

1 Large Neighborhood Beam Search for 2 Domain-Independent Dynamic Programming

3 Ryo Kuroiwa  

4 Department of Mechanical and Industrial Engineering, University of Toronto, 5 King's College Rd,
5 Toronto, ON M5S 3G8, Canada

6 J. Christopher Beck  

7 Department of Mechanical and Industrial Engineering, University of Toronto, 5 King's College Rd,
8 Toronto, ON M5S 3G8, Canada

9 — Abstract —

10 Large neighborhood search (LNS) is an algorithmic framework that removes a part of a solution
11 and performs search in the induced search space to find a better solution. While LNS shows strong
12 performance in constraint programming, little work has combined LNS with state space search. We
13 propose large neighborhood beam search (LNBS), a combination of LNS and state space search.
14 Given a solution path, LNBS removes a partial path between two states and then performs beam
15 search to find a better partial path. We apply LNBS to domain-independent dynamic programming
16 (DIDP), a recently proposed generic framework for combinatorial optimization based on dynamic
17 programming. We empirically show that LNBS finds better quality solutions than a state-of-the-art
18 DIDP solver in five out of nine benchmark problem types with a total of 8570 problem instances. In
19 particular, LNBS shows a significant improvement over the existing state-of-the-art DIDP solver in
20 routing and scheduling problems.

21 **2012 ACM Subject Classification** Computing methodologies → Discrete space search

22 **Keywords and phrases** Large Neighborhood Search, Dynamic Programming, State Space Search,
23 Combinatorial Optimization

24 **Digital Object Identifier** 10.4230/LIPIcs.CP.2023.7

25 **Supplementary Material** *Software (Source Code)*: [https://github.com/domain-independent-dp/
26 didp-rs/releases/tag/lpbs-cp23](https://github.com/domain-independent-dp/didp-rs/releases/tag/lpbs-cp23)

27 **Funding** This research is supported by the Natural Sciences and Engineering Research Council of
28 Canada.

29 **1** Introduction

30 In constraint programming (CP), large neighborhood search (LNS) [34] achieves strong per-
31 formance in solving combinatorial optimization problems such as routing [24] and scheduling
32 problems [27]. LNS is an algorithmic framework that removes a part of a solution and then
33 performs search in the induced partial search space (neighborhood) to find a better solution.
34 Typically, LNS uses tree search to find a better solution in a partial search space, where each
35 search node represents a partial assignment of decision variables, and a solution of a problem
36 corresponds to a leaf node, where all variables are assigned values.

37 Dynamic programming (DP) is a powerful method for multiple combinatorial optimization
38 problems [14, 15], and the hybridization of CP, decision diagrams, and DP is a topic of active
39 research [1, 22, 23, 5, 28, 19, 17]. Recently, domain-independent dynamic programming
40 (DIDP), a model-based paradigm for combinatorial optimization based on DP, has been
41 proposed [25]. In DIDP, a model of a problem is represented by a state transition system. A
42 solution corresponds to a path in a state space graph, where each vertex represents a state
43 and each edge represents a transition between two states. The current state-of-the-art DIDP



© Ryo Kuroiwa and J. Christopher Beck;

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. 7; pp. 7:1–7:22



Leibniz International Proceedings in Informatics

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

44 solver is complete anytime beam search (CABS) [26], which is based on beam search, an
 45 algorithm searches for a path in a state space graph by maintaining a fixed number of states
 46 at a time.

47 In this paper, we propose large neighborhood beam search (LNBS), a combination of
 48 LNS and beam search in a state space graph. LNBS tries to improve a solution path by
 49 removing a partial path between two states and then performing beam search to find a
 50 better partial path. While LNBS has the freedom to select a neighborhood (i.e., a partial
 51 path to remove), we propose a strategy that dynamically adjusts the size of a neighborhood
 52 based on a multi-armed bandit problem. With our strategy, LNBS is complete, i.e., it finds
 53 and proves an optimal solution given enough time, but, of course, it is aimed at problems
 54 where its solution quality is more important than proved optimality. We implement LNBS
 55 for DIDP and empirically evaluate its performance. The experimental results show that
 56 LNBS outperforms CABS in five out of nine benchmark problem types in terms of solution
 57 quality. In addition, LNBS performs better than a commercial CP solver, which uses LNS
 58 [27], in seven problems while CABS is better than CP in six problems. Since LNBS performs
 59 particularly well in routing and scheduling problems, we also investigate the reason for this
 60 performance and gain insight from empirical analysis.

61 **2 Background**

62 We first introduce domain-independent dynamic programming and complete anytime beam
 63 search. Then, we present large neighborhood search (LNS). We also describe LNS with
 64 decision diagrams [18], a recently proposed method for combinatorial optimization that can
 65 be considered a combination of LNS and state space search.

66 **2.1 Domain-Independent Dynamic Programming**

67 A combinatorial optimization problem is to find a set of discrete decisions, e.g., a permutation,
 68 to minimize or maximize an objective function. In dynamic programming (DP), a problem is
 69 recursively formulated by decomposing it into subproblems, represented by *states*, and the
 70 optimal objective value of each subproblem is represented by the *value function*, which maps
 71 a state to a real number.

72 Domain-independent dynamic programming (DIDP) is a model-based paradigm for
 73 combinatorial optimization based on DP [25]. In DIDP, a DP formulation of a combinatorial
 74 optimization problem is defined by a state-transition system in Dynamic Programming
 75 Description Language (DyPDL). A DyPDL model is a seven-tuple $\langle \mathcal{V}, S^0, \mathcal{K}, \mathcal{T}, \mathcal{B}, \mathcal{C}, h \rangle$
 76 consisting of *state variables* \mathcal{V} , the *target state* S^0 , *constants* \mathcal{K} , *transitions* \mathcal{T} , *base cases* \mathcal{B} ,
 77 *state constraints* \mathcal{C} , and the *dual bound* h .

78 A state variable $v \in \mathcal{V}$ has a type of *set*, *element*, or *numeric*. Each set and element
 79 variable v_i is associated with a set of objects $N_i = \{0, \dots, n_i - 1\}$. The domain of a set
 80 variable v_i is 2^{N_i} , and the domain of an element variable v_i is N_i . A numeric variable takes
 81 a real value. A constant in \mathcal{K} is a value independent of the state variables.

82 A state is a complete value assignment to the state variables, and the target state S^0
 83 is a state. We denote the value of a state variable v_i in a state S by $S[v_i]$. The value of a
 84 state S is represented by the value function $V(S)$, and the objective of a DyPDL model is
 85 to compute the value of the target state, $V(S^0)$. We can define dominance between states
 86 based on state variables. If a state S dominates another state S' , denoted by $S' \preceq S$, then
 87 $V(S)$ is equal or better (less/greater for minimization/maximization) than $V(S')$. For the
 88 details of dominance in DyPDL, please refer to previous work [25].

89 A transition $\tau \in \mathcal{T}$ is a four-tuple $\langle \text{eff}_\tau, \text{cost}_\tau, \text{pre}_\tau, \text{forced}_\tau \rangle$ defining how a state S is
 90 transformed to a successor state (a subproblem) $S[\tau]$, and how $V(S)$ is computed based
 91 on the value of $V(S[\tau])$. The set of *effects* $\text{eff}_\tau = \{(v_i, e_{v_i}) \mid v_i \in \mathcal{V}\}$ defines how each
 92 state variable v_i is updated by an *expression* e_{v_i} . The expression e_{v_i} consists of predefined
 93 operations on state variables and constants and returns $S[\tau][v_i]$ given S . For example, for a
 94 set variable U , $U \setminus \{i\}$ is an expression representing removing an element i from U , which
 95 can be used as e_U and results in $S[\tau][U] = S[U] \setminus \{i\}$. The *cost expression* cost_τ is an
 96 expression that takes $V(S[\tau])$ in addition to S and returns a real number $\text{cost}_\tau(V(S[\tau]), S)$.
 97 The preconditions pre_τ are conditions on the state variables, i.e., expressions returning a
 98 binary value \top or \perp . The preconditions define when the transition is *applicable*. For example,
 99 if $\text{pre}_\tau = \{i \in U\}$, then τ is applicable in S iff $i \in S[U]$. The flag forced_τ is a boolean value
 100 indicating whether the transition is a *forced transition*. Let $\mathcal{T}(S)$ be the set of applicable
 101 transitions in state S . If forced transitions are applicable, the first defined one τ is selected,
 102 and all other forced and non-forced transitions are ignored, i.e., $\mathcal{T}(S) = \{\tau\}$. The value of
 103 $V(S)$ is computed by taking the best $\text{cost}_\tau(V(S[\tau]), S)$ over all $\tau \in \mathcal{T}(S)$.

104 Base cases \mathcal{B} are sets of conditions. For any $B \in \mathcal{B}$, if a state S satisfies all conditions
 105 in B , denoted by $S \models B$, then it is called a *base state*, and $V(S)$ is defined non-recursively
 106 by an expression $e_B(S)$. State constraints \mathcal{C} are conditions that must be satisfied by all
 107 states. If one of the state constraints $c \in \mathcal{C}$ is violated, denoted by $S \not\models c$, then $V(S) = \infty$
 108 ($V(S) = -\infty$) for minimization (maximization). The dual bound $h(S)$ is a lower (upper)
 109 bound on $V(S)$ for minimization (maximization).

110 Overall, if the problem is minimization, the DP formulation is defined as follows.

$$111 \quad \text{compute } V(S^0) \tag{1}$$

$$112 \quad \text{s.t. } V(S) = \begin{cases} \min_{\tau \in \mathcal{T}(S)} \text{cost}_\tau(V(S[\tau]), S) & \text{if } S \models \mathcal{C} \wedge \forall B \in \mathcal{B}, S \not\models B \\ e_B(S) & \text{if } S \models \mathcal{C} \wedge \exists B \in \mathcal{B}, S \models B \\ \infty & \text{if } S \not\models \mathcal{C} \end{cases} \tag{2}$$

$$113 \quad V(S) \leq V(S') \quad \text{if } S' \preceq S \tag{3}$$

$$114 \quad V(S) \geq h(S). \tag{4}$$

116 The first line states that the optimal objective value is $V(S^0)$. Equation (2) recursively
 117 defines the value function V . Inequalities (3) and (4) are bounds on the value function. For
 118 maximization, we replace min with max in the first line of Equation (2) and swap \leq and \geq
 119 in Inequalities (3) and (4). A *solution* for the DP formulation is a sequence of transitions
 120 that transforms the target state S^0 into a base state. Concretely, for a sequence of transitions
 121 $x = \langle x_1, \dots, x_n \rangle$, let $S^{i+1} = S^i[x_{i+1}]$ for $i = 0, \dots, n-1$. Then, x is a solution if $x_{i+1} \in \mathcal{T}(S^i)$,
 122 $S^i \models \mathcal{C}$, and $\forall B \in \mathcal{B}, S^i \not\models B$ for $i = 0, \dots, n-1$, $S^n \models \mathcal{C}$, and $\exists B \in \mathcal{B}, S^n \models B$. The solution
 123 is an *optimal solution* if $\text{cost}_{\tau_i}(V(S^{i+1}), S^i) = V(S^i)$ for $i = 0, \dots, n-1$ in addition.

