



**HAL**  
open science

## Variable neighborhood search for graphical model energy minimization

Abdelkader Ouali, David Allouche, Simon de Givry, Samir Loudni, Yahia Lebbah, Lakhdar Loukil, Patrice Boizumault

► **To cite this version:**

Abdelkader Ouali, David Allouche, Simon de Givry, Samir Loudni, Yahia Lebbah, et al. Variable neighborhood search for graphical model energy minimization. *Artificial Intelligence*, 2020, 278, pp.103194. 10.1016/j.artint.2019.103194 . hal-02463467

**HAL Id: hal-02463467**

**<https://normandie-univ.hal.science/hal-02463467>**

Submitted on 21 Dec 2021

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution - NonCommercial 4.0 International License

# Variable Neighborhood Search for Graphical Model Energy Minimization

Abdelkader Ouali<sup>1</sup>, David Allouche<sup>2</sup>, Simon de Givry<sup>2</sup>, Samir Loudni<sup>1</sup>  
Yahia Lebbah<sup>3</sup>, Lakhdar Loukil<sup>3</sup> and Patrice Boizumault<sup>1</sup>

<sup>1</sup> *University of Caen Normandy, CNRS, UMR 6072 GREYC, 14032 Caen, France*

<sup>2</sup> *INRA, MIA Toulouse, UR-875, 31320 Castanet-Tolosan, France*

<sup>3</sup> *University of Oran 1 Ahmed Ben Bella, Lab. LITIO, 31000 Oran, Algeria*

---

## Abstract

Graphical models factorize a global probability distribution/energy function as the product/sum of local functions. A major inference task, known as MAP in Markov Random Fields and MPE in Bayesian Networks, is to find a global assignment of all the variables with maximum a posteriori probability/minimum energy. A usual distinction on MAP solving methods is complete/incomplete, i.e. the ability to prove optimality or not. Most complete methods rely on tree search, while incomplete methods rely on local search. Among them, we study Variable Neighborhood Search (VNS) for graphical models. In this paper, we propose an iterative approach above VNS that uses (partial) tree search inside its local neighborhood exploration. The proposed approach performs several neighborhood explorations of increasing search complexity, by controlling two parameters, the discrepancy limit and the neighborhood size. Thus, optimality of the obtained solutions can be proven when the neighborhood size is maximal and with unbounded tree search. We further propose a parallel version of our method improving its anytime behavior on difficult instances coming from a large graphical model benchmark. Last we experiment on the challenging minimum energy problem found in Computational Protein Design, showing the practical benefit of our parallel version. A solver is available at <https://github.com/toulbar2/toulbar2>.

*Keywords:*

---

## 1. Introduction

2 Probabilistic graphical models [1] are formed by variables linked to each other  
3 by stochastic relationships. They enable to model complex systems with hetero-

*Preprint submitted to Artificial Intelligence Journal*

*September 3, 2019*

4 geneous data and to capture uncertainty. Graphical models have been applied in  
5 a wide range of areas such as image analysis, speech recognition, bioinformatics,  
6 and ecology.

7 We focus on models with discrete variables like Markov Random Field and  
8 Bayesian Network. Our goal is to find a global assignment of all the variables with  
9 maximum a posteriori probability. This optimization task defines an NP-complete  
10 problem [2]. Solving methods can be categorized in two groups: exact and local  
11 search methods. Exact methods rely on tree search, variable elimination, linear  
12 programming, or a combination of them [3, 4, 5]. Graph-cut and message-passing  
13 algorithms like loopy belief propagation and variational approaches [6, 7, 8, 9, 10]  
14 are exact only in some particular cases (*e.g.*, binary image denoising or tree struc-  
15 tured problems). Local search methods are stochastic algorithms like Gibbs sam-  
16 pling, Guided Local Search [11, 12], and Stochastic Greedy Search [13]. Some  
17 of them have theoretical asymptotic proof of convergence, *i.e.*, the optimal solu-  
18 tion is guaranteed to be found if infinite time is available. In practice, they may  
19 exhibit a better anytime behavior than exact methods on large and difficult prob-  
20 lems [14, 12, 13], *i.e.*, they produce better solutions in less time.

21 A few attempts have been done to combine exact and local search methods.  
22 A simple way is to run sequentially a local search algorithm then tree search,  
23 where solutions found by local search will be used as initial upper bounds for  
24 branch and bound exact methods. Another approach is to design a local search  
25 framework where the neighborhood exploration is performed by tree search in  
26 a systematic or non-systematic way as it is done in Large Neighborhood Search  
27 (LNS) [15, 16, 17, 18, 19] and Variable Neighborhood Search (VNS) [20, 21].  
28 VNS/LDS+CP [21] combines a metaheuristic, VNS, with Limited Discrepancy  
29 Search (LDS) [22], a partial tree search method (Section 2). We propose in this  
30 paper an iterative variant of VNS/LDS+CP, called Unified Decomposition Guided  
31 VNS (UDGVNS), adapted to graphical models and able to prove optimality when  
32 the neighborhood size is maximal and with unbounded tree search.

### 33 *Contributions and plan.*

- 34 1. We introduce UDGVNS, a new iterative approach above DGVNS<sup>1</sup> (Decom-  
35 position Guided VNS) unifying complete and incomplete search methods.  
36 UDGVNS restores the completeness of DGVNS by applying successive calls  
37 with an increasing discrepancy limit.

---

<sup>1</sup>DGVNS [23] exploits, within VNS/LDS+CP, structural knowledge coming from tree decom-  
position in order to efficiently guide the exploration of large neighborhoods (Section 2).

- 38 2. We describe a coarse-grained parallel version called UPDGVNS (for Unified  
 39 Parallel DGVNS) allowing asynchronous cooperative execution of UDGVNS  
 40 processes with centralized information exchange as in [24, 25]. As for  
 41 UDGVNS, the parallel release enables one to control the compromise be-  
 42 tween optimality proof and anytime behavior. Compared to UDGVNS, the  
 43 parallel version enables to improve the anytime behavior on difficult in-  
 44 stances.
- 45 3. We propose a new operator denoted `add1/jump` for managing the neigh-  
 46 borhood size  $k$  that exploits the graph of clusters provided by a tree decom-  
 47 position of the problem.
- 48 4. We present an extensive empirical study that includes a wide range of in-  
 49 stances coming from various benchmarks (Cost Function Network (CFN),  
 50 Computer Vision and Pattern Recognition (CVPR), Uncertainty in Artificial  
 51 Intelligence (UAI) 2008 and Probabilistic Inference Challenge (PIC)  
 52 2011) which compares our techniques to state-of-the-art ones. Experimen-  
 53 tal results show that our approaches offer a good compromise between the  
 54 number of problems completely solved compared to the quality of the best  
 55 solution found.
- 56 5. We report experiments on the challenging minimum energy problem in  
 57 Computational Protein Design (CPD). For this aim, we designed new larger  
 58 instances that are well structured and supposed to be more difficult to solve  
 59 than those generated in [26]. We show the practical benefit of our ap-  
 60 proaches compared to TOULBAR2 and FIXBB<sup>2</sup> the simulated annealing al-  
 61 gorithm provided by the Rosetta package for CPD.

62 The paper is organized as follows. Section 2 recalls preliminaries. Section 3  
 63 presents UDGVNS. Section 4 describes our parallel version UPDGVNS. Sections 5-  
 64 6 report experiments we performed. Finally, we conclude and draw some perspec-  
 65 tives.

## 66 2. Preliminaries

### 67 2.1. Graphical Model

68 **Definition 1.** A probabilistic graphical model (or Gibbs distribution) [1] is a  
 69 triplet  $(\mathcal{X}, \mathcal{D}, \mathcal{F})$  with  $\mathcal{X} = \{X_1, \dots, X_n\}$ , a set of  $n$  random variables,  $\mathcal{D} =$   
 70  $\{D_1, \dots, D_n\}$ , a set of finite domains of values of maximum size  $d = \max_{i=1}^n |D_i|$ ,

---

<sup>2</sup>fixed backbone design application

71 and  $\mathcal{F}$ , a set of potential functions. Each variable  $X_i$  takes values in  $D_i$ . An as-  
 72 signment of  $\mathcal{X}$  is a set  $x = (x_1, \dots, x_n)$ , with  $x_i \in D_i$ . The set of all possible  
 73 assignments of  $\mathcal{X}$  is denoted  $\Delta = \prod_{i=1}^n D_i$ . Let  $A = \{D'_1, \dots, D'_n\}$  with  $D'_i \subseteq D_i$   
 74 represent a restricted set of  $\Delta$  called a partial assignment. If  $S$  is a subset of  
 75  $V = \{1, \dots, n\}$ ,  $X_S$ ,  $x_S$  and  $\Delta_S$  are respectively the subset of random variables  
 76  $\{X_i, i \in S\}$ , the assignment  $(x_i, i \in S)$  obtained from  $x$ , and the set of all possi-  
 77 ble assignments of  $X_S$ . Given a set  $\mathcal{S}$  of partitions of  $V$ , the set  $\mathcal{F} = \{f_S\}_{S \in \mathcal{S}}$  of  
 78 maps from  $\Delta_S$  to  $\mathbb{R}^+$  is said to factorize a joint probability distribution  $\mathbb{P}$  iff:

$$\mathbb{P}(x) = \frac{1}{Z} \prod_{f_S \in \mathcal{F}} f_S(x_S) \quad (1)$$

79 where  $Z = \sum_{x \in \Delta} \prod_{f_S \in \mathcal{F}} f_S(x_S)$  is the normalizing constant, also called parti-  
 80 tion function.

Among the various tasks, the *Most Probable Explanation* (MPE) problem is to find the most likely assignment  $x \in \Delta$  to all the variables in  $\mathcal{X}$  maximizing  $\mathbb{P}(x)$ . By taking the opposite of the logarithm of  $\mathbb{P}(x)$ , i.e.,

$$-\log \mathbb{P}(x) = \sum_{f_S \in \mathcal{F}} -\log f_S(x_S) + \log Z = \sum_{f_S \in \mathcal{F}} \varphi(x_S) + \log Z$$

81 we obtain an additive model with  $\varphi(x_S)$  called an *energy function*. Finding a  
 82 solution of minimum energy is equivalent to MPE. In the rest of the paper, we  
 83 consider energy minimization. When  $\varphi(x_S)$  maps to  $\mathbb{N}^+ \cup \{\infty\}$ , the correspond-  
 84 ing deterministic graphical model is called a *Cost Function Network* (CFN) [27].  
 85 Finding a solution of minimum cost is the same as doing energy minimization on  
 86 the equivalent probabilistic model [28].

87 Specific solving methods have been proposed to solve these problems but two  
 88 general approaches can be considered. The first one applies traditional search  
 89 techniques based on backtracking or branch and bound. In the worst case, their  
 90 time complexity is in  $O(d^n)$  while being generally linear in space. The second one  
 91 relies on methods that exploit the notion of decomposition of graphs and which are  
 92 based on Dynamic Programming (DP) (see Section 2.4.1). These methods make it  
 93 possible to guarantee a time complexity in  $O(d^w)$  (where  $w$  is the minimal width  
 94 over all the tree decompositions) but with an exponential space complexity.

## 95 2.2. DFBB and Limited Discrepancy Search

96 Depth-First Branch and Bound (DFBB) methods explore a search tree in a sys-  
 97 tematic way by recursively choosing the next unassigned variable to assign and by

98 choosing a value in its domain for the assignment (the *branch* part) until a better  
 99 solution is found or it can be proved that the subtree rooted at the current search  
 100 node has no better solutions and it can be pruned (the *bound* part). DFBB depends  
 101 on its variable and value ordering heuristics for branching in order to find good  
 102 solutions rapidly and to reduce the size of the search tree to be explored. It also  
 103 depends on its lower bound computation in order to prune the search for mini-  
 104 mization problems. Typically, lower bounds are built by dynamic programming  
 105 with bounded memory, such as mini-buckets heuristic [29], or by solving a lin-  
 106 ear relaxation of the problem or its dual in an exact or approximate way. In the  
 107 experiments, we exploit during search an approximate dual lower bound called  
 108 Existential Directional Arc consistency (EDAC) [30] that performs fast incremen-  
 109 tal problem reformulations with extra domain value pruning. More information  
 110 can be found in [31].

111 Limited Discrepancy Search (LDS) [22] is a heuristic method that explores  
 112 the search tree in a non-systematic way by making a limited number of *wrong*  
 113 decisions w.r.t. its value ordering heuristic. We assume a binary search tree where  
 114 at each search node either the selected variable is assigned to its chosen preferred  
 115 value (left branch) or the value is removed from the domain (right branch). Each  
 116 value removal corresponds to a wrong decision made by the search, it is called  
 117 a *discrepancy*. The number of discrepancies is limited by a parameter denoted  
 118  $\ell$ . See Algorithm 1, where  $\text{lb}(A)$  gives a lower bound on the minimum energy  
 119  $\min_{x \in \prod_{D_i \in A} D_i} -\log \mathbb{P}(x)$  of the partial assignment  $A$ .

120 In order to detect if a complete search has been done, LDS returns true if and  
 121 only if the discrepancy limit is never reached. Otherwise it returns false as soon  
 122 as  $\ell = 0$  (line 1). If it returns true then LDS is equivalent to a complete DFBB.

123 In order to produce better quality solutions as time passes, a simple strategy  
 124 is to iterate LDS with an increasing number of discrepancies  $\ell$  going from  $\ell_{min}$   
 125 to  $\ell_{max}$ . See Algorithm 2 for this Iterative LDS (ILDS) method, where  $+_\ell$  is a  
 126 special sum operator that will be discussed in Section 3. The minimum energy  
 127 and its corresponding solution are provided in global variables  $ub$  and  $x$ . In the  
 128 sequel, we give no initial upper bound ( $ub = \infty$ ).

129 **Proposition 1.**  $\text{ILDS}(0, n(d-1), +, \infty, \{\})$  is a complete method with a worst-  
 130 case time complexity exponential in the number of variables and a linear space  
 131 complexity.

