




1 On Inferring Cumulative Constraints

2 Konstantin Sidorov   

3 Delft University of Technology, The Netherlands

4 — Abstract —

5 Cumulative constraints are central in scheduling with constraint programming, yet propagation is
6 typically performed per constraint, missing multi-resource interactions and causing severe slowdowns
7 on some benchmarks. I present a preprocessing method for inferring additional cumulative constraints
8 that capture such interactions without search-time probing.

9 This approach interprets cumulative constraints as linear inequalities over occupancy vectors and
10 generates valid inequalities by (i) discovering covers, the sets of tasks that cannot run in parallel,
11 (ii) strengthening the cover inequalities for the discovered sets with lifting, and (iii) injecting the
12 resulting constraints back into the scheduling problem instance. Experiments on standard RCPSP
13 and RCPSP/max test suites show that these inferred constraints improve search performance
14 and tighten objective bounds on favorable instances, while exhibiting more stable performance on
15 unfavorable instances than the prior work. Additionally, these experiments discover 25 new lower
16 bounds and five new best solutions; seven of the lower bounds are obtained directly from the inferred
17 constraints.

18 **2012 ACM Subject Classification** Theory of computation → Constraint and logic programming;
19 Theory of computation → Discrete optimization; Mathematics of computing → Integer programming

20 **Keywords and phrases** Constraint programming, scheduling, cumulative, lifting, valid inequalities

21 **Digital Object Identifier** 10.4230/LIPIcs.CP.2026.1

22 **Related Version** *Preprint*: <https://arxiv.org/abs/2602.15635>

23 **Supplementary Material** The experimental results reported in this paper are obtained with:
24 *Software (Lifting implementation and infrastructure)*: <https://doi.org/10.5281/zenodo.20044686>
25 *Dataset (Solver logs and report files)*: <https://doi.org/10.5281/zenodo.18663551>

26 **Funding** *Konstantin Sidorov*: supported by the TU Delft AI Labs program as part of the XAIT lab.

27 **Acknowledgements** I would like to thank Imko Marijnissen for the enlightening discussions that
28 sparked the development of the technique presented in this paper, and the anonymous reviewers for
29 improving the presentation of the approach and empirical results.

30 **1** Introduction

31 Cumulative constraints and their use in scheduling problems are one of the longstanding
32 success stories of constraint programming (CP), serving as a central abstraction and a source
33 of powerful inferences about resource-constrained scheduling problems. A large body of
34 work has demonstrated how powerful propagation on a single cumulative resource can be,
35 including effective detection of overload situations, pruning based on resource envelopes, and
36 increasingly sophisticated explanations and propagator combinations. All these developments
37 have contributed to the state-of-the-art performance of CP solvers on scheduling benchmarks.

38 Nevertheless, most existing techniques in this field reason about one cumulative con-
39 straint at a time. Even when a model contains several resources, propagation is typically
40 performed independently per resource, with indirect interactions only mediated through
41 domain narrowing of start-time variables. However, the recent work by Sidorov et al. [33] has
42 made a first step to expose the limitations of this single-resource viewpoint, including the $3n$
43 *problem* structure, a pathological case when an instance is equivalent to a trivially infeasible



© Konstantin Sidorov;

licensed under Creative Commons License CC-BY 4.0

32nd International Conference on Principles and Practice of Constraint Programming (CP 2026).

Editor: Nicolas Beldiceanu; Article No. 1; pp. 1:1–1:24

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

44 single-machine scheduling problem but is intractable for a conventional lazy clause-generating
 45 solver, a paradigm subsuming many of the state-of-the-art CP codes.

46 While Sidorov et al. have addressed a special (disjunctive) case of this problem with a
 47 *Unite and Lead* (referred to in the following as UNL) approach and managed to discover
 48 large hidden structures in several known research benchmarks, their technique has two
 49 fundamental limitations. First, while it succeeds at discovering disjunctive structures (that is,
 50 the collections of tasks that cannot be run in parallel), it does not support any extensibility
 51 for deriving more general multi-resource interactions. Second, UNL has to explicitly probe
 52 for disjunctive cliques throughout the search, incurring significant overhead even in instances
 53 where no such structure exists. Since the clique search overhead is not offset by stronger
 54 reasoning on instances without a disjunctive structure, the solver performance deteriorates
 55 sharply on such problems. As a result, Sidorov et al. show that, when UNL does not improve
 56 the search, it commonly slows down the solver by more than an order of magnitude.

57 In this paper, I address these limitations by proposing a new algorithmic framework for
 58 inferring cumulative constraints at the root search node. Instead of restricting attention to
 59 disjunctive reasoning, this approach operates over cumulative constraints as *linear inequalities*
 60 on a special representation of start time variables (occupancy vectors), and reasoning across
 61 multiple resources reduces to *generating valid cutting planes* in the same sense as it is
 62 understood in integer programming. This viewpoint subsumes a wide range of multi-
 63 constraint reasoning techniques from earlier work, including the version of UNL restricted to
 64 the root node and classical lower-bound techniques for RCPSPs.

65 Next, I leverage this connection by describing a workflow for inferring cumulative con-
 66 straints based on the idea of *lifting* [27]. Lifting is a general-purpose technique from
 67 mixed-integer programming for tightening the generated cutting planes $\mathbf{a}^T \mathbf{x} \leq b$ by picking a
 68 variable y not mentioned in the inequality and finding the largest β such that $\mathbf{a}^T \mathbf{x} + \beta y \leq b$
 69 is still valid. I apply this idea by (i) discovering sets of tasks C that cannot be run in parallel,
 70 (ii) lifting the constraint of the form “tasks in C cannot consume more than $(|C| - 1)$ units
 71 of resource,” and (iii) introducing the lifted constraints as new cumulative constraints into
 72 the model.

73 I evaluate the method as a preprocessing step for two state-of-the-art CP solvers on a
 74 suite of standard RCPSP benchmarks. The results suggest that the approach described
 75 in this paper successfully retains most of the positive traits of UNL (substantial gains on
 76 instances with hidden structure), but does so with less severe runtime penalties than UNL.
 77 This is helpful for advancing the best-known bounds for the benchmark instances: with the
 78 presented approach, I report the improved schedules for five benchmarks and tighten the
 79 lower bounds for 25 instances, with seven of them not requiring any search to verify the
 80 correctness of the new bound.

81 The remainder of the paper is organized as follows. Section 2 introduces constraint
 82 programming, the **Cumulative** constraint, and the multi-resource scheduling model that
 83 motivate this work. Section 3 establishes the theoretical foundation of the approach: it
 84 reformulates **Cumulative** constraints as linear inequalities over occupancy vectors, recalls
 85 the necessary machinery of covers and lifting from integer programming, and connects valid
 86 inequalities of the occupancy-vector polyhedron to new **Cumulative** constraints via *inference*
 87 *transfer lemma*. Section 4 builds on this foundation by describing the complete lifting-based
 88 procedure for inferring auxiliary **Cumulative** constraints, including cover enumeration, the
 89 lifting workflow, and the quality metric used to select the most restrictive inferred constraints.
 90 I evaluate this approach in Section 5 from two angles: (a) as a preprocessing technique for
 91 adding cumulative constraints to the input instance before running a CP solver, and (b) as a

92 technique for inferring **Cumulative** constraints with unit capacity,¹ by comparing it with
 93 the UNL approach. Section 6 discusses earlier methods for joint reasoning on cumulative
 94 constraints and their connection with the lifting approach. Finally, Section 7 concludes the
 95 narrative and discusses future research directions.

96 **2 Scheduling with Cumulative constraints**

97 **Constraint programming and the Cumulative constraint.** A constraint satisfaction problem
 98 (CSP) consists of a tuple $(\mathcal{X}, \mathcal{C}, \mathcal{D})$ where \mathcal{X} is the set of *variables*, \mathcal{C} is the set of *constraints*
 99 which specify the relations between variables, and \mathcal{D} is the *domain* which specifies for each
 100 variable which values it can take [30]. A *solution* \mathcal{I} is a mapping that maps each variable in
 101 \mathcal{X} to a single value in the domain of that variable in \mathcal{D} which satisfies all of the constraints
 102 in \mathcal{C} .

103 Constraint programming (CP) is a paradigm for solving CSPs; CP solvers enforce
 104 constraints through *propagators*, the functions that each enforce the local consistency of one
 105 constraint by removing values from \mathcal{D} that cannot participate in any solution. CP solvers
 106 interleave propagation with search via a branch-and-bound procedure: after applying the
 107 propagators, the solver makes a decision that creates several subproblems by splitting the
 108 domain of a variable into two or more parts. This process of applying propagators and
 109 making decisions is performed until either a solution \mathcal{I} is found, the problem is found to be
 110 unsatisfiable, or a termination criterion is met.

111 **Cumulative** is a constraint useful for many scheduling problems to model limited renew-
 112 able resources, such as available work-hours. Definition 1 formalizes the notion of a task
 113 and the **Cumulative** constraint, where the variable s_j encodes the start time of task j , and
 114 $(s_j + d_j)$ encodes its finish time. Thus, we implicitly associate each task $j \in T$ with an
 115 interval $[s_j, s_j + d_j)$.

116 Throughout the text, I use the Iverson bracket $\llbracket P \rrbracket$, defined for an arbitrary predicate P
 117 as follows: $\llbracket P \rrbracket = 1$, if P is true, and $\llbracket P \rrbracket = 0$ otherwise.

118 **► Definition 1.** Let T be a set of *tasks*, and let a task $j \in T$ be defined by its start variable
 119 s_j , resource usage $a_j \in \mathbb{Z}_{\geq 0}$, and duration $d_j \in \mathbb{Z}_{\geq 0}$. Finally, let $b \in \mathbb{Z}_{\geq 0}$ be the **capacity**
 120 of the resource. Then the **Cumulative**($\mathbf{s}, \mathbf{a}, \mathbf{d}, b$) constraint is the condition that at any time
 121 point τ the resource usage of intervals $[s_j, s_j + d_j)$ covering τ does not exceed the capacity:

$$122 \quad \forall \tau \in \mathbb{Z} : \sum_{j \in T} a_j \llbracket s_j \leq \tau < s_j + d_j \rrbracket \leq b. \quad (1)$$

123 An important special case of **Cumulative** is where the resource has unit capacity $b = 1$
 124 and each task has resource usage $a_j \in \{0, 1\}$, known as **Disjunctive**. Note that $a_j = 0$
 125 is permitted and represents a task that does not consume the resource modeled by this
 126 constraint, which is helpful for a multi-resource model, since tasks do not have to consume
 127 all the modeled resources.