124 2.2 Complete Anytime Beam Search for DIDP

125 Previous research has shown that a subset of DyPDL models can be solved by cost-algebraic
 126 heuristic search [11], a generalized version of the shortest path algorithm [25]. Multiple DIDP
 127 solvers using cost-algebraic heuristic search algorithms have been proposed [25, 26]. These
 128 solvers perform *state space search*, which finds a path from the target state to a base state in
 129 a *state space graph*, a directed graph where each vertex is a state. In the state space graph,
 130 an edge from state S to $S[\tau]$ exists if $\tau \in \mathcal{T}(S)$. Following the previous work, we assume
 131 that $\text{cost}_\tau(V(S[\tau]), S)$ is expressed as $w_\tau(S) \times V(S[\tau])$ where w_τ is an expression returning

Algorithm 1 Beam Search for DyPDL.

```

1: function BEAMSEARCH( $\bar{f}$ ,  $b$ )
2:    $g(S^0) \leftarrow 0, f(S^0) \leftarrow h(S^0), x(S^0) \leftarrow \langle \rangle$ .
3:    $O \leftarrow \{S^0\}, \bar{x} \leftarrow \text{NULL}, \text{complete} \leftarrow \top$ .
4:   while  $O \neq \emptyset$  and  $\bar{x} = \text{NULL}$  do
5:      $G \leftarrow \emptyset$ . ▷ A set of states in the next layer.
6:     for all  $S \in O$  do
7:       if  $\exists B \in \mathcal{B}$  such that  $S \models B$  then ▷ A base state.
8:         if  $g(S) \times e_B(S) < \bar{f}$  then ▷ A better solution.
9:            $\bar{f} \leftarrow g(S) \times e_B(S), \bar{x} \leftarrow x(S)$ .
10:        else
11:          for all  $\tau \in \mathcal{T}(S) : S[\tau] \models \mathcal{C}$  do
12:            if  $\nexists S' \in G$  such that  $S[\tau] \preceq S'$  and  $g(S) \times w_\tau(S) \geq g(S')$  then
13:               $x(S[\tau]) \leftarrow \langle x(S); \tau \rangle$ .
14:               $g(S[\tau]) \leftarrow g(S) \times w_\tau(S), f(S[\tau]) \leftarrow g(S[\tau]) \times h(S[\tau])$ .
15:              if  $\exists S' \in G$  such that  $S' \preceq S[\tau]$  and  $g(S[\tau]) \leq g(S')$  then
16:                 $G \leftarrow G \setminus \{S'\}$ . ▷ Remove a dominated state.
17:                if  $f(S[\tau]) < \bar{f}$  then ▷ Pruning by the primal bound.
18:                   $G \leftarrow G \cup \{S[\tau]\}$ .
19:             $O \leftarrow \{S \in G \mid f(S) < \bar{f}\}$ .
20:            if  $|O| > b$  then
21:               $O \leftarrow$  the best  $b$  states in  $G$  minimizing  $f$ .
22:               $\text{complete} \leftarrow \perp$ .
23:            if  $O \neq \emptyset$  then
24:               $\text{complete} \leftarrow \perp$ .
25:            return  $\bar{x}, \text{complete}$ .
```

132 a real value and \times is a binary operator satisfying a cost-algebra. In particular, we focus on
 133 nonnegative w_τ and binary operators $+$ or \max , i.e., $\text{cost}_\tau(V(S[\tau]), S) = w_\tau(S) + V(S[\tau])$
 134 or $\text{cost}_\tau(V(S[\tau]), S) = \max\{w_\tau(S), V(S[\tau])\}$. The weight of the edge $(S, S[\tau])$ is defined
 135 as $w_\tau(S)$, and the cost of a path from the target state S^0 , which corresponds to a sequence
 136 of the transitions $x = \langle x_1, \dots, x_n \rangle$, is $\text{cost}_x(S^0) = \prod_{i=0}^{n-1} w_{x_{i+1}}(S^i)$ where $S^{i+1} = S^i[x_{i+1}]$
 137 for $i = 0, \dots, n-1$. As each edge weight is nonnegative, the cost of a path is nonnegative
 138 and non-decreasing in length. In this paper, we focus on minimization while DyPDL and
 139 cost-algebraic heuristic search can handle both minimization and maximization.

140 The state-of-the-art cost-algebraic heuristic search solver is complete anytime beam search
 141 (CABS) [36, 26]. CABS performs beam search, which searches at most b states in the *open*
 142 *list* O at each layer of the state space graph. We show the pseudo-code of beam search in
 143 Algorithm 1. In addition to a DyPDL model and b , beam search takes the primal bound
 144 \bar{f} as an input, which is the best-known objective value and could be infinity. With each
 145 state S , the best path $x(S)$ from S^0 and its cost $g(S) = \text{cost}_{x(S)}(S^0)$ (the *g-value*) are
 146 maintained in lines 13 and 14. Starting from $O = \{S^0\}$, beam search processes a state S in
 147 O . If S is a base state, and the best path to S has a better cost than \bar{f} in line 8, then \bar{f}
 148 and the solution \bar{x} are updated. Otherwise, $S[\tau]$ is added to the candidate set G for each
 149 transition $\tau \in \mathcal{T}(S)$, which is called the *expansion* of S , and we say that S is expanded
 150 (lines 11–18). When expanding $S \in O$, if there exists a state $S' \in G$ that dominates $S[\tau]$
 151 and $g(S') \leq g(S) \times w_\tau(S)$, then $S[\tau]$ is not added to G . In addition to $g(S[\tau])$, the priority

152 $f(S[\tau]) = g(S[\tau]) \times h(S[\tau])$ (the *f-value*) is computed in line 14, which is a lower bound on
 153 the optimal path cost from S^0 to a base state via $S[\tau]$. If $f(S[\tau]) \geq \bar{f}$, the corresponding
 154 path does not lead to an improved solution, so $S[\tau]$ is not added to G in line 18. After
 155 expanding all states, O is updated to the best b states in G according to f in lines 20-21.
 156 This procedure is repeated until a solution whose cost is better than the given primal bound
 157 is found, or no successor states are generated. The variable `complete` maintains whether the
 158 search is complete. If states in G are pruned due to the beam width b (line 22), or O is not
 159 empty when a solution is found (line 24), there may exist a better solution, so the search is
 160 not complete. If `complete` = \top , then \bar{x} is the optimal solution if it is not NULL, or the model
 161 is infeasible if \bar{x} = NULL. CABS performs a sequence of beam search with exponentially
 162 increasing beam width $b = 1, 2, 4, \dots$ using the best objective value found so far as the primal
 163 bound \bar{f} until the optimality of the best solution or the infeasibility is proved.


164 2.3 Large Neighborhood Search

165 Large neighborhood search (LNS) iteratively removes a part of a solution and solves the
 166 resulting subproblem (neighborhood) [34]. A solution for a CP problem is represented as a
 167 complete value assignment to decision variables. Given a solution, LNS removes a subset of
 168 the value assignments and solves the subproblem where the remaining variables are fixed to
 169 the values assigned in the original solution. Typically, a tree search algorithm is used. In
 170 a tree search algorithm, a search node is a partial value assignment to decision variables,
 171 successor nodes are generated by assigning a value to an unassigned variable, and a solution
 172 corresponds to a leaf node, where all variables are assigned values.

173 2.3.1 Large Neighborhood Search with Decision Diagrams

174 For combinatorial optimization, recent work proposed LNS with decision diagrams (DD-LNS)
 175 [18]. Although DD-LNS was not explicitly framed as state space search, we interpret it as a
 176 state space search algorithm. While DD-LNS is independent of DIDP, it also uses the DP
 177 formulation of a problem as input while assuming that the solution has n transitions. Given
 178 a sequence of transitions $\langle x_1, \dots, x_n \rangle$, DD-LNS keeps the first d transitions, $\langle x_1, \dots, x_d \rangle$, and
 179 searches for the remaining $n - d$ transitions. To find such a sequence, DD-LNS constructs
 180 a decision diagram (DD), a directed graph where nodes are partitioned into layers. In the
 181 constructed DD, each vertex corresponds to a state, and each edge corresponds to a transition,
 182 so it is a state space graph. The first layer in the DD contains only the node corresponding to
 183 S^d , where $S^i = S^{i-1} \llbracket x_i \rrbracket$ for $i = 1, \dots, d$. DD-LNS iteratively constructs a layer by applying
 184 transitions to the states in the current layer until reaching layer $n - d + 1$.

185 Because constructing an exact DD is intractable, DD-LNS constructs a restricted DD,
 186 which keeps a subset of states in the exact DD. In each layer, states satisfying certain
 187 conditions are kept, some of which are selected randomly. Let the number of such states be
 188 K . If K is smaller than a parameter W , from the remaining states, DD-LNS also keeps the
 189 best $W - K$ states that minimize a priority function, the rough lower bound (RLB). The
 190 RLB of a state is a lower bound on the cost of a solution path via that state, which is the
 191 same as the *f-value*. If the RLB of a state is larger than the best solution cost, the state is
 192 removed from the DD as it does not lead to a better solution. Therefore, the procedure of
 193 constructing a restricted DD can be considered beam search with randomization. Indeed,
 194 the authors acknowledged that DD-LNS is a hybridization of LNS and beam search [18].
 195 DD-LNS decreases d by 1 if a better solution is not found with d , starting from $d = n - 2$ and
 196 restarting from $d = n - 2$ if $d = 0$. When $d = 0$ and the restricted DD keeps all states except



■ **Figure 1** Partial state space graph induced by the prefix $\langle 1 \rangle$ and the suffix $\langle 4 \rangle$ (highlighted in red) in Example 1. The current partial path is highlighted in blue and an alternative partial path is highlighted in green. Dashed transitions conflict with the suffix (explained in Section 3.2).

197 for those removed based on RLB, then it is the exact DD, so DD-LNS proves the optimality
 198 of the solution. Since DD-LNS can be interpreted as a state space search algorithm and
 199 is designed for combinatorial optimization, we implement it for DIDP and experimentally
 200 compare it with our method.

201 3 Large Neighborhood Beam Search

202 We start with a simple idea of LNS for state space search: given a solution path, $x =$
 203 $\langle x_1, \dots, x_n \rangle$ which connects states $\langle S^0, \dots, S^n \rangle$, we remove a partial path $\langle x_i, \dots, x_{i+d-1} \rangle$
 204 and search for a better partial path from S^{i-1} to S^{i+d-1} . If we find a better solution
 205 $\langle x_1, \dots, x_{i-1}, x'_i, \dots, x'_{i+d-1}, x_{i+d}, \dots, x_n \rangle$, we repeat this procedure with the new solution. While
 206 the overview of the algorithm is simple, there are design choices on how to select a partial
 207 path to remove and how to search for a better partial path. The novelty of our method
 208 compared to existing methods arises from such choices in addition to the fact that it is used
 209 for DIDP. First, we describe the modifications of beam search for DyPDL to search for a
 210 partial path from S^{i-1} to S^{i+d-1} . Then, we propose strategies to select a partial path to
 211 remove.

212 3.1 Beam Search for DyPDL in a Partial State Space Graph

213 We want to find a path from S^{i-1} to S^{i+d-1} instead of from S^0 to a base state. We could
 214 modify line 8 in Algorithm 1 so that it checks if $S = S^{i+d-1}$ instead of $\exists B \in \mathcal{B}, S \models B$.
 215 However, in DyPDL, it may not be desirable as shown in the following example.

