

"Boosting Decision Diagram-Based Branch-and-Bound by Pre-Solving with Aggregate Dynamic Programming"

Coppé, Vianney ; Gillard, Xavier ; Schaus, Pierre

ABSTRACT

Discrete optimization problems expressible as dynamic programs can be solved by branch-and-bound with decision diagrams. This approach dynamically compiles bounded-width decision diagrams to derive both lower and upper bounds on unexplored parts of the search space, until they are all enumerated or discarded. Assuming a minimization problem, relaxed decision diagrams provide lower bounds through state merging while restricted decision diagrams obtain upper bounds by excluding states to limit their size. As the selection of states to merge or delete is done locally, it is very myopic to the global problem structure. In this paper, we propose a novel way to proceed that is based on pre-solving a so-called aggregate version of the problem with a limited number of states. The compiled decision diagram of this aggregate problem is tractable and can fit in memory. It can then be exploited by the original branch-and-bound to generate additional pruning and guide the compilation of restricted decision diagrams toward good solutions. The results of the numerical study we conducted on three combinatorial optimization problems show a clear improvement in the performance of DD-based solvers when blended with the proposed techniques. These results also suggest an approach where the aggregate dynamic programming model could be used in replacement of the relaxed decision diagrams altogether.

CITE THIS VERSION

Coppé, Vianney ; Gillard, Xavier ; Schaus, Pierre. *Boosting Decision Diagram-Based Branch-and-Bound by Pre-Solving with Aggregate Dynamic Programming*. The 29th International Conference on Principles and Practice of Constraint Programming (Toronto, Canada, du 27/08/2023 au 31/08/2023). <http://hdl.handle.net/2078.1/278691>

Le dépôt institutionnel DIAL est destiné au dépôt et à la diffusion de documents scientifiques émanant des membres de l'UCLouvain. Toute utilisation de ce document à des fins lucratives ou commerciales est strictement interdite. L'utilisateur s'engage à respecter les droits d'auteur liés à ce document, principalement le droit à l'intégrité de l'œuvre et le droit à la paternité. La politique complète de copyright est disponible sur la page [Copyright policy](#)

DIAL is an institutional repository for the deposit and dissemination of scientific documents from UCLouvain members. Usage of this document for profit or commercial purposes is strictly prohibited. User agrees to respect copyright about this document, mainly text integrity and source mention. Full content of copyright policy is available at [Copyright policy](#)

1 Boosting Decision Diagram-Based 2 Branch-and-Bound by Pre-Solving with Aggregate 3 Dynamic Programming

4 Vianney Coppé ✉ 

5 UCLouvain, Louvain-la-Neuve, Belgium

6 Xavier Gillard ✉ 

7 UCLouvain, Louvain-la-Neuve, Belgium

8 Pierre Schaus ✉ 

9 UCLouvain, Louvain-la-Neuve, Belgium

10 — Abstract —

11 Discrete optimization problems expressible as dynamic programs can be solved by branch-and-bound
12 with decision diagrams. This approach dynamically compiles bounded-width decision diagrams to
13 derive both lower and upper bounds on unexplored parts of the search space, until they are all
14 enumerated or discarded. Assuming a minimization problem, relaxed decision diagrams provide
15 lower bounds through state merging while restricted decision diagrams obtain upper bounds by
16 excluding states to limit their size. As the selection of states to merge or delete is done locally, it
17 is very myopic to the global problem structure. In this paper, we propose a novel way to proceed
18 that is based on pre-solving a so-called aggregate version of the problem with a limited number of
19 states. The compiled decision diagram of this aggregate problem is tractable and can fit in memory.
20 It can then be exploited by the original branch-and-bound to generate additional pruning and guide
21 the compilation of restricted decision diagrams toward good solutions. The results of the numerical
22 study we conducted on three combinatorial optimization problems show a clear improvement in the
23 performance of DD-based solvers when blended with the proposed techniques. These results also
24 suggest an approach where the aggregate dynamic programming model could be used in replacement
25 of the relaxed decision diagrams altogether.

26 **2012 ACM Subject Classification** Mathematics of computing → Combinatorial optimization

27 **Keywords and phrases** Discrete Optimization, Decision Diagrams, Aggregate Dynamic Programming

28 **Digital Object Identifier** 10.4230/LIPIcs.CP.2023.22

29 **1** Introduction

30 On top of their use for Boolean encodings [27], formal verification [25], model checking
31 [15], computer-aided design [29] and much more, *decision diagrams* (DDs) have recently
32 emerged as a tool for discrete optimization. They provide a compact way to encode a set
33 of solutions to a problem. Still, for large problems, DDs representing the whole solution
34 space – called *exact* DDs – can quickly become intractable to compute. Two variants of
35 DDs can be used instead: *restricted* [10] and *relaxed* [1, 8] DDs that respectively encode a
36 subset and superset of the set of solutions. When compiled based on a *dynamic programming*
37 (DP) model, these approximate DDs allow to compute bounds on the objective function for
38 any subproblem while controlling the size of the DD compiled. Restricted DDs aim to find
39 good admissible solutions by iteratively extending a bounded set of promising candidates
40 while dropping others, in a beam search fashion. On the other hand, relaxed DDs rely
41 on a problem-dependent state merging scheme to maintain an acceptable DD size while
42 preserving all solutions of the problem. In [9], Bergman et al. presented a *branch-and-bound*
43 algorithm solely based on these two ingredients, thus introducing a new general-purpose
44 discrete optimization framework and solver.



© Vianney Coppé, Xavier Gillard and Pierre Schaus;
licensed under Creative Commons License CC-BY 4.0

29th International Conference on Principles and Practice of Constraint Programming (CP 2023).

Editor: Roland H. C. Yap; Article No. 22; pp. 22:1–22:17

Leibniz International Proceedings in Informatics



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

45 In addition to exploiting the compactness of DP models, the main novelty of this
46 approach is its unique way of deriving lower and upper bounds. In the last few years, some
47 algorithmic improvements have been suggested to further strengthen these bounds. Assuming
48 a minimization problem, Gillard et al. [19] showed how user-defined lower bound formulas
49 can be integrated to prune DDs during their compilation and thus concentrate the search
50 on promising parts of the search space. They also proposed a way to compute tighter lower
51 bounds for all nodes contained in a relaxed DD through *local bounds*. Rudich et al. [30]
52 introduced a *peeling* operator that splits a relaxed DD in two: one part containing all paths
53 traversing a selected exact node and the other containing all remaining paths. It allows both
54 to warm-start the compilation of subsequent relaxed DDs and to strengthen the bounds of
55 the nodes inside the relaxed DD on which the peeling has been performed. More recently, [16]
56 generalized the ideas of [19] and introduced the use of a cache storing new thresholds that
57 further enhance the pruning power of the solver. Other factors impacting the quality of the
58 bounds provided by relaxed DDs have been studied, including variable orderings [7, 11, 26]
59 and alternative compilation schemes [24]. Yet, all these approaches rely on a problem-specific
60 state merging operator at the heart of the relaxation, which does not yield tight relaxations
61 for all problems, as our computational experiments show.

62 After covering the necessary background about DD-based optimization, this paper presents
63 an alternate relaxation scheme for deriving good bounds by incorporating ideas from *aggregate*
64 *dynamic programming* [2, 3] to the DD-based discrete optimization framework. The underlying
65 idea of the approach is to deduce information about an original problem instance by creating
66 and solving an aggregate – relaxed – version of it. This is achieved by *aggregating* the
67 states of the DP model as to obtain a much smaller DP state space. If this aggregation is
68 adequately specified, one can compute a lower bound for any original subproblem by finding
69 the optimal solution of its aggregate version. Furthermore, this optimal aggregate solution
70 can be *disaggregated* and transposed in the original problem to find good heuristic solutions.
71 In practice, the aggregation-based lower bounds are used as additional pruning within the
72 compilation of relaxed and restricted DDs. Moreover, aggregate solutions are translated
73 into node selection heuristics to steer the compilation of restricted DDs toward resembling
74 solutions to the original problem, which are thus expected to be good.

75 Throughout the paper, the framework is illustrated on three different combinatorial
76 problems: the Talent Scheduling Problem, the Pigment Sequencing Problem and the Aircraft
77 Landing Problem. They are then used for the experimental evaluation of the framework,
78 the results of which show that the aggregation-based bound brings additional pruning and
79 enables solving more instances. Furthermore, the aggregation-based node selection heuristic
80 improves the quality of the solutions found early in the search and thus contributes to
81 speeding up the overall resolution. Finally, we show that a DD-based solver using only the
82 aggregation-based bound as relaxation performs almost equally well, which is a promising
83 direction for problems for which defining a merging operator is difficult or inefficient.

