

Domain-Independent Dynamic Programming with Constraint Propagation

Imko Marijnissen¹, J. Christopher Beck², Emir Demirović¹, Ryo Kuroiwa^{3,4}

¹Delft University of Technology, The Netherlands

²University of Toronto, Canada

³National Institute of Informatics, Japan

⁴The Graduate University for Advanced Studies, SOKENDAI, Japan

i.c.w.m.marijnissen@tudelft.nl, jcb@mie.utoronto.ca, e.demirovic@tudelft.nl, kuroiwa@nii.ac.jp

Abstract

There are two prevalent model-based paradigms for combinatorial problems: 1) state-based representations, such as heuristic search, dynamic programming (DP), and decision diagrams, and 2) constraint and domain-based representations, such as constraint programming (CP), (mixed-)integer programming, and Boolean satisfiability. In this paper, we bridge the gap between the DP and CP paradigms by integrating constraint propagation into DP, enabling a DP solver to prune states and transitions using constraint propagation. To this end, we implement constraint propagation using a general-purpose CP solver in the Domain-Independent Dynamic Programming framework and evaluate using heuristic search on three combinatorial optimisation problems: Single Machine Scheduling with Time Windows ($1|r_i, \delta_i|\sum w_i T_i$), the Resource-Constrained Project Scheduling Problem (RCPSP), and the Travelling Salesperson Problem with Time Windows (TSPTW). Our evaluation shows that constraint propagation significantly reduces the number of state expansions, causing our approach to solve more instances than a DP solver for $1|r_i, \delta_i|\sum w_i T_i$ and RCPSP, and showing similar improvements for tightly constrained TSPTW instances. The runtime performance indicates that the benefits of propagation outweigh the overhead for constrained instances, but that further work into reducing propagation overhead could improve performance further. Our work is a key step in understanding the value of constraint propagation in DP solvers, providing a model-based approach to integrating DP and CP.

Code — <https://doi.org/10.5281/zenodo.19051392>

Experiments — <https://doi.org/10.5281/zenodo.19051248>

Appendices — <https://arxiv.org/abs/2603.16648>

1 Introduction

Combinatorial optimisation is key in Artificial Intelligence (AI), encompassing problems such as planning, scheduling, and variants of satisfiability. Within this field, Dynamic Programming (DP) and Constraint Programming (CP) provide declarative model-based frameworks such as Domain-Independent Dynamic Programming (Kuroiwa and Beck 2026) and MiniZinc (Nethercote et al. 2007; Stuckey et al.

2014) that enable users to formulate representations of problems that can then be solved by generic solvers.

While solving the same problem, the techniques used by solvers following these paradigms differ. DP solvers are based on state-based problem representations, which facilitate duplicate/dominance detection techniques. While certain DP solvers, such as ddo (Gillard, Schaus, and Coppé 2020) and CODD (Michel and van Hoeve 2024) use decision diagram branch-and-bound (Bergman et al. 2016), in this work, we focus on DP solvers using heuristic search by Kuroiwa and Beck (2026). In contrast, CP solvers focus on using depth-first search and inference techniques to prune the search space by determining which values *cannot* be assigned to variables in a solution.

Previous works introduce constraint propagation in DP approaches, but they either only used problem-specific propagation (Fontaine, Dibangoye, and Solnon 2023; Brita, van der Linden, and Demirović 2025), or did not use heuristic search (Tanaka and Fujikuma 2012). Other, related but orthogonal, work has focused on using dominance and caching (Chu and Stuckey 2015; Chu, de la Banda, and Stuckey 2010; Smith 2005) in a CP solver. To the best of our knowledge, no work has been done on integrating generic constraint propagation in a model-based DP framework.

To address this gap, we create a framework for integrating constraint propagation from CP with DP modelling and implement our approach using general-purpose CP and DP solvers, providing:

1. A dual view of the problem, using a state-based DP view and an integer-based CP view, which provides:
 - Dominance and duplicate detection, and heuristic search using the DP view of the problem.
 - Strong inference techniques for constraints using the CP view of the problem, allowing pruning of states and strengthening of the dual bound.
2. A generic model-based way of integrating heuristic search with constraint propagation, allowing *any* form of constraint propagation rather than *only* relying on domain-specific inference techniques (though these techniques are supported by our framework as well).

We evaluate our framework on three combinatorial problems: Single Machine Scheduling with Time Windows ($1|r_i, \delta_i|\sum w_i T_i$), the Resource-Constrained Project

Scheduling Problem (RCPSp), and the Travelling Salesperson Problem with Time Windows (TSPTW). For these problems, we show that constraint propagation significantly reduces the number of state expansions, enabling our approach to solve more instances than a DP solver alone for $1|r_i, \delta_i| \sum w_i T_i$ and RCPSp, and tightly constrained TSPTW instances. The performance in terms of runtime indicates that the benefits of propagation outweigh the overhead for tightly constrained instances, but that further work into reducing the propagation time is needed. In our in-depth analysis, we vary $1|r_i, \delta_i| \sum w_i T_i$ and TSPTW instance constrainedness, demonstrating that propagation performs exceptionally well when instances are highly constrained.

This work provides a key step in ascertaining the usefulness of constraint propagation in DP solvers using heuristic search, allowing us to better understand the complementary strengths and weaknesses of the approaches.

2 Preliminaries

We introduce the necessary concepts related to dynamic programming, the Domain-Independent Dynamic Programming (DIDP) framework and its Rust interface, heuristic search in DIDP, and constraint programming.

2.1 Domain-Independent Dynamic Programming

Dynamic Programming (DP) is a problem-solving methodology based on state-based representations. *Domain-Independent Dynamic Programming (DIDP)* is a model-based framework for solving combinatorial optimisation problems formulated as DP models. One of its strengths is that DIDP models can be solved using generic solvers, removing the need to implement specialised algorithms for each DP model. It has recently seen success on combinatorial problems, such as single machine scheduling and TSPTW (Kuroiwa and Beck 2026).

A DIDP model is a state-transition system where a solution is created by starting with a *target state* \mathcal{S}^0 and applying a sequence of transitions until a *base state* is reached. A *successor state* $\mathcal{S}_{[\tau]}$ is created from \mathcal{S} by applying a possible transition $\tau \in T(\mathcal{S})$ with weight $w_\tau(\mathcal{S})$. The cost of a solution is the sum of transition weights plus the cost of the base state, where $V(\mathcal{S})$ maps \mathcal{S} to its optimal solution cost and the *dual bound* is a lower (upper) bound on $V(\mathcal{S})$ in minimisation (maximisation) problems. Thus, the goal of a DIDP solver is to compute $V(\mathcal{S}^0)$.

While the DIDP framework allows the user to solve generic DP models, it is limited in its expressivity due to the restrictions of its declarative modelling language, Dynamic Programming Problem Description Language (DyPDL). The *Rust Programmable Interface for DIDP (RPID)* defines a DIDP model in terms of Rust functions, allowing the user to define more complex dual bounds, state constraints and/or state variables, while also providing a significant speedup compared to DIDP (Kuroiwa and Beck 2025).

A RPID model uses three interfaces (called *traits* in Rust) to describe a DP model: 1) **Dp** - `TargetState()` defines the target state, `IsBaseState(\mathcal{S})` determines when \mathcal{S} is a base state with cost `BaseCost(\mathcal{S})`, and `GenSucc(\mathcal{S})` gener-

ates a set of successors $\mathcal{S}_{[\tau]}$ with transition weight $w_\tau(\mathcal{S})$, 2) **Dominance** - `Dominates($\mathcal{S}, \mathcal{S}'$)` determines if \mathcal{S} dominates \mathcal{S}' to prevent exploring redundant states and 3) **Bound** - `Dual(\mathcal{S})` returns the dual bound.