216 ► **Example 1.** Consider the following DP formulation, where $U \subseteq \{0, 1, 2, 3, 4\}$ is a set
 217 variable, $k \in \{0, 1, 2, 3, 4\}$ is an element variable, and c_{lj} for $l, j \in \{0, 1, 2, 3, 4\}$ is a constant.

218 compute $V(\{1, 2, 3, 4\}, 0)$

$$219 \quad V(U, k) = \begin{cases} \min_{j \in U} c_{kj} + V(U \setminus \{j\}, j) & \text{if } U \neq \emptyset \\ 0 & \text{if } U = \emptyset. \end{cases}$$

220

221 Each transition in the DyPDL model has precondition $j \in U$ and effect $(U, U \setminus \{j\})$ for
 222 some j , and so we denote each transition by j . Each solution corresponds to a permutation
 223 of the transitions 1, 2, 3, 4. A solution $\langle 1, 2, 3, 4 \rangle$ connects a sequence of states $\langle (\{1, 2, 3, 4\}, 0),$
 224 $(\{2, 3, 4\}, 1), (\{3, 4\}, 2), (\{4\}, 3), (\emptyset, 4) \rangle$. Consider removing $\langle 2, 3 \rangle$ from the solution. We
 225 visualize the partial state space graph in Figure 1. An algorithm tries to find a path from

■ **Algorithm 2** Beam Search in a partial state space graph.

```

1: function BEAMSEARCHFORPARTIALPATH( $\bar{f}$ ,  $b$ , prefix, suffix)
2:    $\hat{S} \leftarrow S^0 \llbracket \text{prefix} \rrbracket$ ,  $g(\hat{S}) \leftarrow \text{cost}_{\text{prefix}}(S^0)$ ,  $f(\hat{S}) \leftarrow g(\hat{S}) \times h(\hat{S})$ ,  $x(\hat{S}) \leftarrow \text{prefix}$ .
3:    $O \leftarrow \{\hat{S}\}$ ,  $\bar{x} \leftarrow \text{NULL}$ , complete  $\leftarrow \top$ .
4:   while  $O \neq \emptyset$  and  $\bar{x} = \text{NULL}$  do
5:      $G \leftarrow \emptyset$ . ▷ A set of states in the next layer.
6:     for all  $S \in O$  do
7:        $S' \leftarrow S$ , success  $\leftarrow \perp$ .
8:       if  $\exists B \in \mathcal{B}$  and  $S' \models B$  then
9:         success  $\leftarrow \top$ . ▷ A base state, success.
10:      else
11:        for  $\tau \leftarrow \text{suffix}_1, \dots, \text{suffix}_{n-i-d+1}$  do ▷ Rollout of the suffix.
12:          if  $S' \not\models \text{pre}_\tau$  then
13:            break. ▷ Preconditions are not satisfied, fail.
14:             $S' \leftarrow S' \llbracket \tau \rrbracket$ ,  $g(S') \leftarrow g(S') \times w_\tau(S')$ ,  $x(S') \leftarrow \langle x(S'); \tau \rangle$ .
15:            if  $S' \not\models \mathcal{C}$  then
16:              break. ▷ State constraints are not satisfied, fail.
17:            if  $\exists B \in \mathcal{B}$  such that  $S' \models B$  then
18:              success  $\leftarrow \top$ . ▷ A base state, success.
19:              break.
20:            if success then
21:              if  $g(S') \times e_B(S') < \bar{f}$  then ▷ A better solution.
22:                 $\bar{f} \leftarrow g(S') \times e_B(S')$ ,  $\bar{x} \leftarrow x(S')$ .
23:            else
24:              for all  $\tau \in \mathcal{T}(S) : S \llbracket \tau \rrbracket \models \mathcal{C}$  do
25:                if  $\nexists S' \in G$  such that  $S \llbracket \tau \rrbracket \preceq S'$  and  $g(S) \times w_\tau(S) \geq g(S')$  then
26:                   $x(S \llbracket \tau \rrbracket) \leftarrow \langle x(S); \tau \rangle$ .
27:                   $g(S \llbracket \tau \rrbracket) \leftarrow g(S) \times w_\tau(S)$ ,  $f(S \llbracket \tau \rrbracket) \leftarrow g(S \llbracket \tau \rrbracket) \times h(S \llbracket \tau \rrbracket)$ .
28:                  if  $\exists S' \in G$ ,  $S' \preceq S \llbracket \tau \rrbracket \wedge g(S \llbracket \tau \rrbracket) \leq g(S')$  then
29:                     $G \leftarrow G \setminus \{S'\}$ . ▷ Remove a dominated state.
30:                  if  $f(S \llbracket \tau \rrbracket) < \bar{f}$  then ▷ Pruning by the primal bound.
31:                     $G \leftarrow G \cup \{S \llbracket \tau \rrbracket\}$ .
32:                 $O \leftarrow \{S \in G \mid f(S) < \bar{f}\}$ .
33:              if  $|G| > b$  then
34:                 $O \leftarrow$  the best  $b$  states in  $G$  minimizing  $f$ .
35:              complete  $\leftarrow \perp$ .
36:            if  $O \neq \emptyset$  then
37:              complete  $\leftarrow \perp$ .
38:          return  $\bar{x}$ , complete.

```

226 $(\{2, 3, 4\}, 1)$ to $(\{4\}, 3)$. The original one, $\langle 2, 3 \rangle$, is only the path. However, a partial path
 227 $\langle 3, 2 \rangle$ from $(\{2, 3, 4\}, 1)$ to $(\{4\}, 2)$ also results in a valid solution $\langle 1, 3, 2, 4 \rangle$.

228 Considering the above example, instead of focusing on a partial path to a state, we focus on
 229 a partial path to a *suffix* of the solution path. Given a solution path $\langle x_1, \dots, x_n \rangle$, if we remove a
 230 partial path $\langle x_i, \dots, x_{i+d-1} \rangle$, then $\langle x_i, \dots, x_{i-1} \rangle$ is the prefix, and $\langle x_{i+d}, \dots, x_n \rangle$ is the suffix. For
 231 a partial path $\langle x'_i, \dots, x'_{i+d'-1} \rangle$, we want to check if $\langle x_1, \dots, x_{i-1}, x'_i, \dots, x'_{i+d'-1}, x_{i+d}, \dots, x_n \rangle$ is

■ **Algorithm 3** Large Neighborhood Beam Search (LNBS).

Input: initial feasible solution \bar{x} .

Output: solution \bar{x} and if the optimality or infeasibility is proved.

```

1: while the time limit is not reached do
2:    $n \leftarrow |\bar{x}|, \bar{f} \leftarrow \text{cost}_{\bar{x}}(S^0)$ .
3:   Select  $d$  such that  $2 \leq d \leq n$ .
4:   Select  $i$  such that  $1 \leq i \leq n - d + 1$ .
5:   Select beam width  $b$ .
6:    $\text{prefix} \leftarrow \langle x_1, \dots, x_{i-1} \rangle, \text{suffix} \leftarrow \langle x_{i+d}, \dots, x_n \rangle$ .
7:    $x, \text{complete} \leftarrow \text{BEAMSEARCHFORPARTIALPATH}(\bar{f}, b, \text{prefix}, \text{suffix})$ 
8:   if  $x \neq \text{NULL}$  then
9:      $\bar{x} \leftarrow x$ .
10:  if  $i = 0 \wedge d = n \wedge \text{complete}$  then
11:    return  $\bar{x}, \top$ .
12: return  $\bar{x}, \perp$ .

```

232 a valid solution. Therefore, for a state S found by a search algorithm, we perform a rollout
233 of the suffix from S and check if each of resulting states satisfies the state constraints and a
234 base case. We show the modified version of beam search in Algorithm 2. This algorithm
235 takes a prefix prefix and a suffix suffix as input. In line 2, we denote the state resulting from
236 applying the prefix to the target state by $\hat{S} = S^0[\text{prefix}]$ and initialize the open list O with
237 \hat{S} in line 3. In lines 8–22, the algorithm performs a rollout of the suffix from S and checks if
238 it results in a better solution. Other parts are the same as Algorithm 1.

239 We use this modified version of beam search in large neighborhood beam search (LNBS)
240 as shown in Algorithm 3. In line 2, $|\bar{x}|$ denotes the length of the current solution \bar{x} . In
241 lines 3–5, LNBS selects parameters d , i , and b . In line 7, LNBS performs beam search in the
242 neighborhood. If an improving solution is found, LNBS updates the current solution \bar{x} in
243 line 9. If the searched neighborhood is the original search space, i.e., $i = 1$ and $d = n$, and
244 beam search proves the optimality or infeasibility, LNBS terminates in line 11. Therefore, if
245 it is guaranteed to select $i = 1$ and $d = n$ with sufficiently large b given enough time, LNBS
246 is guaranteed to find the optimal solution or prove the infeasibility, i.e., it is complete. CABS
247 can be considered a configuration of LNBS, where $i = 1$, $d = n$, and b increases exponentially.
248 DD-LNS can also be considered a configuration of LNBS, where i ranges from $n - 2$ to 1,
249 d is $n - i + 1$, and b is fixed to W while beam search is extended with the randomization
250 mechanism. We will describe the strategies that we use to select d , i , and b below.

251 3.2 Removing Conflicting Transitions

252 In Example 1, consider finding a partial path from a prefix $\langle 1 \rangle$, which results in state
253 $\hat{S} = (\{2, 3, 4\}, 1)$, to a suffix $\langle 4 \rangle$ using beam search. In \hat{S} , three transitions 2, 3, and 4 are
254 applicable. However, applying transition 4 does not lead to a feasible solution because it
255 is already used in the suffix and cannot be applied twice: it requires $4 \in U$ and removes 4
256 from U , but no other transition adds 4 to U , so applying 4 makes the suffix inapplicable.
257 Generalizing this example, if we know that a transition τ makes a transition τ' in the suffix
258 inapplicable, then we can ignore τ when searching for a partial path. In particular, we focus
259 on the effects of τ that add/remove an element to/from a set variable and the preconditions
260 of τ' that require the element to be/not to be in that set variable.

261 ► **Proposition 2.** *Suppose that a DyPDL model $\langle \mathcal{V}, S^0, \mathcal{K}, \mathcal{T}, \mathcal{B}, \mathcal{C}, h \rangle$ has a set variable $U \in \mathcal{V}$
 262 whose domain is 2^N , where N is a set of objects. There does not exist a solution $\langle x_1, \dots, x_n \rangle$
 263 such that any pair of $\tau = x_i$ and $\tau' = x_j$ for $1 \leq i < j \leq n$ satisfy either of the following
 264 conditions:*

- 265 1. *There exists $k \in N$ such that $(U, U \setminus \{k\}) \in \text{eff}_\tau$, $(k \in U) \in \text{pre}_{\tau'}$, and each $\tau'' \in \mathcal{T} \setminus \{x_i, x_j\}$
 266 does not change U or $(U, U \setminus \{l\}) \in \text{eff}_{\tau''}$ for some $l \in N$.*
- 267 2. *There exists $k \in N$ such that $(U, U \cup \{k\}) \in \text{eff}_\tau$, $(k \notin U) \in \text{pre}_{\tau'}$, and each $\tau'' \in$
 268 $\mathcal{T} \setminus \{x_i, x_j\}$ does not change U or $(U, U \cup \{l\}) \in \text{eff}_{\tau''}$ for some $l \in N$.*

269 **Proof.** For the first case, x_i removes k from U , and no other transition adds k to U . Since
 270 x_j requires k to be in U , once we apply x_i , we cannot apply x_j later. The other case is
 271 proved similarly. ◀

272 Before starting beam search in a neighborhood, we remove a transition τ from the model
 273 if there exists a transition τ' in the suffix such that τ and τ' satisfy one of the conditions in
 274 Proposition 2. Detecting such a pair of transitions is done once at the beginning by checking
 275 the expression trees representing the preconditions and effects of transitions.