84 Although this paper is – to the best of our knowledge – the first to combine aggregate
85 dynamic programming with the DD-based branch-and-bound paradigm proposed by Bergman
86 et al, there has already been some hybridization work to combine discrete optimization with
87 DDs and other methods. For instance, in [12], Cappart et al. propose to use reinforcement
88 learning to guess the variable ordering that should be used to derive the best possible bounds
89 from the compiled approximate DDs. Other attempts combined DDs with Lagrangian
90 relaxation [13, 23] or MIP [5, 22, 31, 32]. On a slightly different note, a method has
91 been proposed where restricted DDs are used to generate good neighborhoods in a *large*
92 *neighborhood search* framework [20].

2 Preliminaries

2.1 Discrete Optimization

A discrete optimization problem \mathcal{P} involves finding the best possible solution x^* from a finite set of feasible solutions $Sol(\mathcal{P}) = D \cap C$. This set is determined by the domain $D = D_0 \times \dots \times D_{n-1}$ from which the variables $x = \langle x_0, \dots, x_{n-1} \rangle$ each take on a value, i.e. $x_j \in D_j$, and by a set of constraints C imposed on the value assignments. The quality of the solutions is evaluated according to an objective function $f(x)$ that must be minimized. Formally, the problem is defined as $\min \{f(x) \mid x \in D \cap C\}$ and any optimal solution x^* must satisfy $x^* \in Sol(\mathcal{P})$ and $\forall x \in Sol(\mathcal{P}) : f(x^*) \leq f(x)$. We describe below three optimization problems that will be utilized in the paper as illustrations for the aggregation-based framework.

Talent Scheduling Problem The *Talent Scheduling Problem* (TalentSched) is a film shoot scheduling problem that considers a set $N = \{0, \dots, n-1\}$ of scenes and a set $A = \{0, \dots, m-1\}$ of actors. Each scene $i \in N$ involves a required set $R_i \subseteq A$ of actors for a duration $D_i \in \mathbb{N}$. Moreover, each actor $k \in M$ has a pay rate C_k and is paid without interruption from their first to their last scheduled scene. The objective of TalentSched is to find a permutation of the scenes that minimizes the total cost of the film shoot.

Pigment Sequencing Problem The *Pigment Sequencing Problem* (PSP) is a single-machine production planning problem that aims to minimize the stocking and changeover costs while satisfying a set of orders. There are different item types $I = \{0, \dots, n-1\}$ with a given stocking cost S_i to pay for each time period between the production and the deadline of an order. For each pair $i, j \in I$ of item types, a changeover cost C_{ij} is incurred whenever the machine switches the production from item type i to j . Finally, the demand matrix Q contains all the orders: $Q_p^i \in \{0, 1\}$ indicates whether there is an order for item type $i \in I$ at time period p with $0 \leq p < H$ and H the time horizon.

Aircraft Landing Problem The *Aircraft Landing Problem* (ALP) requires to schedule the landing of a set of aircrafts $N = \{0, \dots, n-1\}$ on a set of runways $R = \{0, \dots, r-1\}$. The aircrafts have a target T_i and latest L_i landing time. Moreover, the set of aircrafts is partitioned in disjoint sets A_0, \dots, A_{c-1} corresponding to different aircraft classes in $C = \{0, \dots, c-1\}$. For each pair of aircraft classes $a, b \in C$, a minimum separation time $S_{a,b}$ between the landings is given. The goal is to find the schedule that minimizes the total waiting time of the aircrafts – the delay between their target time and scheduled landing time – while respecting their latest landing time.

2.2 Dynamic Programming

Dynamic programming (DP) is a *divide-and-conquer* strategy introduced by Bellman [4] for solving discrete optimization problems with an inherent recursive structure. It works by recursively decomposing the problem in smaller and overlapping subproblems. The cornerstone of the approach is the caching of intermediate results that allows each distinct subproblem to be solved only once. A *DP model* of a discrete optimization problem \mathcal{P} can be defined as a labeled transition system consisting of:

- the *control variables* $x_j \in D_j$ with $j \in \{0, \dots, n-1\}$.
- a set of *state-spaces* $S = \{S_0, \dots, S_n\}$ among which one distinguishes the *initial state* r , the *terminal state* t and the *infeasible state* $\hat{0}$.

22:4 Boosting DD-Based Branch-and-Bound with Aggregate Dynamic Programming

- 135 ■ a set t of *transition functions* s.t. $t_j : S_j \times D_j \rightarrow S_{j+1}$ for $j = 0, \dots, n-1$ taking the
 136 system from one state s^j to the next state s^{j+1} based on the value d assigned to variable
 137 x_j , or to \perp if assigning $x_j = d$ is infeasible. These functions should never allow one to
 138 recover from infeasibility, i.e. $t_j(\hat{0}, d) = \hat{0}$ for any $d \in D_j$.
- 139 ■ a set h of *transition value functions* s.t. $h_j : S_j \times D_j \rightarrow \mathbb{R}$ representing the immediate
 140 reward of assigning some value $d \in D_j$ to the variable x_j for $j = 0, \dots, n-1$.
- 141 ■ a *root value* v_r .
- 142 On that basis, the objective function $f(x)$ of \mathcal{P} is formulated as follows:

$$\begin{aligned} & \text{minimize } f(x) = v_r + \sum_{j=0}^{n-1} h_j(s^j, x_j) \\ & \text{subject to } s^{j+1} = t_j(s^j, x_j), \text{ for all } j = 0, \dots, n-1, \text{ with } x_j \in D_j \\ & \quad s^j \in S_j, j = 0, \dots, n \text{ and } x \in C. \end{aligned} \tag{1}$$

- 144 **TalentSched** A DP model for TalentSched was introduced in [17] that we slightly adapt
 145 here to make it suitable for the relaxation discussed in Section 2.3.1. States of this model
 146 are pairs (M, P) where M and P are disjoint sets of scenes that respectively must or might
 147 still be scheduled. The only case where P is non-empty happens when a state is relaxed.
- 148 ■ Control variables: $x_j \in N$ with $0 \leq j < n$ decides which scene is shot in j -th position.
- 149 ■ State spaces: $S = \{(M, P) \mid M, P \subseteq N, M \cap P = \emptyset\}$. The root state is $r = (N, \emptyset)$ and
 150 the terminal states are of the form (\emptyset, P) .
- 151 ■ Transition functions:

$$t_j(s^j, x_j) = \begin{cases} (s^j.M \setminus \{x_j\}, s^j.P \setminus \{x_j\}) & \text{if } x_j \in s^j.M, \\ (s^j.M \setminus \{x_j\}, s^j.P \setminus \{x_j\}), & \text{if } x_j \in s^j.P \text{ and } |s^j.M| < n - j, \\ \hat{0}, & \text{otherwise.} \end{cases}$$

- 153 A scene from P can only be selected if there are more spots left than scenes in M .
- 154 ■ Transition value functions: let $a(Q) = \cup_{i \in Q} R_i$ be the required set of actors for a set
 155 of scenes Q . Given a state $s = (M, P)$, the set of actors that are guaranteed to be
 156 on-location is computed as $o(s) = a(s.M) \cap a(N \setminus (s.M \cup s.P))$ because they are required
 157 both for a scene that must still be scheduled and for another that is guaranteed to be
 158 scheduled. In the transition value functions, we add all the actors from R_{x_j} to this set
 159 and sum the individual costs: $h_j(s^j, x_j) = D_{x_j} \sum_{k \in o(s^j) \cup R_{x_j}} C_k$.
- 160 ■ Root value: $v_r = 0$.

- 161 **PSP** The PSP was already tackled with a DD-based approach in [16, 20]. We hereby recall
 162 the DP model from [16] that allows the machine to be idle at some time periods. In this
 163 model, the decisions are made backwards – this allows to define transition functions that only
 164 lead to feasible production schedules. If variable x_j decides the type of item to produce at
 165 period j , the reverse variable ordering x_{H-1}, \dots, x_0 is thus used. To simplify the transition
 166 functions, let us denote by P_r^i the time period at which the r -th item of type i must be
 167 delivered, i.e. $P_r^i = \min\{0 \leq q < H \mid \sum_{p=0}^q Q_p^i \geq r\}$ for all $i \in N, 0 \leq r \leq \sum_{0 \leq p < H} Q_p^i$.
 168 Moreover, we define a dummy item type \perp used for idle periods and $N' = N \cup \{\perp\}$.

- 169 States are pairs (i, R) with i the item type that the machine is currently set to produce
 170 and R a vector that gives the remaining number R_i of demands to satisfy for each type i .
- 171 ■ Control variables: $x_j \in N'$ with $0 \leq j < H$ decides the item type to produce at period j .

172 ■ State space: $S = \{s \mid s.i \in N', \forall i \in N, 0 \leq s.R_i \leq \sum_{0 \leq p < H} Q_p^i\}$. The root state is given
 173 by $r = \langle \perp, (\sum_{0 \leq p < H} Q_p^0, \dots, \sum_{0 \leq p < H} Q_p^{n-1}) \rangle$ and the terminal states are of the form
 174 $\langle i, (0, \dots, 0) \rangle$ with $i \in N'$.