Using these functions, Equation (1) recursively defines the value $V(\mathcal{S})$ of a state (for a minimisation problem) as the minimum value of the successor states and the transition weight, where Inequality (2) specifies the dominance relationship and Inequality (3) the dual bound. Thus, solving Equation (1) provides the optimal solution.

$$V(\mathcal{S}) = \begin{cases} \text{BaseCost}(\mathcal{S}) & \text{if IsBaseState}(\mathcal{S}) \\ \min_{(w_\tau(\mathcal{S}), \mathcal{S}_{[\tau]}) \in \text{GenSucc}(\mathcal{S})} w_\tau(\mathcal{S}) + V(\mathcal{S}_{[\tau]}) & \text{else} \end{cases} \quad (1)$$

$$V(\mathcal{S}) \leq V(\mathcal{S}') \quad \text{if Dominates}(\mathcal{S}, \mathcal{S}') \quad (2)$$

$$V(\mathcal{S}) \geq \text{Dual}(\mathcal{S}) \quad (3)$$

Heuristic Search Heuristic search is a strategy for solving problems, using a *heuristic* $h(\mathcal{S})$ to influence which state to *expand*. In this work, we use two heuristic search algorithms implemented in RPID: A^* (Hart, Nilsson, and Raphael 1968) and Complete Anytime Beam Search (CABS) (Zhang 1998). In RPID, both A^* and CABS use the dual bound as $h(\mathcal{S})$ and the cost to a node $g(\mathcal{S})$ to guide the search, expanding the node with the minimum $f(\mathcal{S}) = g(\mathcal{S}) + h(\mathcal{S})$. However, as opposed to A^* , CABS stores at most b states with minimum $f(\mathcal{S})$ in each *layer* (storing at most $\mathcal{O}(n \times b)$ states at a time, where n is the branching factor), increasing b until termination. CABS has shown superior memory use and ability to prove optimality/infeasibility (Kuroiwa and Beck 2026) while also providing intermediate solutions.

Pseudocode for heuristic search in DIDP can be seen in Algorithm 1; it starts with the target state, (heuristically) selecting a state to expand, and *generating* its successors while avoiding dominated/redundant states, iterating this process until a termination condition is met.

Algorithm 1: Simplified Heuristic Search in RPID

```

 $\mathcal{O} \leftarrow \{\text{TargetState}()\}, \mathcal{E} \leftarrow \emptyset, \text{Primal} \leftarrow \infty$ 
while  $\neg \text{ShouldTerminate}()$  do
   $\mathcal{S} \leftarrow \arg \min_{\mathcal{S} \in \mathcal{O}} f(\mathcal{S}), \mathcal{O} \leftarrow \mathcal{O} \setminus \{\mathcal{S}\}$ 
  if IsBaseState( $\mathcal{S}$ ) then
     $\text{Primal} \leftarrow \min\{\text{Primal}, g(\mathcal{S}) + \text{BaseCost}(\mathcal{S})\}$ 
    continue
  for  $(w_\tau(\mathcal{S}), \mathcal{S}_{[\tau]}) \in \text{GenSucc}(\mathcal{S})$  do
     $h(\mathcal{S}_{[\tau]}) \leftarrow \text{Dual}(\mathcal{S}_{[\tau]})$ 
    if  $f(\mathcal{S}_{[\tau]}) \leq \text{Primal} \wedge \nexists \mathcal{S}' \in \mathcal{E} \text{ s.t.}$ 
       $\text{Dominates}(\mathcal{S}', \mathcal{S}_{[\tau]}) \wedge g(\mathcal{S}') \leq g(\mathcal{S}_{[\tau]})$  then
         $\mathcal{O} \leftarrow \text{Insert}(\mathcal{O}, \mathcal{S}_{[\tau]}), \mathcal{E} \leftarrow \mathcal{E} \cup \{\mathcal{S}_{[\tau]}\}$ 

```

2.2 Constraint Programming

Constraint Programming (CP) is a paradigm for solving combinatorial problems in the form of a *constraint satisfaction problem (CSP)*, represented by the tuple $(\mathcal{X}, \mathcal{C}, \mathcal{D})$, where \mathcal{X} is the set of *variables*, \mathcal{C} is the set of *constraints*

specifying the relations between *variables*, and \mathcal{D} is the *domain* which maps $x \in \mathcal{X}$ to its possible values $\mathcal{D}(x) \subseteq \mathcal{Z}$. The goal is to find a *solution* \mathcal{I} : a mapping of each variable $x \in \mathcal{X}$ to a value $v_x \in \mathcal{D}(x)$ which satisfies all constraints.

A CSP can be transformed into a *constraint optimisation problem (COP)* by adding an objective function which maps a solution \mathcal{I} to a value: $o: \mathcal{I} \mapsto \mathcal{Z}$. The goal is then to find a solution that minimises (or maximises) this value.

One strength of CP solvers is the use of inference techniques, which prune values from the domains of variables that cannot be part of a feasible solution. Concretely, constraints are represented in the CP solver by one or more *propagators*. Thus, a propagator can be seen as a function $p: \mathcal{D} \mapsto \mathcal{D}'$ such that $\forall x \in \mathcal{X} : \mathcal{D}(x) \supseteq \mathcal{D}'(x)$, where a propagator is at *fixed-point* if $p(\mathcal{D}) = \mathcal{D}$. We define $\text{LB}(x)$ ($\text{UB}(x)$) as the minimum (maximum) value that can be assigned to $x \in \mathcal{X}$ in \mathcal{D} . Importantly, propagators are applied independently, but multiple propagators can be used together to express complex relationships between variables.

In this work, we use two constraints and corresponding propagation algorithms, which are implemented in CP solvers such as Pumpkin (Flippo et al. 2024), Huub (Dekker et al. 2025), and CP-SAT (Perron and Didier 2025):

The *Disjunctive* constraint (Carlier 1982), states that a set of jobs should not overlap. Each job has a variable starting time and a (variable) duration. The most prevalent propagation algorithm is edge-finding (Vilím 2004), which reasons over whether scheduling a job i as early (late) as possible is not possible due to another set of jobs Ω to infer that i should be processed after (before) Ω .

The *Cumulative* constraint (Aggoun and Beldiceanu 1992), states that the cumulative resource usage of a set of tasks should not exceed the capacity at any time point. Each task has a variable starting time, a duration, and resource usage. The most common propagation algorithm is timetable filtering (Beldiceanu and Carlsson 2002), which reasons about when tasks are *guaranteed* to consume resources to infer that another task *cannot* be processed at these times.

3 Related Work

Combining constraint propagation with dynamic programming was discussed by Hooker and van Hoes (2018), and while there has been work in this area, a generic approach for incorporating constraint propagation with heuristic search in a model-based way has remained elusive.

Fontaine, Dibangoye, and Solnon (2023) apply propagation after each solution to tighten time windows when solving TSPTW using a variation of A^* (Vadlamudi et al. 2012). Additionally, pruning techniques for DP approaches play a key role in reducing the search space for optimal decision tree problems (Brita, van der Linden, and Demirović 2025). While these approaches apply propagation within DP, their approach, unlike ours, relies on problem-specific propagation in a non-model-based approach.