276 3.3 Bandit-Based Depth Selection

277 Selecting the depth of a neighborhood, d , in line 3 of Algorithm 3 is non-trivial. If d
 278 is too small, it is unlikely that an improving solution exists. However, if d is too large,
 279 each neighborhood search takes a long time. We want to select d such that the total cost
 280 improvement is maximized within the time limit.

281 We formulate the depth selection as the budgeted multi-armed bandit problem with
 282 continuous random costs [35]. We have the set of depths $D \subseteq \{2, \dots, n\}$. If we select a depth
 283 $d \in D$ and perform search in line 7 of Algorithm 3, we obtain a new solution x by spending
 284 search time t . As t is assumed to take a value in $[0, 1]$ in the budgeted multi-armed bandit
 285 problem, we divide the actual time by the time limit T . If the cost $\text{cost}_x(S^0)$ is smaller than
 286 the current best solution cost \bar{f} , the reward is $r = (\bar{f} - \text{cost}_x(S^0)) / \bar{f}$. If no better solution is
 287 found, the reward is $r = 0$. We call this process a *round*, and we repeat rounds until reaching
 288 the time limit T . We do not know the reward r and time t before finishing a round, so we
 289 use random variables r_{dk} and t_{dk} representing the reward and time if depth d is used at
 290 round k . Let a be a strategy that selects a depth a_k in the round k . The number of rounds
 291 performed by a by the time limit, K_{aT} , is also a random variable. The objective is to find a
 292 strategy a that maximizes the total expected reward $\mathbb{E}[\sum_{k=1}^{K_{aT}} r_{a_k k}]$.

293 We use Budgeted-UCB [35]. At each round, if some depths in D have not been selected
 294 before, Budgeted-UCB selects one of them. Otherwise, let m_{dk} be the number of rounds
 295 where the depth d is selected up to round $k - 1$, and let \bar{r}_{dk} and \bar{t}_{dk} be the average reward
 296 and search time for d up to round $k - 1$. Budgeted-UCB selects the depth d that maximizes

$$297 \frac{\bar{r}_{dk}}{\bar{t}_{dk}} + \frac{\epsilon_{dk}}{\bar{t}_{dk}} + \frac{\epsilon_{dk}}{\bar{t}_{dk}} \frac{\min\{\bar{r}_{dk} + \epsilon_{dk}, 1\}}{\max\{\bar{t}_{dk} - \epsilon_{dk}, \lambda\}} \quad (5)$$

298 where $\epsilon_{dk} = \sqrt{\frac{2 \log(k-1)}{m_{dk}}}$ and λ is a positive lower bound of the search time of each round.

299 In practice, we initialize $D = \{2, 4, 8, \dots, 2^l, n\}$, where n is the length of the initial feasible
 300 solution and l is the maximum integer such that $2^l < n$. If we get a solution whose length
 301 n' is different from n at round k , we replace n with n' in D using $m_{n',k+1} = m_{n,k+1}$,
 302 $\bar{r}_{n',k+1} = \bar{r}_{n,k+1}$, and $\bar{t}_{n',k+1} = \bar{t}_{n,k+1}$ and ignore depths greater than n' in D . If multiple
 303 depths have not been selected before or have the same value, we select the minimum depth

among them. For λ , we use the time of the first round divided by 10 while there is no guarantee that it is a lower bound. Thus, the theoretical analysis of Budgeted-UCB studied in the original paper [35] does not necessarily apply to our setting. In addition, while Xie et al. [35] assumed that the pairs $\{(r_{dk}, t_{dk})\}_{k=1}^{\infty}$ are i.i.d., we do not have such a guarantee.

3.4 Start Selection

Once LNBS determines the depth d to use, it selects a starting point i , which induces the prefix and the suffix, in line 4 of Algorithm 3. We select i considering the cost change in a neighborhood. Concretely, the cost change by partial path $\langle x_i, \dots, x_{i+d-1} \rangle$ is defined as

$$\delta_{di} = \text{cost}_{\langle x_1, \dots, x_{i+d-1} \rangle}(S^0) - \text{cost}_{\langle x_1, \dots, x_{i-1} \rangle}(S^0). \quad (6)$$

Since the path cost is non-decreasing, $\delta_{di} \geq 0$. We ignore i with $\delta_{di} = 0^1$ and select one uniformly at random from the remaining options.

Our second approach is to select i based on the probability biased by δ_{di} . As we explain in the next subsection, for each d and i , the beam width b is maintained. Since smaller b_{di} leads to a shorter search time, we discount the probability of selecting i by b_{di} . Concretely, given the depth d , we can select the starting point i with the probability

$$p_{di} = \frac{\delta_{di}/b_{di}}{\sum_{j=1}^{n-d+1} \delta_{dj}/b_{dj}}. \quad (7)$$

3.5 Beam Width Selection

Given the depth d and the starting point i , LNBS selects a beam width b in line 5 of Algorithm 3. Here, we use a similar strategy to CABS: for each d and i , we initialize the beam width b_{di} to be 1 and update it to $2b_{di}$ after each round with d and i . If we find an improved solution in line 7, we reset $b_{d'i'} = 1$ only for d' and i' such that $i' > i$ or $i' + d' < i + d$; if $i' \leq i$ and $i' + d' \geq i + d$, the prefix and the suffix for the neighborhood induced by i' and d' do not change, and we know that a better partial path was not found with beam widths smaller than $b_{d'i'}$.² If the neighborhood is exhausted, i.e., `complete` = \top in line 7, we ignore the combination of d and i in lines 3 and 4 until a new solution is found and b_{id} is reset to 1. Since the number of neighborhoods is finite, LNBS eventually exhausts all the neighborhoods and finds the optimal solution, which guarantees completeness.

4 Experimental Evaluation

We compare LNBS with CABS, DD-LNS, CP, and mixed-integer programming (MIP).

4.1 Experimental Settings

As benchmarks, we use the same problems and instances used by previous work [26]: the traveling salesperson problem with time windows (TSPTW), the capacitated vehicle routing problem (CVRP), the multi-commodity pickup and delivery traveling salesperson problem

¹ Theoretically, a better solution may be found with such i if the suffix is not empty because a partial path may change the state from which the suffix is applied, which may change the cost of the suffix.

² Theoretically, a better solution may be found with beam width smaller than $b_{d'i'}$ if the updated primal bound changes the search behavior.

337 (m-PDTSP), the single machine total weighted tardiness problem ($1||\sum s_i T_i$), the talent
 338 scheduling problem (Talent), the simple assembly line balancing problem to minimize the
 339 number of stations (SALBP-1), the bin packing problem (BPP), and the graph-clear problem
 340 (GCP). We use the same DP, CP, and MIP models as the previous work [26].³

341 LNBS, CABS, and DD-LNS are implemented in didp-rs v0.3.2⁴ using Rust 1.65.0. In
 342 LNBS and DD-LNS, CABS is run first to find a feasible solution, and then LNBS and DD-
 343 LNS are run to improve the solution. For DD-LNS, we use $W = 1000$ and $p = 0.1$ following
 344 the original paper [18]. As we described above, LNBS removes conflicting transitions from a
 345 suffix, selects the depth using Budgeted-UCB, and geometrically increases the beam width for
 346 each neighborhood. To select the starting point of a partial path, we consider two approaches,
 347 uniform and biased sampling. While biased sampling achieves better solution quality in
 348 CVRP and m-PDTSP, uniform sampling is better in TSPTW, $1||\sum w_i T_i$, SALBP-1, MOSP,
 349 and GCP (see Appendix A). In what follows, we only show results from uniform sampling.
 350 In Appendix A, we also evaluate the importance of removing conflicting transitions and
 351 Budgeted-UCB. We confirm that these mechanisms significantly improve the solution quality.
 352 A more comprehensive ablation study is left for future work.

353 We implemented the DP models using didppy, a Python interface for didp-rs. We use
 354 IBM ILOG CP Optimizer 22.1.0 for the CP models and Gurobi Optimizer 9.5.0 for the
 355 MIP models. CP Optimizer is known to use LNS [27]. The DP, CP, and MIP models are
 356 implemented in Python 3.10.2. All experiments are run on an Intel Xeon Gold 6148 processor
 357 with a single thread, an 8 GB memory limit, and a time limit of 1800 seconds. For LNBS
 358 and DD-LNS, we take the median of 5 runs.

359 Following the previous work [26], we use the primal gap and the primal integral to measure
 360 the performance [7]. If an algorithm finds a solution x^t with the cost $f(x^t)$ at time t , and
 361 the optimal or best-known solution cost is f^* , the primal gap at t for the algorithm is

$$362 \quad p(t) = \begin{cases} 1 & \text{if } f(x^t) \cdot f^* < 0 \\ 0 & \text{if } f(x^t) = f^* = 0 \\ \frac{|f(x^t) - f^*|}{\max\{|f(x^t)|, |f^*|\}} & \text{otherwise.} \end{cases} \quad (8)$$

363 If the algorithm does not find a solution at time t , then $p(t) = 1$. Let t_1, \dots, t_{l-1} be time
 364 points where a better solution is found, $t_0 = 0$, and $t_l = T$ where T is the time limit. Then,
 365 the primal integral, $P(T)$ is defined as

$$366 \quad P(T) = \sum_{i=1}^l p(t_{i-1}) \cdot (t_i - t_{i-1}). \quad (9)$$

367 We use the primal gap at the time limit, $p(T)$, and the primal integral, $P(T)$, as the measures.

368 4.2 Experimental Results

369 We show the number of instances where the optimality of a solution is proved, the average
 370 primal gap at the time limit, and the average primal integral for each problem in Table 1.
 371 We omit MIP in Table 1 because it is outperformed by CP in the primal gap and the primal
 372 integral for all problem types (see Appendix A). In the number of optimally solved instances,
 373 MIP is the best in CVRP (with 26 problems solved). As reported by previous work [26],

³ <https://github.com/Kurorororo/didp-models>

⁴ <https://didp.ai>

	CP			CABS			DD-LNS			LNBS		
	#	gap	p.i.	#	gap	p.i.	#	gap	p.i.	#	gap	p.i.
TSPTW (340)	47	0.0259	49.0	257	0.0033	9.0	109	0.0100	23.7	241	0.0016	5.7
CVRP (207)	0	0.3174	601.1	6	0.1772	339.5	0	0.2504	461.6	6	0.1640	316.8
m-PDTSP (1178)	1049	0.0122	25.5	1030	0.0023	5.1	459	0.0102	21.7	1029	0.0022	5.0
$1 \sum w_i T_i$ (375)	150	0.0009	3.5	286	0.0346	74.4	100	0.0409	83.5	275	0.0051	13.0
Talent (1000)	0	0.0081	29.3	232	0.0173	38.2	0	0.0602	114.6	232	0.0041	11.1
SALBP-1 (2100)	1584	0.0046	28.4	1799	0.0003	2.2	1507	0.0067	13.5	1682	0.0022	7.3
BPP (1615)	1234	0.0014	7.7	1159	0.0017	6.1	775	0.0190	35.4	1139	0.0021	8.1
MOSP (570)	437	0.0044	13.0	526	0.0000	0.4	353	0.0203	37.5	523	0.0002	0.7
GCP (135)	1	0.0151	44.3	103	0.0000	0.6	3	0.0009	2.7	102	0.0001	0.6
Larger Instances												
m-PDTSP (240)	77	0.1491	285.9	101	0.0694	153.2	79	0.1345	265.7	98	0.0652	146.6
MOSP (760)	0	0.0676	150.6	150	0.0002	4.4	0	0.0402	72.7	148	0.0025	10.4
GCP (50)	0	0.5289	1268.3	0	0.0013	10.8	0	0.0764	137.8	0	0.0038	19.5

■ **Table 1** Summary of the experimental result. ‘#’ is the number of optimally solved instances, ‘gap’ is the average primal gap at the time limit, and ‘p.i.’ is the average primal integral.