132 *Proof.* First, we prove that ILDS returns true if and only if optimality was proven.  
 133 Each iteration does at most  $\ell$  discrepancies along the path from the root search  
 134 node to a terminal node. With a sufficiently large discrepancy limit, LDS never

---

**Algorithm 1: Limited Discrepancy Search algorithm**


---

```

Function LDS ( $\ell, A, ub : In/Out, x : Out$ ) : Boolean
   $left \leftarrow \mathbf{true}; right \leftarrow \mathbf{true};$ 
  if ( $\exists D_i \in A, |D_i| > 1$ ) then
    Choose an unassigned variable  $X_i \in \mathcal{X}$  such that  $|D_i| > 1$ ;
    Choose a value  $x_i \in D_i$ ;
     $A' \leftarrow (A \setminus \{D_i\}) \cup \{\{x_i\}\}$ ;
    if ( $\text{lb}(A') < ub$ ) then
      |  $left \leftarrow \text{LDS}(\ell, A', ub, x)$ ; // left branch
    if ( $\ell > 0$ ) then
      |  $A'' \leftarrow (A \setminus \{D_i\}) \cup \{D_i \setminus \{x_i\}\}$ ;
      | if ( $\text{lb}(A'') < ub$ ) then
      | |  $right \leftarrow \text{LDS}(\ell - 1, A'', ub, x)$ ; // right branch
    else
  1 | return false; // search is incomplete
  else
  2 |  $ub \leftarrow \text{lb}(A), x \leftarrow A$ ; // new solution found
  return  $left \wedge right$ ; // true if both branches are complete

```

---

135 reaches its discrepancy limit ( $\ell = 0$ ) and explores a complete search tree, so both  
 136 LDS and ILDS (line 3) return true. Because we have in the worst case  $n(d - 1)$   
 137 value removals (right branches) to reach a terminal node before assigning all the  
 138 variables, we can set  $\ell_{max} = n(d - 1)$  and at least, the last iteration is complete<sup>3</sup>,  
 139 *i.e.*,  $\text{LDS}(n(d - 1), \mathcal{D}, ub, x)$  is equivalent to a complete DFBB. Here,  $d$  is the  
 140 maximum domain size of all the variables.

141 Another condition for completeness is reached when LDS finds a solution with  
 142 a cost equal to a known lower bound of the problem. In this case, LDS will stop  
 143 branching and ILDS will return true before reaching its last iteration (line 3).

144 For simplicity reasons, let assume variables with Boolean domains ( $d = 2$ ).  
 145 The maximum height  $h$  of the explored search tree is therefore equal to the number  
 146 of variables  $h = n$ . The number of terminal nodes with exactly  $\ell$  discrepancies is  
 147 bounded by  $\binom{h}{\ell}$ . In the worst case, ILDS runs for  $\ell = 0..h$ . The number of termi-  
 148 nal nodes of LDS for  $\ell = h$  is equal to  $\binom{h}{0} + \binom{h}{1} + \binom{h}{2} + \dots + \binom{h}{h} = \sum_{0 \leq k \leq h} \binom{h}{k} = 2^h$ .  
 149 Thus in the worst case, one iteration of LDS has a time complexity in  $\Theta(2^h)$ . By  
 150 doing at most  $h + 1$  iterations (from  $\ell = 0$  to  $\ell_{max} = h$ ), ILDS will explore at  
 151 least  $\Omega(2^h)$  terminal nodes. The asymptotic time complexity of ILDS is therefore

---

<sup>3</sup>In practice,  $\ell$  was less than or equal to 128 for all instances completely solved within 1 hour CPU time limit by LDS and VNS methods in Section 5.

---

**Algorithm 2:** Iterative LDS algorithm

---

```
Function ILDS ( $\ell_{min}, \ell_{max}, +\ell, ub : In/Out, x : Out$ ) : Boolean
   $r \leftarrow 0$ ; // number of discrepancy iterations
   $\ell \leftarrow \ell_{min}$ ; // initial discrepancy limit
  while ( $\ell \leq \ell_{max}$ ) do
     $opt \leftarrow \text{LDS}(\ell, \mathcal{D}, ub, x)$ ;
  3   if ( $opt \vee ub = \text{lb}(\mathcal{D})$ ) then return true;
     $r \leftarrow r + 1$ ;
  4   if ( $\ell < \ell_{max}$ ) then  $\ell \leftarrow \min(\ell_{max}, \ell_{min} + \ell r)$ ;
    else  $\ell \leftarrow \infty$ ;
  return false ;
```

---

152 exponential in the number of variables.

153 Because LDS has a linear space complexity, thanks to its depth-first search  
154 principle as in DFBB, ILDS has also a linear space complexity.  $\square$

155 In [32], a similar stopping condition for optimality proof was presented. Be-  
156 cause heuristics are often less informed near the root of the search tree, it is usually  
157 better to make wrong decisions at the beginning of the search [22, 33]. It favors  
158 exploring new parts of the search tree, possibly finding better solutions that will  
159 prune the remaining part of the search tree already explored at previous iterations.  
160 Our actual implementation of LDS exploits this fact (right branch done before left  
161 branch when the discrepancy limit is not reached).

### 162 2.3. Variable Neighborhood Search

163 VNS [20] is a metaheuristic that uses a finite set of pre-selected neighborhood  
164 structures  $N_k, k = 1, 2, \dots, k_{max}$  to escape from local minima by systematically  
165 changing the neighborhood structure if the current one does not improve the cur-  
166 rent incumbent solution. VNS repeatedly performs three major steps. In the first  
167 one, called *shaking*, a solution  $x'$  is randomly generated in the neighborhoods of  
168  $x$  denoted  $N_k(x)$ . In the second one, a local search method is applied from  $x'$   
169 to obtain a local optimum  $x''$ . In the third one, called *neighborhood change*, if  $x''$   
170 is better, than  $x$  is replaced with  $x''$  and  $k$  is set to 1; otherwise,  $k$  is increased by  
171 one.

172 The use of VNS scheme for solving deterministic graphical models started  
173 with VNS/LDS+CP [21] and improved in DGVNS [23] (see section 2.4). This  
174 approach is related to LNS [15], but it adjusts dynamically the neighborhood size  
175 and exploits a tree decomposition of the constraint graph of the problem when the  
176 search seems to stagnate as in VNS.

177 **2.4. Decomposition Guided Variable Neighborhood Search**

178 Recently, Fontaine *et al.* [23] investigated the incorporation of tree decompo-  
 179 sition in order to efficiently guide the exploration of large neighborhoods. They  
 180 proposed Decomposition Guided VNS (DGVNS), a first local search approach that  
 181 exploits the graph of clusters provided by a tree decomposition of the constraint  
 182 graph of the problem to build relevant neighborhood structures. The next section  
 183 2.4.1 defines formally the constraints tree decomposition. Then we present in sec-  
 184 tion 2.4.2 the construction of the initial solution exploited by DGVNS. In section  
 185 2.4.3, we detail the main DGVNS algorithm, and show in section 2.4.4 how to  
 186 instantiate VNS/LDS+CP algorithm from DGVNS. Finally, we briefly discuss the  
 187 impact of a tree decomposition on the performance of DGVNS.

188 **2.4.1. Tree decomposition**

189 **Definition 2.** A tree decomposition of a connected graphical model  $G$  is a pair  
 190  $(C_T, T)$  where  $T = (I, A)$  is a tree with nodes set  $I$  and edges set  $A$  and  $C_T =$   
 191  $\{C_i \mid i \in I\}$  is a family of subsets of  $\mathcal{X}$ , called clusters, such that: (i)  $\cup_{i \in I} C_i = \mathcal{X}$ ,  
 192 (ii)  $\forall f_S \in \mathcal{F}, \exists C_i \in C_T$  s.t.  $S \subseteq C_i$ , (iii)  $\forall i, j, k \in I$ , if  $j$  is on the path from  $i$   
 193 to  $k$  in  $T$ , then  $C_i \cap C_k \subseteq C_j$ .

194 **Definition 3.** The intersection of two clusters  $C_i$  and  $C_j$  is called a separator, and  
 195 noted  $sep(C_i, C_j)$ .

196 **Definition 4.** A graph of clusters for a tree decomposition  $(C_T, T)$  is an undi-  
 197 rected graph  $G = (C_T, E)$  that has a vertex for each cluster  $C_i \in C_T$ , and there  
 198 is an edge  $(C_i, C_j) \in E$  when  $sep(C_i, C_j) \neq \emptyset$ .

199 The width  $w$  of a tree decomposition  $(C_T, T)$  is equal to  $\max_{C_i \in C_T} |C_i| - 1$ .  
 200 The treewidth  $w^*$  of  $G$  is the minimal width over all the tree decompositions of  $G$ .

201 Several studies have focused on the computation of tree decompositions [34,  
 202 35]. The proposed algorithms can be classified into two approaches: the exact  
 203 algorithms that compute decompositions of optimal width (equal to the treewidth)  
 204 and the heuristic methods, which do not offer a guarantee on optimality. The mo-  
 205 tivation of heuristic approaches is due to the fact that the optimal computation is a  
 206 NP-hard problem [36]. Heuristic approaches dedicated to handle graphical mod-  
 207 els often use triangulation (such as *Minimum Fill-in* (min-fill) [37] and *Maximum*  
 208 *Cardinality Search* (MCS) [38]). These heuristic approaches allow to process  
 209 graphs of several thousand of vertices in reasonable time, but without guaran-  
 210 teeing the quality of the obtained decompositions in terms of deviation from the  
 211 optimum  $w^*$ . Other heuristics and pre- or post-processing rules may be applied to

---

**Algorithm 3: Decomposition Guided VNS algorithm**


---

```

Procedure DGVNS ( $\ell, k_{min}, k_{max}, ub : In/Out, x : Out$ )
  let  $(C_T, T)$  be a tree decomposition of  $(\mathcal{X}, \mathcal{D}, \mathcal{F})$ ;
  LDSr ( $n(d-1), \mathcal{D}, ub, x$ );           // find an initial solution
   $c \leftarrow 1$ ;                       // current cluster index
   $i \leftarrow 0$ ;                       // nb. of successive failed neighborhood sizes
   $k \leftarrow k_{min}$ ;                 // initial neighborhood size
  while ( $k \leq k_{max} \wedge \neg TimeOut$ ) do
5    $A \leftarrow \text{getNeighborhood}(x, C_c, k)$ ;
    $ub' \leftarrow ub$ ;
6   LDSr ( $\ell, A, ub', x'$ );           // neighborhood search
   if ( $ub' < ub$ ) then
7      $x \leftarrow x', ub \leftarrow ub'$ ; // new best solution
8      $i \leftarrow 0, k \leftarrow k_{min}$ ;
   else
      $i \leftarrow i + 1$ ;
9      $k \leftarrow \min(k_{max}, k_{min} + i)$ ;
10     $c \leftarrow 1 + c \bmod |C_T|$ ;     // visit next cluster
Function getNeighborhood ( $x, C, k$ )
  if  $k \geq |\mathcal{X}|$  then
11     $X_{un} \leftarrow \mathcal{X}$ 
  else
12     $Cand \leftarrow \text{CompleteCluster}(C, k)$ ;
13     $X_{un} \leftarrow \text{Random}(Cand, k)$ ; // Random selection of  $k$  conflict
    variables
14     $A \leftarrow \{D_i \mid X_i \in X_{un}\} \cup \{\{x_i\} \mid X_i \in \mathcal{X} \setminus X_{un}\}$ ; // Unassign selected
    variables
  return  $A$ ;

```

---

212 reduce the width of the decomposition [35, 34]. In this paper, we use a heuristic  
 213 approach based on min-fill.

#### 214 2.4.2. Initial solution and restricted LDS

215 VNS relies on an initial solution  $x$ . Without any infinite terms in the problem  
 216 (corresponding to forbidden assignments or hard constraints),  $x$  can be produced  
 217 by a greedy search algorithm such as LDS ( $0, \mathcal{D}, \infty, x$ ). Otherwise, we can ei-  
 218 ther relax the problem (by replacing every infinite energy term by the sum of the  
 219 greatest finite term of each original energy function in the problem) or rely on a  
 220 complete search method. For that, we made a modified version of LDS, called

221 Restricted LDS ( $\text{LDS}^r$ ), that stops immediately after a first solution is found<sup>4</sup>.

222 In particular,  $\text{LDS}^r(n(d-1), \mathcal{D}, \infty, x)$  will either find a solution of finite cost  
223 (*i.e.*, satisfying all the constraints) or prove the problem has no feasible solution.

### 224 2.4.3. DGVNS algorithm

225 Algorithm 3 shows the pseudo-code of DGVNS. It exploits the graph of clusters  
226 provided by a tree decomposition of the constraint graph of the problem to build  
227 relevant neighborhood structures. Let  $\mathcal{X}$  be the set of variables, and let  $N_{k,c}$  be  
228 the neighborhood structure, where  $k$  is the neighborhood size and  $C_c$  is the cluster  
229 where the variables will be selected from.

230 First, an initial solution  $x$  is generated by  $\text{LDS}^r$ , as detailed in section 2.4.2.  
231 Second, to favor moves on regions that are closely linked,  $x$  is partially destroyed  
232 by unassigning a subset of  $k$  variables and an exploration of its (large) neighbor-  
233 hood is performed until the solution is repaired with a new one. To select the  
234 variables to be unassigned, DGVNS uses a neighborhood heuristic based on clus-  
235 ters (see function `getNeighborhood`, line 5): the set of candidate variables  
236  $C_{and}$  to be unassigned are selected in the same cluster  $C_c$ . Indeed, the concept of  
237 cluster embodies this criterion, because of its size (smaller than the original prob-  
238 lem), and by the strong connection of the variables it contains. If ( $k > |C_c|$ ), we  
239 complete the set of candidate variables to be unassigned by adding clusters adja-  
240 cent to  $C_c$  in order to take into account the topology of the graph of clusters. This  
241 treatment is achieved by function `CompleteCluster` ( $C_c, k$ ) (line 12). Third, a  
242 subset of  $k$  variables  $X_{un}$  is randomly selected in  $C_{and}$  (line 13), and then rebuilt  
243 using  $\text{LDS}^r$ . In the particular case where  $k$  is greater than the variables cardinality  
244 (line 11), the whole variables are selected. The neighborhood change in DGVNS  
245 is performed in the same way as in VNS. However, DGVNS considers successively  
246 all the clusters  $C_c$ . This ensures a better diversification by covering a large number  
247 of different regions, and to locate the region containing the global optimum.