Generic propagation in a DP approach has been applied by Tanaka and Fujikuma (2012), making use of the *Disjunctive* for the $|\mathcal{r}_i, \delta_i| \sum w_i T_i$ problem. Their approach uses the Lagrangian relaxation to generate tight

bounds based on a network representation, pruning states by propagating the *Disjunctive* constraint. However, as opposed to our work, they do not use a model-based approach, limiting the transferability of their techniques.

From a CP perspective, the work by Chu, de la Banda, and Stuckey (2010) proposes automatic identification and exploitation of subproblem dominance/equivalence, which can be used to prune states in a CP solver. In a similar vein, Smith (2005) proposes to create a representation of a key based on the (reduced) CP domains for caching. In our work, rather than adjusting the CP solver by adding dominance/equivalence constraints, we take the natural definition of dominance/equivalence in a DP model and incorporate propagation into this framework rather than the other way around, which means that we are able to more easily integrate with heuristic search using dual bounds.

Furthermore, the work of Perron (1999) transfers heuristic search to a CP solver. However, they do not leverage the dominance/duplication detection of the DP model.

Lastly, Fox, Sadeh, and Baykan (1989) proposed a model for problem solving by combining constraint satisfaction and heuristic search. However, the main difference between our work and theirs is that we perform constraint propagation for search in a generic model-based framework.

4 DIDP with Constraint Propagation

We describe our framework for incorporating constraint propagation, and, for three problems, we describe the DP model and the incorporation of constraint propagation.

4.1 Framework

There are many possible architectures for combining solvers. Part of this space has been explored by the development of hybridisations of CP with fields such as Linear Programming (Beck and Refalo 2003; Laborie and Rogerie 2016), Satisfiability (SAT) (Ohrimenko, Stuckey, and Codish 2009; Feydy and Stuckey 2009), and Mixed-Integer Programming (Hooker 2005). Additionally, there is extensive work on hybridisation architectures in the context of SAT Modulo Theories (Nieuwenhuis, Oliveras, and Tinelli 2006). In our initial effort to develop a generic hybridisation of DP and CP, we chose to develop the simplest interface that would allow a DP solver to exploit the main technology of CP: constraint propagation.

To this end, we utilise a generic CP solver implementing constraint propagation to provide information to the DP solver. An overview of the interactions between these two components can be seen in Figure 1. At its core, the CP solver provides two pieces of information: 1) when a state




Figure 1: An overview of the interactions in our framework.

\mathcal{S} is infeasible, i.e., dynamically adding state/transition preconditions, and 2) a dual bound $\text{DualCP}(\mathcal{S})$. It should be noted that this information can be determined by *any* approach (e.g., decision diagram solvers or linear programs).

To improve usability, we *only* modify the RPID model and *not* the search, by replacing $\text{GenSucc}(\mathcal{S})$ in Algorithm 1 with $\text{GenSuccPropagation}(\mathcal{S}, \text{Primal})$ (Algorithm 2), and $\text{Dual}(\mathcal{S}_{\lceil \tau \rceil})$ with $\max\{\text{Dual}(\mathcal{S}_{\lceil \tau \rceil}), \text{DualCP}(\mathcal{S}_{\lceil \tau \rceil}, \mathcal{D}')\}$. Algorithm 2 first creates the CP model, which is propagated to obtain \mathcal{D}' and checked for infeasibility using $\text{IsInfeasible}(\mathcal{S}, \mathcal{D}')$ and $\text{DualCP}(\mathcal{S}, \mathcal{D}')$; one benefit of using CP is that \mathcal{D}' can be used to strengthen the dual bound. Next, successor infeasibility is checked in $\text{IsSuccInfeasible}(\mathcal{S}_{\lceil \tau \rceil}, \mathcal{D}')$; since \mathcal{D}' is valid for any successor $\mathcal{S}_{\lceil \tau \rceil}$, \mathcal{D}' is used in $\text{IsSuccInfeasible}(\mathcal{S}_{\lceil \tau \rceil}, \mathcal{D}')$ to avoid propagation. In practice, we only initialise \mathcal{X} and \mathcal{C} for the target state, and adjust \mathcal{D} to represent every state.

Algorithm 2: $\text{GenSuccPropagation}(\mathcal{S}, \text{Primal})$

```

 $(\mathcal{X}, \mathcal{C}, \mathcal{D}) \leftarrow \text{CPModel}(\mathcal{S})$ 
 $\mathcal{D}' \leftarrow \text{Propagate}(\mathcal{X}, \mathcal{C}, \mathcal{D})$ 
if  $\text{IsInfeasible}(\mathcal{S}, \mathcal{D}') \vee g(\mathcal{S}) + \text{DualCP}(\mathcal{S}, \mathcal{D}') \geq$ 
  Primal then
  return  $\emptyset$ 
Successors  $\leftarrow \emptyset$ 
for  $(w_\tau(\mathcal{S}), \mathcal{S}_{\lceil \tau \rceil}) \in \text{GenSucc}(\mathcal{S})$  do
  if  $\neg \text{IsSuccInfeasible}(\mathcal{S}_{\lceil \tau \rceil}, \mathcal{D}')$  then
    Successors  $\leftarrow \text{Successors} \cup \{(w_\tau(\mathcal{S}), \mathcal{S}_{\lceil \tau \rceil})\}$ 
return Successors

```

4.2 $1|r_i, \delta_i| \sum w_i T_i$

The single-machine scheduling problem $1|r_i, \delta_i| \sum w_i T_i$ is formulated as follows: given are n jobs $\mathcal{J} = \{j_0, \dots, j_{n-1}\}$ which need to be processed on a single machine without overlapping. Each job $i \in \mathcal{J}$ has a duration $p_i \in \mathcal{N}^+$, a release time $r_i \in \mathcal{N}$, a deadline $d_i \in \mathcal{N}$, a weight $w_i \in \mathcal{N}$, and a latest possible finish time $\delta_i \in \mathcal{N}$.

The goal is to assign a start time s_i to each job such that $r_i \leq s_i < s_i + p_i \leq \delta_i$, minimising the weighted tardiness $\sum_{i \in \mathcal{J}} w_i T_i$, where $T_i = \max\{0, s_i + p_i - d_i\}$.

DP Model The DP model (adapted from Abdul-Razaq, Potts, and Van Wassenhove (1990)) schedules one job at a time. We track the unscheduled jobs \mathcal{U} and current time t in $\mathcal{S} = (\mathcal{U}, t)$. Additionally, we define the next available time after scheduling i as $t'(t, i) = \max\{t, r_i\} + p_i$, the set of unscheduled jobs for which the time-window is not violated when scheduled $B = \{i \in \mathcal{U} : t'(t, i) \leq \delta_i\}$, and the successor state $\mathcal{S}_{\lceil i \rceil} = (\mathcal{U} \setminus \{i\}, t'(t, i))$.

The Bellman equation is then defined as:

$$V(\mathcal{U}, t) = \begin{cases} 0 & \text{if } \mathcal{U} = \emptyset \\ \infty & \text{if } \mathcal{U} \neq B \\ \min_{i \in \mathcal{U}} w_i \times T_i + V(\mathcal{S}_{\lceil i \rceil}) & \text{else} \end{cases} \quad (4)$$

$$V(\mathcal{U}, t^a) \leq V(\mathcal{U}, t^b) \quad \text{if } t^a \leq t^b \quad (5)$$

$$V(\mathcal{U}, t) \geq \sum_{i \in \mathcal{U}} w_i \times \max\{0, \max\{r_i, t\} + p_i - d_i\} \quad (6)$$

Equation (4) states that a state is infeasible if there are jobs that cannot be scheduled due to time windows. Inequality (5) states that state a dominates state b if state a schedules the same jobs as state b but does so in less time. Inequality (6) describes a dual bound as the sum of tardinesses if all unscheduled jobs were started at $\max\{r_i, t\}$.