175 ■ Transition functions:

$$176 \quad t_j(s^j, x_j) = \begin{cases} \langle t_j^i(s^j, x_j), t_j^R(s^j, x_j) \rangle, & \text{if } x_j \neq \perp \text{ and } s^j.R_{x_j} > 0 \text{ and } j \leq P_{s^j.R_{x_j}}^{x_j}, \\ \langle t_j^i(s^j, x_j), t_j^R(s^j, x_j) \rangle, & \text{if } x_j = \perp \text{ and } \sum_{i \in N} s^j.R_i < j + 1, \\ \hat{0}, & \text{otherwise.} \end{cases}$$

where

$$t_j^i(s^j, x_j) = \begin{cases} x_j, & \text{if } x_j \neq \perp \\ s^j.i, & \text{otherwise.} \end{cases}$$

$$t_j^R(s^j, x_j) = \begin{cases} (s^j.R_0, \dots, s^j.R_{x_j} - 1, \dots, s^j.R_{n-1}), & \text{if } x_j \neq \perp \\ s^j.R, & \text{otherwise.} \end{cases}$$

177 In the transition function, the first condition ensures that there remains at least one
 178 item to produce for the chosen type and that the current time period j is earlier than its
 179 deadline. The second condition ensures that idle periods cannot be scheduled when the
 180 remaining quantity to produce is equal to the number of periods left.

181 ■ Transition value functions: the changeover and stocking costs are computed as:

$$182 \quad h_j(s^j, x_j) = \left\{ \begin{array}{l} C_{x_j s^j.i}, \quad \text{if } x_j \neq \perp \text{ and } s^j.i \neq \perp \\ 0, \quad \text{otherwise.} \end{array} \right\} + \left\{ \begin{array}{l} S_{x_j} \cdot (j - P_{s^j.R_{x_j}}^{x_j}), \quad \text{if } x_j \neq \perp \\ 0, \quad \text{otherwise.} \end{array} \right\}$$

183 ■ Root value: $v_r = 0$.

184 **ALP** We reproduce here the DP model presented in [28] where states are pairs (Q, ROP) ,
 185 with Q a vector that gives the remaining number of aircrafts of each class to schedule and
 186 ROP a *runway occupation profile*: a vector containing pairs (l, c) that respectively give the
 187 time and aircraft class of the latest landing scheduled on each runway. Similarly to the PSP
 188 modeling, we denote by \perp either a dummy aircraft class or a dummy runway.

189 ■ Control variables: we use pairs of variables $(x_j, y_j) \in (C \times R) \cup \{(\perp, \perp)\}$ with $0 \leq j < n$
 190 that represent the decision to place an aircraft of class x_j on runway y_j , or to schedule
 191 nothing at all in case of (\perp, \perp) .

192 ■ State spaces: $S = \{(Q, ROP) \mid \forall i \in C : Q_i \geq 0, \forall k \in R : ROP_k.l \geq 0, ROP_k.c \in C \cup \{\perp\}\}$.
 193 The root state is $r = (\langle |A_0|, \dots, |A_{c-1}| \rangle, \langle (0, \perp), \dots, (0, \perp) \rangle)$ and the terminal states are
 194 of the form $(\langle 0, \dots, 0 \rangle, ROP)$.

195 ■ Transition functions: if A_i^k gives the aircraft from class i that must be scheduled when
 196 there are k aircrafts left from this class, we can define the function computing the earliest
 197 landing time given a state s , a class x and a runway y :

$$198 \quad E(s, x, y) = \begin{cases} T_{A_x^s.Q_x}, & \text{if } s.ROP_y.l = 0 \text{ and } s.ROP_y.c = \perp, \\ \max(s.ROP_y.l + \min_{i \in C} S_{i,x}, T_{A_x^s.Q_x}), & \text{if } s.ROP_y.l > 0 \text{ and } s.ROP_y.c = \perp, \\ \max(s.ROP_y.l + S_{s.ROP_y.c,x}, T_{A_x^s.Q_x}), & \text{otherwise.} \end{cases}$$

199 This allows us to define the transition functions as:

$$200 \quad t_j(s^j, x_j, y_j) = \begin{cases} \langle t_j^Q(s^j, x_j, y_j), t_j^{ROP}(s^j, x_j, y_j) \rangle, & \text{if } x_j \neq \perp \text{ and } s^j.Q_{x_j} > 0 \\ & \text{and } E(s^j, x_j, y_j) \leq L_{A_{x_j}^{s^j.Q_{x_j}}}, \\ \langle t_j^Q(s^j, x_j, y_j), t_j^{ROP}(s^j, x_j, y_j) \rangle, & \text{if } x_j = \perp \text{ and } \sum_{i \in C} s^j.Q_i = 0, \\ \hat{0}, & \text{otherwise.} \end{cases}$$

201 where

$$\begin{aligned}
 202 \quad t_j^Q(s^j, x_j, y_j) &= \begin{cases} \langle s^j.Q_0, \dots, s^j.Q_{x_j} - 1, \dots, s^j.Q_{c-1} \rangle, & \text{if } x_j \neq \perp \\ s^j.Q, & \text{otherwise.} \end{cases} \\
 t_j^{ROP}(s^j, x_j, y_j) &= \begin{cases} \langle s^j.ROP_0, \dots, (E(s^j, x_j, y_j), x_j), \dots, s^j.ROP_{r-1} \rangle, & \text{if } x_j \neq \perp \\ s^j.ROP, & \text{otherwise.} \end{cases}
 \end{aligned}$$

203 The first condition of the transition function ensures that there remains at least one
 204 aircraft of the chosen class and that its earliest landing time is not greater its latest
 205 landing time. The second condition only allows us to schedule dummy aircrafts when
 206 there are no aircrafts left to schedule.

- Transition value functions: the waiting time of the aircraft is computed as:

$$h_j(s^j, x_j, y_j) = \begin{cases} E(s^j, x_j, y_j) - T_{A_{x_j}}^{s^j.Q_{x_j}}, & \text{if } x_j \neq \perp \\ 0, & \text{otherwise.} \end{cases}$$

- 207 ■ Root value: $v_r = 0$.

208 Because the runways are identical and independent, there are many symmetries in this model.
 209 This can be mitigated by sorting the ROP of every state by increasing latest landing time,
 210 breaking ties according to the previous aircraft class scheduled.

211 2.3 Decision Diagrams

212 When used to manipulate the DP model of a discrete optimization problem \mathcal{P} , DDs are
 213 graphical encodings that represent a set of solutions of the problem. More precisely, a
 214 DD $\mathcal{B} = (U, A, \sigma, l, v)$ is a layered directed acyclic graph composed of a set of nodes U
 215 interconnected by a set of arcs A . Starting from a single node u_r corresponding to a DP state
 216 given by the function $\sigma(u_r)$, the process of iteratively extending a set of partial solutions is
 217 called the *compilation* of a DD and is described by Algorithm 1. Note that the highlighted
 218 portions concern the ingredients introduced in Section 3 and can be ignored for now. The
 219 algorithm begins by initializing a layer L_i that only contains the *root node* u_r , assuming its
 220 state $\sigma(u_r)$ belongs to the i -th stage of the DP model. The subsequent layers of the DD are
 221 then constructed sequentially by applying each valid transition of the DP model to every
 222 node of the last completed layer at lines 8–16. Each layer thus corresponds to a stage of
 223 the DP model and contains a single node for each state reached in order to preserve the
 224 compactness of the model. The arcs $a \in A$ materialize the transitions that exist between
 225 the states of consecutive stages. In particular, the arc $a = (u \xrightarrow{d} u')$ connecting nodes
 226 $u \in L_j, u' \in L_{j+1}$ represents the transition between $\sigma(u)$ and $\sigma(u')$. The decision associated
 227 with this transition is stored by the *label* $l(a) = d \in D_j$ and the transition value is given by
 228 the *arc value* $v(a)$.

229 The algorithm completes when the last layer L_n is generated, constituted by a single
 230 node t called the *terminal* node. The DD thus constructed contains a set of $u_r \rightsquigarrow t$
 231 paths that can be combined with any previously discovered $r \rightsquigarrow u_r$ path, connecting
 232 the *root* of the problem to u_r . Any $r \rightsquigarrow t$ path $p = (a_0, \dots, a_{n-1})$ represents a solution
 233 given by $x(p) = (l(a_0), \dots, l(a_{n-1}))$. The objective value of such solution can also be
 234 retrieved from the sequence of arcs by accumulating their values, and adding the root
 235 value: $v(p) = v_r + \sum_{j=0}^{n-1} v(a_j)$. The set of solutions contained in the DD is denoted as
 236 $Sol(\mathcal{B}) = \{x(p) \mid \exists p : r \rightsquigarrow t, p \in \mathcal{B}\}$. A DD rooted at a node u_r is *exact* if it perfectly
 237 represents the set of solutions of the corresponding subproblem $\mathcal{P}|_{u_r}$, i.e. $Sol(\mathcal{B}) = Sol(\mathcal{P}|_{u_r})$
 238 and $v(p) = f(x(p)), \forall p \in \mathcal{B}$. The best value among the $u_1 \rightsquigarrow u_2$ paths in \mathcal{B} is denoted
 239 $v^*(u_1 \rightsquigarrow u_2 \mid \mathcal{B})$, and in particular $v^*(u \mid \mathcal{B}) = v^*(r \rightsquigarrow u \mid \mathcal{B})$.