248 The core of DGVNS is its reconstruction phase. It relies on a non local solver  
249 combining constraint propagation and Restricted LDS with a fixed discrepancy  
250 to explore the neighborhood of the solution. One advantage of this choice is its  
251 exploration speed that improves the quality profile and allows a more balanced  
252 exploration of the search tree. First, a subset of  $k$  variables are selected in  $\mathcal{X}$ .  
253 Then, a partial assignment  $A$  is generated from the current solution  $x$  by unfixing  
254 the  $k$  selected variables, and then re-built in the best way (line 6).

---

<sup>4</sup>At the end of line 2 of Alg. 1, it stops the recursive LDS procedure and returns false.

255 Let *succ* a successor function<sup>5</sup> and  $N_{k,c}$  the current neighborhood structure: if  
 256  $LDS^r$  finds a (first) solution of better quality  $x'$  in the neighborhood of  $x$  (line 6),  
 257 then  $x'$  becomes the current solution (line 7),  $k$  is reset to  $k_{min}$  (line 8), and the  
 258 next cluster is considered (line 10). Otherwise, DGVNS looks for improvements in  
 259  $N_{(k+1),succ(c)}$ , a neighborhood structure where  $(k + 1)$  variables will be unassigned  
 260 (line 9). In fact, when a local minimum is found in the current neighborhood,  
 261 moving from  $k$  to  $(k + 1)$  will also provide some diversification by enlarging the  
 262 neighborhood size. The search stops when it reaches the maximal neighborhood  
 263 size allowed or a timeout.

#### 264 2.4.4. From DGVNS to VNS/LDS+CP

265 When a tree decomposition of the constraint graph of the problem is not avail-  
 266 able, the constraints can be handled as a single cluster (i.e.  $|C_T| = 1$ ), and DGVNS  
 267 behaves as VNS/LDS+CP algorithm [23]. In this case, the search process com-  
 268 pletes the variables with  $LDS^r$  without taking into account the constraints connec-  
 269 tivity, which is intuitively less efficient than the solving process when  $|C_T| > 1$ .  
 270 More precisely, when the variables  $X_{un}$  are strongly connected through the con-  
 271 straint graph,  $LDS^r$  will efficiently instantiate these variables thanks to constraint  
 272 propagation, called at every LDS search node, where it is well known that its ef-  
 273 fectiveness depends on the connectivity of the variables to be instantiated. This  
 274 idea is cleverly exploited by DGVNS thanks to the concept of cluster provided  
 275 by the tree decomposition of the constraint graph. In VNS/LDS+CP, the neigh-  
 276 borhood heuristic (function `getNeighborhood`) randomly selects  $k$  variables  
 277 to unassign among conflicted ones. Such a heuristic which is mainly based on  
 278 random choices has a major drawback since it does not take advantage of the  
 279 topology of the constraint graph. For instance, it may select unrelated variables  
 280 (i.e., no constraint may be fully unassigned), and all selected variables may also  
 281 have a high degree (i.e., they occur in many constraints). In such a case, it is  
 282 unlikely to rebuild them without violating several constraints, and thus to find a  
 283 solution of better quality than the current one. Nevertheless, as reported by [23],  
 284 VNS/LDS+CP remains efficient on some problems, but DGVNS is much more  
 285 efficient.

#### 286 2.4.5. Synthesis

287 The motivation of exploiting the graph of clusters of a tree decomposition  
 288 within VNS/LDS+CP algorithm is to build pertinent neighborhood structures en-

---

<sup>5</sup> $succ(c) = 1 + c \text{ mod } |C_T|$ .

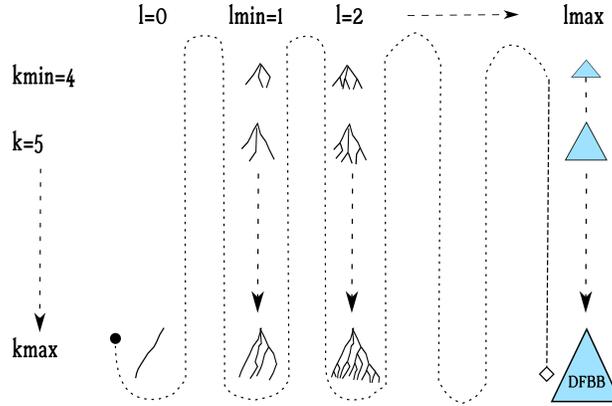


Figure 1: A general overview of UDGvNS algorithm, exploring successively different search trees, starting from an initial greedy search ( $\bullet$ ), and ending to a complete search ( $\diamond$ ).

289 abling a better diversification. Clearly, the quality of the tree decomposition im-  
 290 pacts greatly the performance of DGVNS.

291 In our prior works [23], we have studied the impact of some parameters related  
 292 to topological properties of the tree decomposition: the *width* of a tree decompo-  
 293 sition  $w$ , *separators size*, and *the decomposability* of a problem ( $\frac{w}{n}$ ), estimated by  
 294 the ratio between the width of a tree decomposition and the number of variables.  
 295 The width of a tree decomposition gives a good indication on the size of subprob-  
 296 lems, while separators size provides information about the connectivity between  
 297 clusters and the degree of their overlap.

298 From this study, we showed that DGVNS is very effective on problems that  
 299 decompose into *weakly connected clusters of reasonable size*, i.e. tree decompo-  
 300 sitions characterized by low values of ( $\frac{w}{n}$ ), and by separators of small size (clusters  
 301 that do not overlap heavily), leading to more pertinent neighborhoods.

### 302 3. Unified Decomposition Guided VNS

303 We present UDGvNS in Algorithm 4, an iterative DGVNS method, unifying two  
 304 complete and incomplete search methods. As done by Iterative LDS, UDGvNS re-  
 305 stores the completeness of DGVNS by applying successive calls with an increasing  
 306 discrepancy limit.

#### 307 3.1. UDGvNS algorithm

308 As in the previous VNS algorithms, the initial solution of UDGvNS is obtained  
 309 by LDS<sup>r</sup> (line 15), corresponding to a greedy search with no discrepancy if the

---

**Algorithm 4: Unified Decomposition Guided VNS algorithm**

---

```
Procedure UDGVNS ( $\ell_{min}, \ell_{max}, +\ell, k_{min}, k_{max}, +k, ub : In/Out, x : Out,$   
 $opt : Out$ )  
  let  $(C_T, T)$  be a tree decomposition of  $(\mathcal{X}, \mathcal{D}, \mathcal{F})$  ;  
15   $opt \leftarrow \text{LDS}^r(n(d-1), \mathcal{D}, ub, x)$  ; // find an initial solution  
16  if  $(ub = \text{lb}(\mathcal{D}))$  then  $opt \leftarrow \text{true}$  ;  
     $c \leftarrow 1$  ; // current cluster index  
     $r \leftarrow 0$  ; // number of discrepancy iterations  
     $\ell \leftarrow \ell_{min}$  ; // initial discrepancy limit  
    while  $(\neg opt \wedge \ell \leq \ell_{max})$  do  
       $i \leftarrow 0$  ; // nb. of successive failed neighborhood sizes  
17       $k \leftarrow k_{min}$  ; // initial neighborhood size  
      while  $(\neg opt \wedge k \leq k_{max})$  do  
         $A \leftarrow \text{getNeighborhood}(x, C_c, k)$  ;  
         $ub' \leftarrow ub$  ;  
18         $opt \leftarrow \text{LDS}^r(\ell, A, ub', x')$  ; // neighborhood search  
19        if  $(ub' = \text{lb}(\mathcal{D}))$  then  $opt \leftarrow \text{true}$  ;  
20        else if  $(A \neq \mathcal{D})$  then  $opt \leftarrow \text{false}$  ;  
21        if  $(ub' < ub)$  then  
           $x \leftarrow x', ub \leftarrow ub'$  ; // new best solution  
22           $i \leftarrow 0, k \leftarrow k_{min}$  ;  
23           $r \leftarrow 0, \ell \leftarrow \ell_{min}$  ;  
        else  
           $i \leftarrow i + 1$  ;  
          if  $(k < k_{max})$  then  
             $k \leftarrow \min(k_{max}, k_{min} + k \cdot i)$  ;  
          else  $k \leftarrow \infty$  ;  
24         $c \leftarrow 1 + c \bmod |C_T|$  ; // visit next cluster  
         $r \leftarrow r + 1$  ;  
        if  $(\ell < \ell_{max})$  then  
           $\ell \leftarrow \min(\ell_{max}, \ell_{min} + \ell \cdot r)$  ;  
        else  
           $\ell \leftarrow \infty$   
      // End
```

---

310 problem to be solved has only finite energy terms. Then UDGVNS tries to improve  
311 the current solution by doing several neighborhood explorations of increasing  
312 search complexity, by controlling two parameters, the discrepancy limit ( $\ell$ ) and  
313 the neighborhood size ( $k$ ), as shown in Fig. 1. It starts from a small neighborhood  
314 with a few variables unassigned ( $k = k_{min} = 4$ ). It explores the neighborhood  
315 using LDS with a small discrepancy limit initially set to one ( $\ell = \ell_{min} = 1$ ). The  
316 unassigned variables are selected from a current cluster of a tree decomposition

317 (and its neighbor clusters if needed) as done in DGVNS. If no solution is found  
 318 then UDGVNS increases its neighborhood size until all the variables of the prob-  
 319 lem are included in the neighborhood ( $k = k_{max} = n$ ). If still no solution is found  
 320 then UDGVNS increases its discrepancy limit and resets its neighborhood size to  
 321  $k_{min}$ . The last iteration corresponds to a complete search on the whole problem  
 322 for proving optimality ( $k = k_{max} = n, \ell = \ell_{max} = n(d - 1)$ ). As soon as a  
 323 better solution is found by the current neighborhood search (line 18), UDGVNS  
 324 stops the search in order to reinitialize the two parameters to their minimum value  
 325 (lines 22–23). By doing so, it favors finding next solutions more rapidly, as it  
 326 is faster to explore many small neighborhoods than a larger one, improving the  
 327 anytime behavior of the search.

328 **Proposition 2.** UDGVNS  $(1, n(d-1), +, 1, n, +, \infty, \{\}, opt)$  is a complete method  
 329 with a worst-case time complexity exponential in the number of variables and a  
 330 linear space complexity.

331 *Proof.* For UDGVNS, optimality can be proven in two cases: (i) when the current  
 332 neighborhood corresponds to the whole problem (condition falsified at line 20,  
 333 since  $(A \neq \mathcal{D})$  is false) and the discrepancy value is greater than or equal to the  
 334 maximum number of right branches (checked as before during LDS searches at  
 335 lines 15,18), and thus DGVNS behaves as an exhaustive search, or (ii) by exam-  
 336 ining the bounds at the root node (line 16) and after each neighborhood search  
 337 (line 19). In this case, the search space is implicitly explored by the algorithm,  
 338 therefore optimality is proven. Case (i) is always reached when  $k = k_{max} = n$   
 339 and  $\ell = \ell_{max} = n(d - 1)$ , corresponding here to the last iteration of the two *While*  
 340 loops of UDGVNS. Notice that in this case ( $k = k_{max}$ ) and ( $\ell = \ell_{max}$ ), all the vari-  
 341 ables are selected candidates (see line 11) to be explored exhaustively by an LDS  
 342 stopping at the first solution better than the current bound. If the current bound is  
 343 not optimal, the search restarts until reaching optimality. In practice, optimality  
 344 proofs are often produced at smaller  $\ell$ , but still for  $k = n$  (or before if condition  
 345 at line 19 becomes true).

346 Assuming a complete search tree over Boolean variables ( $d = 2$ ), the worst-  
 347 case time complexity of the initial  $LDS^r$  at line 15 is in  $\Theta(2^n)$ . The inner  $LDS^r$   
 348 at line 18 with  $k$  variables and  $\ell$  discrepancies has a search tree with maximum  
 349 height  $h = k$ , assuming variables with Boolean domains. In the worst case, its  
 350 asymptotic time complexity is in  $\Theta(2^h)$  for  $\ell = h = k$  (see Proposition 1). The  
 351 number of  $LDS^r$  calls depends on the problem upper bound. Each time a strictly  
 352 better upper bound is found (line 21),  $LDS^r$  is stopped and we reset  $k$  and  $\ell$  to  
 353 their minimum value (lines 22–23). When the energy functions map to  $\mathbb{N}^+ \cup \{\infty\}$

354 as in Cost Function Network (CFN), there will be a finite number of upper bound  
355 improvements. Moreover if  $+_\ell$  and  $+_k$  are strictly increasing functions, then there  
356 is a finite number of  $\text{LDS}^r$  calls until  $\ell = \ell_{max}$  and  $k = k_{max}$ , and no better  
357 solution exists (otherwise,  $\ell = \ell_{min}$  and  $k = k_{min}$ , and the search will continue).  
358 Thus, UDGVNS terminates and returns  $opt = \mathbf{true}$  if and only if it exists a feasible  
359 optimal solution. It has the same exponential time and linear space complexities  
360 as ILDS.  $\square$

### 361 3.2. Strategies for managing parameters of UDGVNS

362 UDGVNS has to control the evolution of two parameters. For each parameter,  
363  $\ell$  and  $k$ , we have tested three updating rules: increase by one at each iteration  
364 ( $+_{\ell/k} = +$ ), multiply by two at each iteration ( $a +_{\ell/k} b = \text{mult}2(a, b) = a \times 2^b$ ),  
365 and apply a Luby strategy [39] ( $a +_{\ell/k} b = \text{Luby}(a, b) = a \times \text{luby}(1 + b)$ )<sup>6</sup>.