128 This special case is significant because many inference procedures that are intractable
 129 for **Cumulative** constraints can be efficiently executed for **Disjunctive** constraints. For
 130 example, determining satisfiability of a single **Cumulative** constraint is NP-complete [2],
 131 but determining satisfiability of a single **Disjunctive** constraint can be done in polynomial
 132 time [37, Section 3].

¹ Also known as disjunctive or no-overlap constraints.

133 **Scheduling problems with multiple Cumulative constraints** Many practical scheduling
 134 problems require reasoning about several resources simultaneously. A canonical example is
 135 the *resource-constrained project scheduling problem* (RCPSP) [8], where a set of n tasks T
 136 must be scheduled—that is, each task j has to get a starting time s_j —subject to m renewable
 137 resources and precedence constraints $s_k - s_j \geq d_j$ for some pairs of tasks (j, k) . Each task
 138 $j \in T$ has a fixed duration $d_j \in \mathbb{Z}_{\geq 0}$ and consumes $a_{ij} \in \mathbb{Z}_{\geq 0}$ units of resource i while active.
 139 Each resource i has a capacity $b_i \in \mathbb{Z}_{\geq 0}$, and the objective is to minimize the makespan –
 140 the time elapsed from the earliest task start² to the latest task finish, $\max_{j \in T}(s_j + d_j)$. The
 141 RCPSP/max variant additionally admits difference constraints of the form $s_k - s_j \geq \gamma_{jk}$
 142 between pairs of tasks, generalizing precedence constraints.

143 A common trait of such problems is that their CP models contain the following collection
 144 of m Cumulative constraints over n variables s encoding the start times of tasks in T :

$$145 \quad \text{Cumulative}(s, \mathbf{a}_i, \mathbf{d}, b_i) \quad (i = 1, \dots, m). \quad (2)$$

146 This part of the CP model is defined by the duration vector $\mathbf{d} \in \mathbb{Z}_{\geq 0}^n$ shared between
 147 constraints, the capacity vector $\mathbf{b} \in \mathbb{Z}_{\geq 0}^m$ containing the capacity of each Cumulative, with b_i
 148 being the capacity of the i -th constraint, and the resource usage matrix $A \in \mathbb{Z}_{\geq 0}^{m \times n}$, where the
 149 i -th row \mathbf{a}_i contains the consumptions on the i -th constraint and the j -th column contains
 150 the consumptions of the j -th task on all constraints.

151 In a conventional CP solver, each of these constraints is handled by an independent
 152 propagator, with interactions between resources mediated only indirectly through domain
 153 narrowing of the start-time variables s . On the other hand, the joint modeling of resource
 154 constraints offers a richer inference structure, as stronger conclusions can be drawn by
 155 reasoning collectively; the relevant ideas in the related work are discussed in Section 6.

156 **3 Cumulative constraints as linear inequalities**

157 This section is organized as follows. Subsection 3.1 introduces occupancy vectors and
 158 establishes that a collection of Cumulative constraints can be restated as a single matrix
 159 inequality over those vectors – this reformulation is the key conceptual bridge between CP
 160 and integer programming that the rest of the paper builds on. Subsection 3.2 recalls the
 161 notions of covers, cover inequalities, and lifting from the theory of knapsack polyhedra [27,
 162 16]; these are standard tools from integer programming presented here for completeness.
 163 Finally, Subsection 3.3 states and proves the *inference transfer lemma* – the main theoretical
 164 contribution of this paper – which establishes that any valid inequality of the occupancy-
 165 vector polyhedron can be introduced into the CP model as a new Cumulative constraint,
 166 and illustrates the complete workflow on an example.

167 **3.1 Polyhedral view via occupancy vectors**

168 We start by introducing a representation of the start-time variables that we need in order to
 169 reason about linear inequalities. Intuitively, we change the variable representation from an
 170 integer start time s_j to a binary vector that records at each integer time point whether a
 171 task j is active.

² We assume without loss of generality that tasks are available from time 0, and therefore $\min_j s_j = 0$.

172 ▶ **Definition 2.** Given a variable x of a CSP and an integer d , from now on referred to as
 173 *time and duration*, the **occupancy vector** $\langle x \mid d \rangle$ is a binary vector indexed by integer
 174 time points $t \in \mathbb{Z}$ with a finite support of d consecutive ones: $\langle x \mid d \rangle_t := \llbracket x \leq t < x + d \rrbracket$.

175 ▶ **Note 3.** The occupancy vector $\langle x \mid d \rangle$ is structurally identical to the `pulse`($x, x + d, 1$)
 176 function of CP Optimizer [21]: both encode the presence of a unit-height pulse over $[x, x + d)$.
 177 The key difference is that `pulse` is used as a building block for cumulative resource expressions
 178 evaluated at a single time point, whereas $\langle x \mid d \rangle$ is treated as a vector admitting linear-
 179 algebraic operations, in particular, the matrix-vector product in Lemma 4.

180 The key point behind this definition is that a system of **Cumulative** constraints from
 181 Equation 2 can be restated as a system of *linear* inequalities over occupancy vectors:

182 ▶ **Lemma 4** (polyhedral view). An assignment of n variables \mathbf{s} satisfies m constraints
 183 `Cumulative`($\mathbf{s}, \mathbf{a}_i, \mathbf{d}, b_i$) for $i = 1, \dots, m$ if and only if the occupancy vectors of \mathbf{s} satisfy

$$184 \quad A \begin{pmatrix} \langle s_1 \mid d_1 \rangle \\ \vdots \\ \langle s_n \mid d_n \rangle \end{pmatrix} \leq (\mathbf{b} \quad \cdots \quad \mathbf{b}), \quad (3)$$

185 where the right-hand side denotes the matrix with every column equal to \mathbf{b} , indexed by time
 186 points $t \in \mathbb{Z}$.

187 **Proof.** (\Rightarrow) Observe that the elements of the left-hand side and right-hand side matrices in
 188 row r and column t are equal to $\sum_j a_{rj} \langle s_j \mid d_j \rangle_t$ and b_r respectively. Consider an arbitrary
 189 pair of indices r, t ; we need to show that $\sum_j a_{rj} \langle s_j \mid d_j \rangle_t \leq b_r$, and by Definition 2, this is
 190 equivalent to

$$191 \quad \sum_j a_{rj} \llbracket s_j \leq t < s_j + d_j \rrbracket \leq b_r.$$

192 But this is implied by Definition 1 for the constraint `Cumulative`($\mathbf{s}, \mathbf{a}_r, \mathbf{d}, b_r$), which is
 193 satisfied by the lemma statement.

194 (\Leftarrow) To show that `Cumulative`($\mathbf{s}, \mathbf{a}_r, \mathbf{d}, b_r$) holds for any r given Equation 3, take an
 195 arbitrary time point t and write Equation 1 for it as $\sum_j a_{rj} \llbracket s_j \leq t < s_j + d_j \rrbracket \leq b_r$. However,
 196 by definition of the matrix product, this is exactly the inequality between the elements of
 197 matrices in Equation 3 at row r , column t . ◀

198 This reformulation is similar to the time-resource decomposition [31] in the sense that it
 199 linearizes a **Cumulative** constraint, but the main difference is that Equation 3 introduces
 200 *one matrix* inequality, which is possible through the use of occupancy vectors. Conversely,
 201 the time-difference decomposition introduces one constraint per time point, as well as a
 202 collection of constraints that enforce Definition 2. While using Equation 3 directly does
 203 not provide any advantage for solving CSPs with **Cumulative** constraints, it provides a new
 204 viewpoint for such models: they can be seen as models with *polyhedral* constraints over
 205 occupancy vectors, with the coefficients of inequalities mapping one-to-one to the arguments
 206 of **Cumulative** constraints.

207 Crucially for this work, this reformulation raises the following question: what *other* linear
 208 inequalities are valid for any occupancy vectors satisfying Equation 3? The following section
 209 introduces the concepts of polyhedral geometry that we need to answer this question.

210 **3.2 Covers and lifting**

211 The occupancy vectors satisfying Equation 3 form a set of the form $\mathcal{P}(A, \mathbf{b}) := \{\mathbf{x} \in \{0, 1\}^n : A\mathbf{x} \leq \mathbf{b}\}$; we now develop the tools needed to reason about valid inequalities for such sets.
 212
 213 In general, A and \mathbf{b} do not contain *all* inequalities that hold on $\mathcal{P}(A, \mathbf{b})$; this gap is filled by
 214 the following definition:

215 **► Definition 5.** *An inequality $\boldsymbol{\pi}^T \mathbf{x} \leq \pi_0$, also written as $(\boldsymbol{\pi}, \pi_0) \in \mathbb{R}_{\geq 0}^{n+1}$, is called a **valid***
 216 ***inequality** for $\mathcal{P}(A, \mathbf{b})$ if it holds for any point in $\mathcal{P}(A, \mathbf{b})$.*

217 Discovering valid inequalities of a given polyhedron $\mathcal{P}(A, \mathbf{b})$ is itself a theoretically rich
 218 topic with major practical implications for integer programming solvers. For the purposes
 219 of this work, we only need a few core notions; I refer the interested reader to survey the
 220 literature on knapsack polyhedra [16] for a more systematic exposition. Consider first a
 221 simple family of valid inequalities:

222 **► Definition 6.** *Given a polyhedron $\mathcal{P}(A, \mathbf{b})$ with a single constraint $\mathbf{a}^T \mathbf{x} \leq b$, we say that*
 223 *$C \subseteq \{1, \dots, n\}$ is a **cover** if $\sum_{i \in C} a_i > b$.*

224 **► Proposition 7.** *A **cover inequality** $\sum_{i \in C} x_i \leq |C| - 1$ is a valid inequality for $\mathcal{P}(A, \mathbf{b})$ if*
 225 *C is a cover.*

226 Next, we need a general-purpose procedure for increasing the left-hand side of a valid
 227 inequality $\boldsymbol{\pi}^T \mathbf{x} \leq \pi_0$ —which may or may not be a cover inequality—known as *lifting* [27].
 228 This procedure maintains the set C of variables that have already been considered for adding
 229 to the inequality, starting with $C \leftarrow \{i : \pi_i \neq 0\}$, and repeating the following steps until
 230 $C = \{1, \dots, n\}$:

- 231 1. Choose an unused variable $i \in \{1, \dots, n\} \setminus C$ or terminate if none exist.
- 232 2. Solve an auxiliary subproblem over the variables in C

$$233 \quad v^*(\boldsymbol{\pi}, \pi_0, C, i; A, \mathbf{b}) := \begin{array}{ll} \text{maximize} & \sum_{c \in C} \pi_c x_c \\ \mathbf{x} \in \{0, 1\}^{|C|} & \\ \text{subject to} & \sum_{c \in C} a_{j,c} x_c \leq b_j - a_{j,i} \quad (j = 1, \dots, m) \end{array} \quad (4)$$

234 that corresponds to maximizing the left-hand side of the current inequality $\boldsymbol{\pi}^T \mathbf{x}$ when
 235 the variable x_i under consideration is set to 1.