■ **Algorithm 1** Compilation of DD \mathcal{B} rooted at node u_r with maximum width W .

```

1:  $i \leftarrow$  index of the layer containing  $u_r$ 
2:  $L_i \leftarrow \{u_r\}$ 
3:  $\tilde{P} \leftarrow \Delta(\tilde{p})$  with  $\tilde{p}$  the optimal solution for  $\pi(\sigma(u_r))$  // retrieve disaggregate solution
4: for  $j = i$  to  $n - 1$  do
5:   if  $|L_j| > W$  then
6:     restrict or relax the layer to get  $W$  nodes with Algorithm 2
7:    $L_{j+1} \leftarrow \emptyset$ 
8:   for all  $u \in L_j$  do
9:      $v_{rlb}(\sigma(u)) \leftarrow \max \{v_{rlb}(\sigma(u)), v_{agg}(\pi(\sigma(u)))\}$  // inject aggregation-based bound
10:    if  $v^*(u | \mathcal{B}) + v_{rlb}(\sigma(u)) \geq \bar{v}$  then // rough lower bound pruning w.r.t. incumbent
11:      continue
12:    for all  $d \in D_j$  do
13:      create node  $u'$  with state  $\sigma(u') = t_j(\sigma(u), d)$  or retrieve it from  $L_{j+1}$ 
14:      create arc  $a = (u \xrightarrow{d} u')$  with  $v(a) = h_j(\sigma(u), d)$  and  $l(a) = d$ 
15:       $score(a) \leftarrow 1$  if  $l(a) \in \tilde{P}_j$ , 0 otherwise
16:      add  $u'$  to  $L_{j+1}$  and add  $a$  to  $A$ 

```

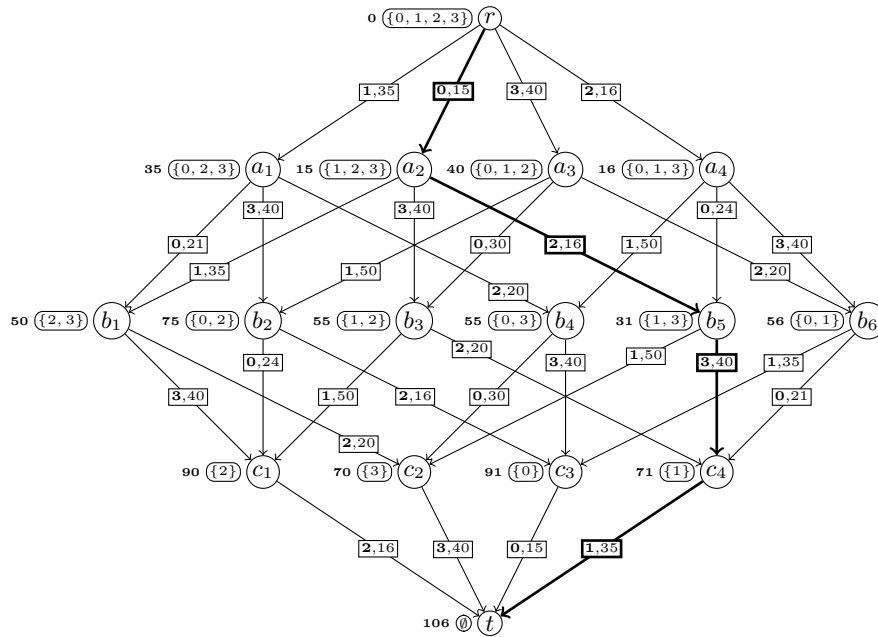
240 ▶ **Example 1.** Let us define a TalentSched instance with 4 scenes with durations $D =$
 241 $\langle 3, 5, 2, 4 \rangle$ and 4 actors with pay rates $C = \langle 10, 20, 30, 40 \rangle$. The actor requirements for each
 242 scene are given by $R = \langle \{0, 3\}, \{0, 1, 3\}, \{0, 2, 3\}, \{0, 1, 2, 3\} \rangle$. Figure 1 shows the exact DD
 243 compiled for this instance with the DP model recalled in Section 2.2. Note that for each
 244 state $s = (M, P)$ corresponding to a node in the DD, we only show the set M since P is
 245 always empty in exact nodes. An optimal solution of the problem is $\langle 0, 2, 3, 1 \rangle$, which gives
 246 an objective value of 106.

247 As the reader might have guessed, the compilation of an exact DD for a combinatorial
 248 optimization problem suffers from the curse of dimensionality as much as the corresponding
 249 DP model. This is why DD-based discrete optimization rarely relies on exact DDs but rather
 250 on *restricted* and *relaxed* DDs. These two variants follow two distinct compilation schemes
 251 that allow to maintain the number of nodes of each layer – called the *width* – under a given
 252 parameter W . In Algorithm 1, this logic is performed at line 5 where the width of the current
 253 layer is compared with W . If needed, the layer is then either restricted or relaxed at line 6
 254 by calling Algorithm 2.

255 2.3.1 Approximate Decision Diagrams

256 As stated by Algorithm 2, restricted DDs simply remove surplus nodes from the layer until it
 257 is reduced to W nodes. A heuristic is used to evaluate the nodes and drop the least promising
 258 ones. Restricted DDs thus generate a subset of the solutions of the corresponding problem,
 259 i.e. $Sol(\bar{\mathcal{B}}) \subseteq Sol(\mathcal{P})$ and $v(p) = f(x(p)), \forall p \in \bar{\mathcal{B}}$ for a restricted DD $\bar{\mathcal{B}}$. They thus provide
 260 upper bounds on the objective value.

261 As opposed to restricted DDs, a relaxed DD $\underline{\mathcal{B}}$ yields lower bounds by representing
 262 a superset of the solutions of the corresponding problem: $Sol(\underline{\mathcal{B}}) \supseteq Sol(\mathcal{P})$ and $v(p) \leq$
 263 $f(x(p)), \forall p \in \underline{\mathcal{B}}$. This is achieved through a problem-specific *state merging* operator $\oplus(\sigma(\mathcal{M}))$
 264 that defines an approximate representation that includes all states $\sigma(\mathcal{M}) = \{\sigma(u) \mid u \in \mathcal{M}\}$
 265 corresponding to the merged nodes \mathcal{M} and preserves all their outgoing transitions, although



■ **Figure 1** The exact DD for the TalentSched instance given in Example 1. Nodes are annotated with their state and the best prefix value. Arcs are labeled with the associated decision in bold and transition value. The arcs constituting one of the optimal solutions are highlighted in bold.

■ **Algorithm 2** Restriction or relaxation of layer L_j with maximum width W .

```

1: while  $|L_j| > W$  do
2:    $\mathcal{M} \leftarrow$  select nodes from  $L_j$  according to their score
3:    $L_j \leftarrow L_j \setminus \mathcal{M}$ 
4:   create node  $\mu$  with state  $\sigma(\mu) = \oplus(\sigma(\mathcal{M}))$  and add it to  $L_j$  // for relaxation only
5:   for all  $u \in \mathcal{M}$  and arc  $a = (u' \xrightarrow{d} u)$  incident to  $u$  do
6:     replace  $a$  by  $a' = (u' \xrightarrow{d} \mu)$  and set  $v(a') = \Gamma_{\mathcal{M}}(v(a), u)$ 

```

266 it may also introduce infeasible transitions. In Algorithm 2, a *meta*-node is created for this
267 merged state at line 4 and the arcs pointing to the deleted nodes are redirected to this
268 merged node at line 6. The operator $\Gamma_{\mathcal{M}}$ permits to adjust the value of these arcs if needed.
269 In all three formulations given below, this operator is the identity function.

270 **TalentSched** The merging operator is defined by $\oplus(\mathcal{M}) = (\oplus_M(\mathcal{M}), \oplus_P(\mathcal{M}))$ where
271 $\oplus_M(\mathcal{M}) = \bigcap_{s \in \mathcal{M}} s.M$ and $\oplus_P(\mathcal{M}) = (\bigcup_{s \in \mathcal{M}} s.M \cup s.P) \setminus (\bigcap_{s \in \mathcal{M}} s.M)$. The definition of
272 $\oplus_P(\mathcal{M})$ ensures that the resulting set of scenes that might be scheduled contains any scene
273 that must or might be scheduled in any of the states, except those that still must be scheduled
274 for all states.