366 Operator  $+_k$  controls the compromise between intensification and diversifica-  
367 tion. The goal of the Luby strategy is to exponentially increase the number of  
368 small neighborhoods explored compared to the number of larger ones. Whereas  
369 classical VNS algorithms will get stuck on large problems<sup>7</sup>, trying to diversify  
370 the search by exploring larger neighborhoods, VNS using Luby will spend more  
371 time on small neighborhoods in order to locally improve the current solution, fa-  
372 voring intensification. By adding randomness on variable and/or value ordering  
373 heuristics<sup>8</sup> used by  $\text{LDS}^r$ , it is possible to find a better solution even when the dis-  
374 crepancy limit decreases when applying the Luby strategy. The  $\text{mult}2$  strategy  
375 reduces the number of neighborhood explorations at a given discrepancy limit, in  
376 order to try larger discrepancy limits more rapidly. If the problem is solvable by a  
377 complete search within the time limit, it will also speed-up the optimality proof.

378 Operator  $+_\ell$  controls the compromise between incomplete and complete search.  
379 Using a fast growing strategy emphasis completeness whereas a slow growth  
380 should favor anytime behavior. The  $\text{mult}2$  strategy tends to favor a non-decreasing  
381 worst-case complexity of the successive neighborhood searches, especially when  
382 going from  $\ell$  with  $k = k_{max}$  to  $2\ell$  with  $k = k_{min}$  (for sufficiently large  $k_{min}$ ).

383 We noticed that it is worthwhile to cover all the variables by the union of the  
384 explored neighborhoods in order to not miss some important variables. We tested

---

<sup>6</sup>Recall  $\text{luby}(i) = \{1, 1, 2, 1, 1, 2, 4, 1, 1, 2, 1, 1, 2, 4, 8, \dots\}$ .

<sup>7</sup>Although it is possible to add a limit on the number of backtracks per neighborhood search as it is done in Large Neighborhood Search methods [15].

<sup>8</sup>Adaptive heuristics such as weighted degree heuristic [40] will also modify the variable ordering from one search to another.

385 a fourth strategy for  $k$  which consists in a slow increment (by  $+1$ ) at the beginning  
 386 until  $k = \max_{i \in I} (|C_i|) + |C_T| - 1$  then it *jumps* directly to  $k = k_{max}$ . This ensures  
 387 that  $k$  grows slowly until the largest cluster has been totally explored by at least  
 388 one neighborhood search. Then, when  $k = k_{max} = |\mathcal{X}|$ , UDGVNS does a restart,  
 389 looking for an improved starting solution, using  $LDS^r$  applied on the whole prob-  
 390 lem. If it fails to find a better solution, or prove optimality, a larger discrepancy  
 391 can be selected and UDGVNS continues its intensification process starting with a  
 392 small neighborhood size (line 17).

#### 393 4. A Parallel Version of UDGVNS

394 This section describes how UDGVNS has been parallelized. We called the  
 395 resulting algorithm Unified Parallel DGVNS (UPDGVNS). Section 4.1 provides a  
 396 general overview of the parallel version. The UPDGVNS algorithm is detailed in  
 397 section 4.2. A more detailed discussion about UPDGVNS properties is given in  
 398 section 4.3.

##### 399 4.1. UPDGVNS in nutshell

400 The parallel version relies on a master/worker model and exploits the UDGVNS  
 401 framework to control the compromise between optimality proof and anytime be-  
 402 havior. UPDGVNS enhances the optimization process of UDGVNS by enabling  
 403 a better diversification. More precisely, UPDGVNS uses the master process as a  
 404 diversification component to explore the search space on the global scale, while  
 405 using the worker process as an intensification component to exploit the search  
 406 space on the neighborhood provided by the master. Figure 2 provides an overview  
 407 of UPDGVNS method based on the following three main components:

- 408 • a *master* process, on the left side, which determines the neighborhoods to  
 409 be explored and updates the global solution  $S$  at each iteration.
- 410 • a set of asynchronous *worker* processes, on the right side, which explore  
 411 independently the parts of the search space assigned by the master process.
- 412 • an interaction model based on an asynchronous communication between the  
 413 master and the workers, where the master process controls the communica-  
 414 tion over the entire processes.

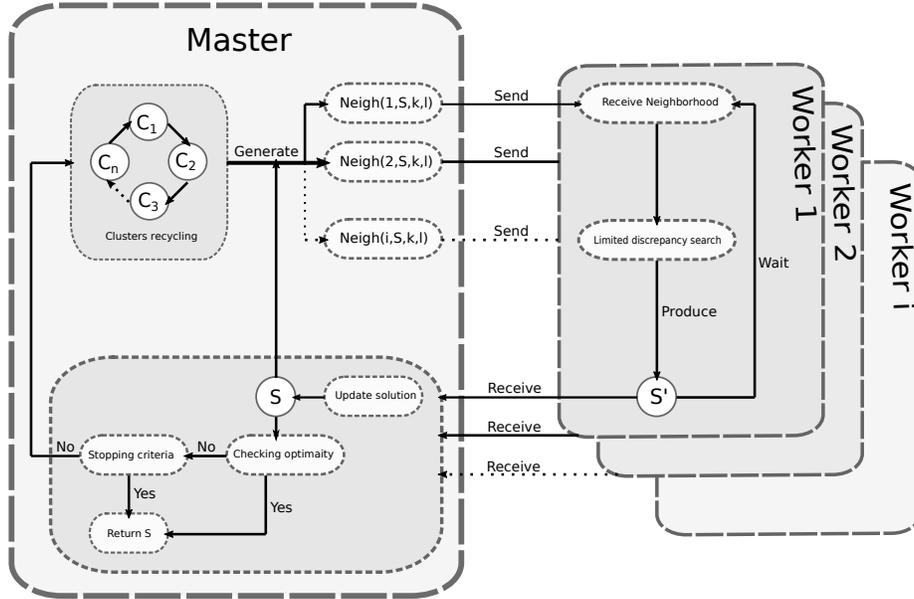


Figure 2: A general overview of UPDGVNS algorithm. The set  $\{C_1, C_2, \dots, C_n\}$  corresponds to the  $n$  clusters provided by the tree decomposition. The total number of the available workers is denoted by  $i$ . The master's global solution is denoted by  $S$ . Local solutions found by workers are denoted by  $S'$ .

415 **4.2. UPDGVNS algorithm**

416 Algorithms 5 and 6 depict the pseudo-code of UPDGVNS in more details. Let  
 417  $\mathcal{P}$  be a data structure allowing to manage a list of parameters used by each worker  
 418 process for the exploration of the neighborhood of a solution  $x$  (i.e.,  $i, k, r, \ell, cl,$   
 419  $x, ub, opt$ ). Initially, the master (see Algorithm 5) initiates the search by launch-  
 420 ing (at line 26)  $npr$ <sup>9</sup> worker processes in parallel with the same initial solution  
 421 (line 25). This is done by initializing the different parameters for the neighbor-  
 422 hood exploration and sending them to each worker  $p$ . Each worker process obtains  
 423 from the master a copy of the current best solution  $x$ , the index  $c$  of the cluster  
 424 to be processed and performs destroy and repair operations on its local copy (see  
 425 Algorithm 6). As soon as a new solution  $x'$  is found by a worker  $p$ , it is sent to  
 426 the master as well as its status (i.e. flag  $opt$ ) by checking whether  $x'$  is proven op-  
 427 timal, by setting flag  $opt$  to *true*. In the second while loop (at line 27), the master  
 428 waits for new solutions found by each worker process. Like UDGVNS, the master

<sup>9</sup>Worker processes are ranked from 1 to  $npr$ , while the master is ranked zero.

---

**Algorithm 5: Master algorithm for Unified Parallel DGVNS algorithm**

---

```
Procedure master ( $npr, \ell_{min}, \ell_{max}, +\ell, k_{min}, k_{max}, +k, ub : In/Out, x : Out,$   
 $opt : Out$ )  
  let  $(C_T, T)$  be a tree decomposition of  $(\mathcal{X}, \mathcal{D}, \mathcal{F})$  ;  
25  $opt \leftarrow \text{LDS}^r(n(d-1), \mathcal{D}, ub, x)$  ; // find an initial solution  
  if  $(ub = \text{lb}(\mathcal{D}))$  then  $opt \leftarrow \text{true}$  ;  
  if  $(\neg opt)$  then  
     $c \leftarrow 1$  ; // current cluster index  
26 for each worker  $p = 1, \dots, npr$  do  
  // initial parameters of neighborhood exploration  
   $\mathcal{P}[p].ub \leftarrow ub, \mathcal{P}[p].x \leftarrow x, \mathcal{P}[p].cl \leftarrow c, \mathcal{P}[p].\ell \leftarrow \ell_{min}, \mathcal{P}[p].k \leftarrow k_{min}$   
  ;  
   $\mathcal{P}[p].r \leftarrow 0$  ; // number of discrepancy iterations  
   $\mathcal{P}[p].i \leftarrow 0$  ; // nb. of succ. failed neighb. sizes  
  Send( $p, \mathcal{P}[p]$ ) ;  
   $c \leftarrow 1 + c \bmod |C_T|$  ; // visit next cluster  
27 while  $(\neg opt \wedge \neg \text{TimeOut})$  do  
  Receive( $p, \mathcal{P}[p]$ ) ; // wait a new solution from worker  $p$   
28 if  $(\mathcal{P}[p].opt)$  then  $opt \leftarrow \text{true}$  // optimality proof check ;  
  if  $(\mathcal{P}[p].k \geq k_{max} \wedge \mathcal{P}[p].\ell \geq \ell_{max})$  then  $\text{TimeOut} \leftarrow \text{true}$  ;  
  // update parameters of neighborhood exploration  
   $\mathcal{P}[p].cl \leftarrow c$  ;  
  if  $(\mathcal{P}[p].ub < ub)$  then  
     $ub \leftarrow \mathcal{P}[p].ub, x \leftarrow \mathcal{P}[p].x$  ; // new best solution  
     $\mathcal{P}[p].i \leftarrow 0, \mathcal{P}[p].k \leftarrow k_{min}$  ;  
     $\mathcal{P}[p].r \leftarrow 0, \mathcal{P}[p].\ell \leftarrow \ell_{min}$  ;  
  else  
     $\mathcal{P}[p].ub \leftarrow ub, \mathcal{P}[p].x \leftarrow x$  ;  
     $\mathcal{P}[p].i \leftarrow \mathcal{P}[p].i + 1$  ;  
    if  $(\mathcal{P}[p].k < k_{max})$  then  
       $\mathcal{P}[p].k \leftarrow \min(k_{max}, k_{min} + k \mathcal{P}[p].i)$  ;  
    else  
       $\mathcal{P}[p].i \leftarrow 0, \mathcal{P}[p].k \leftarrow k_{min}$  ;  
       $\mathcal{P}[p].r \leftarrow \mathcal{P}[p].r + 1$  ;  
      if  $(\mathcal{P}[p].\ell < \ell_{max})$  then  
         $\mathcal{P}[p].\ell \leftarrow \min(\ell_{max}, \ell_{min} + \ell \mathcal{P}[p].r)$  ;  
     $c \leftarrow 1 + c \bmod |C_T|$  ;  
29 if  $(\neg opt \wedge \neg \text{TimeOut})$  then Send( $p, \mathcal{P}[p]$ ) ;
```

---

429 controls how the discrepancy limit  $\ell$  and the neighborhood size  $k$  evolve during  
430 successive explorations and updates the shared global best solution  $x$  according to  
431 solutions sent by the workers. However, contrary to UDGVNS, whenever  $k$  reaches  
432  $k_{max}$  and the discrepancy value is greater than or equal to  $\ell_{max}$  or the internal flag

---

**Algorithm 6:** Worker algorithm for UPDGVNS algorithm

---

```
Procedure worker ()  
    while ( $\neg$ TimeOut) do  
        Receive(0, P);  
         $A \leftarrow$  getNeighborhood( $P.x, C_{P.cl}, P.k$ ) ;  
         $P.opt \leftarrow$  LDSr ( $P.l, A, P.ub, P.x$ ) ;  
        if ( $P.ub = \text{lb}(\mathcal{D})$ ) then  $P.opt \leftarrow$  true;  
        else if ( $A \neq \mathcal{D}$ ) then  $P.opt \leftarrow$  false;  
        Send(0, P);
```

---

433  $opt$  of the worker is *true* (line 28), the whole process stops and the master returns  
434 the (optimal) solution. Otherwise, if the time limit is not reached, the workers  
435 that are ready to restart a new exploration are re-launched starting from the best  
436 available overall solution on the next clusters (line 29).

#### 437 4.3. UPDGVNS properties

438 This section discusses in depth the key properties of UPDGVNS that contribute  
439 to its success, namely diversification and workload distribution between workers.

##### 440 4.3.1. Diversification in UPDGVNS

441 Regarding the parallelization scheme, diversification is ensured in three ways:

- 442 1. the parallel exploration of different clusters provides a form of diversifica-  
443 tion to UPDGVNS method by exploring independently different parts of the  
444 search space. Moreover, the unassigned variables are selected from a cur-  
445 rent cluster of a tree decomposition (and its neighbor clusters if needed) in  
446 a random way as done in DGVNS algorithm.
- 447 2. the control of the size of the neighborhood ( $k$ ) and the discrepancy limit ( $\ell$ )  
448 is done per worker in an asynchronous and independent manner. This allows  
449 to relaunch each worker asking for a next cluster with diverse parameters  
450 and different initial partial assignments, thus enhancing the diversification  
451 scheme of UPDGVNS.
- 452 3. to prevent redundant searches among different workers, i.e., case where two  
453 or more identical<sup>10</sup> neighborhoods are explored, we add randomness in the

---

<sup>10</sup>This may arise when all the variables of the problem are included in the neighborhood ( $k = k_{max} = n$ ).

454 variable ordering heuristic used in  $LDS^r$  algorithm. More precisely, our  
455 heuristic for variable ordering first selects the variable having the lowest  
456 ratio domain cardinality divided by weighted degree ( $dom/wdeg$ ), break-  
457 ing ties by selecting one variable randomly. This leads to variations in the  
458 exploration of the search tree performed by each worker.

459 All these features provide a high level of diversification by exploring different  
460 regions in parallel and allow to decrease the probability that different workers  
461 perform the same exploration of the search space even for  $k = k_{max} = n$ .