- 236 3. Set $\pi_i \leftarrow \pi_0 - v^*(\boldsymbol{\pi}, \pi_0, C, i; A, \mathbf{b})$ and $C \leftarrow C \cup \{i\}$, where $v^*(\boldsymbol{\pi}, \pi_0, C, i; A, \mathbf{b})$ is the
 237 maximum value of the objective of the problem given by Equation 4.

238 The lifting machinery applies to any polyhedron $\mathcal{P}(A, \mathbf{b})$ with nonnegative integer coeffi-
 239 cients. In particular, it applies to the occupancy-vector polyhedron arising from Equation 3
 240 – and crucially, any valid inequality discovered for that polyhedron can be translated directly
 241 back into a **Cumulative** constraint. The next section states this correspondence formally.

242 **3.3 From valid inequalities to Cumulative constraints**

243 First, we prove the main result of Section 3 that connects the scheduling problems with
 244 **Cumulative** constraints with the theory of valid inequalities:

245 **► Lemma 8** (inference transfer). *Any assignment of n variables \mathbf{s} satisfying m constraints*
 246 ***Cumulative**($\mathbf{s}, \mathbf{a}_i, \mathbf{d}, b_i$) for $i = 1, \dots, m$ also satisfies **Cumulative**($\mathbf{s}, \boldsymbol{\pi}, \mathbf{d}, \pi_0$) as long as*
 247 *($\boldsymbol{\pi}, \pi_0$) is a valid inequality for the polyhedron $\mathcal{P}(A, \mathbf{b})$.*

248 **Proof.** We start by applying Lemma 4 to rewrite the implication we need to prove in a
 249 polyhedral view: $AS \leq (\mathbf{b} \ \cdots \ \mathbf{b}) \implies \boldsymbol{\pi}^T S \leq (\pi_0 \ \cdots \ \pi_0)$, where

$$250 \quad S := \begin{pmatrix} \langle s_1 \mid d_1 \rangle \\ \vdots \\ \langle s_n \mid d_n \rangle \end{pmatrix}.$$

251 If this implication is not true, then there is a satisfying assignment of variables \mathbf{s} such
 252 that $AS \leq (\mathbf{b} \ \cdots \ \mathbf{b})$ holds but $\boldsymbol{\pi}^T S \leq (\pi_0 \ \cdots \ \pi_0)$ does not. Let t be the time point
 253 where the latter inequality is not true, that is, $\sum_{j=1}^n \pi_j x_j > \pi_0$ for $x_j = \langle s_j \mid d_j \rangle_t$.

254 We thus have constructed a vector $\mathbf{x} \in \{0, 1\}^n$ such that $A\mathbf{x} \leq \mathbf{b}$ and $\boldsymbol{\pi}^T \mathbf{x} > \pi_0$. However,
 255 that contradicts the definition of the valid inequality, since \mathbf{x} is a point in the polyhedron
 256 $\mathcal{P}(A, \mathbf{b})$ for which the valid inequality $\boldsymbol{\pi}^T \mathbf{x} \leq \pi_0$ does not hold. ◀

257 The inference transfer lemma serves a foundational role in the approach presented in this
 258 paper, as it describes the inferred **Cumulative** constraints in terms of another well-studied
 259 object: valid inequalities of a polyhedron. On the other hand, Lemma 8 does not *prescribe*
 260 any approach for constructing valid inequalities. However, the previous subsection introduced
 261 such a procedure in the form of lifting. To see how the concepts introduced in this section
 262 – occupancy vectors, valid inequalities, and lifting – tie together, consider the following
 263 example:

264 ▶ **Example 9.** Consider a CSP on variables $\mathbf{s} = (s_1, s_2, s_3, s_4)$ with a constraint

$$265 \quad \text{Cumulative}(\mathbf{s}, (5, 3, 2, 4), \mathbf{d}, 7) \tag{5}$$

266 for some vector of durations \mathbf{d} . In polyhedral view, this constraint can be written as

$$267 \quad 5\langle s_1 \mid d_1 \rangle + 3\langle s_2 \mid d_2 \rangle + 2\langle s_3 \mid d_3 \rangle + 4\langle s_4 \mid d_4 \rangle \leq 7.$$

268 To infer a new **Cumulative** constraint, we use Lemma 8 to look for valid inequalities of
 269 the following polyhedron:

$$270 \quad \{(x_1, x_2, x_3, x_4) \in \{0, 1\}^4 : 5x_1 + 3x_2 + 2x_3 + 4x_4 \leq 7\}. \tag{6}$$

271 Observe that $C = \{2, 3, 4\}$ is the *cover* for Equation 6 ($3 + 2 + 4 = 9 > 7$), and it
 272 corresponds to the *cover inequality* $x_2 + x_3 + x_4 \leq 2$ forbidding the simultaneous choice of
 273 those three variables.

274 Since this is a valid inequality, we can improve it by applying the lifting procedure; in this
 275 example, we lift the only remaining variable x_1 . Consider a constraint $\alpha x_1 + x_2 + x_3 + x_4 \leq 2$
 276 for arbitrary $\alpha \geq 0$ and observe that it holds for any solution satisfying Equation 6 with
 277 $x_1 = 0$ regardless of α . Suppose that $x_1 = 1$, and rewrite the lifted constraint and Equation 6
 278 as follows:

$$279 \quad \underbrace{3x_2 + 2x_3 + 4x_4 \leq 2}_{\text{Equation 6}} \implies \underbrace{\alpha \leq 2 - (x_2 + x_3 + x_4)}_{\text{lifted constraint for } x_1=1}. \tag{7}$$

280 Since we look for the tightest valid constraint, we can achieve this by choosing the largest α
 281 such that Equation 7 is still true for any $x_2, x_3, x_4 \in \{0, 1\}$. Specifically, we achieve this by
 282 setting

$$283 \quad \alpha \leftarrow 2 - \max\{x_2 + x_3 + x_4 \mid 3x_2 + 2x_3 + 4x_4 \leq 2, x_2, x_3, x_4 \in \{0, 1\}\}, \tag{8}$$

284 which also happens to be the value of the lifting subproblem v^* in Equation 4 for the values
 285 of $(A, \mathbf{b}, C, \boldsymbol{\pi}, \pi_0)$ used in this example. To solve Equation 8, observe that $x_3 = 1, x_2 =$
 286 $x_4 = 0$ is a feasible solution, but setting any two variables to one is not possible under the
 287 $3x_2 + 2x_3 + 4x_4 \leq 2$ constraint. Therefore, setting $\alpha \leftarrow 2 - 1 = 1$ retains the validity of the
 288 inequality, and we obtain the lifted inequality $x_1 + x_2 + x_3 + x_4 \leq 2$.

289 Going back to the original problem (inferring a new **Cumulative** constraint), we can
 290 summarize the reasoning above as follows: the constraint in Equation 5 implies

$$291 \quad \text{Cumulative}(\mathbf{x}, (1, 1, 1, 1), \mathbf{d}, 2).$$

292 In other words, we have discovered that *at most two tasks can run in parallel* in any feasible
 293 assignment. \lrcorner

294 Before proceeding, it is important to point out that the discovery of valid inequalities,
 295 while similar to its integer programming counterpart, has a subtle difference from it: the
 296 goal in the integer programming case is usually not to merely produce new inequalities, but
 297 to *separate* a non-integer solution \mathbf{x}^* from the linear programming relaxation by a valid
 298 inequality [17]; this, in turn, suggests starting with a cover violated by \mathbf{x}^* . In the absence
 299 of such guidance, some other strategy for enumerating covers is required. The next section
 300 describes these steps—cover generation, lifting, and constraint introduction—in more detail.

301 **4** Inferring Cumulative constraints by lifting

302 This section introduces an approach for inferring auxiliary **Cumulative** constraints from
 303 m constraints $\text{Cumulative}(\mathbf{s}, \mathbf{a}_i, \mathbf{d}, b_i)$, $i = 1, \dots, m$. To employ lifting in the absence of a
 304 guiding non-solution, I propose to:

- 305 1. Find a collection \mathcal{C} of at most N_{cover} covers for the constraints $A\mathbf{x} \leq \mathbf{b}, \mathbf{x} \in \{0, 1\}^n$.
- 306 2. For each cover $C \in \mathcal{C}$, lift the cover inequality $\sum_{c \in C} x_c \leq |C| - 1$ with respect to $\mathcal{P}(A, \mathbf{b})$
 307 as described in Subsection 3.2.
- 308 3. Among the lifted **Cumulative** constraints, add a limited number N_{out} of them, giving
 309 preference to the most restrictive constraints.