Constraint Propagation The variables \mathcal{X} and domains \mathcal{D} created in $\text{CPModel}(\mathcal{S})$ are described in Definition (8); for each unscheduled job i , a CP variable for the start time is created, which starts after r_i and all scheduled jobs.

The CP constraints \mathcal{C} are then defined by Constraint (7), which uses the Disjunctive to ensure that no jobs overlap.

$$\text{Disjunctive}([s_i \mid i \in \mathcal{U}], [p_i \mid i \in \mathcal{U}]) \quad (7)$$

$$s_i \in [\max\{r_i, t\}, \delta_i - p_i] \quad \forall i \in \mathcal{U} \quad (8)$$

After $\text{Propagate}(\mathcal{X}, \mathcal{C}, \mathcal{D})$, we use the reduced domains \mathcal{D}' in $\text{DualCP}(\mathcal{S}, \mathcal{D}')$ as shown in Inequality (9), which states a dual bound as the sum of tardinesses of all unscheduled jobs if started at $\text{LB}(s_i)$. Additionally, Equation (10) describes $\text{IsSuccInfeasible}(\mathcal{S}_{\lceil \tau \rceil}, \mathcal{D}')$, where a successor is infeasible if τ cannot be scheduled at $t'(t, \tau)$.

$$V(\mathcal{U}, t) \geq \sum_{i \in \mathcal{U}} w_i \times \max\{0, \text{LB}(s_i) + p_i - d_i\} \quad (9)$$

$$V(\mathcal{S}_{\lceil \tau \rceil}) = \infty \quad \text{if } t'(t, \tau) \notin \mathcal{D}'(s_\tau) \quad (10)$$

4.3 RCPSP

The Resource-Constrained Project Scheduling Problem can be formulated as follows: given are n tasks $\mathcal{T} = \{t_0, \dots, t_{n-1}\}$ and a set of resources \mathcal{R} , where each resource $r \in \mathcal{R}$ has a capacity $C_r \in \mathcal{N}^+$. Each task has a duration $p_i \in \mathcal{N}^+$, and a resource usage $u_{ir} \in \mathcal{N}$ per resource $r \in \mathcal{R}$.

The goal is to assign a start time $s_i \in \mathcal{N}$ to each task such that 1) the cumulative resource usage never exceeds *any* resource capacity, and 2) a set of precedence constraints is respected, as shown in Inequalities (11)-(12) (where $H = \sum_{i \in \mathcal{T}} p_i$). We define the predecessors of task i as Pre_i . The objective is to minimise the makespan: $\max_{i \in \mathcal{T}} s_i + p_i$.

$$\sum_{i \in \mathcal{T} : s_i \leq t \wedge s_i + p_i \geq t} u_{ir} \leq C_r \quad \forall t \in [0, H], r \in \mathcal{R} \quad (11)$$

$$s_i + p_i \leq s_j \quad \forall (i, j) \in \mathcal{P} \quad (12)$$

DP Model We introduce a DP model based on scheduling one task at a time. We keep track of the partial schedule PS mapping each scheduled task j to its scheduled start time PS_j (and each unscheduled task to \perp), and the current time t in $\mathcal{S} = (PS, t)$. We define the set of scheduled tasks $S = \{i \mid i \in \mathcal{T} : PS_i \neq \perp\}$ (and $\mathcal{U} = \mathcal{T} \setminus S$).

Equation (13) defines $t'(t, \tau)$ as the earliest time where 1) there is no resource conflict when scheduling τ and 2) all of the predecessors of τ have finished. We then define the successor state as $\mathcal{S}_{\lceil \tau \rceil} = ((PS \setminus \{(\tau, \perp)\}) \cup \{(\tau, t'(t, \tau))\}, t'(t, \tau))$; we also define $\text{Makespan}(PS, t) = \max\{\max_{i \in S} \{PS_i + p_i\}, \max_{i \in \mathcal{U}} \{t + p_i\}\}$, and $\text{W}(\mathcal{S}, \tau) = (\text{Makespan}(\mathcal{S}_{\lceil \tau \rceil}, t'(t, \tau)) - \text{Makespan}(\mathcal{S}, t))$.

$$t'(t, \tau) = \begin{cases} \min h \in [t, H] \text{ s.t. } \nexists r \in \mathcal{R} : \\ u_{\tau r} + \sum_{i \in S : s_i \leq h \wedge s_i + p_i > h} u_{ir} > C_r \\ \wedge \text{Pre}_\tau \subseteq \{j \mid j \in S \wedge s_j + p_j \leq h\} \end{cases} \quad (13)$$

The Bellman equation is then defined as:

$$V(PS, t) = \begin{cases} 0 & \text{if } S = \mathcal{T} \\ \min_{i \in \mathcal{U}: Pre_i \subseteq S} w(S, \tau) + V(\mathcal{S}_{[\tau]}) & \text{if } S \neq \mathcal{T} \end{cases} \quad (14)$$

$$V(\mathcal{S}^a) \leq V(\mathcal{S}^b) \quad \text{if } \begin{cases} \mathcal{S}^a = \mathcal{S}^b \wedge t^a \leq t^b \wedge \\ \forall i \in S : \max(PS_j^a, PS_j^b) \\ + p_j > t \rightarrow PS_i^a \leq PS_i^b \end{cases} \quad (15)$$

$$V(\mathcal{S}_{[\tau]}) \leq V(\mathcal{S}_{[\tau']}) \quad \text{if } t'(t, \tau) + p_\tau \leq t'(t, \tau') \quad (16)$$

$$V(PS, t) \geq \text{Length}(\text{CriticalPath}(\mathcal{P}, \mathcal{U})) \quad (17)$$

$$V(PS, t) \geq \max_{r \in \mathcal{R}} \left[\frac{\sum_{i \in \mathcal{U}} u_{ir} \times p_i}{C_r} \right] \quad (18)$$

Inequality (15) describes dominance between states a and b ; if they have scheduled the same tasks but a has scheduled all tasks sooner than b , then a dominates b . Similarly, Inequality (16) describes a transition dominance using the left-shift rule (Demeulemeester and Herroelen 1992); if there are two transitions $\tau, \tau' \in \mathbb{T}(S)$, then $\mathcal{S}_{[\tau]}$ dominates $\mathcal{S}_{[\tau']}$ if we can complete task τ *before* starting τ' . Inequality (17) states that $V(S)$ is larger than the length of the longest path through the precedence graph (i.e., the critical path). Inequality (18) states an energy-based dual-bound.

Constraint Propagation The variables \mathcal{X} and domains \mathcal{D} created in $\text{CPModel}(S)$ are described in Definitions (22)-(24); for each task i , a CP variable s_i for the start time is created, and a CP variable o for the objective is created.