#### 462 4.3.2. Workload distribution in UPDGVNS

463 In our parallel algorithm, we do not decompose the whole search space into a  
464 partition of subproblems but rather explore different randomly-selected subprob-  
465 lems which may overlap. First, each subproblem is related to a particular cluster  
466 in the tree decomposition and each time a worker asks for a job it gets the next  
467 (current) cluster  $c$  in a global round-robin fashion (Algorithm 6) and explores the  
468 assigned cluster independently starting from the best overall available solution.  
469 So, in terms of workload distribution between workers, this remains more or less  
470 balanced. Moreover, since the proof of optimality is performed sequentially, i.e.,  
471 by a single worker who explores the entire search tree, this does not induce unbal-  
472 anced workloads among workers.

473 Second, the way the size of the neighborhood ( $k$ ) and the discrepancy limit  
474 ( $\ell$ ) evolve is done per worker in an asynchronous and independent manner. This  
475 can lead to situation where one worker may search on a small subproblem, while  
476 another worker on a larger subproblem. If the larger subproblem is not defined  
477 over the whole problem, the balance on subproblems is quickly established with  
478 the rest of the workers once the current worker finishes its exploration. This is  
479 achieved by the master which updates parameters of the neighborhood exploration  
480 accordingly (i.e., neighborhood size and discrepancy limit) in order to relaunch  
481 the worker on the next cluster. At the end, the first worker finishing its search  
482 on the whole problem with no discrepancy limit will report optimality and stop  
483 UPDGVNS.

## 484 5. UAI Evaluation Results

### 485 5.1. Benchmarks description

486 We performed experiments on probabilistic and deterministic graphical mod-  
487 els coming from a large multi-paradigm benchmark repository [28]<sup>11</sup>. Among the  
488 3016 available instances, we selected all the instances that were used in previous  
489 Uncertainty in Artificial Intelligence (UAI) competitions, in image analysis, or  
490 in CFN optimization. It includes 319 instances from *UAI 2008 Evaluation* and  
491 Probabilistic Inference Challenge (PIC) 2011<sup>12</sup>, 1461 instances from the Com-  
492 puter Vision and Pattern Recognition (CVPR) OpenGM2 benchmark<sup>13</sup> [41], and  
493 281 instances from the Cost Function Library<sup>14</sup>. In order to have a fair compari-  
494 son between solvers, we preprocessed all the instances by polynomial-time prob-  
495 lem reformulations and simplifications that remove variables (using bounded and  
496 functional variable elimination [42]), values (using *dead-end elimination* [43]),  
497 and fixed-value potentials, after an initial lower bound computation by *Equiva-  
498 lence Preserving Transformations* [31] (enforcing *Virtual Arc Consistency (VAC)*  
499 as a message-passing algorithm). The resulting instances are smaller while pre-  
500 serving the same optimum. We used TOULBAR2 with options `-A -z=2` for this  
501 preprocessing step. We kept 1669 non-trivial instances (with more than one vari-  
502 able) for the experiments. The number of variables  $n$  ranges from 4 (CVPR-  
503 GeomSurf-3-gm13) to 48,566 (CVPR-ColorSeg-8-crops-small) with mean value  
504  $n \approx 403.4$  (instead of 1,316.5 before preprocessing). For DGVNS methods, we  
505 built tree decompositions using *min-fill* heuristic. Because the number of clus-  
506 ters  $m = |C_T|$  can be very large ( $m \approx 256.7$ ), we merged any pair of connected  
507 clusters  $(C_i, C_j)$  when the separator size is too large compared to the individual  
508 cluster sizes ( $|sep(C_i, C_j)| > 0.7 \min(|C_i|, |C_j|)$ ), resulting in  $m \approx 19.9$  and  
509 mean treewidth  $\max_{i \in T}(|C_i|) \approx 76.9$  (instead of 49.6 without merging). In or-  
510 der to experiment sequential and parallel methods on the most difficult instances,  
511 we selected a subset of instances unsolved in 1 hour by all our DFBB, LDS, and  
512 VNS algorithms. We took at most twenty instances per problem category (avoid-  
513 ing over-representation issues by some categories), resulting in a selection of 114

---

<sup>11</sup>[genoweb.toulouse.inra.fr/~degivry/evalgm](http://genoweb.toulouse.inra.fr/~degivry/evalgm)

<sup>12</sup>[graphmod.ics.uci.edu/uai08/Evaluation/Report/Benchmarks](http://graphmod.ics.uci.edu/uai08/Evaluation/Report/Benchmarks), [www.cs.huji.ac.il/project/PASCAL](http://www.cs.huji.ac.il/project/PASCAL)

<sup>13</sup>[hci.iwr.uni-heidelberg.de/opengm2](http://hci.iwr.uni-heidelberg.de/opengm2)

<sup>14</sup>[costfunction.org/benchmark](http://costfunction.org/benchmark)

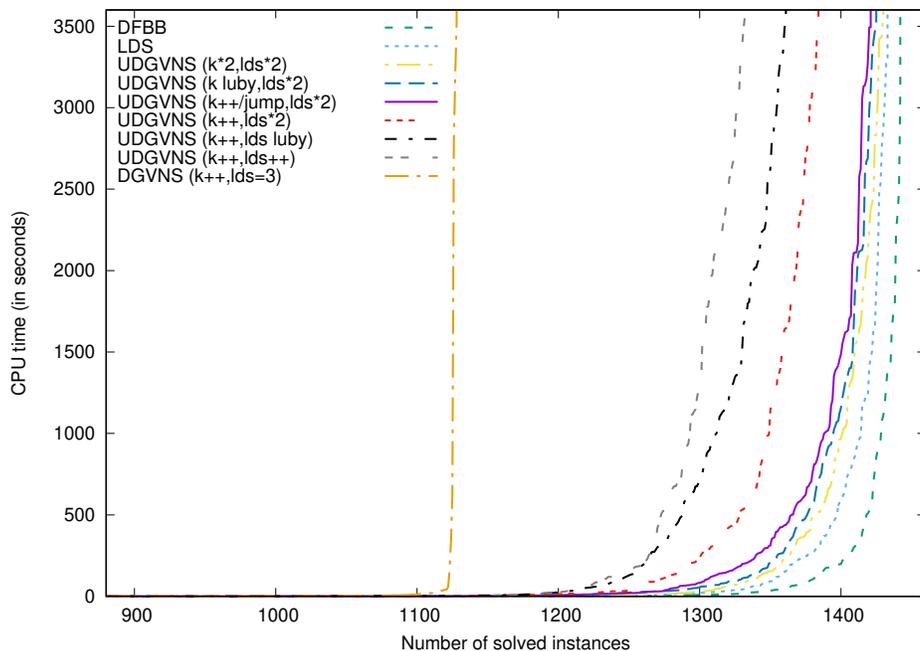


Figure 3: Number of instances solved by our approach as times passes on a restricted benchmark set (Methods are sorted as timeout limit (3600s)).

514 difficult instances<sup>15</sup>.

## 515 5.2. Experimental protocol

516 LDS<sup>r</sup> employs a randomized (for breaking ties) dynamic variable ordering  
 517 heuristic<sup>16</sup>. Its value ordering heuristic chooses the EAC value as the preferred  
 518 value and lower bounds are deduced by enforcing EDAC, as explained in [30]. In  
 519 the following, we set  $k_{min} = 4$ ,  $k_{max} = n = |\mathcal{X}|$ ,  $\ell_{min} = 1$ , and  $\ell_{max} = n(d - 1)$ .  
 520 DFBB corresponds to UDGVNS  $(\infty, \infty, +, \infty, \infty, +)$ , LDS corresponds to  
 521 UDGVNS  $(1, \infty, \text{mult}2, |\mathcal{X}|, |\mathcal{X}|, +)$ , DGVNS to UDGVNS  $(3, 3, +, k_{min}, k_{max}, +)$ .  
 522 These methods and their parallelization based on MPI (UPDGVNS) have been im-  
 523 plemented into the new version 1.0.0 of TOULBAR2<sup>17</sup>.

<sup>15</sup>UAI DBN, Grid, Linkage, ObjectDetection, CVPR ChineseChars, ColorSeg-8, InPainting-4, InPainting-8, ProteinInteraction, and CFN CELAR, ProteinDesign, SPOT5, Warehouse categories.

<sup>16</sup>Weighted degree heuristic as defined in [40].

<sup>17</sup><https://github.com/toulbar2/toulbar2>

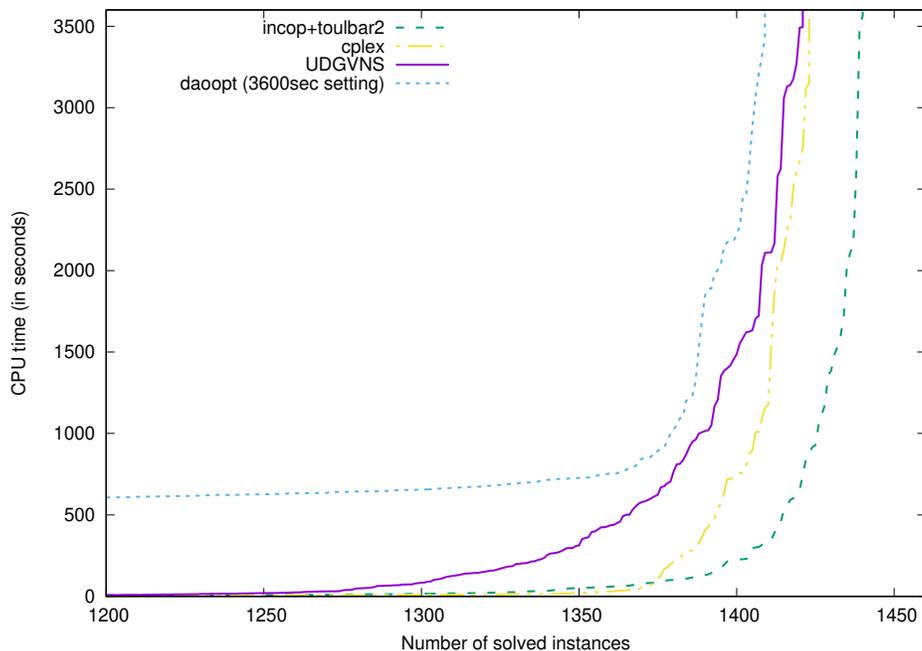


Figure 4: Number of instances solved by each method as time passes (UDGVNS =  $\text{UDGVNS}(k \text{ add1/jump}, \ell \text{ mult}2)$ ).

524 We compared with state-of-the-art exact solvers. DAOOPT<sup>18</sup> won PIC 2011. It  
525 has a time-bounded initial phase of lower bound computation based on *Message*  
526 *Passing Linear Programming* algorithm [8, 9] and mini-bucket elimination [29]  
527 with iterative *min-fill* heuristic, further improved by Join Graph Linear Program-  
528 ming [44]. It also finds good initial upper bounds using LDS (with discrepancy  
529 limit set to 1) and stochastic local search  $GLS^+$  [12]. We used the standalone code  
530 of DAOOPT version 0.99.7g-UAI12 (with option settings `-mplp=2000 -mplps=60`  
531 `-slsX=20 -slsT=10 -t 30000 -orderTime=180 -jglp=1000 -jglps=60 -i 35 -m`  
532 `4096 -match -y -lds=1` for 3600-second time limit). We tested three parameter set-  
533 tings as suggested in [45], controlling the time spent to compute initial lower and  
534 upper bounds. In the 3600sec setting, SLS is run 20 times with 10 seconds per run.  
535 The best solution found is used as an initial upper bound for an AND/OR exhaus-  
536 tive tree search. We compared also with an older version of TOULBAR2, namely

<sup>18</sup><https://github.com/lotten/daoopt>

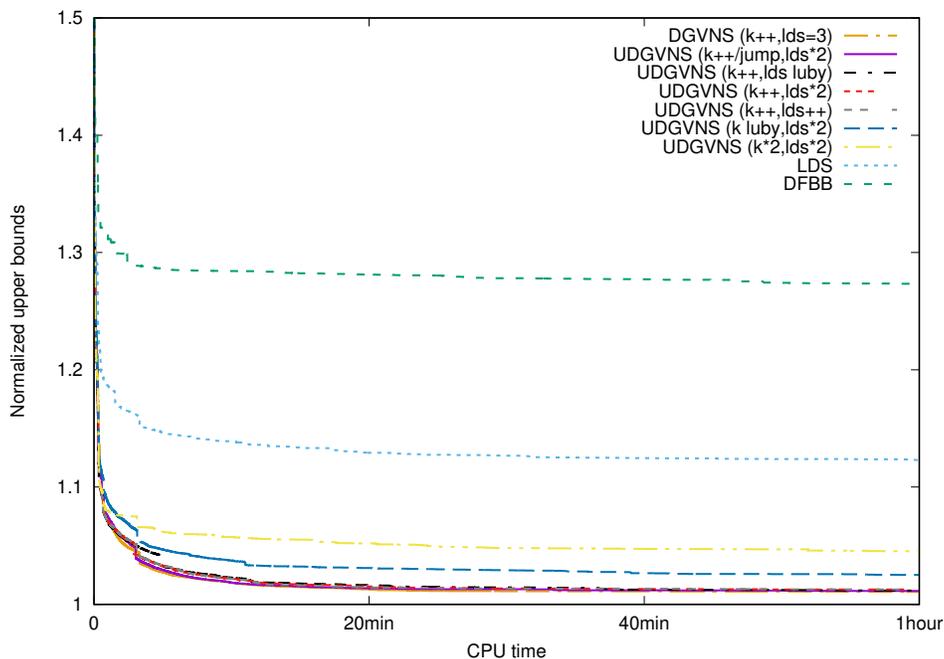


Figure 5: Average evolution of normalized upper bounds for UDGVNS versus DFBB and LDS algorithms on 114 difficult instances.