310 To rank constraints by restrictiveness in Steps 1 and 3, I use the following quality metric:

311 **► Definition 10.** *Given a constraint $\text{Cumulative}(\mathbf{s}, \boldsymbol{\pi}, \mathbf{d}, \pi_0)$, its **capacity bound** is the*
 312 *ratio of total resource usage over the time horizon to its capacity per time unit: $\mathcal{L}(\mathbf{a}, \mathbf{d}, b) :=$*
 313
$$\frac{\sum_i d_i \pi_i}{\pi_0}.$$

314 The relevance of this quantity, as pointed out in the early RCPSP work [34, p. 255] [18,
 315 Section 2], is that it gives a lower bound on the duration spanned by the tasks:

316 **► Lemma 11.** *For any satisfying assignment of a constraint $\text{Cumulative}(\mathbf{s}, \boldsymbol{\pi}, \mathbf{d}, \pi_0)$, let*
 317 $\text{EST} := \min_i s_i$ *and* $\text{LFT} := \max_i (s_i + d_i)$. *Then* $\text{LFT} - \text{EST} \geq \mathcal{L}(\boldsymbol{\pi}, \mathbf{d}, \pi_0)$.

318 The complete procedure is given by Algorithm 1 (cover enumeration) and Algorithm 2
 319 (lifting and constraint selection). Aside from the main workflow—find a cover, lift it, score
 320 the resulting constraint—the implementation employs two additional optimizations. First,
 321 it discards any valid inequality $(\boldsymbol{\pi}, \pi_0)$ dominated by a model constraint, in the sense that
 322 $\boldsymbol{\pi} \leq \mathbf{a}_i$ and $\pi_0 \geq b_i$ for some i . Second, after lifting a cover C to a constraint with support
 323 $C' \supseteq C$, all subsets of C' of cardinality $|C|$ are marked as covers not requiring further lifting;
 324 the following example illustrates why this is beneficial.

325 ▶ **Example 12.** Suppose that Algorithm 2 has lifted a cover $C_0 = \{1, 2\}$ to a constraint

$$326 \quad \underbrace{\langle x_1 \mid d_1 \rangle + \langle x_2 \mid d_2 \rangle}_{\text{cover inequality}} + \underbrace{\langle x_3 \mid d_3 \rangle + \cdots + \langle x_{100} \mid d_{100} \rangle}_{\text{lifted terms}} \leq 1. \quad (9)$$

327 In other words, lifting has discovered that no two tasks among the set $C' = \{1, 2, \dots, 100\}$
 328 can run in parallel, which translates to an update $Q \leftarrow Q \cup (C', 2)$.

329 Now, consider later iterations of the loop over covers C that visit other binary covers
 330 within C' ; since any two tasks form a cover, this happens for $\binom{100}{2} - 1 = 4949$ iterations unless
 331 some task pairs were discarded before the lifting. Observe that the condition for skipping
 332 holds for *any two tasks* in C' since $(C', 2) \in Q$, $C' \supseteq C$, and $|C| \geq 2$; that means that none
 333 of those iterations invokes the v^* function, and skipping ≈ 5000 redundant calls to v^* is a
 334 performance gain, because (a) solving lifting subproblems is the bottleneck of Algorithm 2,
 335 and (b) lifting on C' is likely to re-discover Equation 9 instead of a new constraint; for
 336 example, it is guaranteed to do so if it lifts the variables in C' before others.

Algorithm 1 Enumeration of covers.

Input: Polyhedron $\mathcal{P}(A, \mathbf{b})$, durations \mathbf{d} , cap N_{cover} .

Output: A set of covers \mathcal{C} .

C1. [Short covers.] For each resource r , scan all task pairs $\{i, j\}$. If $a_{ri} + a_{rj} > b_r$, add $\{i, j\}$ as a binary cover. Otherwise, if another task k satisfies $a_{rk} > b_r - (a_{ri} + a_{rj})$, add the cover $\{i, j, k^*\}$, where k^* is a task with the longest duration among the ones satisfying $a_{rk} > b_r - (a_{ri} + a_{rj})$.

C2. [Filter.] Retain the N_{cover} short covers with the highest capacity bound $\mathcal{L}(\cdot, \mathbf{d}, \cdot)$, discarding the rest.

C3. [Long covers.] For each resource r and each consumption value v appearing on \mathbf{a}_r , let $B := \{i : a_{ri} = v\}$ and let k be the smallest integer with $kv > b_r$. If $|B| \geq k$, add as covers the k longest and the k shortest tasks in B .

337

Algorithm 2 Lifting-based inference of auxiliary Cumulative constraints.

Input: Polyhedron $\mathcal{P}(A, \mathbf{b})$ on n variables, durations \mathbf{d} , start-time variables \mathbf{s} .

Input: Covers \mathcal{C} , cap N_{out} , lifting subproblem call budget N_{calls} .

Output: A set of at most N_{out} valid Cumulative constraints.

L1. [Initialize.] Set $Q \leftarrow \emptyset$ (the set of known cover supersets) and $O \leftarrow \emptyset$ (the output constraints). Set the remaining call budget $\rho \leftarrow N_{\text{calls}}$.

L2. [Process next cover.] Take the next cover $C \in \mathcal{C}$. If $\exists (C', k) \in Q$ with $C' \supseteq C$ and $k \leq |C|$, skip to **L2**. Otherwise, set $\pi_i \leftarrow \llbracket i \in C \rrbracket$ for $i = 1, \dots, n$, $\pi_0 \leftarrow |C| - 1$, and $I \leftarrow \{1, \dots, n\} \setminus C$.

L3. [Lift next variable.] If $I = \emptyset$ or $\rho = 0$, go to **L4**. Otherwise choose $i \leftarrow \arg \min_{i \in I} d_i$, solve the lifting subproblem $v^*(\boldsymbol{\pi}, \pi_0, C, i; A, \mathbf{b})$, set $\pi_i \leftarrow \pi_0 - v^*(\boldsymbol{\pi}, \pi_0, C, i; A, \mathbf{b})$, update $I \leftarrow I \setminus \{i\}$, $C \leftarrow C \cup \{i\}$, and $\rho \leftarrow \rho - 1$. Repeat.

L4. [Record.] Update $Q \leftarrow Q \cup (\{i : \pi_i = 1\}, |C_0|)$ where C_0 is the original cover. If no existing constraint r satisfies $\mathbf{a}_r \geq \boldsymbol{\pi}$ and $b_r \leq \pi_0$, add $\text{Cumulative}(\mathbf{s}, \boldsymbol{\pi}, \mathbf{d}, \pi_0)$ to O . If \mathcal{C} is not exhausted, go to **L2**.

L5. [Select.] Return the N_{out} constraints in O with the highest value of $\mathcal{L}(\cdot, \mathbf{d}, \cdot)$.

338

339 **5 Experimental evaluation**

340 In this section, I evaluate the performance of the presented approach on a collection of
 341 research benchmarks. The implementation of the approach, together with the infrastructure
 342 for running the experiments and processing their results, is available in the supplementary
 343 material. The modeling part is based on the CPMpy library [12]; all lifting subproblems
 344 (Equation 4) are solved with Gurobi 12 [13]. I ran the experiments on DelftBlue [9], with
 345 each run of an instance being allocated a single core of an Intel Xeon E5-6248R 24C 3.0GHz
 346 processor and 4000 MB of RAM with a time limit of one hour.

347 I use the benchmark collection previously used by Sidorov et al. [32] to evaluate the UNL
 348 approach for the same two problem models (RCPSP and RCPSP/max) with Cumulative
 349 constraints and difference constraints of the form $x_j - x_i \geq \gamma_{ij}$ for RCPSP/max and
 350 $x_j - x_i \geq d_i$ for RCPSP, both having the latest completion time of all tasks (makespan) as
 351 the minimization objective. More specifically, I use 736 RCPSP instances from MiniZinc
 352 benchmarks [36, 25, 3, 6, 20, 19] and 349 RCPSP/max benchmarks distributed by PSPLIB [19]
 353 through C, D, UBO, and SM suites. All data files are available in the supplementary materials.

354 Given a minimization objective \mathcal{O} , each of the solvers is run with one of the two search
 355 directions. In the **primal search**,³ the solver generates a series of problems with an extra
 356 assumption $[\mathcal{O} \leq o_n]$ for decreasing values of o_n . In the **dual search**,⁴ the solver generates
 357 a series of problems with an extra assumption $[\mathcal{O} \leq o_n]$ for increasing values of o_n .

358 Each of the runs involves one of the following CP solvers: Pumpkin⁵ [11] and CP-SAT [28].
 359 The choice of the solvers is dictated by two factors: (a) both of the mentioned codes are
 360 advanced, state-of-the-art solvers performing well at the recent editions of MiniZinc Challenge,
 361 and (b) both of them support both the primal and the dual search directions. In the baseline
 362 runs, the solver is run on the input CSP as is, whereas the runs with the lifting first execute
 363 Algorithm 1 with $N_{\text{cover}} = 100$, then Algorithm 2 with $N_{\text{out}} = 5$, and proceed with the solver
 364 execution on the augmented CSP in the remaining time.

365 All CP solver configurations use default presolve settings: CP-SAT performs structural
 366 simplifications without inferring new cumulative constraints, and its lifting-based cut gen-
 367 eration operates on energy overflows in the linear programming relaxation rather than on
 368 occupancy vectors. Pumpkin performs no presolve. Therefore, I do not expect overlap with
 369 the inferences produced by the presented approach.

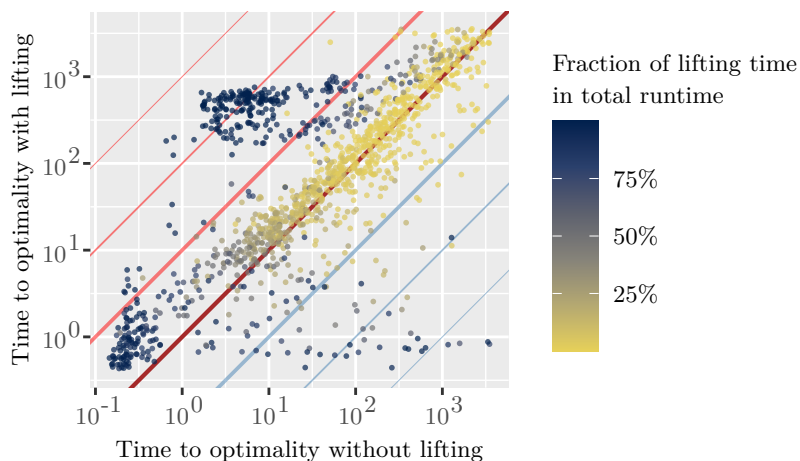
370 I budget the preprocessing by limiting the number of candidate covers considered (N_{cover})
 371 and the number of added inferred constraints (N_{out}), rather than by imposing a wall-time
 372 cutoff, since wall-time on shared HPC systems is not deterministic. All solvers were executed
 373 with the free search option.