The CP constraints \mathcal{C} are defined as: Constraint (19) specifies the `Cumulative` which ensures the resource capacity is not exceeded (where $St = [s_i \mid i \in \mathcal{U}]$, $P = [p_i \mid i \in \mathcal{U}]$, and $U_r = [u_{ir} \mid i \in \mathcal{U}]$), Constraint (20) ensures that precedences hold, and Constraint (21) defines the makespan.

$$\text{Cumulative}(St, P, U_r, C_r) \quad \forall r \in \mathcal{R} \quad (19)$$

$$s_i + p_i \leq s_j \quad \forall (i, j) \in \mathcal{P} \quad (20)$$

$$s_i + p_i \leq o \leq \text{Primal} \quad \forall i \in \mathcal{U} \quad (21)$$

$$s_j = PS_j \quad \forall j \in S \quad (22)$$

$$s_i \in [t, H - p_i] \quad \forall i \in \mathcal{U} \quad (23)$$

$$o \in [0, H] \quad (24)$$

After `Propagate`($\mathcal{X}, \mathcal{C}, \mathcal{D}$), we use the bounds \mathcal{D}' in $\text{DualCP}(S, \mathcal{D}')$ as described in 1) Inequality (25), which specifies a lower-bound on the makespan based on the earliest completion time (Vilfm 2009), and 2) Inequality (26), which states that the makespan of a state is at least as large as the maximum latest finish time of the unscheduled tasks. Additionally, Equation (27) describes `IsSuccInfeasible`($\mathcal{S}_{[\tau]}, \mathcal{D}'$), where a successor is not feasible if task τ cannot be scheduled at time $t'(t, \tau)$.

$$V(S) \geq \max_{r \in \mathcal{R}} \max_{\Omega \subseteq \mathcal{T}} \left[\frac{C_r \min_{i \in \Omega} \text{LB}(s_i) + \sum_{i \in \Omega} (u_{ir} \times p_i)}{C_r} \right] \quad (25)$$

$$V(PS, t) \geq \max_{i \in \mathcal{U}} \{ \text{LB}(s_i) + p_i \} \quad (26)$$

$$V(\mathcal{S}_{[\tau]}) = \infty \quad \text{if } t'(t, \tau) \notin \mathcal{D}'(s_\tau) \quad (27)$$

4.4 TSPTW

The Travelling Salesperson Problem with Time Windows can be formulated as follows: given are n locations $\mathcal{C} = \{l_0, \dots, l_{n-1}\}$ which need to be visited, with a travel time $c_{ij} \in \mathcal{N}$ between locations i and j . Each location i needs to be visited within $[r_i, \delta_i]$; if the salesperson arrives at location i before r_i then they wait until time r_i . We define c_{ij}^* as the shortest travel time between location i and j .

The goal is to determine the order in which to visit the locations, starting and ending at the depot l_0 while visiting each location besides the depot exactly once, such that the total travel time is minimised (excluding waiting times).

DP Model The DP model (Kuroiwa and Beck 2026) is based on visiting one location at a time. We track the unvisited locations \mathcal{U} , current location l , and current time t in $\mathcal{S} = (\mathcal{U}, l, t)$. Additionally, we define $t'(i, j) = \max\{t + c_{ij}, r_j\}$, $\mathcal{S}_{[j]} = (\mathcal{U} \setminus \{j\}, j, t'(l, j))$, $O = \{j \mid j \in \mathcal{U} : t + c_{ij}^* > \delta_j\}$, $\text{MinTo}(l) = \min_{i \in \mathcal{C}: i \neq j} c_{il}$, $\text{MinFrom}(l) = \min_{i \in \mathcal{C}: i \neq j} c_{li}$, and $R_i = \{j \in \mathcal{U} \mid t + c_{ij} \leq \delta_j\}$.

The Bellman equation is then defined as:

$$V(S) = \begin{cases} \infty & \text{if } O \neq \emptyset \\ c_{l_0} & \text{if } \mathcal{U} = \emptyset \\ \min_{j \in R_l} c_{lj} + V(\mathcal{S}_{[j]}) & \text{else} \end{cases} \quad (28)$$

$$V(\mathcal{U}, l, t_a) \leq V(\mathcal{U}, l, t_b) \quad \text{if } t_a \leq t_b \quad (29)$$

$$V(\mathcal{U}, l, t) \geq \begin{cases} \text{MinTo}(l_0) + \sum_{i \in \mathcal{U}} \text{MinTo}(i) \\ \text{MinFrom}(l) + \sum_{i \in \mathcal{U}} \text{MinFrom}(i) \end{cases} \quad (30)$$

Inequality (29) states that state a dominates state b if 1) they have visited the same locations, 2) are at the same current location, and 3) the current time of state a is earlier than the current time of state b . Inequality (30) uses the minimum travel time *to* an unvisited location i from another location and the minimum travel time *from* an unvisited location i to another location to calculate a dual bound on \mathcal{S} .

Constraint Propagation The variables \mathcal{X} and domains \mathcal{D} created in $\text{CPModel}(S)$ are described in Definitions (33)-(35); for each unvisited location i , a CP variable s_i for the arrival time and time to next location p_i is created, where the bounds on p_i can be tightened if $\exists j \in \mathcal{U} \setminus \{i\} : \text{LB}(s_j) \geq \text{UB}(s_i)$, signifying that the depot cannot be visited from i . Additionally, a CP variable o for the objective is created.

The constraints \mathcal{C} , a relaxation of TSPTW, are defined as 1) Constraint (31) stating that no locations can be visited simultaneously, and 2) Constraint (32) defining the objective.

$$\text{Disjunctive}([s_i \mid i \in \mathcal{U} \cup \{l\}], [p_i \mid i \in \mathcal{U} \cup \{l\}]) \quad (31)$$

$$o = \sum_{i \in \mathcal{C}} p_i \leq \text{Primal} \quad (32)$$

$$s_i \in [\max\{t, r_i\}, \delta_i] \quad \forall i \in \mathcal{U} \cup \{l\} \quad (33)$$

$$p_i \in \{c_{ij} \mid j \in (\mathcal{U} \cup \{l_0\}) \setminus \{i\} \wedge c_{ij} \neq \infty\} \quad \forall i \in \mathcal{U} \cup \{l\} \quad (34)$$

$$o \in [0, \max_{i \in \mathcal{U}} \delta_i + c_{i_0}] \quad (35)$$

After `Propagate`($\mathcal{X}, \mathcal{C}, \mathcal{D}$), we use the bounds \mathcal{D}' in $\text{DualCP}(S, \mathcal{D}')$ as described in Inequality (36), which states

that a dual bound is the sum of the lower bounds of the travel times inferred by constraint propagation. Additionally, Equation (37) describes $\text{IsSuccInfeasible}(\mathcal{S}_{[\tau]}, \mathcal{D}')$, where visiting τ next is infeasible if its arrival time or the travel time from the previous location is not possible.

$$V(\mathcal{U}, l, t) \geq \sum_{i \in \mathcal{U} \cup l} \text{LB}(p_i) \quad (36)$$

$$V(\mathcal{S}_{[\tau]}) = \infty \quad \text{if } t'(l, \tau) \notin \mathcal{D}'(s_\tau) \vee c_{l\tau} \notin \mathcal{D}'(p_\tau) \quad (37)$$

5 Experimentation

Our aim is to empirically show the impact of constraint propagation on the number of state expansions and runtime when integrated with a DP solver. For all three problems, we show that constraint propagation reduces the search space. For $1|r_i, \delta_i| \sum w_i T_i$ and TSPTW, we show that it is especially effective for tightly constrained instances.