537 INCOP+TOULBAR2<sup>19</sup> [28] won the *UAI 2010 Evaluation* at 20-minute time limit.  
 538 INCOP+TOULBAR2 takes a starting solution from the best result of three runs of  
 539 the *IDWalk* [46] local search algorithm (100,000 local moves per run). It is fol-  
 540 lowed by an exhaustive hybrid best-first search [5]. IBM ILOG CPLEX 12.7.0.0  
 541 (using parameters EPAGAP, EPGAP, and EPINT set to zero to avoid premature  
 542 stop) was reported as being very competitive on some image analysis [41] and  
 543 Markov Random Field problems [28]. CPLEX explores its search tree using best-  
 544 first search. It applies several heuristics methods to find good solutions before  
 545 and during the search. We also compared with message-passing algorithms: LIB-  
 546 DAI<sup>20</sup> [47], winner of the *UAI 2010 Evaluation* at 20sec. and 1hour time limits,  
 547 MPLP2<sup>21</sup> [8, 9], and TRW-S<sup>22</sup> [6]. Note that LIBDAI and TRW-S are applied on

<sup>19</sup>[www.inra.fr/mia/T/toulbar2](http://www.inra.fr/mia/T/toulbar2) version 0.9.8 with parameters *-i-dee-hbfs*.

<sup>20</sup>[bitbucket.org/jorism/libdai.git](https://bitbucket.org/jorism/libdai.git) version 0.3.2 using UAI 2010 settings.

<sup>21</sup>[cs.nyu.edu/~dsontag/code/README\\_v2.html](http://cs.nyu.edu/~dsontag/code/README_v2.html) using  $2 \cdot 10^{-7}$  int gap thres.

<sup>22</sup>[github.com/opengm/opengm](https://github.com/opengm/opengm) v2.3.5 with TRW-S v1.3 stopping after 100,000 iterations or  $10^{-5}$  gap thres.

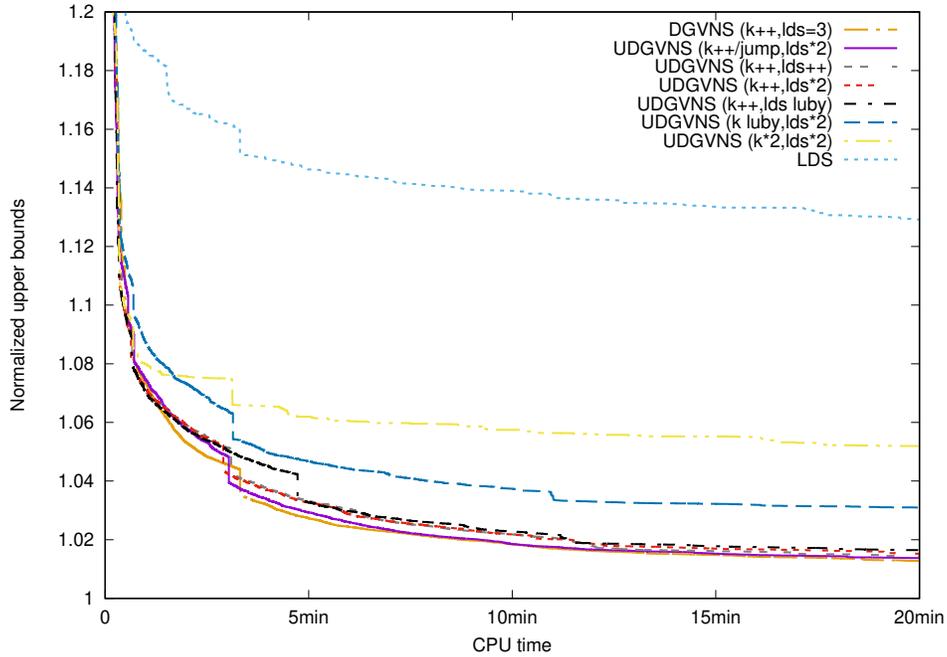


Figure 6: Any time upper bound zoom for UDGVNS versus LDS.

548 the original instances rather than the preprocessed ones as we found they produced  
 549 better results without applying VAC first. All solvers read problems in *uai* tabular  
 550 format, except CPLEX which uses the local polytope formulation (called support  
 551 encoding in [28]). All computations were performed on a cluster of 48-core AMD  
 552 Opteron 6176 at 2.3 GHz and 384 GB of RAM with a 1-hour CPU time limit<sup>23</sup>.

### 553 5.3. Experimental results

#### 554 5.3.1. Optimality results

555 The efficiency of DFBB, LDS, and VNS methods to prove optimality is shown  
 556 in the cactus plot of Figure 3. DFBB was slightly more efficient than LDS and  
 557 solved 1442 (resp. 1433) instances among 1669 in 1-hour time limit. They are  
 558 followed by three UDGVNS strategies with  $(k \text{ mult } 2, \ell \text{ mult } 2)$  (1430 solved),  
 559  $(k \text{ Luby}, \ell \text{ mult } 2)$  (1425 solved) and  $(k \text{ add1/jump}, \ell \text{ mult } 2)$  (1421 solved),  
 560 remaining very close in terms of performance. Next, a set of three less-and-less ef-

<sup>23</sup>Using parameter `-pe parallel_smp min(2x, 30)` on a SUN Grid Engine for a method exploiting  $x$  core(s) to ensure half-load of the nodes on the cluster.

561 efficient UDGVNS strategies rise: ( $k$  add1,  $\ell$  mult2) (1384 solved), ( $k$  add1,  $\ell$  Luby)  
 562 (1361 solved) and ( $k$  add1,  $\ell$  add1) (1333 solved), showing the importance of  
 563 faster discrepancy increase to speed up optimality proofs. The worst strategy  
 564 here was using a fixed discrepancy level ( $\ell = 3$  as originally proposed in [23])  
 565 which solved 1128 instances. Figure 4 compares the cactus plot of UDGVNS ver-  
 566 sus CPLEX, DAOOPT, and INCOP+TOULBAR2. DAOOPT (3600sec setting) solved  
 567 1409 instances, CPLEX solved 1423, and INCOP+TOULBAR2 1440 instances.

### 568 5.3.2. Anytime upper bound profiles

569 In order to summarize the evolution of upper bounds as time passes, we took a  
 570 subset of 114 difficult instances, unsolved in 1 hour by our DFBB, LDS, and VNS  
 571 methods (whereas CPLEX could solve 17 of these instances). Specifically, for each  
 572 instance  $I$  we normalize all energies as follows: the best, potentially suboptimal  
 573 solution found by any algorithm is 1, the worst solution is 2. This normalization  
 574 is invariant to translation and scaling. Figure 5 shows the upper bound behavior  
 575 for different VNS strategies compared to DFBB and LDS. Figure 6 reports an  
 576 anytime upper bound zoom. The ranking of best methods is the opposite of the  
 577 cactus plot order, except for ( $k$  add1/jump,  $\ell$  mult2) which comes in second  
 578 position. According to details in Figure 6, the  $\ell = 3$  strategy got the best upper  
 579 bounds in average, but still very close to the other VNS strategies, except may-  
 580 be ( $k$  Luby,  $\ell$  mult2) and ( $k$  mult2,  $\ell$  mult2). We conclude that our new  
 581 iterative UDGVNS method (especially ( $k$  add1/jump,  $\ell$  mult2)) offers a good  
 582 compromise between anytime behavior and optimality proof. These results also  
 583 show that variable neighborhood search is by far superior to classical systematic  
 584 DFBB or non-systematic LDS tree search methods, improving by more than 20%  
 585 (resp. 10%) the quality of the solutions.

### 586 5.3.3. Comparing UDGVNS with state-of-the-art methods

587 In the following figures, we assume UDGVNS with ( $k$  add1/jump,  $\ell$  mult2)  
 588 strategy. Figure 7 compares UDGVNS with state-of-the-art methods. Message-  
 589 passing algorithms like TRW-S and LIBDAI gave the worst results. They could  
 590 not find any solution for 20 (resp. 19) instances (mostly in UAI/Linkage and  
 591 CFN/SPOT5 categories, both containing hard constraints). The same problem oc-  
 592 curred for MPLP2 on 5 instances (SPOT5), but it obtained much better results on  
 593 the remaining instances. CPLEX ran out of memory on two instances without pro-  
 594 ducing any solution due to the heavy local polytope encoding (CFN/Warehouse/capb,  
 595 capmq5). All other methods got better results in average and produced at least one  
 596 solution per instance. According to its initial phase setting, DAOOPT provides dif-

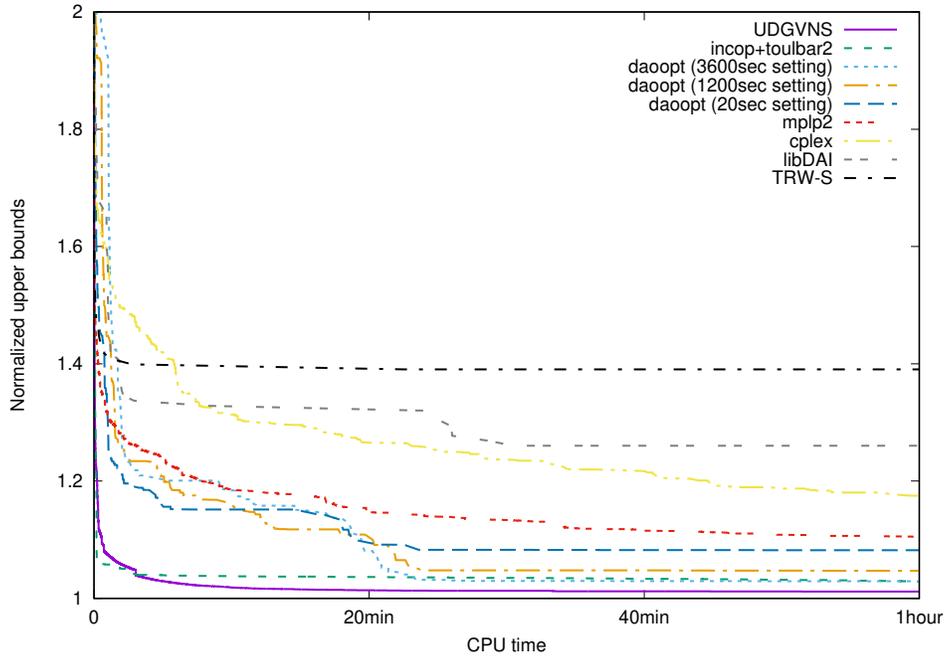


Figure 7: Comparing the anytime behavior of UDGVNS against state-of-the-art methods.

597 ferent anytime behaviors, very close to the best solutions in 1 hour. UDGVNS  
 598 performed the best, slightly better than INCOP+TOULBAR2, improving by 1.7%  
 599 (resp. 2.3%) on average after 1 hour (20 min).

#### 600 5.3.4. Parallelization

601 Finally, in order to evaluate the impact of core numbers, we consider the anytime  
 602 upper bound behavior of the parallel release: UPDGVNS using ( $k$  add1/jump,  
 603  $\ell$  mult2) with  $\ell_{min} = 3$ , taken from the best strategies enlightened by UDGVNS.  
 604 We made a comparison with CPLEX using 10 and 30 cores. Figure 8 shows that  
 605 CPLEX with 10 or 30 cores exhibits better anytime behavior than CPLEX using 1  
 606 core, but still being far from the other competitors (30 cores gave solutions 10%  
 607 higher than UPDGVNS after 1 hour). We could not compare with the parallel  
 608 version of DAOOPT as it is based on a different cluster engine (*condor*) and it does  
 609 not parallelize its initial phase.

610 Figure 9 shows that UPDGVNS (with 10 or 30 cores) provides slightly better  
 611 upper bounds than UDGVNS (using 1 core) in less than 20 min. The results seems  
 612 to be, in average, poorly sensitive to the cores number, due to the fact that the

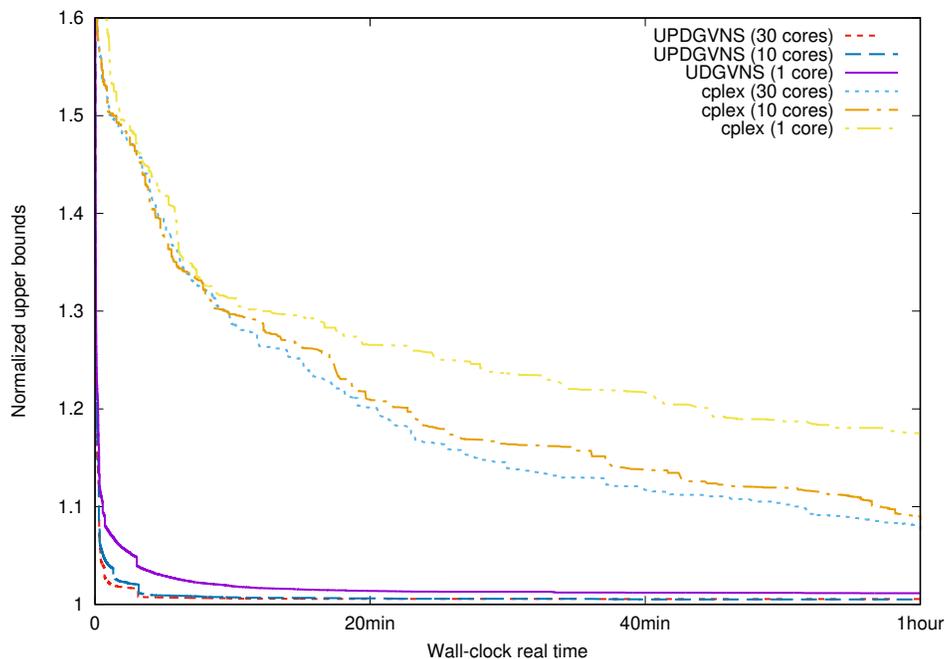


Figure 8: Anytime upper bound with 1, 10 and 30 processors respectively for CPLEX and UPDGVNS (UPDGVNS = UPDGVNS( $k$  add1/jump,  $\ell$  mult2)).

613 10-core curve is extremely close from the 30-core one. Figure 9 also compares  
 614 the anytime upper bound quality with those provided by single-core INCOP lo-  
 615 cal solver followed by a hybrid best-first search in TOULBAR2 and by DAOOPT  
 616 with options tuned for the 1200-second UAI 2014 challenge. The 10 and 30-core  
 617 UPDGVNS curves converge quickly in less than 2min. INCOP+TOULBAR2 quickly  
 618 drops out around 1 min and never reaches the same quality level. UDGVNS con-  
 619 verges slower but still going down after 20 min. DAOOPT (1200sec setting) gave  
 620 results 10% in average worse than UPDGVNS with either 10 and 30 processors.

