassigned; the canonical mapping leaves k-th pigeon unassigned, which falsifies ω_P . On the other hand, if $\omega_C = \neg \llbracket x_k = u \rrbracket \lor \neg \llbracket x_k = v \rrbracket$ is a unique value clause, then $\omega_P = \bar{p}_{u,k} \lor \bar{p}_{v,k}$ is true for any critical assignment and is thus a tautology.

Next, consider the case when ω_C is a *constraint* clause. By Lemma 25, we can 768 assume without loss of generality that $\omega_C \in \mathcal{P}$ is a constraint clause with respect to 769 All different $(X \cup Y)$. Thus, if ω_C is false, then any assignment of $X \cup Y$ assigns the 770 same values to some two variables in this set. By Hall's theorem, that means that there 771 is a set of variables $Q \subseteq X \cup Y$ with indices V and a set of values $K \subset [1, 3n - 1]$ such 772 that $|K| < |V| \le 2n$ and assigning all variables in Q to values in K falsifies ω_C . Con-773 versely, that means that assigning some variable in Q to a value outside of K, equivalently, in 774 $K = [1, 3n-1] \setminus K$ satisfies ω_C . We have encoded a condition $\bigvee_{v \in V} (x_v \notin K) = \bigvee_{v \in V} (x_v \in K)$ 775 implying ω_C , which after applying the canonical mapping becomes $\omega_P = \bigvee_{k \in \bar{K}, v \in V} p_{k,v}$ and 776 can be directly added to the proof, since it coincides with $\sigma(K, V)$ from Equation 7. 777

Last, suppose that $\omega_C = \omega''_C \diamond \omega''_C$ is a resolution of some previously introduced propositional 778 clauses ω'_C, ω''_C . Let ω'_P and ω''_P be the corresponding clauses in \mathcal{P}' ; suppose first that neither 779 of those clauses was skipped. In that case, we observe the encoding ω_P of ω_C is critically 780 implied by ω'_P and ω''_P : otherwise, there would be a critical assignment α satisfying both ω'_P 781 and ω_P'' but falsifying ω_P , and applying the canonical mapping to it yields an assignment 782 \mathcal{I} satisfying ω'_{C} and ω'_{C} but falsifying ω_{C} , contradicting the soundness of the resolution 783 step. In case when one (or both) of the original clauses is a tautology, the same reasoning 784 holds, because resolving with a tautological clause does not increase the set of falsifying 785 assignments. 786

B Exponential lower bound on pigeonhole proofs

⁷⁸⁸ We start with an auxiliary fact that establishes a lower bound on the width of a valid proof:

⁷⁸⁹ ► Lemma 26. Let $n \in \mathbb{N}_{\geq 1}$, $0 < \gamma \leq 1$, and consider a sequence \mathcal{P} such that any clause in ⁷⁹⁰ it is either a pigeon clause of PHP_{(2+γ)n}, a critical implication of the previous clauses, or ⁷⁹¹ a subproblem clause $\sigma(K, V)$, $1 \leq |K| < |V| \leq 2n$. Then \mathcal{P} has a clause with at least $2\gamma n^2$ ⁷⁹² literals.

⁷⁹³ **Proof.** Given a clause ω , let Complexity(ω) be the smallest number of pigeon clauses that ⁷⁹⁴ imply ω on all critical assignments; in particular, Complexity(ω) = 1 for all pigeon clauses ⁷⁹⁵ $\omega \in \text{PHP}_{(2+\gamma)n}$, Complexity(\bot) = $(2 + \gamma)n$, and Complexity($\sigma(K, V)$) = |V|.