374 MIP is good at finding an optimal solution for small instances, while it fails to find a feasible
375 solution for larger instances, which results in poor performance in the primal gap and the
376 primal integral on average. In other problems, LNBS solves more instances optimally than
377 MIP except for BPP, where LNBS solves 1139 and MIP solves 1157.

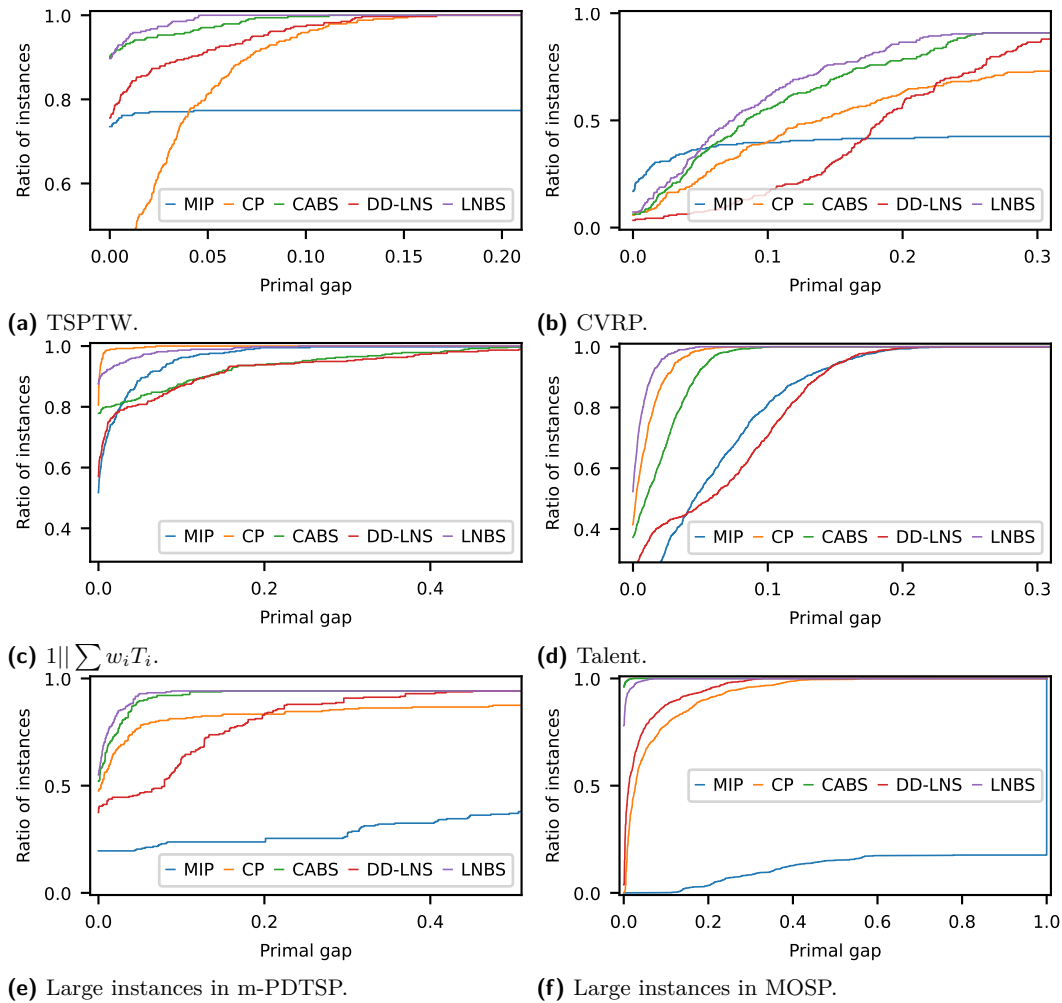
378 Compared to CABS, LNBS achieves the better primal gap and the primal integral in
379 TSPTW, CVRP, m-PDTSP, $1||\sum w_i T_i$, and Talent. We show the distribution of the primal
380 gap in TSPTW, CVRP, $1||\sum w_i T_i$, and Talent in Figure 2. The primal integral has a similar
381 trend to the primal gap in these problems (see Appendix A). In contrast, CABS is better
382 than LNBS in SALBP-1, BPP, MOSP, and GCP. These results are consistent with the
383 observation that LNS is effective for routing and scheduling problems in CP [24, 27]. In the
384 routing problems (TSPTW, CVRP, and m-PDTSP), CABS is already better than CP, and
385 LNBS is even better than CABS. In $1||\sum w_i T_i$, while LNBS shows a significant improvement
386 from CABS (0.0051 from 0.0346) outperforming MIP (0.0188), CP is still the best. However,
387 in Talent, LNBS outperforms CP while CABS does not. Overall, LNBS is better than CP in
388 seven problems while CABS is better than CP in six problems.

389 In TSPTW, the difference in the primal gap between LNBS and CABS (0.0016 and
390 0.0033) seems small, but it is because they achieve almost the same primal gap in many
391 instances: they optimally solve all 135 instances in the Dumas set [10] and achieve almost the
392 same average primal gap in the AFG set [2], which has 50 instances, and GendreauDumas
393 Extended set [16], which has 130 instances. However, in the OhlmannThomas set [31], which
394 has 25 instances, no instance is optimally solved, and LNBS shows a significant improvement
395 in the primal gap (0.0184 from 0.0395).

396 In the number of optimally solved instances, CABS is equal to or better than LNBS in all
397 problems. While CABS always searches the entire state space graph, LNBS searches multiple
398 neighborhoods, and the entire state space graph is just one of them. Nevertheless, LNBS
399 proves the optimality of more than 93% of instances that are optimally solved by CABS.

400 DD-LNS performs worse than CABS and LNBS in all the problems. Previous work has
401 reported that DD-LNS is effective for TSPTW [18]. Note however that the DD-LNS and
402 LNBS results in Table 1 are not the results reported by Gillard and Schaus. To validate that
403 our implementation and experimental settings do not handicap DD-LNS, we compare the
404 results of DD-LNS with those of the original paper in TSPTW.⁵ Our DD-LNS implementation

⁵ https://github.com/xgillard/ijcai_22_DDLNS/blob/main/results/tsptw/ddlms/results_w1000_



■ **Figure 2** Distribution of the primal gap at the time limit. Higher and left is better.

405 finds a better solution than the original in all instances with the time limit of 600 seconds.
 406 This difference is likely due to the difference between the DP models used by us (from
 407 Kuroiwa and Beck [26]) and Gillard and Schaus. In TSPTW, a solution is a tour that starts
 408 from a depot, visits each customer j within the time window $[a_j, b_j]$, and returns to the
 409 depot, and an optimal solution minimizes the total travel time. In both DP models, state
 410 variables are the set of unvisited customers U , the current customer i , and the current time
 411 t , and each transition corresponds to visiting a customer or the depot. Each DP model has a
 412 dual bound (called RLB by Gillard and Schaus), a lower bound on the optimal solution cost.
 413 While Gillard and Schaus use a dual bound based on a minimum spanning tree, Kuroiwa and
 414 Beck use a simpler one based on the minimum travel time between customers. In addition,
 415 Kuroiwa and Beck use information that was not considered by Gillard and Schaus. First,
 416 they use dominance between states based on the current time: a state S dominates another
 417 state S' if $S[U] = S'[U]$, $S[i] = S'[i]$, and $S[t] \leq S'[t]$. Furthermore, since the time to visit
 418 customer j is underestimated by $t + c_{ij}^*$, where c_{ij}^* is the shortest travel time from i to j , they

t600.txt

419 define state constraints $\forall j \in U, t + c_{ij}^* \leq b_j$. The dominance and the state constraints are
 420 useful to prune states, which potentially explains the performance gap. Another difference
 421 is whether considering the time window constraints at the depot or not. In the benchmark
 422 instances used above⁶ (and in Kuroiwa and Beck [26]), a time window $[a_0, b_0]$ is defined for
 423 the depot. Gillard and Schaus explicitly model a required return to the depot within $[a_0, b_0]$
 424 while Kuroiwa and Beck do not. However, in the benchmark instances, $b_0 \geq b_j + c_{j0}$ holds
 425 for all customers j , where c_{j0} is the travel time from j to the depot. Thus, if all customers
 426 are visited within the time windows, the depot can be reached within the time window, and
 427 explicit modeling of the depot return window is unnecessary.

428 4.3 Larger Instances

429 LNBS performs better than CABS in m-PDTSP and worse in MOSP and GCP, but the
 430 difference in the average primal gap is small. To evaluate the difference more clearly, we use
 431 larger instances for these problems.

432 In m-PDTSP, a vehicle visits all nodes in a graph, picks up some commodities at some
 433 nodes, and delivers them to others. Each commodity has a weight and the total weight of
 434 commodities that a vehicle can carry is limited by the capacity. In the benchmark set for
 435 m-PDTSP, three types of instances are used: Class 1, Class 2, and Class 3 [21], and Class 1
 436 instances are generated from instances of the sequential ordering problem (SOP) [3]. We
 437 generate larger Class 1 instances by using 30 SOP instances in TSPLIB⁷ that were not used
 438 by the previous work. The original instances have at most 47 nodes, and the new instances
 439 have 42 to 378 nodes. We use the same methods as the previous work [21] with the maximum
 440 weight $q \in \{1, 5\}$ and the capacity $Q \in \{5q, 10q, 20q, 100q\}$, resulting in 240 instance in total.

441 In MOSP, an instance is represented by a matrix, and the original set uses at most
 442 125×125 matrices. We add instances using 150×150 to 1000×1000 matrices [8, 12].

443 In GCP, an instance is represented by a graph. The original instance set uses random and
 444 random planar graphs with 20, 30, and 40 nodes. We generate 50 instances using random
 445 graphs with 100 and 200 nodes following the method used by previous work [29].