621 The trends observed over all instances are quite similar to those obtained on  
 622 selected instances for each family. The only exception are the Pedigree instances.  
 623 Table 1 gives the solving time to find and prove optimality on UAI-Linkage cate-  
 624 gory for U(P) DGVNS, CPLEX and DAOOPT (in parenthesis, unnormalized upper  
 625 bound found after 1 hour). We report DAOOPT time from [48], obtained on a clus-  
 626 ter of dual 2.67 GHz Intel Xeon X5650 6-core CPUs and 24 GB of RAM. We  
 627 can see that UDGVNS and UPDGVNS (with 10 or 30 cores) are clearly dominated  
 628 by CPLEX which exhibits better results. DAOOPT remains competitive on these

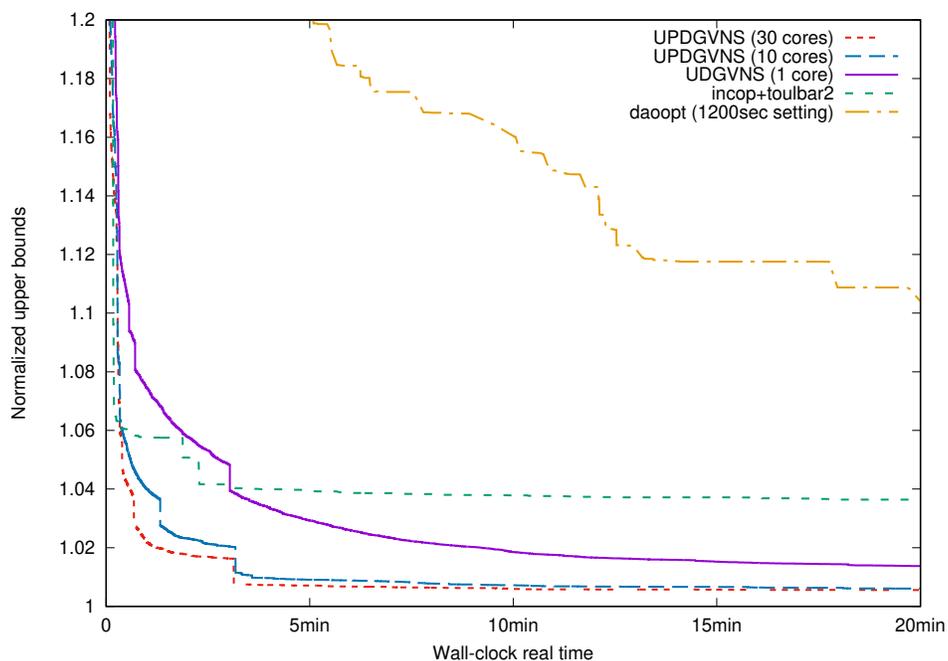


Figure 9: Comparing the anytime behavior of parallel UPDGVNS.

629 instances but still being far from CPLEX in terms of CPU times.

## 630 6. Computational Protein Design

631 In Computational Protein Design (CPD), the challenge of identifying a protein  
 632 that performs a given task is defined as the combinatorial optimization of a com-  
 633 plex pairwise energy function over amino acid sequences and 3D geometry [49].  
 634 This holds great interest for medicine, synthetic biology, nanotechnologies and  
 635 biotechnologies [50][51][52]. We used the CPD problem as a difficult benchmark  
 636 to test our tree-decomposition based methods. For that, we generated 21 large in-  
 637 stances with small treewidth selected from the Protein Data Bank<sup>24</sup> (PDB). These  
 638 instances have been selected on the basis of 3D criteria described in a supplemen-  
 639 tary material<sup>25</sup>. The instances contain from  $n = 107$  up to 292 variables with a

<sup>24</sup>[www.pdb.org](http://www.pdb.org)

<sup>25</sup>[genoweb.toulouse.inra.fr/~degivry/evalgm/CFN/ProteinDesignUAI2017](http://genoweb.toulouse.inra.fr/~degivry/evalgm/CFN/ProteinDesignUAI2017)

Linkage ( <i>optimum / worst solution</i> )	pedigree19 (4625/21439)	pedigree31 (5258/166553)	pedigree44 (6651/104904)	pedigree51 (6406/629929)
CPLEX (1 core)	790	59.3	6.35	36.23
CPLEX (10 cores)	191	9.00	2.48	9.43
CPLEX (30 cores)	<b>75</b>	<b>7.17</b>	<b>2.69</b>	<b>5.34</b>
DAOPT (1 core)	375,110	16,238	95,830	101,788
DAOPT (20 cores)	27,281	1,055	6,739	6,406
DAOPT (100 cores)	7,492	201	1,799	1,578
UDGVNS (1 core)	- (4949)	- (5258)	- (6722)	- (6406)
UPDGVNS (10 cores)	- (4762)	3,341	- (6651)	- (6406)
UPDGVNS (30 cores)	- (4626)	1,775	- (6651)	- (6406)

Table 1: Best CPU time (in seconds) for sequential versions and best wall-clock time for multiple-core ones to find and prove optimality on Pedigree instances. A “-” means no proof of optimality in 1 hour (except DAOPT with no time limit) (in parenthesis, unnormalized upper bound found after 1 hour).

640 maximum domain size from  $d = 383$  to 450, and between 1,623 and 6,208 bi-  
641 nary cost functions. The min-fill treewidth ranges from  $w = 21$  to 68, resulting in  
642 small ratios of treewidth per number of variables, from  $\frac{w}{n} = 0.16$  to 0.34.

### 643 6.1. Experimental Protocol

644 We compared UDGVNS and UPDGVNS with VNS/LDS+CP and Fixed Back-  
645 Bone (a CPD dedicated algorithm provided by the Rosetta package). We also  
646 compared with TOULBAR2. We tested TOULBAR2 using Virtual Arc Consis-  
647 tency [31] in pre-processing and hybrid best-first search with tree decomposition  
648 (BTD-HBFS) using min-fill heuristic. Dead end elimination is turned off, accord-  
649 ing to [26] (the resulting command line is: TOULBAR2 -B=1 -A -DEE: -O=-3).  
650 The TOULBAR2 experiments use a 24-hour CPU time limit.

651 Concerning the VNS methods, in order to evaluate variability due to the ran-  
652 dom selection of neighborhoods, a set of 10 runs per instance with different  
653 seeds has been performed with a time limit of 1-hour per run. For the par-  
654 allel strategy, the number of processes  $n_{pr}$  is set to 96 (i.e. maximum num-  
655 ber of available processors). For UDGVNS and UPDGVNS, following the results  
656 observed in Section 5, we considered the following two settings for operators  
657  $+_k$  and  $+_\ell$ : ( $k$  add1,  $\ell = 3$ ) which yields the best anytime performances and  
658 ( $k$  add1 / jump,  $\ell$  mult2) which offers the best compromise between both any-  
659 time performances and optimality proof.  $k_{min}$ ,  $k_{max}$  and  $\ell_{min}$  have been respec-  
660 tively set to 4,  $n$  (the total number of variables) and 3 (they correspond to the same

Instance	Time(s)		
	UDGVNS	UPDGVNS	TOULBAR2
5dbl	1,828.27	791.16	783.18
3r8q	-	-	41,700.10
4bxp	-	-	4,261.67
1f00	-	-	9,749.00
1xaw	-	-	2,917.04
5e10	839.52	196.43	1,171.98
2gee	-	-	9,795.59
5e0z	416.12	172.96	999.66
3lf9	-	-	2,960.64
5eqz	-	-	41,813.00

Table 2: Comparing  $\text{UDGVNS}(k \text{ add1/jump}, \ell \text{ mult}2)$ ,  $\text{UPDGVNS}(npr, k \text{ add1/jump}, \ell \text{ mult}2)$  and TOULBAR2 in terms of optimality results within the 24-hours time limit. A ‘-’ indicates that the corresponding solver failed to prove optimality.

661 parameter settings as those described in [23]). All computations were performed  
662 on a cluster of 96-core AMD Opteron 6174 at 2.2 GHz and 256 GB of RAM.

## 663 6.2. Experimental Results

664 As for UAI instances, we evaluated the effectiveness of UDGVNS and UPDGVNS  
665 on CPD instances in terms of optimality proof (cf. Section 6.2.1) and solution  
666 quality vs. CPU time (cf. Section 6.2.2). We also study in Section 6.2.3 the  
667 impact of varying the number of processes for the parallel release UPDGVNS.

### 668 6.2.1. Optimality results

669 Our first set of experiments aims at evaluating the efficiency of UDGVNS and  
670 UPDGVNS in terms of optimality proof comparing with TOULBAR2. We used the  
671 setting  $(k \text{ add1/jump}, \ell \text{ mult}2)$  with  $\ell_{min} = 3$ , taken from the best strategies  
672 enlightened by UDGVNS. Table 2 reports the CPU-times required by UDGVNS,  
673 UPDGVNS and TOULBAR2 to find and prove the optimum within the 3600-second  
674 time limit (24 hours for TOULBAR2). As we can see, TOULBAR2 clearly outper-  
675 forms both VNS methods: UDGVNS and UPDGVNS were able to prove optimal-  
676 ity only on 3 instances among the 9 ones closed by TOULBAR2. However, on  
677 two instances (5e10 and 5e0z), UDGVNS is up to 2.40 times faster than TOUL-  
678 BAR2 while UPDGVNS increases speeds by up to 5.96 times, despite the fact that  
679 BTD-HBFS benefits from the lower bounds reported by HBFS in individual clus-  
680 ters to improve its anytime behavior and from the global pruning lower bounds of

681 BTD. This greatly improves the overall performance of TOULBAR2 compared to  
682 VNS/LDS+CP.

### 683 6.2.2. CPU time and solution quality results

684 Our second set of experiments aims at evaluating VNS capability with respect  
685 to finding the optimal solution or a solution of better quality on instances for which  
686 optimal solutions are unavailable. For this aim, we selected ( $k$  add1,  $\ell = 3$ ) as  
687 setting for operators  $+_k$  and  $+\ell$ .

688 Table 3 shows a comparative evaluation of VNS methods with FIXBB and  
689 TOULBAR2. For each instance and each VNS method, we report the number of  
690 successful runs to reach the optimum (or the best known solution for unsolved  
691 instances) within a 3600-second time limit, the average CPU times (in seconds)  
692 over the 10 runs (for unsuccessful runs, the CPU time is set to the time limit)  $\pm$   
693 the standard deviation. The energy gap  $\Delta E$  between the best VNS solution and the  
694 two external references FIXBB and TOULBAR2 are also given. For TOULBAR2,  
695 reported CPU times correspond to times to find an optimal solution (for solved  
696 instances) or a best one (for unsolved ones) within the 24-hour time limit.

697 *A) VNS methods vs. FIXBB.* Rosetta Modeling suite is one of the most popular  
698 software package used in the CPD field. It provides a Monte Carlo based Simu-  
699 lated Annealing algorithm called FIXBB. In this work, the best solutions exhibited  
700 by 1000 FIXBB cycles performed on each CPD instance have been used as base-  
701 line to compare solution quality of the solutions provided by VNS methods when  
702 TOULBAR2 BTD-HBFS fails to solve the instance.

703 The FIXBB CPU times are two orders of magnitude higher than the 1h time  
704 limit imposed for VNS evaluation. They are not reported as they exceed 100 hours  
705 in sequential mode, even if FIXBB cycles are independent and thus are easy to  
706 parallelize. As we can see in Table 3 (see column (4)  $\Delta E$ ), the solution quality  
707 of FIXBB is in all cases inferior to the best solution found by the VNS methods.  
708 The energetic gap  $\Delta E$  between FIXBB solution and the best VNS overall solution  
709 ranges between +0.16 and +5.20 Rosetta Energy Unit (R.E.U). As shown in [26]  
710 such a level of energy difference can strongly impact the designed protein solu-  
711 tion (i.e. the corresponding sequences of the two methods can be far in terms of  
712 hamming distance).