⁷⁹⁶ Observe that a proof contains a clause $\hat{\omega}$ of complexity greater than γn and not greater ⁷⁹⁷ than 2n. To verify this, let $V^* \leq 2n$ be the largest cardinality of a set V used in a subproblem ⁷⁹⁸ clause. If $V^* > \gamma n$, then the subproblem clause that achieves this bound also satisfies the ⁷⁹⁹ complexity bound. Otherwise, assume that this is not the case and that all proof clauses

12:22 Unite and Lead: Finding Disjunctive Cliques for Scheduling Problems

have complexity that is either at most γn or more than 2n. Observe that if ω is critically implied by ω' and ω'' , then Complexity(ω) \leq Complexity(ω') + Complexity(ω''), since the implication relation is transitive. As the formula clauses have complexity 1 and subproblem clauses have complexity at most γn , then *all* their resolvents ω have to have the complexity of at most γn : otherwise, Complexity(ω) $\leq \gamma n + \gamma n \leq 2n$ violates our assumption. But then the same bound has to hold for the resolvents of those resolvents, and so on until the empty clause, which has complexity ($2 + \gamma$)n > 2n, which is a contradiction.

We complete the proof by demonstrating that this clause satisfies the lemma conditions. Let $\hat{\Omega}$ be the set of formula clauses that implies $\hat{\omega}$ such that $\gamma n < |\hat{\Omega}| \le 2n$; observe that $\hat{\omega}$ has at least $|\hat{\Omega}|((2 + \gamma)n - |\hat{\Omega}|)$ literals [4, Lemma 1], and this expression achieves the smallest value $2\gamma n^2$ at the endpoints of the interval on $|\hat{\Omega}|$.

We now derive an exponential bound on the proof length of PHP_{3n} with extra clauses:

▶ Lemma 27. Let $n \in \mathbb{N}_{\geq 1}$, and consider a sequence of clauses \mathcal{P} such that any clause is either a pigeon clause of PHP_{3n}, a resolution of the previous clauses, or a subproblem clause $\sigma(K, V)$ with $1 \leq |K| < |V| \leq 2n$. Then \mathcal{P} has at least $2^{0.08n}$ clauses for sufficiently large n.

Proof. Suppose this is not the case, and we discovered a proof of length $L < 2^{0.08n}$. For the rest of the proof, we say that the clause is *long* if it has at least n^2 literals. By pigeonhole principle, there is a variable $x_{i,j}$ that is contained in at least $\frac{Ln^2}{3n(3n-1)} \ge \frac{1}{9}L$ long clauses. Observe that setting $x_{i,j} = 1$ with $x_{k,j} = x_{i,\ell} = 0$ for all $k \ne i, \ell \ne j$ and discarding all satisfied clauses from both the proof and the formula yields a proof of a PHP_{3n-1} with the same assumptions on subproblem clauses but having at most $\frac{8}{9}L$ long clauses.

Repeating this variable elimination procedure for $T = \lceil \log_{9/8} L \rceil$ steps yields a proof of PHP_{3n-T} with at most $\left(\frac{8}{9}\right)^T L < 1$ long clause, or, in other words, *no* long clauses. On the other hand, we have a pigeonhole formula with $3n - T = 3n - \lceil 0.08n \times \log_{9/8} 2 \rceil \ge \left(2 + \frac{1}{2}\right) n$ variables and a proof with no clauses of length at least n^2 . However, we can now rewrite $3n - T = \hat{\gamma}n$ with $\hat{\gamma} \ge 1/2$ and apply Lemma 26, which implies that the proof in question has to have a clause with at least $2\hat{\gamma}n^2 \ge n^2$ literals. That is a contradiction, therefore, \mathcal{P} has to have length of at least $2^{0.08n}$.

The derivation of the intractability theorem bound is a compilation of the results above:

Proof of Theorem 6. Let \mathcal{P}' be the pigeonhole formula proof produced from \mathcal{P} from Lemma 23 with $|\mathcal{P}| \geq |\mathcal{P}'|$. By Lemma 27, we can rewrite the bound further as $|\mathcal{P}| = \Omega(2^{0.08n})$, and the theorem bound follows from the $1.05 < 2^{0.08}$ inequality.

832 C Explanations

Assume without loss of generality that $\text{EST}_i \leq \text{EST}_j$, as otherwise the tasks *i* and *j* can be swapped. Then we generate the explanations $\mathcal{E}_{sn}(s_i, s_j, p_i, p_j)$ as prescribed by Table 1.