374 To evaluate the solver performance when it did not conclude optimality, I measure the
 375 rate of progress of a solver towards the best-known bound with the following metrics. First,
 376 if $M(t)$ is the lowest makespan discovered at time $t \in [0, T]$, and M^* is the lowest discovered
 377 makespan for this problem, then the **primal integral** is $\int_0^T \frac{M(t) - M^*}{M(t)}$ and measures how
 378 fast the solver progresses towards good solutions [5]. Conversely, if $B(t)$ is the highest lower
 379 bound discovered at time $t \in [0, T]$, and B^* is the highest discovered lower bound for this
 380 problem, the **dual integral** is defined as $\int_0^T \frac{B^* - B(t)}{B^*}$.

³ Also known as the linear SAT-UNSAT search.

⁴ Also known as the linear UNSAT-SAT search or destructive lower bound.

⁵ The version distributed in supplementary materials outputs intermediate lower bounds.



■ **Figure 1** Impact of using lifting on time to optimality. Every point corresponds to a pair of runs on the same instances, using the same underlying solver with the same search direction. Diagonal lines are evenly spaced and correspond to a tenfold relative change between the durations.

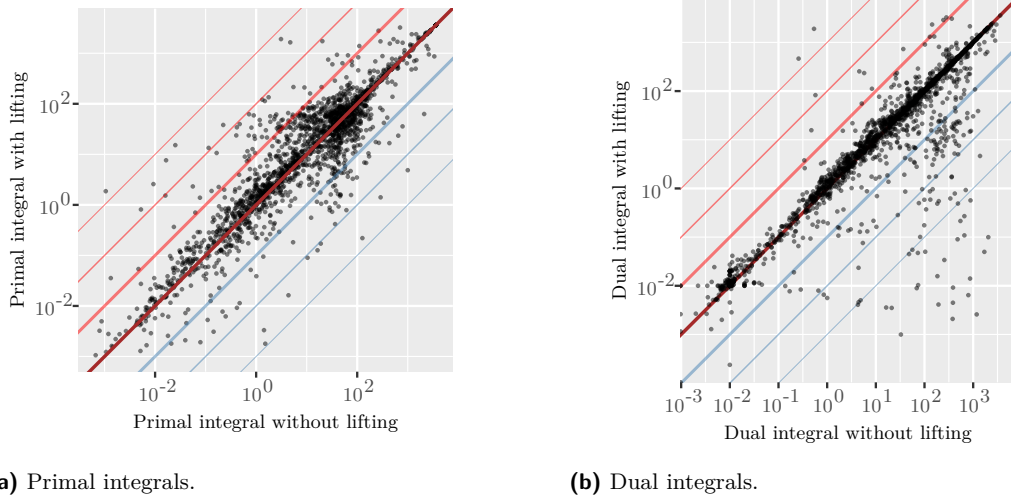
381 5.1 Evaluation against direct CP solver application

382 First, I report the evaluation of the lifting approach for the instances where both the run
 383 performing lifting and its counterpart that does not lift (while having the same settings
 384 otherwise) have proven optimality. As shown on Figure 1, lifting can cut down the solving
 385 time by large factors, sometimes from taking dozens of minutes to sub-second times. On
 386 the other hand, most of the instances with large (tenfold or more) slowdowns correspond
 387 to the cases where the lifting time dominates the solving time; more specifically, those are
 388 UBO500 and UBO1000 instances⁶ for which the number of lifting subproblems exceeds
 389 10^4 . Appendix C evaluates for these families a variant of the lifting configuration with an
 390 additional termination criterion of 2×10^4 lifting subproblems, showing that it eliminates
 391 the majority of preprocessing-dominated slowdowns at the cost of occasionally weaker lifted
 392 constraints.

393 However, most of the instances were not solved to optimality with either approach, which
 394 means that making progress with further comparisons is more viable with primal and dual
 395 integrals as metrics. I proceed with examining this comparison on Figure 2b, which supports
 396 the thesis that lifting is a helpful preprocessing step for well-suited instances: it reduces the
 397 search effort on many instances, while typically not incurring prohibitive search overhead
 398 if it backfires. On the flip side (Figure 2a), this is not true for primal integrals: in other
 399 words, adding new constraints does not help discover good solutions as consistently, but can
 400 considerably hinder this process.

401 These results already show an improvement over the UNL, since Sidorov et al. mention a
 402 much larger variance for dual integrals, especially on the degradation side. For example, they
 403 report that their approach degrades the search by at least an order of magnitude on many
 404 instances, whereas in the current evaluation, this is an *exceptionally* rare event happening
 405 only in 19 pairs of runs. I investigate this in more detail in Table 1, which shows that
 406 observing the degradation by at least a factor of 2 or larger is around the bottom 5% of

⁶ All instances from other collections take fewer than 2×10^4 lifting subproblems to terminate.



■ **Figure 2** Comparison of integrals between the baseline and the lifting approach.

■ **Table 1** Percentiles of the observed distribution of the ratio of dual integrals between baseline runs and the corresponding runs with lifting across problem types; values larger than one indicate an advantage for lifting.


Problem type	2.5%	5%	25%	50%	75%	95%	97.5%
RCPSP	0.33	0.48	0.85	1.00	1.05	10.7	47.4
RCPSP/max	0.41	0.50	0.77	0.97	1.41	51.0	292

407 observed integral ratios, whereas the top 5% of the runs correspond to $10\times$ improvements
 408 for RCPSP benchmarks and to $50\times$ improvements for RCPSP/max benchmarks.

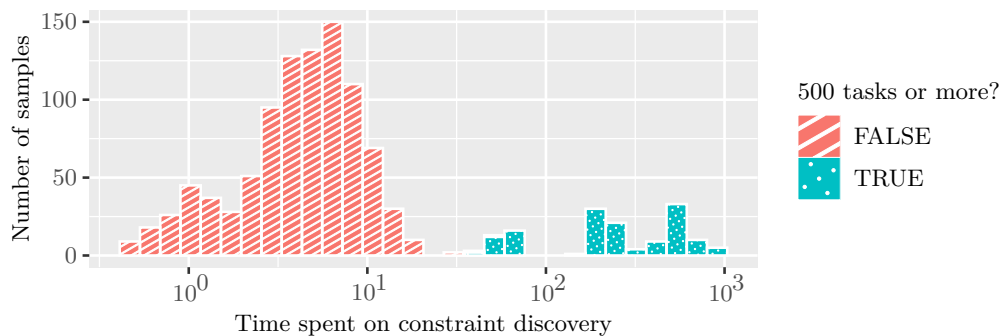
409 I also observe that in many high-performing runs of the lifting approach, the underlying
 410 constraints – or, at least, the one yielding the highest capacity bound as introduced in
 411 Definition 10 – commonly happened to be a **Disjunctive** constraint, that is, to have the
 412 right-hand side equal to one. In a similar vein, I investigate how the influence of lifting on the
 413 search depends on the constraint with the highest capacity bound, which is summarized in
 414 Figure 3. Indeed, the lifting runs that ended up on a **Disjunctive** constraint as the best one
 415 are responsible for the biggest relative gains; however, the runs that landed on a **Cumulative**
 416 constraint still routinely produce large relative improvements, even if more modest. That
 417 said, discovering general **Cumulative** constraints is only helpful when the capacity of the
 418 discovered constraint is not large; conversely, in the runs where the best constraint inferred
 419 through lifting had a capacity larger than eight units, the CP solvers have not been able to
 420 convert it into a significantly faster search.

421 I present the runtime distribution of the preprocessing phase of the lifting approach (that
 422 is, Algorithm 1 and Algorithm 2 but not the CP solver) in Figure 4. As can be seen on the
 423 histogram, under the budget imposed on lifting in this evaluation, all the lifting-specific work
 424 takes a few seconds in the majority of instances; on the other hand, the runs where lifting
 425 took more than a minute are covered exclusively by two collections with more than 500 tasks
 426 per instance, that is, UBO500 and UBO1000.

427 Last, I report the updated state-of-the-art bounds for these benchmarks discovered with
 428 the lifting approach in Appendix A, together with the criterion I used to choose which bounds



■ **Figure 3** Ratios of dual integrals between baseline runs and the corresponding runs with lifting across the right-hand side values of the best inferred constraint; values larger than one indicate an advantage for lifting.



■ **Figure 4** Distribution of time required to run the inference of `Cumulative` constraints.


429 to report. While many of the new bounds have been discovered during search, some are
 430 computed by taking the largest capacity bound among lifted constraints, thereby making
 431 them verifiable by inspecting the constraint (rather than retracing the reasoning made by a
 432 solver).⁷ While many such search-less bounds are encoded by `Disjunctive` constraints (and
 433 therefore could, in principle, have also been proven by UNL without any search), there is a
 434 bound for instance #6 from the UBO1000 collection that improves upon the previously best
 435 known bound and is produced by a `Cumulative` with the capacity of three units. I also note
 436 that no less than twelve of the instances in `Pack` and `Pack-d` collections [6, 20] can be closed
 437 immediately by using one of the lifted cumulative constraints, with the capacities varying
 438 between one and three.

439 5.2 Evaluation against UNL

440 I now proceed with the direct comparison with UNL; to this end, I run the following two
 441 solver configurations on each of the benchmark instances:

- 442 ■ UNL with dual search in the same configuration as in Sidorov et al.
- 443 ■ Pumpkin with dual search, prefaced with lifting in the same configuration as in the

⁷ See also the supplementary material for solver-independent verification of the search-less bounds.



(a) Time to optimality for the instances solved to optimality with both approaches.

(b) Dual integrals for all evaluated instances.

■ **Figure 5** Comparison between the lifting approach and UNL.

444 previous subsection, *but only considering covers with two elements*. This additional
 445 restriction is introduced to ensure that the lifting does not gain any advantage from
 446 discovering non-Disjunctive constraints.