275 **PSP** A valid merging operator is $\oplus(\mathcal{M}) = (\perp, \langle \min_{s \in \mathcal{M}} s.R_0, \dots, \min_{s \in \mathcal{M}} s.R_{n-1} \rangle)$. The
276 configuration of the machine is always reset to the dummy item type \perp as there is little chance
277 that merged states agree on it. For each item type, the remaining number of demands is
278 computed by taking the minimum value among all merged states, meaning that any demand
279 satisfied by at least one state is considered satisfied in the merged state.

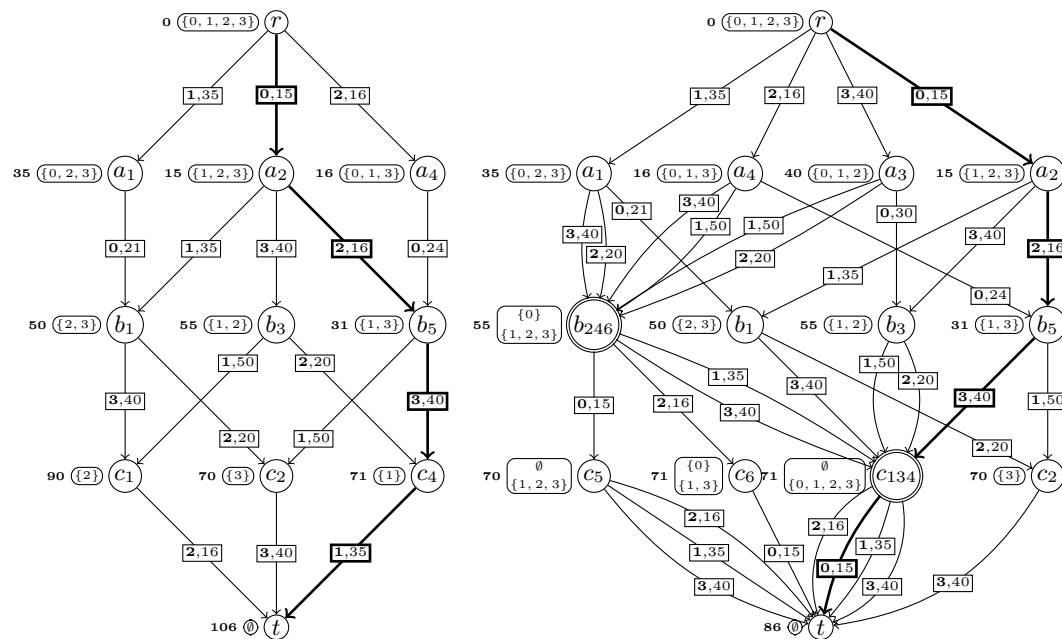


Figure 2 Respectively on the left and the right, a restricted and relaxed DD for the TalentSched instance given in Example 1, compiled with W set to 3 and 4. Merged nodes are circled twice.

280 **ALP** The merging operator is again defined separately for each component of the states:
 281 $\oplus(\mathcal{M}) = (\oplus_Q(\mathcal{M}), \oplus_{ROP}(\mathcal{M}))$. First, the minimum remaining quantity of aircrafts for
 282 each class is stored in the merged state: $\oplus_Q(\mathcal{M}) = \langle \min_{s \in \mathcal{M}} s.Q_0, \dots, \min_{s \in \mathcal{M}} s.Q_{c-1} \rangle$. For
 283 the ROP, the minimum latest landing time on each runway is kept and the last aircraft classes
 284 scheduled are reset to \perp : $\oplus_{ROP}(\mathcal{M}) = \langle (\min_{s \in \mathcal{M}} s.ROP_0.l, \perp), \dots, (\min_{s \in \mathcal{M}} s.ROP_{r-1}.l, \perp) \rangle$.

285 **Example 2.** Figure 2 shows approximate DDs for the TalentSched instance introduced
 286 in Example 1. Despite having a maximum width of 3, the best solution contained in the
 287 restricted DD is the optimal solution previously found. With a maximum width of 4, the
 288 relaxed DD provides a global lower bound of 86. The path corresponding to this lower bound
 289 is given by the assignment $\langle 0, 2, 3, 0 \rangle$, which is infeasible because scene 0 is scheduled twice.

2.3.2 Branch-and-Bound

291 In [9], a branch-and-bound algorithm based only on restricted and relaxed DDs was introduced.
 292 It maintains a queue of open nodes that represent the set of subproblems that remain to
 293 process. For each of them, a restricted DD is compiled in an attempt to improve the
 294 incumbent solution. Then, a relaxed DD is constructed in order to both decompose the given
 295 subproblem into even smaller ones and to compute a lower bound for each of them. These
 296 nodes are then added to the branch-and-bound queue for further exploration, unless the
 297 lower bound permits their direct elimination. Ultimately, the optimality of the best solution
 298 discovered during the search is confirmed once the queue has been emptied.

299 2.3.3 Rough Lower Bound

300 The *rough lower bound* (RLB) [19] is an additional optional modeling component that can
 301 be specified to speed up the resolution of any optimization problem. For any node u , the
 302 RLB gives a lower bound on the best value one can obtain when solving the corresponding
 303 subproblem $\sigma(u)$, i.e. $\underline{v}_{rlb}(\sigma(u)) \leq v^*(u \rightsquigarrow t \mid \mathcal{B})$ with \mathcal{B} the exact DD for the problem. It is
 304 used at line 10 of Algorithm 1 to filter nodes *a priori* by comparing this lower bound with
 305 the incumbent value \bar{v} . Since the RLB is computed for each node of the approximate DDs
 306 compiled throughout the branch-and-bound, it needs to be computationally cheap.

307 The RLB has the potential both to focus the compilation of restricted DDs on promising
 308 parts of the search space and to strengthen the bounds obtained through relaxed DDs.
 309 Furthermore, the branch-and-bound algorithm uses the RLB to make pruning decisions, if it
 310 happens to be tighter than the bound obtained with relaxed DDs.

311 **Example problems** In our computational experiments, we use the lower bound given by
 312 Theorem 1 in [17] for TalentSched and the same RLB as in [20] for the PSP. We do not detail
 313 them in this article for the sake of conciseness.

314 3 Aggregate Dynamic Programming for Decision Diagrams

315 As stated in the introduction, optimizations techniques based on DP and DDs can prove
 316 highly effective [6, 13, 14, 18, 19]. In some cases, however, the state space of the DP models
 317 is simply too large and the bounds derived from restricted and relaxed DDs are of little
 318 to no use. This can be imputed either to the node selection heuristic or to the relaxation
 319 scheme. The *MinLP* heuristic traditionally used favors keeping nodes with the best prefix
 320 values. This locally-optimal selection policy may result in the elimination of all nodes that
 321 lead to the optimal solution, or even to any feasible solution, particularly in cases of highly
 322 constrained problems. In the latter case, the compilation of a restricted DD is a pure waste
 323 of time: no feasible solution is found at the end of the compilation, and not even a bound on
 324 the objective value can be exploited to reduce the optimality gap. The same phenomenon
 325 is detrimental to the usefulness of compiled relaxed DDs whose bounds might be of low
 326 quality when the node selection heuristic is oblivious to the global structure of the problem.
 327 Indeed, the merging operator yields a loose representation when applied to an arbitrary set
 328 of nodes for most problems. In the absence of a perfect heuristic, this situation will occur
 329 under certain conditions. It inspired our pursuit of a more globally-focused approach that
 330 could enhance the usefulness of the compiled DDs. This section presents a framework for
 331 integrating *aggregate dynamic programming* ideas with DD-based optimization that aims to
 332 address some of these shortcomings. Instead of relaxing the original problem by reasoning
 333 on merged states, it proposes to use problem instance and state aggregation operators that
 334 yield a simpler and relaxed version of the problem, which can be solved exactly. Solutions
 335 of the aggregated problem can provide bounds that capture the global problem structure,
 336 as well as guidance for the compilation of restricted DDs. This section details the role and
 337 meaning of the components of the framework one by one.

338 3.1 Preprocessing: Problem Instance Aggregation

339 The goal of this preprocessing step is to create an aggregate and simpler problem instance by
 340 reducing one or more dimensions of the problem. The *instance aggregation operator* Π must
 341 be defined such that the aggregate problem instance $\mathcal{P}' = \Pi(\mathcal{P})$ is a relaxation of the original

342 problem instance \mathcal{P} . In practice, assuming the problem reasons over a set of *elements*, a
 343 clustering algorithm can be used to create clusters of such elements. Then, the aggregate
 344 problem instance can be obtained by considering *aggregate* elements that encompass all
 345 elements in a given cluster and by adapting the instance data accordingly. Formally, if a
 346 set E of elements is clustered into K clusters, we define two mapping functions: $\Phi : E \rightarrow$
 347 $\{0, \dots, K - 1\}$ that gives the cluster for each original element and $\Phi^{-1} : \{0, \dots, K - 1\} \rightarrow 2^E$
 348 that gives the set of original elements for a given cluster.