5.1 Experimentation Setup

The experiments are run single threaded on an Intel Xeon Gold 6248R 24C 3.0GHz processor (Delft High Performance Computing Centre 2024) with a limit of 30 minutes and 16GB of memory. We use the following benchmarks:


- $1|r_i, \delta_i| \sum w_i T_i$ - We generate 900 instances according to Davari et al. (2016). The instances contain 50 jobs with duration $p_i \in [1, 10]$, release date $r_i \in [0, \tau P]$ (where P is the sum of durations), deadline $d_i \in [r_i + p_i, r_i + p_i + \rho P]$, latest finish time $\delta_i \in [d_i, d_i + \phi P]$, and weight $w_i \in [1, 10]$. We generate 10 instances for each combination of $\tau \in \{0, 0.2, 0.4, 0.6, 0.8, 1\}$, $\rho \in \{0.05, 0.25, 0.5\}$, $\phi \in \{0.9, 1.05, 1.2, 1.35, 1.5\}$.
- **RCPSP** - We make use of the 480 J90 instances from PSPLIB (Kolisch and Sprecher 1997).
- **TSPTW** - We consider instances used by López-Ibáñez et al. (2013), excluding too easy/difficult sets or sets with fractional distances; resulting in 130 instances by Gendreau et al. (1998) and 50 instances by Ascheuer (1996).

We use Pumpkin (Flippo et al. 2024) for propagation, and CABS and A* of RPID 0.3.1 for the DP-based approaches. The code and models are in the supplements. We evaluate the following approaches (only considering generic model-based approaches, not problem-specific ones, to gauge the benefit of propagation in *general-purpose* DP solvers):

- **A*/CABS** - The RPID model using A* or CABS.
- **A*/CABS+CP** - The RPID model using A* or CABS *with* constraint propagation; the propagators are executed once per state. We also considered fixed-point propagation, but the computational cost was too high to be able to compete with other methods.
- **ORT** - Uses MiniZinc 2.9.2 with OR-Tools CP-SAT 9.12 as solver (using free search, causing the solver to make use of a portfolio approach); included in the comparison as a reference state-of-the-art constraint-based solver. The CP model descriptions are in Appendix B.



(a) $1|r_i, \delta_i| \sum w_i T_i$ instances solved compared to state expansions. Our approach is able to solve the most instances per expansion.



(b) $1|r_i, \delta_i| \sum w_i T_i$ instances solved compared to time. Our approach solves the most instances.

Figure 2: Number of $1|r_i, \delta_i| \sum w_i T_i$ instances solved compared to the number of state expansions and time.

Experimentation Summary Overall, constraint propagation significantly prunes the search space for all problems. For $1|r_i, \delta_i| \sum w_i T_i$, RCPSP, and tightly constrained TSPTW problems, our approach solves more instances than the DP approach while using considerably fewer states. The runtime performance indicates that further work on reducing propagation time is warranted.

5.2 $1|r_i, \delta_i| \sum w_i T_i$

Starting with $1|r_i, \delta_i| \sum w_i T_i$, Figure 2a shows that our approach solves significantly more instances using fewer states than the baseline versions, exhibiting the effectiveness of constraint propagation. This effectiveness is confirmed by Appendix A.1, which shows that our approach guides the search to better solutions in fewer expansions. While A*+CP initially solves more instances, it plateaus sooner than CABS+CP, ultimately solving fewer instances.

Looking at the number of instances solved over time, Figure 2b shows that the best performing approach is CABS+CP, which proves optimality for three more instances and infeasibility (at the target state) on thirteen more than CABS. Interestingly, CABS+CP only overtakes CABS in the number of instances solved after 500 seconds, showing the impact of propagation overhead. Moreover, the lack of instances solved by OR-Tools indicates that it is our combination of DP *with propagation* that works well.

Parameter Analysis An overview of the number of instances solved across values of ϕ (where a lower value




Figure 3: Number of $1|r_i, \delta_i| \sum w_i T_i$ instances solved over ϕ . The marked parts are instances proven infeasible and the unmarked parts are instances proven optimal. Our approach solves the most instances when they are tightly constrained.

means that the instances have tighter latest finish times), can be seen in Figure 3. It shows that CABS+CP outperforms CABS on constrained instances where ϕ is small(er), indicating that the constraint propagation prunes well when instances are tightly constrained. Furthermore, the number of solved instances decreases as ϕ increases for all approaches, but the rate of decrease differs. Specifically, when $\phi \leq 1.05$, CABS+CP outperforms CABS in terms of instances proven infeasible *and* optimal, but for instances where $\phi \geq 1.2$, CABS is equal to or outperforms CABS+CP.

5.3 RCPSP

Considering RCPSP, Figure 4a shows that constraint propagation increases the number of instances solved per state expansion considerably for CABS, exhibiting that constraint propagation is crucial when using a state-based approach for RCPSP. Appendix A.2 indicates that constraint propagation provides the most benefit when proving optimality.


Additionally, Figure 4b shows that CABS+CP solves more instances per second than CABS alone. Since A^* (+CP) rarely finds primal bounds, which are key for pruning, it falls behind CABS(+CP). OR-Tools is the best-performing solver, possibly due to its depth-first backtracking search being well-suited for the problem. Nonetheless, A^* /CABS+CP provide better primal and/or dual bounds than OR-Tools on 70 instances (as shown in Figure 11 of Appendix A.2).

5.4 TSPTW


Turning to TSPTW, Figure 5a shows that constraint propagation *decreases* the number of solved instances due to a lack of pruning. This observation is demonstrated in the number of solved instances per state expansion of both A^* / A^* +CP and CABS/CABS+CP overlapping, likely due to instances being too loosely constrained to make inferences. This observation is further corroborated by Appendix A.3.

This effect is also reflected by the number of instances solved over time in Figure 5b, which shows that adding constraint propagation results in an increase in solving time. However, the worst-performing approach is OR-Tools.

Parameter Analysis To determine the impact of instance constrainedness, we investigate parameterised in-




(a) RCPSP instances solved compared to state expansions. Our approach solves significantly more instances per state expansion.




(b) RCPSP instances solved compared to time; our approach outperforms the DP baseline, but the best-performing approach is CP.

Figure 4: Number of RCPSP instances solved compared to the number of state expansions and time.



(a) TSPTW instances solved compared to state expansions; the line for A^* +CP fully overlaps with A^* . Our approach expands slightly fewer states than the DP approach.



(b) TSPTW instances solved compared to time. The DP approach solves the most instances.

Figure 5: Number of TSPTW instances solved compared to the number of state expansions and time.




Figure 6: Number of solved TSPTW Rifki and Solnon (2025) instances across α . The marked parts are instances proven infeasible and the unmarked parts are instances proven optimal. The lower plot uses an extended set of instances to show $1.0 \leq \alpha \leq 1.5$. Our approach solves the most instances when they are tightly constrained.

stances (not part of the previous set) introduced by Rifki and Solnon (2025) where $n = 31$. For these instances, a low α indicates that many locations need to be visited relative to the horizon, while a low β means that time windows are tight. We consider two sets: 1) a general set of 135 instances where we select three instances per combination of $\alpha \in \{1.0, 1.5, 2, 2.5, 3\}$ and $\beta \in \{0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1\}$, and 2) a set of 540 instances where we select ten instances per combination of $\alpha \in \{1.0, 1.1, 1.2, 1.3, 1.4, 1.5\}$ and the same values of β .