446 As shown in Table 1, LNBS clearly outperforms CABS in m-PDTSP in the primal gap and
 447 the primal integral, but CABS is better in MOSP and GCP. We also show the distribution
 448 of the primal gap in m-PDTSP and MOSP in Figure 2. The primal integral has a similar
 449 tendency to the primal gap, and the result for GCP is qualitatively similar to that of MOSP.


450 4.4 Analysis of Problem Characteristics

451 In routing problems, a solution is a route visiting all nodes in a graph, and its cost is the
 452 length of the route. In the DP models for these problems, each transition corresponds to
 453 visiting one node, and the cost of a partial path increases when a transition is applied. We
 454 expect that different partial solutions tend to have different costs, and it is relatively easy to
 455 find a better partial path; because the path costs are diverse, unless the current partial path
 456 is optimal, better partial paths are included in a partial state space graph with high density.
 457 In such a case, beam search is likely to find a better partial path although it searches in a
 458 fraction of the partial state space graph restricted by the beam width.

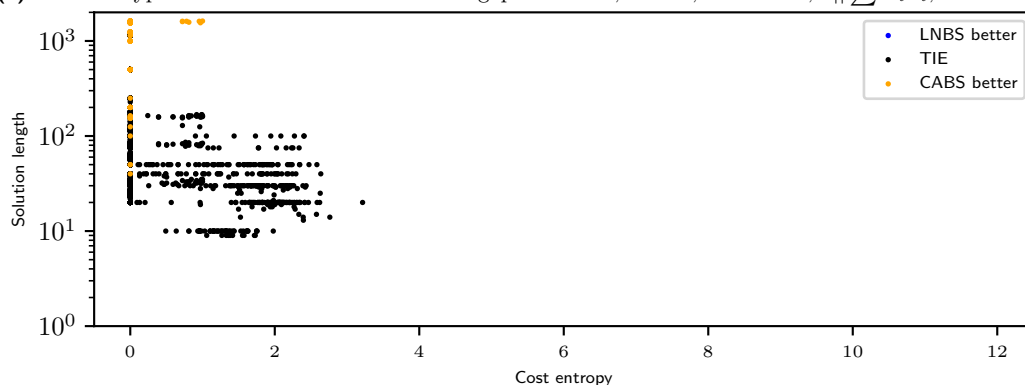
459 In contrast, in SALBP-1 and BPP, the problem is to pack weighted items into capacitated
 460 bins while minimizing the number of bins. In the DP models, each transition packs one item

⁶ <https://lopez-ibanez.eu/tsptw-instances>

⁷ <http://comopt.ifl.uni-heidelberg.de/software/TSPLIB95/sop/>



(a) Problem types where LNBS has lower mean gap: TSPTW, CVRP, m-PDTSP, $1||\sum w_i T_i$, and Talent.



(b) Problem types where CABS has lower mean gap: SALBP-1, BPP, MOSP, and GCP.

Figure 3 Entropy of the cost distribution over partial paths vs. the solution length in each problem instance. ‘LNBS better’ means LNBS finds a better solution, ‘TIE’ means that LNBS and CABS achieve the same solution cost, and ‘CABS better’ means that CABS finds a better solution.

461 into a bin, and the cost increases only when a new bin is opened. In the DP models for
 462 MOSP and GCP, the cost is computed by taking the maximum weights of edges in a path,
 463 and it does not increase unless a new edge has a higher weight than the current maximum.
 464 Therefore, we expect that many partial paths tend to have the same cost in these problems,
 465 making it difficult to improve a solution by searching only a partial state space graph.

466 Based on these observations, we hypothesize that LNBS tends to perform better than
 467 CABS when path costs in a partial state space graph are diverse. To test this hypothesis, we
 468 evaluate the diversity of costs in a partial state space graph using entropy in information theory.
 469 Given a solution for a DyPDL model $x = \langle x_1, \dots, x_n \rangle$, let $Y_{di}(x)$ be the set of solution paths
 470 whose prefix is $\langle x_1, \dots, x_{i-1} \rangle$ and the suffix is $\langle x_{i+d}, \dots, x_n \rangle$. Let $C = \{\text{cost}_y(S^0) \mid y \in Y_{di}(x)\}$
 471 be the set of the path costs. Then, the entropy of the path costs is defined as follows:

$$472 \quad H(Y_{di}(x)) = - \sum_{c \in C} \frac{|\{y \in Y_{di}(x) \mid \text{cost}_y(S^0) = c\}|}{|Y_{di}(x)|} \log_2 \frac{|\{y \in Y_{di}(x) \mid \text{cost}_y(S^0) = c\}|}{|Y_{di}(x)|}. \quad (10)$$

473 As this value gets larger, the cost distribution becomes more diverse, and we expect that
 474 LNBS will perform better than CABS. However, even if the entropy is large, if the problem
 475 itself is easy, both CABS and LNBS will find optimal or near-optimal solutions. To consider
 476 such cases, we also evaluate the length of the initial solution found by CABS.

477 We evaluate entropy and the length of the initial solution found for each problem instance

478 used in Section 4.2. We first run CABS until it finds a feasible solution and record the length
 479 of the solution. Then, we remove the first eight transitions from the solution and enumerate
 480 all feasible prefixes of the solution, i.e., we use $d = 8$ and $i = 1$. In Figure 3, we show a scatter
 481 plot of entropy and the solution length divided into two plots to emphasize the differences
 482 between the problem types where LNBS has a lower primal gap on average (Figure 3a) and
 483 those where CABS is better (Figure 3b). With low entropy, CABS dominates. For higher
 484 entropy, the solution length begins to play a factor: for short solutions, CABS and LNBS
 485 perform equally but for longer solutions, LNBS tends to perform better. Indeed, for the
 486 problems where CABS performs better on average (Figure 3b), the entropy is quite low (less
 487 than 3.5). This result suggests that the entropy of the cost distribution over partial paths
 488 is related to the performance of LNBS. Since this analysis is based on path costs, it is not
 489 directly applicable to LNS with tree search, where a solution corresponds to a leaf node.
 490 However, if we consider the factors of neighborhood size and cost distribution over leaf nodes,
 491 we may be able to apply this analysis to LNS for CP and MIP.

492 **5 Related Work**

493 As we discussed, LNBS can be considered a generalization of DD-LNS [18], which combines
 494 DDs and LNS. DP is closely related to DDs [23], and DDs have been actively used for
 495 combinatorial optimization [9]. For example, DDs are used to obtain bounds on the optimal
 496 objective value [4, 33], and heuristics based on DDs have been proposed [6]. Moreover, ddo, a
 497 general-purpose DD solver for combinatorial optimization has been developed [5, 19]. In CP,
 498 DDs are used for constraint propagation [1, 22]. Recently, HADDOCK, a modeling language
 499 of a DD based on a state transition system, was proposed for CP [17].

500 In state space search, there exist several methods that improve a solution path by searching
 501 in a partial state space graph, but they were not framed as LNS. In classical planning, plan
 502 neighborhood graph search (PNGS) first constructs a partial state space graph by performing
 503 local search from each state in a solution path and then finds the shortest path in the graph
 504 [30]. In sliding tile puzzles, iterative tunneling search with A* (ITSA*) iteratively expands a
 505 partial state space graph, which includes states close to a given path, and finds the shortest
 506 path in that graph [13]. Unlike the above two algorithms, Joint and local path A* (LPA*)
 507 [32] try to find a better partial path between two states in a given path using A* [20]. While
 508 Joint and LPA* fix the length of a partial path to remove and deterministically select a
 509 neighborhood, LNBS dynamically adjusts them and uses beam search instead of A*.

510 **6 Conclusion**

511 We proposed large neighborhood beam search (LNBS), a state space search algorithm based
 512 on large neighborhood search (LNS) and beam search for domain-independent dynamic
 513 programming (DIDP). Our configuration of LNBS exploits the multi-armed bandit problem
 514 and random sampling to select a neighborhood. We proved that LNBS is complete. LNBS
 515 finds better quality solutions on average than the state-of-the-art DIDP solver, complete
 516 anytime beam search (CABS), in five out of the nine benchmark problems. In particular,
 517 LNBS performs well in routing and scheduling problems, and our analysis suggests that this
 518 performance is related to the diversity of the cost distribution over partial paths. A deeper
 519 investigation of the characteristics of the problems that make LNS effective in state space
 520 search and tree search is an interesting direction for future work. Based on such analysis,
 521 developing better configurations for LNBS may also be possible.

522 — References —

- 523 1 H. R. Andersen, T. Hadzic, J. N. Hooker, and P. Tiedemann. A constraint store based on
524 multivalued decision diagrams. In *Principles and Practice of Constraint Programming – CP*
525 *2007*, pages 118–132, 2007. doi:10.1007/978-3-540-74970-7_11.
- 526 2 Norbert Ascheuer. *Hamiltonian path problems in the on-line optimization of flexible manufac-*
527 *turing systems*. PhD thesis, Technische Universität Berlin, 1995.
- 528 3 Norbert Ascheuer, Michael Jünger, and Gerhard Reinelt. A branch & cut algorithm for
529 the asymmetric traveling salesman problem with precedence constraint. *Computational*
530 *Optimization and Applications*, 17:25–42, 2000. doi:10.1023/A:1008779125567.
- 531 4 David Bergman, Andre A. Cire, Willem-Jan van Hoeve, and J. N. Hooker. Optimization
532 bounds from binary decision diagrams. *INFORMS Journal on Computing*, 26(2):253–268,
533 2014. doi:10.1287/ijoc.2013.0561.
- 534 5 David Bergman, Andre A. Cire, Willem Jan Van Hoeve, and J. N. Hooker. Discrete optimization
535 with decision diagrams. *INFORMS Journal on Computing*, 28(1):47–66, 12 2016. doi:
536 10.1287/ijoc.2015.0648.
- 537 6 David Bergman, Andre A. Cire, Willem-Jan van Hoeve, and Tallys Yunes. Bdd-based
538 heuristics for binary optimization. *Journal of Heuristics*, 20(2):211–234, 2014. doi:10.1007/
539 s10732-014-9238-1.
- 540 7 Timo Berthold. Measuring the impact of primal heuristics. *Operations Research Letters*,
541 41:611–614, 2013. doi:10.1016/j.orl.2013.08.007.
- 542 8 Marco Antonio Moreira De Carvalho and Nei Yoshihiro Soma. A breadth-first search applied to
543 the minimization of the open stacks. *Journal of the Operational Research Society*, 66:936–946,
544 6 2015. doi:10.1057/jors.2014.60.
- 545 9 Margarita P. Castro, Andre A. Cire, and J. Christopher Beck. Decision diagrams for discrete
546 optimization: A survey of recent advances. *INFORMS Journal on Computing*, 34(4):2271–2295,
547 2022. doi:10.1287/ijoc.2022.1170.
- 548 10 Yvan Dumas, Jacques Desrosiers, Eric Gelinat, and Marius M Solomon. An optimal algorithm
549 for the traveling salesman problem with time windows. *Operations Research*, 43(2):367–371,
550 1995. doi:10.1287/opre.43.2.367.
- 551 11 Stefan Edelkamp, Shahid Jabbar, and Alberto Lluch Lafuente. Cost-algebraic heuristic
552 search. In *Proceedings of the 20th National Conference on Artificial Intelligence (AAAI)*, pages
553 1362–1367, 2005.
- 554 12 Rafael de Magalhães Dias Frinhani, Marco Antonio Moreira de Carvalho, and Nei Yoshihiro
555 Soma. A pagerank-based heuristic for the minimization of open stacks problem. *PLoS ONE*,
556 13(8):1–24, 2018. doi:10.1371/journal.pone.0203076.
- 557 13 David A Furcy. ITSA*: Iterative tunneling search with A*. In *Proceedings of AAAI Workshop*
558 *on Heuristic Search, Memory-Based Heuristics and Their Applications*, pages 21–26, 2006.
- 559 14 Maria Garcia de la Banda and Peter J. Stuckey. Dynamic programming to minimize the
560 maximum number of open stacks. *INFORMS Journal on Computing*, 19(4):607–617, 2007.
561 doi:10.1287/ijoc.1060.0205.
- 562 15 Maria Garcia de la Banda, Peter J. Stuckey, and Geoffrey Chu. Solving talent scheduling
563 with dynamic programming. *INFORMS Journal on Computing*, 23(1):120–137, 2011. doi:
564 10.1287/ijoc.1090.0378.
- 565 16 Michel Gendreau, Alain Hertz, Gilbert Laporte, and Mihnea Stan. A generalized insertion
566 heuristic for the traveling salesman problem with time windows. *Operations Research*, 46(3):330–
567 346, 1998. doi:10.1287/opre.46.3.330.
- 568 17 Rebecca Gentzel, Laurent Michel, and W.-J. van Hoeve. HADDOCK: A language and architec-
569 ture for decision diagram compilation. In Helmut Simonis, editor, *Principles and Practice of*
570 *Constraint Programming – CP 2020*, pages 531–547, 2020. doi:10.1007/978-3-030-58475-7_
571 31.