447 As can be seen from Figure 5, the proposed approach outperforms UNL with respect to
 448 time to optimality on the majority of the applicable instances (Figure 5a) and has varying
 449 performance with respect to the dual integral (Figure 5b). This is explained by the fact that
 450 the lifting approach explicitly introduces Disjunctive constraints into the model, which are
 451 then used in specialized—and highly optimized—propagation algorithms. In contrast, UNL
 452 maintains all the facts around the disjointness of task pairs in a single graph and applies
 453 pruning for some of the cliques of that graph.

454 These plots also suggest that the lifting approach does not fully subsume UNL, as for some
 455 instances, lifting yields a notably worse performance. Interestingly, while many instances
 456 favorable for UNL are adequately explained by reasons unrelated to the methodology of
 457 this paper (such as the implementation of the preprocessing stage or discrepancies between
 458 the underlying solver versions), this is not true in all cases. For example, instance #40
 459 from the UBO50 collection is solved to optimality for ≈ 30 seconds with UNL and for ≈ 8
 460 minutes with lifting, with UNL also doing $90\times$ fewer constraint propagations until optimality.
 461 The reason for this is that the lifting approach chooses N_{out} constraints with high capacity
 462 bounds (which in this case is the sum of durations) that also happen to share many variables,
 463 which translates into an insufficient extraction of the disjointness information; this reasoning
 464 is made more specific in Appendix B.

465 6 Related work

466 The need to reason simultaneously about multiple resources has been recognized across
 467 optimization communities throughout the history of earlier developments. For instance, early
 468 exact approaches for RCPSp consistently employed global lower-bounding techniques that
 469 used the entire problem instance structure, such as the lower bounds mentioned by Klein
 470 and Scholl [18] based on solving relaxations to node packing (LB4) and parallel machine

471 scheduling (LB5) problems. The presented approach can be seen as a generalization of these
472 techniques in the following sense: Algorithm 2 with an appropriate choice of the input covers
473 and an ordering of lifted variables can yield the same bounds as LB4 or LB5 implementations
474 by letting a CP solver propagate the lifted `Cumulative` constraints. However, unlike using
475 LB4 or LB5 directly, the new `Cumulative` remains known to a CP solver as one of the
476 constraints rather than as an exogenous relaxation step, with any propagation algorithms for
477 `Cumulative` applicable to it.

478 Within the CP community, the limitations of the single-resource viewpoint were also
479 acknowledged for a long time. Beldiceanu and Carlsson [4] introduced the `Cumulatives`
480 constraint to capture multiple resources in a single constraint, which was later improved by
481 faster filtering algorithms [22, 23]. While these works highlighted the need for multi-resource
482 interaction, their focus was not on discovering new inference strategies but on running
483 existing strategies faster, by replacing what they call the *ping-pong* effect, where propagation
484 for one `Cumulative` triggers a propagation for another, with a single pass accounting for all
485 `Cumulative` constraints. Conversely, the lifting approach can not only draw conclusions from
486 several `Cumulative` constraints at once but can also draw conclusions unreachable (without
487 an exponential blowup) by reasoning over them in isolation.

488 Other efforts to mitigate the weaknesses of the single-resource viewpoint have focused
489 on extracting stronger inferences from individual constraints by deriving non-dominated
490 implied constraints. Carlier and Néron [7] investigated the systematic derivation of such
491 constraints for a single `Cumulative` by searching for one-to-one mappings from original
492 resource consumptions to a reduced set of demands that preserve infeasibility while maximizing
493 pruning. This logic was further refined by Baptiste and Bonifas [1] by examining the dual
494 linear program of a preemptive relaxation of `Cumulative` and introducing a mechanism for
495 assigning a higher consumption level to a fixed number of tasks within a demand group. While
496 these techniques strengthen the reasoning of a single resource, they remain limited to intra-
497 resource analysis. This work instead uses lifting as a tool for *inter*-resource reasoning; rather
498 than refining a single `Cumulative` in isolation, the presented work derives global inequalities
499 that capture the collective contention of tasks across the entire problem structure.

500 The earlier work by Sidorov et al. [33] implements a whole-problem reasoning scheme
501 without merging constraints, opting instead to discover and communicate disjointness between
502 tasks across resources. While UNL can theoretically emulate the node-packing bound (LB4)
503 of Klein and Scholl, it lacks the expressive power to capture variable non-unit demands (LB5).
504 Also, it relies on managing auxiliary variables during search, which can be both a strength
505 (by discovering a disjointness conditional to a search state) and a hindrance (if there is no
506 useful disjointness structure). Conversely, the presented work compiles these interactions
507 into conventional `Cumulative` constraints before initiating any search by a CP solver.

508 The core mechanism that supports the contributions of this paper—lifting—originated in
509 the integer programming community as a method for tightening cutting planes for knapsack
510 constraints. First introduced by Padberg [27], the technique iteratively strengthens an
511 inequality by introducing variables not yet in the support. Later refinements, such as the
512 ones proposed by Zemel [40], demonstrated that all sequential lifting coefficients could be
513 computed efficiently via dynamic programming. Currently, lifting is a standard part of the
514 modern MILP solver practice, as exemplified by the SCIP optimization suite [15], which
515 extensively employs lifting procedures within its knapsack constraint handler.

516 In the context of scheduling, polyhedral methods have been originally applied to single-
517 machine problems [29, 10]. For the general RCPSP, early polyhedral approaches [26] focused
518 on conflict variables (indicating a precedence within a pair of tasks) rather than start

519 times, limiting their direct applicability in standard CP models. More recently, approaches
 520 combining CP propagators with LP relaxations have been explored, with CP-SAT [28], one of
 521 the state-of-the-art CP solvers, integrating scheduling cuts as special separation procedures
 522 for tightening LP relaxations. While scheduling cuts are constructed with a procedure
 523 similar to the one proposed here (choose a cover and lift it), those cuts are derived from
 524 *energy*⁸ overflows; on the other hand, lifting as presented in this paper operates directly on
 525 the *demands* via occupancy vectors, allowing us to view **Cumulative** constraints as linear
 526 inequalities and apply standard lifting machinery directly to the start-time representation.

527 Collectively, these works support the thesis that the single-resource viewpoint is insufficient
 528 for complex scheduling benchmarks. While integer programming and specialized branch-
 529 and-bound approaches have long exploited whole-problem structures, CP has historically
 530 struggled to integrate these insights without sacrificing the modularity of propagators. By
 531 establishing a formal link between **Cumulative** constraints and polyhedral geometry, this
 532 work offers a pragmatic path to bring these powerful global inferences into CP solvers.

533 **7** Conclusions

534 This paper presents a novel approach for aggregating **Cumulative** constraints by (i) reformulating
 535 them as linear constraints in the occupancy-vector representation, (ii) discovering
 536 sets of tasks that cannot be run jointly, and (iii) lifting the linear constraint blocking this
 537 group of tasks. The experimental evaluation shows that this approach can, at its best, be
 538 substantially helpful—or, at its worst, not as obstructive as UNL—for proving objective
 539 bounds, although those advantages do not translate as well to discovering new high-quality
 540 solutions.

541 One possible direction for future work is a tighter coupling with the search loop, for
 542 instance, by adding **Cumulative** constraints not *before* the search but *during* the search. The
 543 core technical problem here is that a **Cumulative** lifted from a conflicting assignment does
 544 not necessarily propagate after backtracking, as observed in earlier work on linear constraint
 545 learning in CP [24]. However, exploiting this procedure in a more heuristic way (e.g., during
 546 restarts) could be helpful to leverage the information learned during search.

547 Last, the running assumption in this work is that both durations and resource consump-
 548 tions are constant. While this is valid for RCPSP and RCPSP/max, this is not true for
 549 several other scheduling problem formulations studied previously, most notably *multi-mode*
 550 scheduling problems [14]. However, extending this approach to multi-mode RCPSP—where
 551 task durations and demands are variable—will require handling the non-linearities resulting
 552 from the products of occupancy vectors and resource consumptions.

553 ——— References ———

- 554 1 Philippe Baptiste and Nicolas Bonifas. Redundant cumulative constraints to compute pree-
 555 mptive bounds. *Discrete applied mathematics*, 234:168–177, 2018. doi:10.1016/j.dam.2017.
 556 05.001.
- 557 2 Philippe Baptiste, Claude Le Pape, and Wim Nuijten. Satisfiability tests and time-bound
 558 adjustments for cumulative scheduling problems. *Annals of operations research*, 92:305–333,
 559 1999. doi:10.1023/a:1018995000688.

⁸ The product of duration and resource consumption.