The common idea behind the listed explanations is to lift the bounds in a way that ensures the absence of overlap when one of the two predicates is true. For example, if $p_i = [s_i \le v_i]$ and $p_j = [s_j \ne v_j]$, then the first predicate $[s_j \ge v_i + d_i]$ in the explanation states the following fact: 'if p_i is true then s_j starts only after task *i* has ended'. Similarly, the second predicate $[s_i \le v_j - d_i + 1]$ in the explanation states that task *i* should end before task *j* starts; we add an extra unit of time since we know that $[s_j \ne v_j] \implies [s_j \ge v_j + 1]$.

Atomic predicate p_i	Atomic predicate p_j	Explanation $\mathcal{E}_{sn}(s_i, s_j, p_i, p_j)$
$[\![s_i \le v_i]\!]$	$[\![s_j \ge v_j]\!]$	$\llbracket s_j \ge v_i + d_i \rrbracket \land \llbracket s_i \le v_j - d_i \rrbracket$
$[\![s_i \le v_i]\!]$	$[\![s_j \neq v_j]\!]$	$[[s_j \ge v_i + d_i]] \land [[s_i \le v_j - d_i + 1]]$
$[\![s_i \le v_i]\!]$	$[\![s_j = v_j]\!]$	$\llbracket s_j \ge v_i + d_i \rrbracket \land \llbracket s_i \le v_j - d_i \rrbracket$
$[\![s_i \neq v_i]\!]$	$[\![s_j \ge v_j]\!]$	$\llbracket s_j \ge v_i + d_i - 1 \rrbracket \land \llbracket s_i \le v_j - d_i \rrbracket$
$[\![s_i \neq v_i]\!]$	$[\![s_j \neq v_j]\!]$	$[[s_j \ge v_i + d_i - 1]] \land [[s_i \le v_j - d_i + 1]]$
$[\![s_i \neq v_i]\!]$	$\llbracket s_j = v_j \rrbracket$	$\llbracket s_j \ge v_i + d_i - 1 \rrbracket \land \llbracket s_i \le v_j - d_i \rrbracket$
$[\![s_i = v_i]\!]$	$[\![s_j \ge v_j]\!]$	$\llbracket s_j \ge v_i + d_i \rrbracket \land \llbracket s_i \le v_j - d_i \rrbracket$
$[\![s_i = v_i]\!]$	$\llbracket s_j \neq v_j \rrbracket$	$[[s_j \ge v_i + d_i]] \land [[s_i \le v_j - d_i + 1]]$
$[\![s_i = v_i]\!]$	$\llbracket s_j = v_j \rrbracket$	$\llbracket s_j \ge v_i + d_i \rrbracket \land \llbracket s_i \le v_j - d_i \rrbracket$

Table 1 List of explanations $\mathcal{E}_{sn}(s_i, s_j, p_i, p_j)$ produced for each pair of atomic constraints p_i, p_j .

D Novel bounds

We report the bounds discovered by our approach if they are *both* better than the previous reported bounds and are not directly reproducible without using our approach. More precisely, we report bounds that are simultaneously (a) tighter than the bounds reported in the previous sources known to us that used the same benchmarks [25, 40, 39, 35], and (b) either tighter than any bound derived by non-disjunctive approaches (CP-SAT and baseline Pumpkin) or matches it but was derived at least ten times faster than with any other approach.

Novel upper bounds (makespans) are reported in Table 2, and novel lower bounds are reported in Table 3; the same data is available as supplementary materials. All durations reported in both tables are in MM:SS format; Table 3 additionally reports closed instances and bounds derived with mining. To indicate the remaining optimality gap, we also state the best known lower bound in Table 2 and the best known makespan in Table 3; in either case, that bound is the tightest among the previously reported values and the values discovered by non-disjunctive approaches.

Problem	Collection	#	Ref. objective	Our objective	Time	Best bound
RCPSP/max	С	61	378	374	07:03	345
$\operatorname{RCPSP}/\operatorname{max}$	\mathbf{C}	69	380	371	08:47	356
RCPSP	J90	5-4	103	102	01:02	101
RCPSP	J120	27-7	125	124	08:21	122
RCPSP	J120	46-10	188	187	57:34	184
RCPSP	J120	9-4	87	86	06:40	85

Table 2 Novel upper bounds derived with our approach.