- 572 18 Xavier Gillard and Pierre Schaus. Large neighborhood search with decision diagrams. In
573 *Proceedings of the 21st International Joint Conference on Artificial Intelligence, IJCAI-22*,
574 pages 4754–4760, 2022. doi:10.24963/ijcai.2022/659.
- 575 19 Xavier Gillard, Pierre Schaus, and Vianney Coppé. Ddo, a generic and efficient framework
576 for mdd-based optimization. In *Proceedings of the 29th International Joint Conference on*
577 *Artificial Intelligence, IJCAI-20*, pages 5243–5245, 2020. doi:10.24963/ijcai.2020/757.
- 578 20 Peter E. Hart, Nils J. Nilsson, and Bertram Raphael. A formal basis for the heuristic
579 determination of minimum cost paths. *IEEE Transactions on Systems Science and Cybernetics*,
580 4(2):100–107, 1968. doi:10.1109/TSSC.1968.300136.
- 581 21 Hipólito Hernández-Pérez and Juan José Salazar-González. The multi-commodity one-to-one
582 pickup-and-delivery traveling salesman problem. *European Journal of Operational Research*,
583 196:987–995, 8 2009. doi:10.1016/j.ejor.2008.05.009.
- 584 22 Samid Hoda, Willem-Jan van Hoes, and J. N. Hooker. A systematic approach to MDD-based
585 constraint programming. In *Principles and Practice of Constraint Programming – CP 2010*,
586 pages 266–280, 2010.
- 587 23 John N. Hooker. Decision diagrams and dynamic programming. In Carla Gomes and Meinolf
588 Sellmann, editors, *Integration of AI and OR Techniques in Constraint Programming for*
589 *Combinatorial Optimization Problems*, pages 94–110, 2013.
- 590 24 Siddhartha Jain and Pascal Van Hentenryck. Large neighborhood search for dial-a-ride
591 problems. In *Principles and Practice of Constraint Programming – CP 2011*, pages 400–413,
592 2011.
- 593 25 Ryo Kuroiwa and J. Christopher Beck. Domain-independent dynamic programming: Generic
594 state space search for combinatorial optimization. In *Proceedings of the 33rd International*
595 *Conference on Automated Planning and Scheduling (ICAPS)*, 2023. doi:10.1609/icaps.
596 v33i1.27200.
- 597 26 Ryo Kuroiwa and J. Christopher Beck. Solving domain-independent dynamic programming
598 problems with anytime heuristic search. In *Proceedings of the 33rd International Conference*
599 *on Automated Planning and Scheduling (ICAPS)*, 2023. doi:10.1609/icaps.v33i1.27201.
- 600 27 Philippe Laborie, Jérôme Rogerie, Paul Shaw, and Petr Vilím. IBM ILOG CP optimizer for
601 scheduling. *Constraints*, 23(2):210–250, 2018. doi:10.1007/s10601-018-9281-x.
- 602 28 Shu Lin, Na Meng, and Wenxin Li. Optimizing constraint solving via dynamic programming.
603 In *Proceedings of the 28th International Joint Conference on Artificial Intelligence, IJCAI-19*,
604 pages 1146–1154, 7 2019. doi:10.24963/ijcai.2019/160.
- 605 29 Michael Morin, Margarita P. Castro, Kyle E.C. Booth, Tony T. Tran, Chang Liu, and
606 J. Christopher Beck. Intruder alert! Optimization models for solving the mobile robot
607 graph-clear problem. *Constraints*, 23(3):335–354, 2018. doi:10.1007/s10601-018-9288-3.
- 608 30 Hootan Nakhost and Martin Müller. Action elimination and plan neighborhood graph
609 search: Two algorithms for plan improvement. In *Proceedings of the 20th International*
610 *Conference on Automated Planning and Scheduling (ICAPS)*, pages 121–128, 2010. doi:
611 10.1609/icaps.v20i1.13402.
- 612 31 Jeffrey W. Ohlmann and Barrett W. Thomas. A compressed-annealing heuristic for the
613 traveling salesman problem with time windows. *INFORMS Journal on Computing*, 19(1):80–
614 90, 2007. doi:10.1287/ijoc.1050.0145.
- 615 32 Daniel Ratner and Ira Pohl. Joint and LPA*: Combination of approximation and search. In
616 *Proceedings of the fifth National Conference on Artificial Intelligence (AAAI)*., pages 173–177,
617 1986. doi:10.5555/2887770.2887798.
- 618 33 Isaac Rudich, Quentin Cappart, and Louis-Martin Rousseau. Peel-And-Bound: Generating
619 Stronger Relaxed Bounds with Multivalued Decision Diagrams. In *Principles and Practice of*
620 *Constraint Programming – CP 2022*, pages 35:1–35:20, 2022. doi:10.4230/LIPIcs.CP.2022.
621 35.

- 622 **34** Paul Shaw. Using constraint programming and local search methods to solve vehicle routing
 623 problems. In *Principles and Practice of Constraint Programming – CP98*, volume 1520, pages
 624 417–431, 1998. doi:10.1007/3-540-49481-2_30.
- 625 **35** Yingce Xia, Xu-Dong Zhang, Nenghai Yu, Geoffrey Holmes, and Yan Liu. Budgeted bandit
 626 problems with continuous random costs. In *Proceedings of the Seventh Asian Conference on*
 627 *Machine Learning*, pages 317–332, 2015.
- 628 **36** Weixiong Zhang. Complete anytime beam search. In *Proceedings of the Fifteenth National*
 629 *Tenth Conference on Artificial Intelligence/Innovative Applications of Artificial Intelligence*
 630 *(AAAI-98/IAAI-98)*, pages 425–430, 1998.

631 **A Additional Experimental Results**

	CP			MIP			LNBS			LNBS/bias		
	#	gap	p.i.	#	gap.	p.i.	#	gap.	p.i.	#	gap.	p.i.
TSPTW (340)	46	0.0275	52.3	227	0.2268	484.1	241	0.0016	5.7	242	0.0019	6.2
CVRP (207)	0	0.3174	601.1	26	0.5845	1157.4	6	0.1640	316.8	6	0.1633	314.8
m-PDTSP (1178)	1050	0.0121	25.4	945	0.0858	180.0	1029	0.0022	5.0	1029	0.0021	4.7
$1 \sum w_i T_i$ (375)	150	0.0003	2.4	109	0.0188	75.6	275	0.0051	13.0	275	0.0056	14.5
Talent (1000)	7	0.0072	27.6	0	0.0573	152.6	232	0.0041	11.1	231	0.0042	11.1
SALBP-1 (2100)	1584	0.0046	28.4	1357	0.3447	634.6	1682	0.0022	7.3	1675	0.0022	7.5
BPP (1615)	1234	0.0014	7.7	1157	0.0385	85.9	1139	0.0021	8.1	1129	0.0021	9.0
MOSP (570)	437	0.0044	13.0	224	0.0394	100.4	523	0.0002	0.7	523	0.0002	0.7
GCP (135)	1	0.0151	44.3	23	0.1102	311.9	102	0.0001	0.6	102	0.0001	0.7
Larger Instances												
m-PDTSP (240)	77	0.1481	284.1	47	0.5811	1096.8	98	0.0652	146.6	97	0.0647	147.0
MOSP (760)	0	0.0675	150.4	0	0.8806	1599.4	148	0.0025	10.4	148	0.0027	10.7
GCP (50)	0	0.5287	1268.1	0	0.5306	977.8	0	0.0038	19.5	0	0.0061	21.2

■ **Table 2** Comparison of CP, MIP, and two configurations of LNBS. ‘#’ is the number of optimally solved instances, ‘gap’ is the average primal gap at the time limit, and ‘p.i.’ is the average primal integral.