- 560 3 Philippe Baptiste and Claude Le Pape. Constraint propagation and decomposition techniques
561 for highly disjunctive and highly cumulative project scheduling problems. *Constraints: an*
562 *international journal*, 5(1/2):119–139, 2000. doi:10.1023/a:1009822502231.
- 563 4 Nicolas Beldiceanu and Mats Carlsson. A new multi-resource cumulatives constraint with
564 negative heights. In *Principles and practice of constraint programming - CP 2002*, volume
565 2470 of *Lecture Notes in Computer Science*, pages 63–79, Berlin, Heidelberg, 2002. Springer
566 Berlin Heidelberg. doi:10.1007/3-540-46135-3_5.
- 567 5 Timo Berthold. Measuring the impact of primal heuristics. *Operations research letters*,
568 41(6):611–614, 2013. doi:10.1016/j.orl.2013.08.007.
- 569 6 Jacques Carlier and Emmanuel Néron. On linear lower bounds for the resource constrained
570 project scheduling problem. *European journal of operational research*, 149(2):314–324, 2003.
571 doi:10.1016/S0377-2217(02)00763-4.
- 572 7 Jacques Carlier and Emmanuel Néron. Computing redundant resources for the resource
573 constrained project scheduling problem. *European journal of operational research*, 176(3):1452–
574 1463, 2007. doi:10.1016/j.ejor.2005.09.034.
- 575 8 Edward W Davis. Resource allocation in project network models — a survey. *Journal of*
576 *Industrial Engineering*, 17(4):177–188, 1966. ISSN: 1045-0165.
- 577 9 Delft High Performance Computing Centre. DelftBlue, 2024.
- 578 10 Martin E Dyer and Laurence A Wolsey. Formulating the single machine sequencing problem
579 with release dates as a mixed integer program. *Discrete Applied Mathematics*, 26(2-3):255–270,
580 1990. doi:10.1016/0166-218x(90)90104-k.
- 581 11 Maarten Flippo, Konstantin Sidorov, Imko Marijnissen, Jeff Smits, and Emir Demirović.
582 A multi-stage proof logging framework to certify the correctness of CP solvers. In *30th*
583 *international conference on principles and practice of constraint programming (CP 2024)*,
584 pages 11:1–11:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2024. doi:10.4230/
585 LIPICS.CP.2024.11.
- 586 12 Tias Guns. Increasing modeling language convenience with a universal n-dimensional array,
587 CPython as python-embedded example. In *Proceedings of the 18th workshop on constraint*
588 *modelling and reformulation at CP (ModRef 2019)*, volume 19, 2019.
- 589 13 Gurobi Optimization, LLC. *Gurobi optimizer reference manual*, 2024.
- 590 14 Sönke Hartmann and Andreas Drexl. Project scheduling with multiple modes: A comparison
591 of exact algorithms. *Networks. An International Journal*, 32(4):283–297, 1998. doi:10.1002/
592 (sici)1097-0037(199812)32:4<283::aid-net5>3.0.co;2-i.
- 593 15 Christopher Hojny, Mathieu Besançon, Ksenia Bestuzheva, Sander Borst, João Dionísio,
594 Johannes Ehls, Leon Eifler, Mohammed Ghannam, Ambros Gleixner, Adrian Göß, Alexander
595 Hoen, Jacob von Holly-Ponientzietz, Rolf van der Hulst, Dominik Kamp, Thorsten Koch,
596 Kevin Kofler, Jurgen Lentz, Marco Lübbecke, Stephen J Maher, Paul Matti Meinhold, Gioni
597 Mexi, Til Mohr, Erik Mühmer, Krunal Kishor Patel, Marc E Pfetsch, Sebastian Pokutta,
598 Chantal Reinartz Groba, Felipe Serrano, Yuji Shinano, Mark Turner, Stefan Vigerske, Matthias
599 Walter, Dieter Weninger, and Liding Xu. The SCIP optimization suite 10.0. *arXiv [math.OC]*,
600 2025. doi:10.48550/arXiv.2511.18580.
- 601 16 Christopher Hojny, Tristan Gally, Oliver Habeck, Hendrik Lüthen, Frederic Matter, Marc E
602 Pfetsch, and Andreas Schmitt. Knapsack polytopes: a survey. *Annals of operations research*,
603 292(1):469–517, 2020. doi:10.1007/s10479-019-03380-2.
- 604 17 Konstantinos Kaparis and Adam N Letchford. Separation algorithms for 0-1 knapsack polytopes.
605 *Mathematical programming*, 124(1-2):69–91, 2010. doi:10.1007/s10107-010-0359-5.
- 606 18 Robert Klein and Armin Scholl. Computing lower bounds by destructive improvement: An
607 application to resource-constrained project scheduling. *European journal of operational research*,
608 112(2):322–346, 1999. doi:10.1016/S0377-2217(97)00442-6.
- 609 19 Rainer Kolisch and Arno Sprecher. PSPLIB - a project scheduling problem library: OR
610 Software - ORSEP Operations Research Software Exchange Program. *European Journal of*
611 *Operational Research*, 96(1):205–216, 1997. doi:10.1016/S0377-2217(96)00170-1.

- 612 20 Oumar Koné, Christian Artigues, Pierre Lopez, and Marcel Mongeau. Event-based MILP
613 models for resource-constrained project scheduling problems. *Computers & operations research*,
614 38(1):3–13, 2011. doi:10.1016/j.cor.2009.12.011.
- 615 21 Philippe Laborie, Jérôme Rogerie, Paul Shaw, and Petr Vilím. IBM ILOG CP optimizer
616 for scheduling. *Constraints: An International Journal*, 23(2):210–250, 2018. doi:10.1007/
617 s10601-018-9281-x.
- 618 22 Arnaud Letort, Mats Carlsson, and Nicolas Beldiceanu. A synchronized sweep algorithm
619 for the k-dimensional cumulative constraint. In *Integration of AI and OR techniques in
620 constraint programming for combinatorial optimization problems*, volume 7874 of *Lecture Notes
621 in Computer Science*, pages 144–159, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg.
622 doi:10.1007/978-3-642-38171-3_10.
- 623 23 Arnaud Letort, Mats Carlsson, and Nicolas Beldiceanu. Synchronized sweep algorithms for
624 scalable scheduling constraints. *Constraints: An International Journal*, 20(2):183–234, 2015.
625 doi:10.1007/s10601-014-9172-8.
- 626 24 Robert Nieuwenhuis. The IntSat method for integer linear programming. In *Principles and
627 practice of constraint programming*, volume 8656 of *Lecture notes in computer science*, pages 574–
628 589, Cham, 2014. Springer International Publishing. doi:10.1007/978-3-319-10428-7_42.
- 629 25 Ramón Alvarez-Valdés Olaguibel and José Manuel Tamarit Goerlich. Heuristic algorithms for
630 resource-constrained project scheduling: A review and an empirical analysis. In *Advances in
631 project scheduling*, volume 9 of *Studies in Production and Engineering Economics*, chapter 5,
632 pages 113–134. Elsevier, North Holland, 1989. doi:10.1016/b978-0-444-87358-3.50009-2.
- 633 26 Ramón Alvarez-Valdés Olaguibel and Josémanuel Tamarit Goerlich. The project scheduling
634 polyhedron: Dimension, facets and lifting theorems. *European Journal of Operational Research*,
635 67(2):204–220, 1993. doi:10.1016/0377-2217(93)90062-r.
- 636 27 Manfred W Padberg. Technical note—a note on zero-one programming. *Operations research*,
637 23(4):833–837, 1975. doi:10.1287/opre.23.4.833.
- 638 28 Laurent Perron and Frédéric Didier. CP-SAT. URL: [https://developers.google.com/
639 optimization/cp/cp_solver/](https://developers.google.com/optimization/cp/cp_solver/).
- 640 29 Maurice Queyranne. Structure of a simple scheduling polyhedron. *Mathematical Programming*,
641 58(1-3):263–285, 1993. doi:10.1007/bf01581271.
- 642 30 Francesca Rossi, Peter van Beek, and Toby Walsh, editors. *Handbook of constraint programming*.
643 Foundations of artificial intelligence. Elsevier Science, London, England, 2006. doi:10.1016/
644 s1574-6526(06)x8001-x.
- 645 31 Andreas Schutt, Thibaut Feydy, Peter J Stuckey, and Mark G Wallace. Why cumulative
646 decomposition is not as bad as it sounds. In *Principles and practice of constraint programming
647 - CP 2009*, Lecture Notes in Computer Science, pages 746–761, Berlin, Heidelberg, 2009.
648 Springer Berlin Heidelberg. doi:10.1007/978-3-642-04244-7_58.
- 649 32 Konstantin Sidorov, Imko Marijnissen, and Emir Demirović. Data collection for “Unite and
650 lead: Finding disjunctive cliques for scheduling problems”, June 2025. doi:10.5281/zenodo.
651 15624416.
- 652 33 Konstantin Sidorov, Imko Marijnissen, and Emir Demirović. Unite and lead: Finding disjunct-
653 ive cliques for scheduling problems. In Maria Garcia de la Banda, editor, *31st international
654 conference on principles and practice of constraint programming (CP 2025)*, volume 340 of *Leib-
655 niz International Proceedings in Informatics (LIPIcs)*, pages 35:1–35:24, Dagstuhl, Germany,
656 2025. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.CP.2025.35.
- 657 34 Joel P Stinson, Edward W Davis, and Basheer M Khumawala. Multiple resource-constrained
658 scheduling using branch and bound. *AIIE Transactions*, 10(3):252–259, 1978. doi:10.1080/
659 0569557808975212.
- 660 35 Peter J Stuckey. RCPSP. <https://people.eng.unimelb.edu.au/pstuckey/rcpsp/>. Accessed:
661 2025-3-29.
- 662 36 Peter J Stuckey, Thibaut Feydy, Andreas Schutt, Guido Tack, and Julien Fischer. The MiniZinc
663 Challenge 2008–2013. *AI magazine*, 35(2):55–60, 2014. doi:10.1609/aimag.v35i2.2539.

- 664 37 Petr Vilím. $O(n \log n)$ filtering algorithms for unary resource constraint. In *Integration of*
 665 *AI and OR techniques in constraint programming for combinatorial optimization problems*,
 666 volume 3011 of *Lecture Notes in Computer Science*, pages 335–347, Berlin, Heidelberg, 2004.
 667 Springer Berlin Heidelberg. doi:10.1007/978-3-540-24664-0_23.
- 668 38 Petr Vilím. Timetable edge finding filtering algorithm for discrete cumulative resources. In
 669 *Integration of AI and OR techniques in constraint programming for combinatorial optimization*
 670 *problems*, volume 6697 of *Lecture Notes in Computer Science*, pages 230–245, Berlin, Heidelberg,
 671 2011. Springer-Verlag. doi:10.1007/978-3-642-21311-3_22.
- 672 39 Petr Vilím, Philippe Laborie, and Paul Shaw. Failure-directed search for constraint-based
 673 scheduling. In *Integration of AI and OR techniques in constraint programming*, volume 9075
 674 of *Lecture Notes in Computer Science*, pages 437–453, Barcelona, Spain, 2015. Springer Cham.
 675 doi:10.1007/978-3-319-18008-3_30.
- 676 40 Eitan Zemel. Easily computable facets of the knapsack polytope. *Mathematics of Operations*
 677 *Research*, 14(4):760–764, 1989. doi:10.1287/moor.14.4.760.

678 **A Novel bounds**

679 I report the bounds discovered by the lifting approach if they are *both* better than the
 680 previously reported bounds and are not directly reproducible without lifting. More precisely,
 681 I report bounds that are simultaneously (a) tighter than the bounds reported in the previous
 682 sources known to us that used the same benchmarks [19, 39, 38, 35, 33], and (b) either
 683 tighter than any bound derived without preprocessing or matches it but was derived at least
 684 ten times faster than with any other approach.