349 **TalentSched** In [17], it is proved that there always exists an optimal solution to the
 350 problem in which scenes with the same set of actors are scheduled together. This gives us the
 351 opportunity to aggregate the problem by creating K clusters of scenes that require a similar set
 352 of actors, which is plausible to occur in real film shoots. Scenes belonging to the same clusters
 353 can then be aggregated by taking the intersection of their actor requirements and adding up
 354 their durations. Formally, we write $\Pi(\mathcal{P} = (N, A, R, D, C)) = (\Pi_N(N), A, \Pi_R(R), \Pi_D(D), C)$
 355 with $\Pi_N(N) = \{0, \dots, K - 1\}$. The aggregate actor requirements are computed as $\Pi_R(R) =$
 356 R' with $R'_i = \cap_{j \in \Phi^{-1}(i)} R_j$ for all $i \in \Pi_N(N)$ and the aggregate durations as $\Pi_D(D) = D'$
 357 with $D'_i = \sum_{j \in \Phi^{-1}(i)} D_j$ for all $i \in \Pi_N(N)$.

358 **PSP** The number of item types considered in a PSP instance dramatically impacts the
 359 size of the state space – for instance, the case with only one item type can be solved
 360 greedily. Therefore, and because it is not unlikely that the machine will produce several
 361 sets of similar items, we propose to cluster item types that have similar stocking and
 362 changeover costs. The instance aggregation operator is thus $\Pi(\mathcal{P} = (I, S, C, H, Q)) =$
 363 $(\Pi_I(I), \Pi_S(S), \Pi_C(C), H, \Pi_Q(Q))$, where the aggregate set of item types is given by $\Pi_I(I) =$
 364 $\{0, \dots, K - 1\}$. Their stocking costs are computed as $\Pi_S(S) = S'$ with $S'_k = \min_{i \in \Phi^{-1}(k)} S_i$
 365 for all $k \in \Pi_I(I)$ and the changeover costs as $\Pi_C(C) = C'$ with $C'_{kl} = \min_{i \in \Phi^{-1}(k), j \in \Phi^{-1}(l)} C_{ij}$
 366 for all $k, l \in \Pi_I(I)$. The aggregate demand matrix is defined as $\Pi_Q(Q) = Q'$ with $Q'_p{}^k =$
 367 $\sum_{i \in \Phi^{-1}(k)} Q_p^i$. However, as the demand matrix is only supposed to contain unit demands,
 368 one must redistribute surplus demands in Q' to the left.

369 **ALP** Similarly to the item types of the PSP, the aircraft classes can be aggregated to
 370 reduce the complexity of the problem. We thus propose to cluster them based on their
 371 minimum separation time with other classes and define the instance aggregation operator as
 372 $\Pi(\mathcal{P} = (N, R, C, A, S, T, L)) = (N, R, \Pi_C(C), \Pi_A(A), \Pi_S(S), T, \Pi_L(L))$. The set of aircrafts,
 373 their target landing time and the number of runways is kept. The aggregate set of classes is
 374 given by $\Pi_C(C) = \{0, \dots, K - 1\}$ and their corresponding set of aircrafts is computed as
 375 $\Pi_A(A) = A'$ with $A'_i = \cup_{j \in \Phi^{-1}(i)} A_j$ for all $i \in \Pi_C(C)$. The smallest separation times between
 376 aggregate classes are kept, as formalized by $\Pi_S(S) = S'$ with $S'_{kl} = \min_{i \in \Phi^{-1}(k), j \in \Phi^{-1}(l)} S_{i,j}$
 377 for all $k, l \in \Pi_C(C)$. Finally, the aggregation operator adapts the latest landing times of all
 378 the aircrafts so that any aircraft with a given target landing time has a greater latest landing
 379 time than all other aircrafts of the same class with a smaller target landing time: $\Pi_L(L) = L'$
 380 with $L'_i = \max \{L_j \mid \Phi(i) = \Phi(j), T_i \leq T_j\}$ for all $i \in A$. This property is assumed to hold
 381 for the original problem instance, and must be preserved so that aircrafts from the same
 382 class can be scheduled sequentially in the DP model.

383 ► **Example 3.** Let us apply the problem instance aggregation to our running example by
 384 creating $K = 2$ aggregate scenes. Assuming the following clustering is found: $\Phi(0) = 0, \Phi(1) =$
 385 $1, \Phi(2) = 0, \Phi(3) = 1$ or equivalently $\Phi^{-1}(0) = \{0, 2\}, \Phi^{-1}(1) = \{1, 3\}$. We thus compute

386 the aggregate scene durations as: $D' = \langle D_0 + D_2, D_1 + D_3 \rangle = \langle 5, 9 \rangle$ and the aggregate actor
 387 requirements as: $R' = \langle \{0, 3\} \cap \{0, 2, 3\}, \{0, 1, 3\} \cap \{0, 1, 2, 3\} \rangle = \langle \{0, 3\}, \{0, 1, 3\} \rangle$.

388 3.2 State Aggregation and Lower Bound

389 A second mapping function accompanies the problem instance aggregation operator: the
 390 *state aggregation operator* $\pi : S \rightarrow S'$ that projects each state of the state space S of the
 391 original problem in the aggregate state space S' . The role of this operator is to translate each
 392 original state to its aggregate version by adapting the state information to fit the aggregate
 393 problem data. Let us denote by \mathcal{B} and \mathcal{B}' the exact DD for problem \mathcal{P} and $\Pi(\mathcal{P})$, respectively.
 394 If the aggregation operators Π and π are defined such that $v^*(u \rightsquigarrow t \mid \mathcal{B}) \geq v^*(u' \rightsquigarrow t' \mid \mathcal{B}')$
 395 for all $u \in \mathcal{B}, u' \in \mathcal{B}'$ with $\pi(\sigma(u)) = \sigma(u')$ and $\pi(\sigma(t)) = \sigma(t')$, then $v^*(u' \rightsquigarrow t' \mid \mathcal{B}')$
 396 can be used as a lower bound in the original problem, which we will denote by $v_{agg}(\pi(\sigma(u)))$.

397 Assuming the aggregate problem can be pre-solved exactly and the solution of each
 398 subproblem is stored, this aggregation-based lower bound can be retrieved very quickly. One
 399 way to exploit it is to incorporate it in the RLB as shown at line 9 of Algorithm 1 so that it
 400 is used as often as possible. Another possibility would be to use the aggregate state space to
 401 replace the state merging scheme in relaxed DDs. Once a layer with greater width than W is
 402 reached, all the states contained in the nodes of the layer could be mapped to the aggregate
 403 state space to pursue the compilation in a lower dimensional space.

404 **TalentSched** The state compression operator for TalentSched is somewhat complex because
 405 we can only map to states where complete aggregate scenes have yet to be scheduled. As a
 406 result, if a state s contains scenes in $s.P$ that can optionally be scheduled, we map it to a
 407 dummy aggregated state. The same logic is applied when $s.M$ only contains a subset of the
 408 scenes that compose an aggregate scene.

$$409 \quad \pi(s) = \begin{cases} (\emptyset, \emptyset), & \text{if } s.P \neq \emptyset, \\ (\emptyset, \emptyset), & \text{if } \exists i \in \Pi_N(N) : (\Phi^{-1}(i) \cap s.M) \neq \emptyset \wedge \Phi^{-1}(i) \not\subseteq s.M, \\ (M', \emptyset), & \text{otherwise, with } M' = \{i \in \Pi_N(N) \mid \Phi^{-1}(i) \subseteq s.M\}. \end{cases}$$

410 **PSP** If we extend the definition of Φ such that $\Phi(\perp) = \perp$, the state aggregation operator
 411 can be defined as $\pi(s) = (\Phi(s.i), R)$ with $R_i = \sum_{j \in \Phi^{-1}(i)} s.R_j$ for all $i \in \Pi_I(I)$. The item
 412 type is projected to its corresponding aggregate type, and the remaining number of items to
 413 produce for each type is separately accumulated within each cluster.

414 **ALP** Again, assuming $\Phi(\perp) = \perp$, the state aggregation operator is defined by $\pi(s) =$
 415 (Q', ROP') with the remaining quantities of aircrafts aggregated as $Q'_i = \sum_{j \in \Phi^{-1}(i)} s.Q_j$
 416 for all $i \in \Pi_C(C)$. For the ROP, one only needs to adapt the class of the last aircraft scheduled
 417 on each runway $ROP'_i = (s.ROP_0.l, \Phi(s.ROP_0.c))$ for all $i \in R$.

418 If lower bounds for original states are obtained only by pre-solving the aggregate problem,
 419 it is unlikely that the solution of an aggregate subproblem mapped with the state aggregation
 420 operator will be available, since the aggregate separation times between aircraft classes lead to
 421 very different landing times. However, a lower bound for an aggregate state $s^1 = (Q^1, ROP^1)$
 422 can be provided by the solution of any state $s^2 = (Q^2, ROP^2)$ such that $Q^1 = Q^2$ and
 423 $ROP_i^1.c = ROP_i^2.c$ and $ROP_i^1.l \geq ROP_i^2.l$ for all $i \in R$.