Figure 6, shows that for $\alpha < 1.2$, CABS+CP outperforms all other approaches. Examining Figure 7, shows that, when $1.0 \leq \alpha \leq 1.5$, the number of instances solved per state expansion is significantly increased compared to Figure 5a. This effect can be explained by CABS+CP reducing the number of states by orders of magnitude compared to CABS (Figure 8), where especially infeasible instances benefit from constraint propagation. Overall, the DP-based methods solve fewer instances as α increases, but the rate of decrease differs, causing CABS to outperform CABS+CP when $\alpha \geq 1.2$. Interestingly, OR-Tools proves optimality on the most instances, likely due to 1) lower memory usage, and 2) stronger propagation (using Circuit (Laurière 1978)).

6 Conclusions and Future Work

We have presented a framework that integrates constraint propagation into a DP solver, and implemented it in the modelling of DIDP. Our approach combines the strengths of heuristic guidance and dominance/duplicate detection of DP with the inference techniques of CP. Our evaluation on three combinatorial optimisation problems shows that adding con-




Figure 7: Number of solved TSPTW Rifki and Solnon (2025) instances ($1.0 \leq \alpha \leq 1.5$) over state expansions. Our approach solves the most instances per expansion.




Figure 8: Comparison of state expansions on TSPTW Rifki and Solnon (2025) instances ($1.0 \leq \alpha \leq 1.5$). Our approach reduces the number of expansions by orders of magnitude.

straint propagation significantly increases the number of instances solved per state expansion, solving more instances than a DP solver alone for $1|r_i, \delta_i| \sum w_i T_i$, RCPSP, and tightly constrained TSPTW instances. As for runtime, the benefit of propagation outweighs the cost for tightly constrained instances, but work into reducing propagation time could improve performance further. In our in-depth analysis, we vary $1|r_i, \delta_i| \sum w_i T_i$ and TSPTW instance constrainedness, showing that propagation performs well when instances are highly constrained. Our work is a key step in understanding the value of constraint propagation in DP solvers, providing a model-based integration of DP and CP.

Our interface is generic, allowing future work to investigate the integration of DP solvers with other techniques besides CP. Another strength of our framework is the simplicity of the interface; similar to SAT Modulo Theories architecture, the interface could be enriched to investigate the impact of additional information on the DP solver. Another direction is to explore what is represented in the CP model, e.g., the DP model could be a relaxation, while the CP model contains the absent elements of the DP model. Finally, we experimented with reducing propagation overhead but obtained inconclusive results, warranting further work into, e.g., reducing redundant work introduced by jumping between states, or by determining *when* to propagate.

Acknowledgements

Imko Marijnissen is supported by the NWO/OCW, as part of the Quantum Software Consortium programme (project number 024.003.037 / 3368).

Ryo Kuroiwa is supported by JSPS KAKENHI grant number JP25K24378.

References

- Abdul-Razaq, T.; Potts, C.; and Van Wassenhove, L. 1990. A survey of algorithms for the single machine total weighted tardiness scheduling problem. *Discrete Applied Mathematics*, 26(2): 235–253.
- Aggoun, A.; and Beldiceanu, N. 1992. Extending CHIP in order to solve complex scheduling and placement problems. In Delahaye, J.; Devienne, P.; Mathieu, P.; and Yim, P., eds., *JFPL'92, 1^{ères} Journées Francophones de Programmation Logique, 25-27 Mai 1992, Lille, France*, 51.
- Ascheuer, N. 1996. *Hamiltonian Path Problems in the On-line Optimization of Flexible Manufacturing Systems*. Ph.D. thesis, Konrad-Zuse-Zentrum für Informationstechnik Berlin.
- Beck, J. C.; and Refalo, P. 2003. A Hybrid Approach to Scheduling with Earliness and Tardiness Costs. *Ann. Oper. Res.*, 118(1-4): 49–71.
- Beldiceanu, N.; and Carlsson, M. 2002. A New Multi-resource cumulatives Constraint with Negative Heights. In Hentenryck, P. V., ed., *Principles and Practice of Constraint Programming - CP 2002, 8th International Conference, CP 2002, Ithaca, NY, USA, September 9-13, 2002, Proceedings*, volume 2470 of *Lecture Notes in Computer Science*, 63–79. Springer.
- Bergman, D.; Ciré, A. A.; van Hoeve, W.; and Hooker, J. N. 2016. Discrete Optimization with Decision Diagrams. *INFORMS J. Comput.*, 28(1): 47–66.
- Brita, C. E.; van der Linden, J. G. M.; and Demirović, E. 2025. Optimal Classification Trees for Continuous Feature Data Using Dynamic Programming with Branch-and-Bound. In Walsh, T.; Shah, J.; and Kolter, Z., eds., *AAAI-25, Sponsored by the Association for the Advancement of Artificial Intelligence, February 25 - March 4, 2025, Philadelphia, PA, USA*, 11131–11139. AAAI Press.
- Carlier, J. 1982. The one-machine sequencing problem. *European Journal of Operational Research*, 11(1): 42–47. Third EURO IV Special Issue.
- Chu, G.; de la Banda, M. G.; and Stuckey, P. J. 2010. Automatically Exploiting Subproblem Equivalence in Constraint Programming. In Lodi, A.; Milano, M.; and Toth, P., eds., *Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems, 7th International Conference, CPAIOR 2010, Bologna, Italy, June 14-18, 2010. Proceedings*, volume 6140 of *Lecture Notes in Computer Science*, 71–86. Springer.
- Chu, G.; and Stuckey, P. J. 2015. Dominance breaking constraints. *Constraints An Int. J.*, 20(2): 155–182.
- Davari, M.; Demeulemeester, E.; Leus, R.; and Nobibon, F. T. 2016. Exact algorithms for single-machine scheduling with time windows and precedence constraints. *J. Sched.*, 19(3): 309–334.
- Dekker, J. J.; Ignatiev, A.; Stuckey, P. J.; and Zhong, A. Z. 2025. Towards Modern and Modular SAT for LCG. In de la Banda, M. G., ed., *31st International Conference on Principles and Practice of Constraint Programming (CP 2025)*, volume 340 of *Leibniz International Proceedings in Informatics (LIPIcs)*, 42:1–42:12. Dagstuhl, Germany: Schloss Dagstuhl – Leibniz-Zentrum für Informatik. ISBN 978-3-95977-380-5.
- Delft High Performance Computing Centre. 2024. Delft-Blue Supercomputer (Phase 2). <https://www.tudelft.nl/dhpc/ark:/44463/DelftBluePhase2>.
- Demeulemeester, E.; and Herroelen, W. 1992. A Branch-and-Bound Procedure for the Multiple Resource-Constrained Project Scheduling Problem. *Management Science*, 38(12): 1803–1818.
- Feydy, T.; and Stuckey, P. J. 2009. Lazy Clause Generation Reengineered. In Gent, I. P., ed., *Principles and Practice of Constraint Programming - CP 2009, 15th International Conference, CP 2009, Lisbon, Portugal, September 20-24, 2009, Proceedings*, volume 5732 of *Lecture Notes in Computer Science*, 352–366. Springer.
- Flippo, M.; Sidorov, K.; Marijnissen, I.; Smits, J.; and Demirović, E. 2024. A Multi-Stage Proof Logging Framework to Certify the Correctness of CP Solvers. In Shaw, P., ed., *30th International Conference on Principles and Practice of Constraint Programming (CP 2024)*, volume 307 of *Leibniz International Proceedings in Informatics (LIPIcs)*, 11:1–11:20. Dagstuhl, Germany: Schloss Dagstuhl – Leibniz-Zentrum für Informatik. ISBN 978-3-95977-336-2.
- Fontaine, R.; Dibangoye, J.; and Solnon, C. 2023. Exact and anytime approach for solving the time dependent traveling salesman problem with time windows. *Eur. J. Oper. Res.*, 311(3): 833–844.
- Fox, M. S.; Sadeh, N. M.; and Baykan, C. A. 1989. Constrained Heuristic Search. In Sridharan, N. S., ed., *Proceedings of the 11th International Joint Conference on Artificial Intelligence, Detroit, MI, USA, August 1989*, 309–315. Morgan Kaufmann.
- Gendreau, M.; Hertz, A.; Laporte, G.; and Stan, M. 1998. A Generalized Insertion Heuristic for the Traveling Salesman Problem with Time Windows. *Oper. Res.*, 46(3): 330–335.
- Gillard, X.; Schaus, P.; and Coppé, V. 2020. Ddo, a Generic and Efficient Framework for MDD-Based Optimization. In Bessiere, C., ed., *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020*, 5243–5245. ijcai.org.
- Hart, P. E.; Nilsson, N. J.; and Raphael, B. 1968. A Formal Basis for the Heuristic Determination of Minimum Cost Paths. *IEEE Trans. Syst. Sci. Cybern.*, 4(2): 100–107.
- Hooker, J. N. 2005. A Hybrid Method for the Planning and Scheduling. *Constraints An Int. J.*, 10(4): 385–401.
- Hooker, J. N.; and van Hoeve, W. J. 2018. Constraint programming and operations research. *Constraints An Int. J.*, 23(2): 172–195.