685 Novel upper bounds (makespans) are reported in Table 2, and novel lower bounds are
 686 reported in Table 3; the same data is available as supplementary materials. All bounds
 687 are reported for RCPSP/max benchmarks; all reported durations are in MM:SS format. To
 688 indicate the remaining optimality gap, I also state the best known lower bound in Table 2
 689 and the best known makespan in Table 3; in either case, that bound is the tightest among
 690 the previously reported values and the values discovered without preprocessing. Additionally,
 691 Table 4 reports the lower bounds that can be verified by an appropriate lifted constraint.

692 As noted in the main text, using **Disjunctive** constraints, when possible, typically yields
 693 the largest gains, compared to non-unit **Cumulative** constraints. A natural hypothesis that
 694 can be proposed from this is that the new results should be credited to the discovery of a
 695 non-trivial collection of **Disjunctive** constraints, rather than to the lifting approach for
 696 deriving general **Cumulative** constraints. To address this, I rerun the reported instances with
 697 the lifting approach in the same configurations as in Subsection 5.1, but with a restriction
 698 on using only binary covers (as in Subsection 5.2), thereby ensuring that the lifting approach
 699 only adds **Disjunctive** constraints. The instances where this **Disjunctive**-restricted lifting
 700 has failed to discover the same objective values and lower bounds are highlighted in Table 2
 701 and Table 3 respectively. As can be seen, in approximately half of the record-setting instances
 702 (eleven out of 23), it was essential to add **Cumulative** constraints with non-unit capacity to
 703 derive novel lower bounds; in a similar vein, only one of the five new best solutions has been
 704 discovered without adding non-unit **Cumulative** constraints.

705 **B Disjointness structure of instance #40 from UBO50**

706 As mentioned in the main text, UNL outperforms the lifting approach on this instance;
 707 this appendix provides a more specific explanation for this. Figure 6 demonstrates the
 708 complement of the disjointness graph of that instance, with vertices corresponding to each of

■ **Table 2** Novel upper bounds derived with lifting. Highlighted rows report solutions that were not reproduced using only `Disjunctive` constraints.

Collection	#	Ref. objective	New objective	Time	Best bound
C	63	366	363	43:50	347
C	67	350	349	23:56	346
UBO100	8	385	383	41:56	376
UBO100	32	434	432	26:59	414
UBO200	4	893	838	29:25	605

709 the fifty tasks that are scheduled in this instance, and edges connect pairs $\{u, v\}$ of vertices
 710 that can be executed in parallel (that is, $a_{i,u} + a_{i,v} \leq b_i$ for all i). Vertices are grouped as
 711 follows:

712 **Universal** These are the `blue` vertices with solid fill that correspond to all tasks shared
 713 between **all** constraints added by lifting.

714 **Occasional** These are the `orange` vertices with hatching that correspond to tasks that are
 715 mentioned in **some but not all** constraints added by lifting.


716 **Void** These are the `green` vertices with dotted fill that correspond to tasks that are mentioned
 717 in **none** of the lifted constraints.

718 The important insight in the structure of this graph is that it is (a) very sparse, with
 719 a random pair of tasks having 87% probability of being disjoint, and (b) the void vertices
 720 are loosely constrained by the rest of the graph; to justify the latter, I observe that 235 out
 721 of 262 maximal independent sets of this graph involve a void vertex. This suggests that
 722 UNL is able to exploit all of the graph structure, which opens the inferences about the tasks
 723 corresponding to the void vertices that are not available to the CP solver after the lifting
 724 procedure concludes.


725 On the other hand, this shows the insufficiency of the selection step based solely on a
 726 quality metric of the `Cumulative` constraint. In this example, the lifting procedure chooses
 727 many independent sets (equivalently, `Disjunctive` constraints) that contain all the universal
 728 vertices and have a large total duration of tasks. Taken on a per-constraint basis, this is a
 729 good decision, as this set covers 40% of the problem variables; unfortunately, none of the
 730 void vertices can be added to this set. I believe that resolving the issue demonstrated by this
 731 instance would require a more elaborate selection procedure that rewards the *overall* coverage
 732 of the space of known conflicts, rather than the quality of each of the added constraints in
 733 isolation.

734 **C** Effect of a call budget on large RCPSP/max instances

735 The experimental evaluation in Section 5 uses two budget parameters to control preprocessing
 736 effort: the number of candidate covers $N_{\text{cover}} = 100$ and the number of added constraints
 737 $N_{\text{out}} = 5$. As noted in Section 5, instances from the UBO500 and UBO1000 collections
 738 generate significantly more lifting subproblems than other families under this configuration,
 739 with some instances exceeding 10^5 subproblems. This appendix evaluates the effect of
 740 an additional termination criterion that caps the total number of lifting subproblems at
 741 2×10^4 , applied to all UBO500 and UBO1000 instances with all other parameters unchanged.
 742 From now on, the configuration with this termination criterion is referred to as *budgeted*



■ **Figure 6** The complement of the disjointness graph of instance #40 from UBO50; vertex color encodes the number of lifted constraints which mention the corresponding task, and the area occupied by a vertex is proportional to the task duration.



(a) Comparison with unbudgeted configuration.

(b) Comparison with budgeted configuration.

■ **Figure 7** Impact of using various lifting configurations on time to optimality. Every point corresponds to a pair of runs on the same instances from UBO500 or UBO1000, using the same underlying solver with the same search direction. Diagonal lines are evenly spaced and correspond to a tenfold relative change between the durations.

743 configuration, and the lifting configuration used in the main text is referred to as *unbudgeted*
744 configuration.

745 First, I report the comparison of the time to optimality with baseline, budgeted, and
746 unbudgeted configurations. As can be seen from Figure 7, budgeting reduces the pathological
747 behavior pointed out in a similar comparison in the main text, as the instances where lifting
748 dominates the runtime now take less time to terminate (because the lifting procedure takes
749 less time). As for the rest of the instances, their performance profile has not substantially
750 changed; the reason for that is that the most advantageous instances for lifting are the ones
751 that were not solved to optimality, which are not reported in this comparison.

752 Budgeting better controls the preprocessing *overhead* on these instances; however, it
753 also does not fully overturn the *gains* obtained from lifting. Table 5 reports the quantiles
754 of the dual integral ratios; as can be seen, the higher end of the distribution became less
755 pronounced, but the rest of the quantiles (until 95%) have not meaningfully changed.

1:22 On Inferring Cumulative Constraints

756 Table 6 reports the instances from Table 3 where the budgeted and unbounded configura-
757 tions produce different bounds after accounting for the search-less bounds from Table 4. Of
758 the 25 novel lower bounds reported in the main text, only three are affected: the budgeted
759 configuration discovers a search bound of 1 260 for instance #4 of UBO500, exceeding the
760 search-less bound of 1 259 that was the best result under the unbudgeted configuration,
761 while the bounds for UBO500 #6 and UBO1000 #68 regress. The remaining bounds are
762 unaffected, confirming that the call budget has a negligible impact on the reported results.

763 One search-less lower bound (#68 from UBO1000) from Table 4 regresses under the
764 budgeted configuration from 2 536 to 2 519, confirming that the call budget occasionally
765 terminates lifting before a stronger constraint is discovered. This also impacts the eventual
766 search performance: the budgeted configuration terminates at the bound 2 537, falling short
767 of the 2 550 lower bound discovered by the configuration from the main text.

768 Together, these results substantiate the claim in Subsection 5.1 that a call budget
769 eliminates the majority of preprocessing-dominated slowdowns at the cost of occasionally
770 weaker lifted constraints.

■ **Table 3** Novel lower bounds derived with lifting. Highlighted rows report bounds that were not reproduced using only **Disjunctive** constraints.

Collection	#	Ref. bound	New bound	Time	Best objective
C	63	343	348	43:08	363
C	66	340	347	49:16	368
UBO100	70	375	387	43:43	408
UBO200	2	682	731	46:25	813
UBO200	3	482	542	30:47	906
UBO200	5	499	559	57:28	767
UBO200	6	538	569	38:29	765
UBO200	8	505	583	48:53	911
UBO200	32	753	794	51:45	863
UBO200	33	784	785	56:14	834
UBO200	34	595	673	35:16	774
UBO200	65	728	747	35:22	814
UBO200	70	724	823	57:43	877
UBO500	3	1159	1260	43:30	1808
UBO500	6	1202	1350	01:05	2035
UBO500	8	1175	1208	00:56	2212
UBO500	38	1258	1398	01:12	2575
UBO1000	2	2321	2518	32:11	5479
UBO1000	6	2311	2329	04:12	3508
UBO1000	9	2274	2414	03:55	4231
UBO1000	10	2367	2441	02:30	3702
UBO1000	35	2382	2563	03:09	4737
UBO1000	68	2478	2550	03:12	4688

1:24 On Inferring Cumulative Constraints

■ **Table 4** Novel *search-less* lower bounds derived with lifting. The bounds marked with an asterisk have not been improved by a CP solver.

Collection	#	Ref. bound	New bound	Objective	Capacity
UBO200	3	482	531	906	1
UBO200	4	514	583*	838	1
UBO200	5	499	533	767	1
UBO200	6	538	558	765	1
UBO200	8	505	559	911	1
UBO500	3	1159	1260*	1808	1
UBO500	4	1230	1259	2774	1
UBO500	6	1202	1345	2035	1
UBO500	8	1175	1328*	2212	1
UBO500	38	1258	1448*	2575	1
UBO1000	2	2321	2515	5479	1
UBO1000	6	2311	2315	3508	3
UBO1000	9	2274	2418*	4231	1
UBO1000	10	2367	2482*	3702	1
UBO1000	35	2382	2568*	4737	1
UBO1000	68	2478	2536	4688	1

■ **Table 5** Percentiles of the observed distribution of the ratio of dual integrals between baseline runs and the corresponding runs with lifting; values larger than one indicate an advantage for lifting.

Lifting configuration	2.5%	5%	25%	50%	75%	95%	97.5%
Unbudgeted	0.40	0.51	0.80	0.98	1.46	33.2	117
Budgeted	0.38	0.54	0.84	0.98	1.20	30.0	79.3

■ **Table 6** Changes of novel lower bounds discovered by budgeted runs in comparison with unbudgeted runs.

Collection	#	Unbudgeted	Budgeted	Change direction
UBO500	4	1259	1260	↗
UBO500	6	1350	1316	↘
UBO1000	68	2550	2537	↘