424 **► Example 4.** Let us compute the aggregation-based lower bound for the root state of
 425 the running example $r = (\{0, 1, 2, 3\}, \emptyset)$ given its aggregate version $\pi(r) = (\{0, 1\}, \emptyset)$ and

426 the clustering performed in Example 3. The aggregate version is trivial to solve since the
 427 objective function is symmetrical and there are only two scenes to schedule. We thus have
 428 $\underline{v}_{agg}(r) = D'_0 \times (C_0 + C_3) + D'_1 \times (C_0 + C_1 + C_3) = 5 \times (1 + 4) + 9 \times (1 + 2 + 4) = 88$, which
 429 is a slightly better lower bound than the one obtained with the relaxed DD of Example 2.

430 3.3 Solution Disaggregation and Node Selection Heuristic

431 In order to exploit the solution of the aggregate version of a subproblem to find good heuristic
 432 solutions for the original subproblem, we need to specify the correspondence between decisions
 433 in the aggregate problem with decisions in the original problem. We therefore define a last
 434 modeling component, called the *decision disaggregation operator* $\delta(d) : D'_k \rightarrow 2^{D_i} \times \dots \times 2^{D_j}$
 435 that maps the instantiation of a variable x'_k in the aggregate problem to a vector of possible
 436 corresponding assignments for variables x_i, \dots, x_j in the original problem.

437 Finally, we define the *path disaggregation operator* that transforms a sequence of decisions
 438 in the aggregate problem to a sequence of sets of possible decisions in the original problem:
 439 $\Delta(p = (a_k, \dots, a_{n'-1})) = \delta(l(a_k)) \cdot \dots \cdot \delta(l(a_{n'-1}))$ where n' is the supposed number of
 440 aggregate variables and \cdot denotes the *concatenation* of two vectors. Using this operator,
 441 we can compute a *score* for each decision made during the compilation of restricted DDs.
 442 At line 3 of Algorithm 1, we first retrieve the optimal value assignment of the aggregate
 443 subproblem and apply the path disaggregation operator on it. Then, a binary *score* is
 444 attributed to each arc at line 15, depending on its compatibility with the disaggregated
 445 solution. At line 2 of Algorithm 2, the maximum score obtained along any path up to each
 446 node can then be used to order nodes from most to least promising, favoring nodes with
 447 incoming paths that are highly compatible with the disaggregated solution. By doing so, the
 448 width of restricted DDs is controlled in the same way as before, enabling the preference of
 449 solutions even when no feasible solution with the maximum possible score is available.

450 **TalentSched** Each aggregate scene corresponds to a set of original scenes, we thus need to
 451 map each aggregate decision to a sequence of original decisions: $\delta(i) = V$ where $V_j = \Phi^{-1}(i)$
 452 for all $0 \leq j < |\Phi^{-1}(i)|$. It corresponds to any of the scenes from the cluster i , duplicated
 453 $|\Phi^{-1}(i)|$ times so that they are all scheduled one after another, preferably.


454 **PSP** The operator is much simpler to define for the PSP, since each decision concerns the
 455 production of one unit of a chosen aggregate item type. It can thus be interpreted as the
 456 decision of producing one unit of any item type in the corresponding cluster: $\delta(i) = \langle \Phi^{-1}(i) \rangle$.

457 **ALP** The only difference with the PSP is that decisions also contain the runway on which
 458 the aircraft is scheduled to land, which remains the same: $\delta(a, r) = \langle \{(a', r) \mid a' \in \Phi^{-1}(a)\} \rangle$.

459 ► **Example 5.** As computed in Example 4, the schedule $\langle 0, 1 \rangle$ is optimal for the aggregate
 460 problem. By disaggregating this solution, we get $\langle \{0, 2\}, \{0, 2\}, \{1, 3\}, \{1, 3\} \rangle$. We can notice
 461 that the optimal schedule $\langle 0, 2, 3, 1 \rangle$ found in Example 1 is compatible with the disaggregated
 462 solution and would thus be favored by the aggregation-based node selection heuristic.

463 4 Computational Experiments

464 The impact of the aggregation-based bounds and heuristics was evaluated experimentally by
 465 extending the generic DD-based solver DDO [21] and injecting the modeling of the three
 466 discrete optimization problems presented throughout the paper. The version of DDO used




■ **Figure 3** Number of instances solved over time for each configuration and problem studied.

467 includes the improvements introduced in [16, 19]. For each problem, random instances were
 468 generated with the following main parameters:


- 469 ■ TalentSched: number of scenes $n \in \{20, 22, 24, 26, 28\}$, number of actors $m \in \{10, 15\}$
 470 and average fraction of actors required for each scene $\rho \in \{0.3, 0.4\}$.
- 471 ■ PSP: number of item types $n = 10$, horizon $H \in \{100, 150, 200\}$ and fraction of time
 472 periods with a demand $\rho \in \{0.9, 0.95, 1\}$.
- 473 ■ ALP: number of aircrafts $n \in \{25, 50, 75, 100\}$, number of runways $r \in \{1, 2, 3, 4\}$, number
 474 of aircraft classes $c = 4$ and mean inter-arrival time $40/r$ for generating the target landing
 475 times according to a Poisson arrival process.

476 Furthermore, the instance generation tries to emulate an increasing number of groups of actor
 477 requirements, item types and aircraft classes that lend themselves more or less to aggregation.
 478 Each instance was presolved in its aggregate state space after aggregating its data according
 479 to k -means clustering for PSP and ALP and a custom hierarchical clustering for TalentSched
 480 that tries to maximize the remaining costs induced by the actor requirements. TalentSched
 481 instances can be presolved exactly with 20 aggregate scenes and PSP instances similarly with
 482 5 aggregate item types. On the other hand, not all ALP instances reduced to 2 aggregate
 483 aircraft have a reasonable number of states so we employ a relaxed DD with maximum width
 484 10000 for the presolving part instead. Note that the present approach does not compete
 485 with the state-of-the-art for TalentSched as it lacks much of the custom symmetry-breaking
 486 logic introduced in [17] and similarly for ALP regarding the dominance-breaking constraints
 487 presented in [28]. Six different configurations were created by combining the default DD-based
 488 solved DDO on one hand and a version using only restricted DDs and no relaxed DDs,
 489 denoted rDDO, on the other hand, with the aggregation-based bounds (AggB) and heuristics
 490 (AggH). Ten minutes were allotted for each configuration to solve each instance.

491 Figure 3 presents the cumulative number of instances solved with respect to the solving
 492 time. For TalentSched, it appears that any configuration of rDDO performs better than any
 493 of DDO. This suggests that the bounds provided by the relaxed DDs are looser than the RLB
 494 while being more expensive to compute. It confirms our intuition that the state merging
 495 scheme yields bounds with a limited impact for some problems, probably because the state
 496 information gets very dilute when many states are merged together. In this case, the RLB
 497 computation is also quite involved – see [17]. Still, adding the AggB and the AggH to either
 498 configurations improves the results by a small margin, although not that significant. This
 499 can be contrasted with the results obtained for the two other problems, which show a clear
 500 improvement when the AggB and the AggH are added to either configurations. Furthermore,
 501 in cases where rDDO alone yields the worst results, incorporating AggB leads to results that
 502 are similar to or better than those achieved by DDO. Combining it with the AggH performs
 503 better than DDO in both cases and almost equally well than DDO+AggB+AggH.



■ **Figure 4** Comparison of the end gap obtained for each instance by DDO and DDO+AggB+AggH.



■ **Figure 5** Comparison of the value of the first solution found by DDO and DDO+AggB+AggH, and of the iteration at which the solution is found for ALP.

504 The impact of the AggB and the AggH can also be measured in terms of end gap $\frac{UB-LB}{UB}$.
 505 Figure 4 compares the end gap obtained for each instance by DDO and DDO+AggB+AggH.
 506 It shows that except for a few instances, DDO+AggB+AggH is always closer to terminating
 507 the search than DDO, especially for PSP. To validate the relevance of the AggH, we also
 508 compare the value of the first solution found by DDO and DDO+AggB+AggH on Figure 5(a).
 509 For TalentSched and PSP, the quality of the first solution is always better when using the
 510 AggH. However, there is no clear trend for the ALP. Unlike TalentSched and PSP, for which
 511 a solution is always found at the first iteration, the landing time windows of ALP make it
 512 difficult to find a feasible solution. This explains both the end gaps close to one in Figure 4
 513 and the ∞ values in Figure 5(a), which represent the absence of a feasible solution. We thus
 514 compare on Figure 5(b) the iteration at which the first solution is found. We observe that
 515 DDO+AggB+AggH finds a feasible solution much earlier than DDO in most cases. This
 516 showcases well the benefits of a node selection heuristic with a more global awareness.