- Kolisch, R.; and Sprecher, A. 1997. PSPLIB - A project scheduling problem library: OR Software - ORSEP Operations Research Software Exchange Program. *European Journal of Operational Research*, 96(1): 205–216.
- Kuroiwa, R.; and Beck, J. C. 2025. RPID: Rust Programmable Interface for Domain-Independent Dynamic Programming. In de la Banda, M. G., ed., *31st International Conference on Principles and Practice of Constraint Programming, CP 2025, August 10-15, 2025, Glasgow, Scotland*, volume 340 of *LIPICs*, 23:1–23:21. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.
- Kuroiwa, R.; and Beck, J. C. 2026. Domain-independent dynamic programming. *Artificial Intelligence*, 354: 104506.
- Laborie, P.; and Rogerie, J. 2016. Temporal linear relaxation in IBM ILOG CP Optimizer. *J. Sched.*, 19(4): 391–400.
- Laurière, J. 1978. A Language and a Program for Stating and Solving Combinatorial Problems. *Artif. Intell.*, 10(1): 29–127.
- López-Ibáñez, M.; Blum, C.; Ohlmann, J. W.; and Thomas, B. W. 2013. The travelling salesman problem with time windows: Adapting algorithms from travel-time to makespan optimization. *Appl. Soft Comput.*, 13(9): 3806–3815.
- Michel, L.; and van Hoeve, W. 2024. CODD: A Decision Diagram-Based Solver for Combinatorial Optimization. In Endriss, U.; Melo, F. S.; Bach, K.; Diz, A. J. B.; Alonso-Moral, J. M.; Barro, S.; and Heintz, F., eds., *ECAI 2024 - 27th European Conference on Artificial Intelligence, 19-24 October 2024, Santiago de Compostela, Spain - Including 13th Conference on Prestigious Applications of Intelligent Systems (PAIS 2024)*, volume 392 of *Frontiers in Artificial Intelligence and Applications*, 4240–4247. IOS Press.
- Nethercote, N.; Stuckey, P. J.; Becket, R.; Brand, S.; Duck, G. J.; and Tack, G. 2007. MiniZinc: Towards a Standard CP Modelling Language. In Bessiere, C., ed., *Principles and Practice of Constraint Programming - CP 2007, 13th International Conference, CP 2007, Providence, RI, USA, September 23-27, 2007, Proceedings*, volume 4741 of *Lecture Notes in Computer Science*, 529–543. Springer.
- Nieuwenhuis, R.; Oliveras, A.; and Tinelli, C. 2006. Solving SAT and SAT Modulo Theories: From an abstract Davis-Putnam-Logemann-Loveland procedure to DPLL(T). *J. ACM*, 53(6): 937–977.
- Ohrimenko, O.; Stuckey, P. J.; and Codish, M. 2009. Propagation via lazy clause generation. *Constraints An Int. J.*, 14(3): 357–391.
- Perron, L. 1999. Search Procedures and Parallelism in Constraint Programming. In Jaffar, J., ed., *Principles and Practice of Constraint Programming - CP'99, 5th International Conference, Alexandria, Virginia, USA, October 11-14, 1999, Proceedings*, volume 1713 of *Lecture Notes in Computer Science*, 346–360. Springer.
- Perron, L.; and Didier, F. 2025. CP-SAT.
- Rifki, O.; and Solnon, C. 2025. On the Phase Transition of the Euclidean Travelling Salesman Problem with Time Windows. *J. Artif. Intell. Res.*, 82: 2167–2188.
- Smith, B. M. 2005. Caching Search States in Permutation Problems. In van Beek, P., ed., *Principles and Practice of Constraint Programming - CP 2005, 11th International Conference, CP 2005, Sitges, Spain, October 1-5, 2005, Proceedings*, volume 3709 of *Lecture Notes in Computer Science*, 637–651. Springer.
- Stuckey, P. J.; Feydy, T.; Schutt, A.; Tack, G.; and Fischer, J. 2014. The MiniZinc Challenge 2008-2013. *AI Mag.*, 35(2): 55–60.
- Tanaka, S.; and Fujikuma, S. 2012. A dynamic-programming-based exact algorithm for general single-machine scheduling with machine idle time. *J. Sched.*, 15(3): 347–361.
- Vadlamudi, S. G.; Gaurav, P.; Aine, S.; and Chakrabarti, P. P. 2012. Anytime Column Search. In Thielscher, M.; and Zhang, D., eds., *AI 2012: Advances in Artificial Intelligence - 25th Australasian Joint Conference, Sydney, Australia, December 4-7, 2012. Proceedings*, volume 7691 of *Lecture Notes in Computer Science*, 254–265. Springer.
- Vilím, P. 2004. O(n log n) Filtering Algorithms for Unary Resource Constraint. In Régim, J.; and Rueher, M., eds., *Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems, First International Conference, CPAIOR 2004, Nice, France, April 20-22, 2004, Proceedings*, volume 3011 of *Lecture Notes in Computer Science*, 335–347. Springer.
- Vilím, P. 2009. Edge Finding Filtering Algorithm for Discrete Cumulative Resources in $O(kn \log n)$. In Gent, I. P., ed., *Principles and Practice of Constraint Programming - CP 2009, 15th International Conference, CP 2009, Lisbon, Portugal, September 20-24, 2009, Proceedings*, volume 5732 of *Lecture Notes in Computer Science*, 802–816. Springer.
- Zhang, W. 1998. Complete Anytime Beam Search. In Mostow, J.; and Rich, C., eds., *Proceedings of the Fifteenth National Conference on Artificial Intelligence and Tenth Innovative Applications of Artificial Intelligence Conference, AAAI 98, IAAI 98, July 26-30, 1998, Madison, Wisconsin, USA*, 425–430. AAAI Press / The MIT Press.