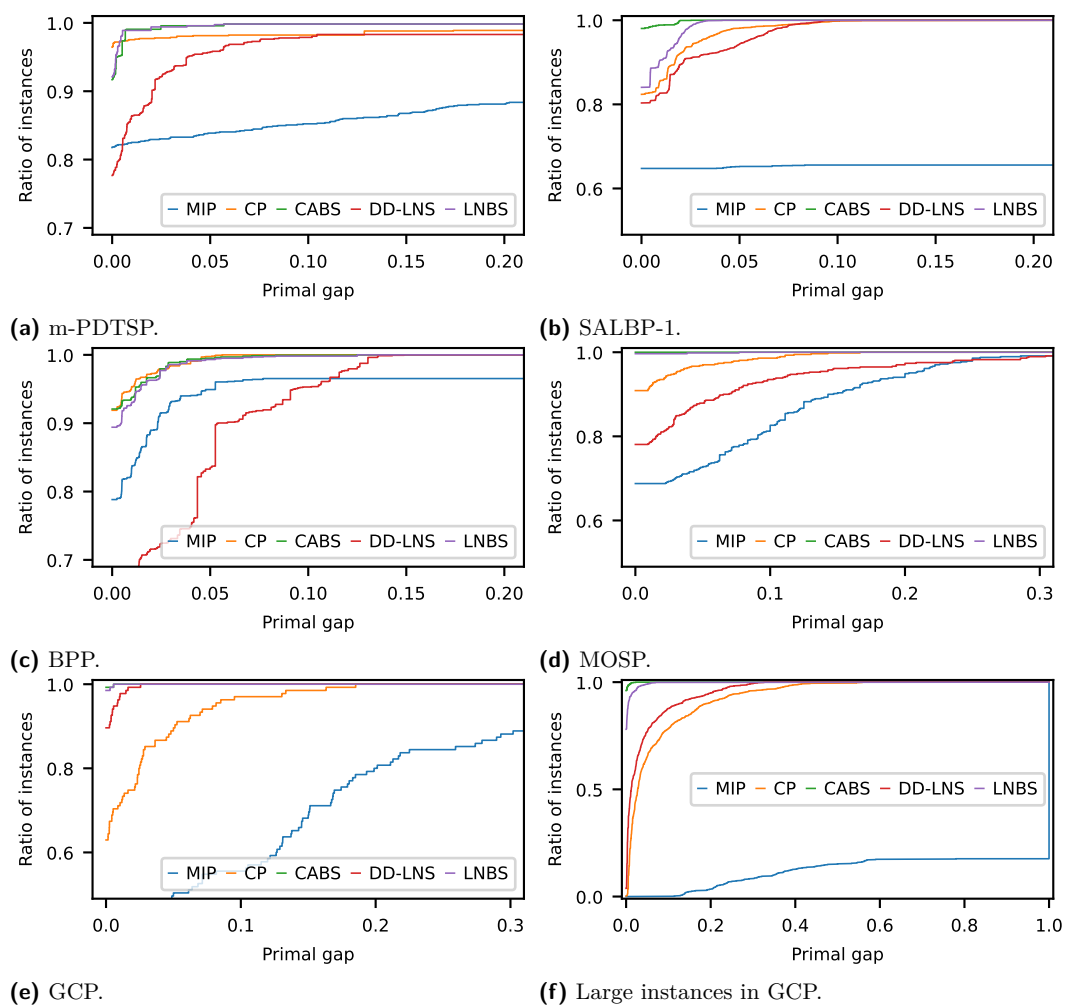
	LNBS			No removing conflicts			No Budgeted-UCB		
	#	gap.	p.i.	#	gap.	p.i.	#	gap.	p.i.
TSPTW (340)	241	0.0016	5.7	234	0.0035	10.6	256	0.0034	9.4
CVRP (207)	6	0.1640	316.8	5	0.1682	322.5	6	0.1767	338.7
$1 \sum w_i T_i$ (375)	275	0.0051	13.0	268	0.0224	56.0	287	0.0340	74.1

■ **Table 3** Comparison of LNBS variants. ‘No removing conflicts’ does not remove conflicting transitions in the suffix. ‘No Budgeted-UCB’ selects the depth uniformly at random instead of using Budgeted-UCB. ‘#’ is the number of optimally solved instances, ‘gap’ is the average primal gap at the time limit, and ‘p.i.’ is the average primal integral.

632 We show the result of MIP in Table 2. MIP solves more instances than CP in TSPTW,
 633 CVRP, and GCP, but CP is better in the primal gap and the primal integral.

634 In addition, we compare two configurations of LNBS in Table 2. One selects the starting
 635 point of a partial path uniformly at random (LNBS), and another selects it according to the
 636 probability distribution biased by partial path costs in Equation (7) (LNBS/bias). LNBS/bias
 637 solves one more instance in TSPTW and outperforms LNBS in CVRP and m-PDTSP in the
 638 primal gap.

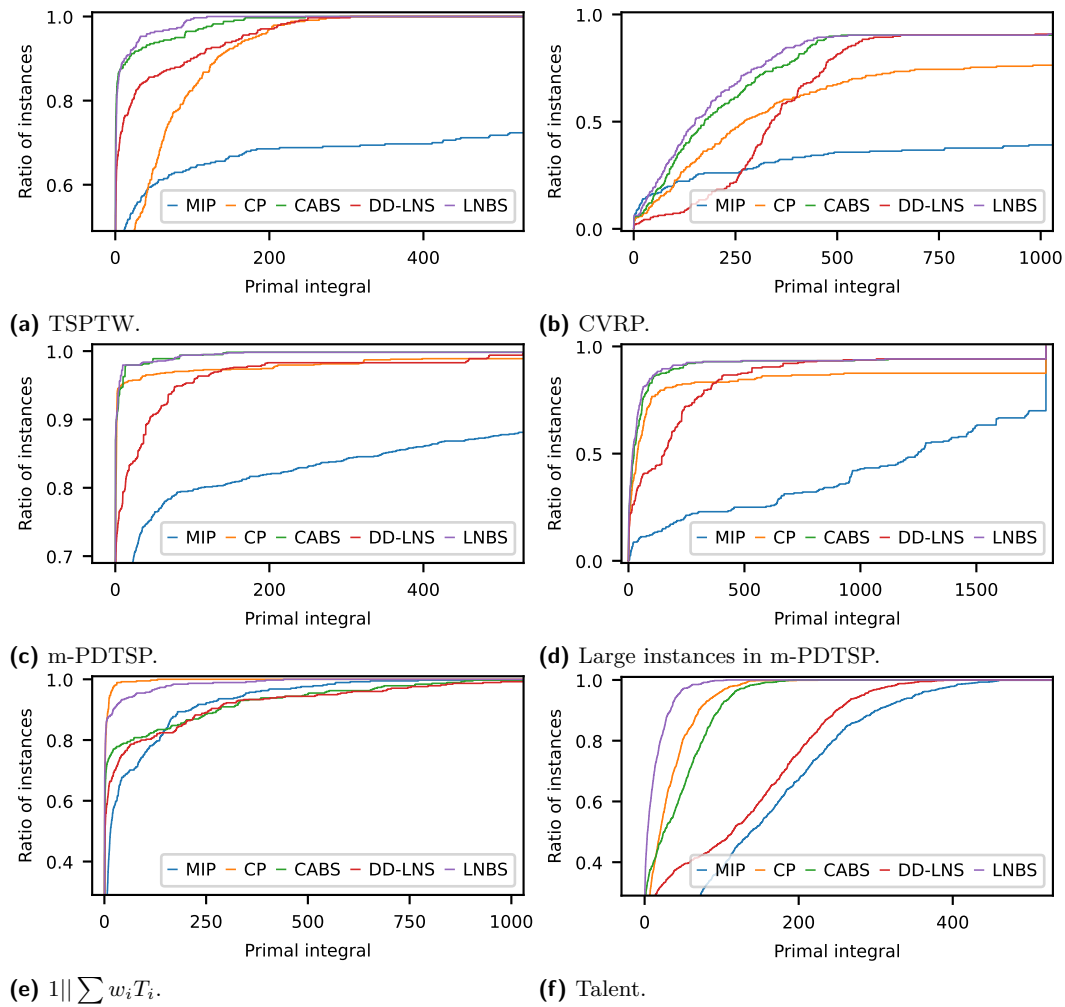
639 We evaluate the importance of other components of LNBS using a subset of the problems:
 640 TSPTW, CVRP, and $1||\sum w_i T_i$. In Table 3, we compare two variants of LNBS, where
 641 conflicting transitions in a suffix are not removed (‘No removing conflicts’), and the depth
 642 is selected uniformly at random instead of Budgeted-UCB (‘No Budgeted-UCB’). The two



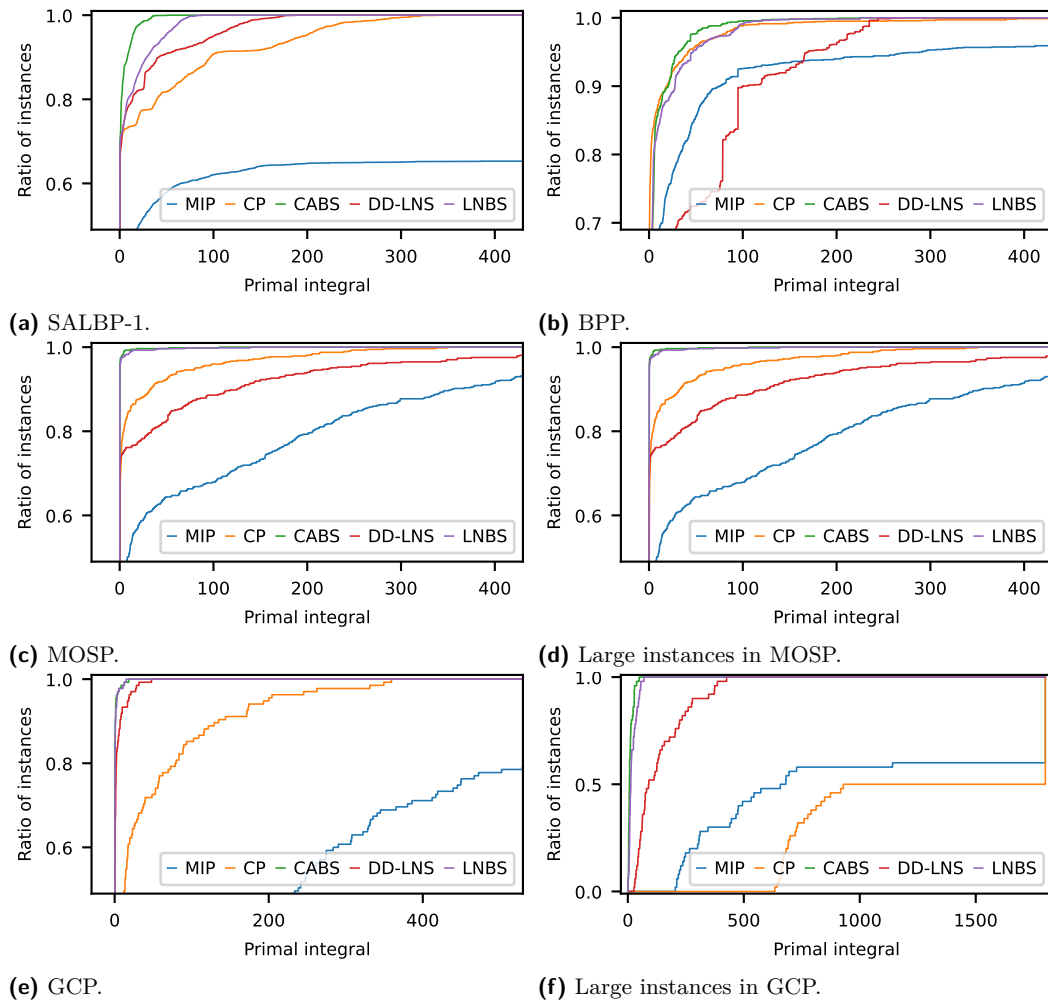
■ **Figure 4** Distribution of the primal gap at the time limit. Higher and left is better.

643 variants perform worse in terms of the primal gap and the primal integral. While ‘No
 644 Budgeted-UCB’ solves more instances to optimality, it is because the largest depth, which
 645 makes LNBS the same as CABS, is more likely to be selected by uniform sampling. Indeed,
 646 the primal gap and the primal integral of ‘No Budgeted-UCB’ are close to those of CABS.

647 In Figure 4, we present the distribution of the primal gap over instances in m-PDTSP,
 648 SALBP-1, BPP, MOSP, Graph-Clear, and large instances in GCP. LNBS is slightly better in
 649 m-PDTSP, and CABS is better in SALBP-1, BPP, and large instances in GCP. Figures 5
 650 and 6 show the distribution of the primal integral. The tendency is similar to that of the
 651 primal gap.



■ **Figure 5** Distribution of the primal integral at the time limit in the problems where LNBS is better. Higher and left is better.



■ **Figure 6** Distribution of the primal integral at the time limit in the problems where CABS is better. Higher and left is better.