517 **5 Conclusion**

518 This paper explained how ideas from aggregate dynamic programming can be incorporated
 519 in DD-based optimization solvers. We proposed to derive lower bounds and node selection
 520 heuristics from a pre-solved aggregate version of the original problem at hand, and explained
 521 how these can be seamlessly added to the DD-based optimization framework. Computational
 522 experiments on three different problems showed that they provide lower bounds that further
 523 strengthen the current approach, and that could even be used as a replacement for relaxed
 524 DDs in some cases. Furthermore, the aggregation-based node selection heuristics were shown
 525 very valuable as they manage to steer the compilation of relaxed DDs toward better solutions
 526 earlier in the search. When applying this idea to a highly constrained problem, the heuristics
 527 proved to quickly lead to feasible solutions that were hard to find otherwise. These results
 528 suggest that aggregation-based bounds and heuristics capture global problem structures well,
 529 as opposed to the greedy *MinLP* heuristic traditionally used to compile approximate DDs.

530 — References —

- 531 1 Henrik Reif Andersen, Tarik Hadzic, John N Hooker, and Peter Tiedemann. A constraint
532 store based on multivalued decision diagrams. In *International Conference on Principles and*
533 *Practice of Constraint Programming*, pages 118–132. Springer, 2007.
- 534 2 Sven Axsäter. State aggregation in dynamic programming—an application to scheduling of
535 independent jobs on parallel processors. *Operations Research Letters*, 2(4):171–176, 1983.
- 536 3 James C Bean, John R Birge, and Robert L Smith. Aggregation in dynamic programming.
537 *Operations Research*, 35(2):215–220, 1987.
- 538 4 Richard Bellman. The theory of dynamic programming. *Bulletin of the American Mathemat-*
539 *ical Society*, 60(6):503–515, 11 1954. URL: [https://projecteuclid.org/443/euclid.bams/](https://projecteuclid.org/443/euclid.bams/1183519147)
540 [1183519147](https://projecteuclid.org/443/euclid.bams/1183519147).
- 541 5 David Bergman and Andre A. Cire. On finding the optimal bdd relaxation. In Domenico
542 Salvagnin and Michele Lombardi, editors, *Integration of AI and OR Techniques in Constraint*
543 *Programming*, volume 10335 of *LNCS*, pages 41–50. Springer, 2017.
- 544 6 David Bergman, Andre A Cire, Ashish Sabharwal, Horst Samulowitz, Vijay Saraswat, and
545 Willem-Jan van Hoeve. Parallel combinatorial optimization with decision diagrams. In *Integ-*
546 *ration of AI and OR Techniques in Constraint Programming: 11th International Conference,*
547 *CPAIOR 2014, Cork, Ireland, May 19-23, 2014. Proceedings 11*, pages 351–367. Springer,
548 2014.
- 549 7 David Bergman, Andre A Cire, Willem-Jan van Hoeve, and John N Hooker. Variable ordering
550 for the application of bdds to the maximum independent set problem. In *International*
551 *conference on integration of artificial intelligence (AI) and operations research (OR) techniques*
552 *in constraint programming*, pages 34–49. Springer, 2012.
- 553 8 David Bergman, Andre A Cire, Willem-Jan van Hoeve, and John N Hooker. Optimization
554 bounds from binary decision diagrams. *INFORMS Journal on Computing*, 26(2):253–268,
555 2014.
- 556 9 David Bergman, Andre A Cire, Willem-Jan van Hoeve, and John N Hooker. Discrete
557 optimization with decision diagrams. *INFORMS Journal on Computing*, 28(1):47–66, 2016.
- 558 10 David Bergman, Andre A Cire, Willem-Jan van Hoeve, and Tallys Yunes. Bdd-based heuristics
559 for binary optimization. *Journal of Heuristics*, 20(2):211–234, 2014.
- 560 11 Quentin Cappart, David Bergman, Louis-Martin Rousseau, Isabeau Prémont-Schwarz, and
561 Augustin Parjadis. Improving variable orderings of approximate decision diagrams using
562 reinforcement learning. *INFORMS Journal on Computing*, 34(5):2552–2570, 2022.
- 563 12 Quentin Cappart, Emmanuel Goutier, David Bergman, and Louis-Martin Rousseau. Improv-
564 ing optimization bounds using machine learning: Decision diagrams meet deep reinforcement
565 learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages
566 1443–1451, 2019.
- 567 13 Margarita P Castro, Andre A Cire, and J Christopher Beck. An mdd-based lagrangian
568 approach to the multicommodity pickup-and-delivery tsp. *INFORMS Journal on Computing*,
569 32(2):263–278, 2020.
- 570 14 Margarita P Castro, Chiara Piacentini, Andre Augusto Cire, and J Christopher Beck. Solving
571 delete free planning with relaxed decision diagram based heuristics. *Journal of Artificial*
572 *Intelligence Research*, 67:607–651, 2020.
- 573 15 Edmund M Clarke, Orna Grumberg, and David E Long. Model checking and abstraction.
574 *ACM transactions on Programming Languages and Systems (TOPLAS)*, 16(5):1512–1542,
575 1994.
- 576 16 Vianney Coppé, Xavier Gillard, and Pierre Schaus. Decision diagram-based branch-and-bound
577 with caching for dominance and suboptimality detection, 2023. [arXiv:2211.13118](https://arxiv.org/abs/2211.13118).
- 578 17 Maria Garcia de la Banda, Peter J Stuckey, and Geoffrey Chu. Solving talent scheduling with
579 dynamic programming. *INFORMS Journal on Computing*, 23(1):120–137, 2011.

- 580 18 Xavier Gillard. *Discrete optimization with decision diagrams: design of a generic solver,*
581 *improved bounding techniques, and discovery of good feasible solutions with large neighborhood*
582 *search*. PhD thesis, UCL-Université Catholique de Louvain, 2022.
- 583 19 Xavier Gillard, Vianney Coppé, Pierre Schaus, and André Augusto Cire. Improving the filtering
584 of branch-and-bound mdd solver. In *International Conference on Integration of Constraint*
585 *Programming, Artificial Intelligence, and Operations Research*, pages 231–247. Springer, 2021.
- 586 20 Xavier Gillard and Pierre Schaus. Large neighborhood search with decision diagrams. In
587 *International Joint Conference on Artificial Intelligence*, 2022.
- 588 21 Xavier Gillard, Pierre Schaus, and Vianney Coppé. Ddo, a generic and efficient framework
589 for mdd-based optimization. In *Proceedings of the Twenty-Ninth International Conference on*
590 *International Joint Conferences on Artificial Intelligence*, pages 5243–5245, 2021.
- 591 22 Jaime E Gonzalez, Andre A Cire, Andrea Lodi, and Louis-Martin Rousseau. Integrated
592 integer programming and decision diagram search tree with an application to the maximum
593 independent set problem. *Constraints*, pages 1–24, 2020.
- 594 23 John N. Hooker. Improved job sequencing bounds from decision diagrams. In Thomas Schiex
595 and Simon de Givry, editors, *Principles and Practice of Constraint Programming*, volume
596 11802 of *LNCS*, pages 268–283. Springer, 2019.
- 597 24 Matthias Horn, Johannes Maschler, Günther R Raidl, and Elina Rönnberg. A*-based construc-
598 tion of decision diagrams for a prize-collecting scheduling problem. *Computers & Operations*
599 *Research*, 126:105125, 2021.
- 600 25 Alan J. Hu. *Techniques for efficient formal verification using binary decision diagrams*. PhD
601 thesis, Stanford University, Department of Computer Science, 1995.
- 602 26 Anthony Karahalios and Willem-Jan van Hoeve. Variable ordering for decision diagrams: A
603 portfolio approach. *Constraints*, 27(1):116–133, 2022.
- 604 27 C.-Y. Lee. Representation of switching circuits by binary-decision programs. *The Bell System*
605 *Technical Journal*, 38(4):985–999, 1959.
- 606 28 Alexander Lieder, Dirk Briskorn, and Raik Stolletz. A dynamic programming approach for
607 the aircraft landing problem with aircraft classes. *European Journal of Operational Research*,
608 243(1):61–69, 2015.
- 609 29 Shin-ichi Minato. *Binary decision diagrams and applications for VLSI CAD*, volume 342.
610 Springer Science & Business Media, 1995.
- 611 30 Isaac Rudich, Quentin Cappart, and Louis-Martin Rousseau. Peel-And-Bound: Generating
612 Stronger Relaxed Bounds with Multivalued Decision Diagrams. In Christine Solnon, editor,
613 *28th International Conference on Principles and Practice of Constraint Programming (CP*
614 *2022)*, volume 235 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 35:1–
615 35:20. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022.
- 616 31 Christian Tjandraatmadja. *Decision Diagram Relaxations for Integer Programming*. PhD
617 thesis, Carnegie Mellon University Tepper School of Business, 2018.
- 618 32 Christian Tjandraatmadja and Willem-Jan van Hoeve. Target cuts from relaxed decision
619 diagrams. *INFORMS Journal on Computing*, 31(2):285–301, 2019. doi:10.1287/ijoc.2018.
620 0830.