E Optimality proof for the RCPSP/max instance #64 of collection J30

855

856

In this appendix, we support the optimality claim on instance # 64 from the J30 collection by showing that the tasks with indices $\Omega = \{1, \ldots, 9, 11, \ldots, 30\}$ are pairwise disjoint, which implies optimality since $\sum_{t \in \Omega} d_t = 169$. This instance has five resources, each with the capacity of five units.

12:24 Unite and Lead: Finding Disjunctive Cliques for Scheduling Problems

Problem	Collection	#	Ref. bound	Our bound	Time	Best objective
RCPSP/max	J30	64	141	169*	< 1s	169
RCPSP/max	J30	65	144	162^{*}	< 1s	162
RCPSP/max	J30	151	142	157^{*}	< 1s	157
RCPSP/max	J30	153	163	176^{*}	< 1s	176
RCPSP/max	J30	155	125	154^{*}	< 1s	154
RCPSP/max	UBO50	10	154	186^{*}	00:05	186
RCPSP/max	UBO100	4	303	365	44:07	376
RCPSP/max	UBO100	7	281	373	51:47	395
RCPSP/max	UBO100	8	364	376	29:06	385
$\operatorname{RCPSP}/\operatorname{max}$	UBO100	32	353	414	53:52	434
RCPSP/max	UBO100	33	328	397	47:04	406
RCPSP/max	UBO200	35	610	747	57:27	823
$\operatorname{RCPSP}/\operatorname{max}$	UBO200	62	621	702	56:46	796
RCPSP/max	UBO500	36	$1 \ 285$	$1\ 423$	26:37	$1 \ 908$
$\operatorname{RCPSP}/\operatorname{max}$	UBO500	61	1 296	1 533	59:50	$1 \ 944$
RCPSP/max	UBO500	64	1 329	1 489	55:08	1 932
RCPSP	J120	46-9	157	159	18:58	166
RCPSP	Pack-d	2	745	$746^{+}_{$	25:39	747
RCPSP	Pack-d	3	624	625^{+*}	13:54	625
RCPSP	Pack-d	47	2 740	$2 742^{+}$	52:35	2745

Table 3 Novel lower bounds derived with our approach. Rows with an asterisk highlight closed instances; rows with a dagger highlight objective values derived with enabled mining.

Table 4 Disjointness justifications for every pair of tasks in $\Omega_1 \cup \Omega_2 \cup \Omega_3$.

	Ω_1	Ω_2	Ω_3
$\overline{\Omega_1}$	#2: 4 + 4	#5: 2 + 4	#1: 2 + 4
Ω_2		#5: 4 + 4	#5: 2 + 4
Ω_3			#1: 4 + 4

To decrease the number of considered pairs, we partition the claimed clique into six parts and record for each of them the lowest amount of each resource consumed by a task:

⁸⁶³ 1. $\Omega_1 = \{6, 11, 21, 24\}$; lowest resource consumption is (2, 4, 1, 2, 2).

2. $\Omega_2 = \{18, 23, 28\};$ lowest resource consumption is (2, 1, 2, 2, 4).

3. $\Omega_3 = \{30\}$; lowest resource consumption is (4, 1, 3, 2, 1).

4. $\Omega_4 = \{1, 2, 3, 4, 5, 8, 9, 12, 13, 14, 15, 16, 17, 19, 20, 25, 29\};$ lowest resource consumption is (1, 1, 1, 4, 1).

5. $\Omega_5 = \{7\}$; lowest resource consumption is (3, 5, 1, 1, 2).

6. $\Omega_6 = \{22, 26, 27\};$ lowest resource consumption is (1, 1, 5, 1, 2).

The last two groups consume 100% of either resource #2 or resource #3; since all other tasks consume some amount of all resources, any pair of tasks involving a task in $\Omega_5 \cup \Omega_6$ is disjoint. The same holds for any pair of tasks involving Ω_4 and any task in $\Omega_1 \cup \cdots \cup \Omega_4$ due to the overflow of resource #4. Thus, it remains to show that any pair of tasks in $\Omega_1 \cup \Omega_2 \cup \Omega_3$ is disjoint. Table 4 handles the six cases for all remaining pairs of task groups.