713 *B) VNS methods vs. TOULBAR2.* The comparison between best solutions found  
714 by VNS methods and TOULBAR2 shows that, excepted VNS/LDS+CP method,  
715 both UDGVNS (except on 1f00) and UPDGVNS provide in all case the same or even

Instance	$ncl$	VNS/LDS+CP		UDGVNS		UPDGVNS		FIXBB		TOULBAR2		Speed-up	
		Succ.	Time (s)	Succ.	Time (s)	Succ.	Time (s)	$\Delta E$	Time (s)	$\Delta E$	Time (s)	(1/2)	(1/3)
<b>5dbl</b>	87	10/10	2,762±67	10/10	249±7	10/10	62±7	0.28	783	0	11.09	44.27	3.98
5jdd	168	0/10	-	10/10	1,611±20	10/10	243±13	4.08	20,662	0.04	-	-	6.63
<b>3r8q</b>	157	0/10	-	10/10	1,246±292	10/10	206±8	4.19	12,762	0	-	-	6.03
<b>4bxp</b>	108	10/10	1,214±32	10/10	907±75	10/10	120±19	0.26	2,966	0	1.33	10.12	7.57
1f00	177	0/10	-	0/10	-	10/10	271±11	4.39	9,749	0	-	-	-
2x8x	131	10/10	3,312±112	4/10	2,731±1,065	10/10	215±11	4.16	69,213	3.61	1.21	15.41	12.71
<b>1xaw</b>	66	0/10	-	2/10	3,585±35	10/10	167±39	2.73	2,804	0	-	-	21.50
<b>5e10</b>	74	10/10	1,667±86	10/10	193±6	10/10	30±2	0.27	1,172	0	8.63	55.78	6.45
1dvo	82	10/10	940±23	10/10	352±5	10/10	75±2	2.90	34,143	0.18	2.67	12.47	4.66
lytq	67	10/10	2,235±211	10/10	292±1	10/10	87±1	1.67	17,064	0.31	7.65	25.57	3.33
2af5	140	0/10	-	10/10	1,029±19	10/10	281±13	4.37	86,029	0.60	-	-	3.65
1ng2	86	10/10	1,066±70	10/10	583±46	10/10	115±4	1.14	38,731	5.93	1.82	9.27	5.07
3sz7	79	0/10	-	10/10	627±54	10/10	144±11	3.11	82,626	0.54	-	-	4.35
<b>2gee</b>	110	10/10	1,648±8	10/10	961±5	10/10	149±27	1.68	5,021	0	1.71	11.07	6.45
<b>5e0z</b>	73	10/10	622±16	10/10	310±2	10/10	45±8	0.16	999	0	2.0	13.90	6.92
lyz7	87	10/10	2,149±10	10/10	2,081±38	10/10	133±4	2.91	83,817	3.21	1.03	16.18	15.67
3e3v	91	0/10	-	10/10	867±207	10/10	111±10	2.57	81,575	0.15	-	-	7.81
<b>3if9</b>	72	10/10	1,636±32	10/10	167±1	10/10	56±2	2.41	2,667	0	9.81	29.06	2.96
lis1	107	10/10	3,179±52	10/10	444±27	10/10	213±13	3.46	63,832	0.42	7.15	14.91	2.08
<b>5eqz</b>	89	10/10	1,850±49	10/10	528±8	10/10	95±3	2.20	12,768	0	3.50	19.44	5.54
4uos	118	10/10	2,305±988	1/10	3,380±661	10/10	167±2	5.21	58,590	17.87	0.68	13.79	20.23
-: TimeOut				<b>(1):</b> VNS/LDS+CP( $k$ add1, $\ell = 3$ )	<b>(2):</b> UDGVNS ( $k$ add1, $\ell = 3$ )	<b>(3):</b> UPDGVNS ( $npr$ , $k$ add1, $\ell = 3$ )							

Table 3: Comparative evaluation on CPD instances. In bold instances completely solved by TOULBAR2.  $ncl$  is the number of clusters.  $\Delta E$  is in Rosetta energy unit. It has been obtained by the difference of solution costs divided by the cost shift used during modeling, in this case  $10^8$ .

Instance	Speed-up		
	(TOULBAR2/VNS/LDS+CP)	(TOULBAR2/UDGVNS)	(TOULBAR2/UPDGVNS)
5dbl	0.28	3.14	12.55
3r8q	-	33.46	202.07
4bxp	3.51	4.69	35.56
1f00	-	17.02	35.93
1xaw	-	2.92	17.49
5e10	0.70	6.07	39.22
2gee	5.94	10.19	65.82
5e0z	1.60	3.22	22.36
3lf9	-	17.75	52.58
5eqz	22.61	79.18	439.48

Table 4: Comparing speed-ups of VNS methods with ( $k$  add1,  $\ell = 3$ ) strategy to obtain the best solution computed by TOULBAR2 within the 3600-seconds time limit on solved CPD instances. A '-' indicates that the corresponding solver was not able to compute a solution of equal/better quality than TOULBAR2.

716 better solution than TOULBAR2 (see column (5)  $\Delta E$  in Table 3). On 11 instances  
717 unsolved by TOULBAR2, UDG VNS and UPDGVNS always obtain solutions of bet-  
718 ter quality. The energetic gap  $\Delta E$  in the worst case reaches 17.86 R.U.E. Con-  
719 cerning the number of successful runs reported over the 10 runs, VNS/LDS+CP  
720 seems less robust respectively than UDG VNS and the parallel release UPDGVNS.  
721 This last one provides in all cases the best solution over all. Table 3 also com-  
722 pares the VNS methods in terms of speedups. We observe that speedup values are  
723 fluctuating from one instance to another, very likely due to the tree decomposition  
724 resulting from the 3D shape of each instance. For VNS/LDS+CP and UDG VNS,  
725 it range between 0.68 and 11.09 over the 14 instances solved by both methods.  
726 As expected, when tree decomposition and parallelization are used, not only the  
727 speed of resolution increases but the reliability too (speedup values between 9.27  
728 and 55.78). Moreover, the comparison between UDG VNS and UPDGVNS shows  
729 significant accelerations (between 2.08 and 21.50), thus confirming the practical  
730 interest of parallelization in addition to the exploitation of problem decomposi-  
731 tion.

732 *C) Comparing anytime performances of VNS methods.* We have also investigated  
733 the anytime performances of three VNS methods by reporting the CPU-times re-  
734 quired within the time limit of 1 hour to reach a solution of equal quality computed  
735 by TOULBAR2 within the 3600-seconds time limit. For solved instances, accord-  
736 ing to details in Table 4, both UDG VNS and UPDGVNS find optimal solutions more

Instance	Time (s)			Speed-up		
	VNS/IDS+CP	UDGVNS	UPDGVNS	(TOULBAR2/VNS/IDS+CP)	(TOULBAR2/UDGVNS)	(TOULBAR2/UPDGVNS)
5jdd	-	1,611.39±20	235.52±11	-	12.82	87.73
2x8x	2,070.49±62	883.50±56	176.48±2	33.42	78.33	392.18
1dvo	885.47±21	264.56±3	68.83±2	38.55	129.05	496.04
1ytq	2,194.63±210	291.93±1	85.19±1	7.77	58.45	200.30
2af5	-	815.62±13	254.68±10	-	105.47	337.79
1ng2	530.56±3	304.94±2	83.88±3	73.00	127.01	461.74
3sz7	3,029.44±566	565.29±40	125.62±7	27.27	146.16	657.74
1yz7	970.24±4	419.47±3	115.22±2	86.38	199.81	727.45
3e3v	2,332.23±46	288.46±2	93.44±2	34.97	282.79	873.01
1isl	2,803.27±45	437.44±28	208.47±12	22.77	145.92	306.19
4uos	460.73±3	420.18±2	153.11±1	127.16	139.43	382.66

Table 5: Comparing CPU times spent by VNS methods with ( $k = \text{add1}, \ell = 3$ ) strategy to obtain the best solution computed by TOULBAR2 within the 3600-seconds time limit on unsolved CPD instances. A '-' indicates that the corresponding solver was not able to compute a solution of equal/better quality than TOULBAR2.

Instance	UDGVNS		UPDGVNS(10)		UPDGVNS(30)		UPDGVNS(96)		Speed-up				
	Succ.	Time (s)	Succ.	Time (s)	Succ.	Time (s)	succ.	Time (s)	(1/2)	(1/3)	(1/4)	(2/4)	(3/4)
5dbl	10/10	249±7	10/10	127±23	10/10	88±12	10/10	62±7	1.96	2.83	3.98	2.02	1.40
5jdd	10/10	1,611±20	10/10	603±67	10/10	342±22	10/10	243±13	2.67	4.71	6.63	2.47	1.40
3r8q	10/10	1,246±292	10/10	508±61	10/10	301±28	10/10	206±8	2.45	4.13	6.03	2.46	1.46
4bxp	10/10	907±75	10/10	301±39	10/10	186±26	10/10	120±19	3.01	4.88	7.57	2.50	1.55
1f00	0/10	-	10/10	589±75	10/10	361±47	10/10	271±11	-	-	-	-	1.32
2x8x	4/10	2,731±1,065	10/10	600±198	10/10	322±44	10/10	215±11	4.55	8.49	12.71	2.79	1.49
1xaw	2/10	3,585±35	10/10	434±162	10/10	219±56	10/10	167±39	8.25	16.37	21.50	2.60	1.31
5e10	10/10	193±6	10/10	59±15	10/10	35±6	10/10	30±2	3.26	5.54	6.45	1.97	1.16
1dvo	10/10	352±5	10/10	134±18	10/10	102±6	10/10	75±2	2.61	3.44	4.66	1.78	1.35
1ytiq	10/10	292±1	10/10	148±12	10/10	102±5	10/10	87±1	1.97	2.87	3.33	1.68	1.16
2af5	10/10	1,029±19	10/10	529±47	10/10	362±20	10/10	281±13	1.94	2.84	3.65	1.87	1.28
1ng2	10/10	583±46	10/10	223±30	10/10	160±20	10/10	115±4	2.61	3.64	5.07	1.94	1.39
3sz7	10/10	627±54	10/10	392±39	10/10	221±20	10/10	144±11	1.60	2.83	4.35	2.71	1.53
2gee	10/10	961±5	10/10	252±28	10/10	203±44	10/10	149±27	3.81	4.72	6.45	1.69	1.36
5e0z	10/10	310±2	10/10	93±19	10/10	65±10	10/10	45±8	3.33	4.74	6.92	2.07	1.45
1yz7	10/10	2,081±38	10/10	242±19	10/10	174±24	10/10	133±4	8.59	11.98	15.67	1.82	1.30
3e3v	10/10	867±207	10/10	265±42	10/10	170±19	10/10	111±10	3.26	5.08	7.81	2.39	1.53
3lf9	10/10	167±1	10/10	100±17	10/10	71±4	10/10	56±2	1.66	2.36	2.96	1.78	1.25
1is1	10/10	444±27	10/10	454±87	10/10	317±30	10/10	213±13	.97	1.40	2.08	2.12	1.48
5eqz	10/10	528±8	10/10	187±22	10/10	120±4	10/10	95±3	2.82	4.40	5.54	1.96	1.26
4uos	1/10	3,380±661	10/10	217±18	10/10	189±6	10/10	167±2	15.58	17.91	20.23	1.29	1.12
-: TimeOut					(1): UDGVNS	(2): UDGVNS (10)	(3): UPDGVNS (30)	(4): UPDGVNS (96)					

Table 6: Impact of the number of processes on the parallelization.

737 quickly than TOULBAR2. For UDGVNS, speedup values range from 2.92 to 79.18  
738 with a mean of 17.76 over all the solved instances. For UPDGVNS, the ratio in  
739 terms of speedup is greatly amplified (between 12.55 and 439.48 with a mean of  
740 92.30 over all the solved instances). On the 11 opened instances, results (sum-  
741 marized in Table 5) show a clear ordering in terms of CPU times across different  
742 solvers, from TOULBAR2, VNS/LDS+CP, UDGVNS, and UPDGVNS. The speedup  
743 values are significantly improved, in particular with UPDGVNS (between 87.73  
744 and 873.01 with a mean of 447, 53 over all the unsolved instances). These re-  
745 sults confirm the superiority of VNS methods in terms of anytime performance as  
746 compared to TOULBAR2.

747 One putative explanation of the observed performance ordering between VNS  
748 methods may be the problem exploration coverage with the picked neighborhoods  
749 during search. Indeed, respectively with 1, 034 and 597 the average total number  
750 of neighborhoods explored during search for the 21 CPD instances is in aver-  
751 age 1.73 higher in VNS/LDS+CP than UDGVNS. Accordingly, tree decomposition  
752 picks more pertinent neighborhoods than VNS/LDS+CP and seems to increase the  
753 probability for a full problem coverage, which can be explained by the decreasing  
754 of possible combination as consequence of the partition in clusters of variables.  
755 Besides, parallelization is one way to increase the coverage probability, because  
756 it is a simple way to increase the number of processed neighborhoods. Conse-  
757 quently, in practice, with 3, 320 neighborhoods in average, UPDGVNS explores in  
758 5.56 times more subproblems than the corresponding sequential version. This fact  
759 can be an explanation of the good quality of the observed results.

### 760 6.2.3. *Parallelization*

761 For the last experiment, we analyzed the performance of our parallel algorithm  
762 by measuring the speed-up on varying the number of processes. We consider the  
763 number of successful runs as well as the average CPU times  $\pm$  the standard devia-  
764 tion (over the 10 runs) of the parallel release UPDGVNS (using  $k$  add1 with  $\ell = 3$ )  
765 to reach the optimum (or the best found solution for unsolved instances) within a  
766 time limit of 1-hour per run. We compare with UDGVNS with ( $k$  add1,  $\ell = 3$ ).  
767 We set  $n_{pr}$  to 10, 30 and 96 respectively (including the master process).

768 As can be seen from Table 6, UPDGVNS with 10, 30 or 96 processors, are  
769 able to obtain better results faster than UDGVNS (using one core). Table 6 also  
770 reports the speed-up values for different number of processes. Comparing to 10-  
771 processes, the improvements to UDGVNS yield a speed-up of 1.6 to 15.58 (3.85  
772 on average). Moreover, with the increase of the number of processes the gains  
773 in terms of CPU times are remarkably amplified. In the case of 30 processes the

774 speed-up is 5.75 on average and for 96 processes it is 7.68 compared to the results  
775 of UDGVNS. These results show that our parallel release on CPD instances is less  
776 sensitive to the communication overhead when increasing the number of processes  
777 and do not impact the overall efficiency of our approach.

## 778 **7. Conclusion and Perspectives**

779 In this paper we proposed a unified view of VNS methods including various  
780 LDS and neighborhood evolution strategies. Experiments performed on difficult  
781 instances, coming from a large graphical model benchmark, showed that our hy-  
782 brid method has a much better anytime behavior than existing tree search methods  
783 and still being competitive for proving optimality. UDGVNS takes advantage of the  
784 good convergence properties of DGVNS and proves optimality in many cases. On  
785 structured or unstructured problems of large sizes, like CPD, UDGVNS obtains  
786 solutions of (very) good quality, thus outperforming the state-of-the-art FIXBB  
787 Rosetta Modeling software package used in the CPD field. We further proposed a  
788 parallel version of our method improving its anytime behavior. It remains as future  
789 work to manage dynamically the tree-decomposition associated to the instance to  
790 solve. Another promising research direction is to use machine learning techniques  
791 to identify the best decompositions to be used for the practical solving [53].

## 792 **Acknowledgments**

793 This work has been partially funded by the french Agence nationale de la  
794 Recherche, reference ANR-16-C40-0028. We are grateful to the genotoul bioin-  
795 formatics platform Toulouse Midi-Pyrenees (Bioinfo Genotoul) and the high per-  
796 formance center of Cerist-Algiers in Algeria for providing computing and storage  
797 resources. We thank Mathieu Fontaine for his contribution to the code of DGVNS.

## 798 **Acronyms**

799 **BTD-HBFS** Backtrack Tree Decomposition - Hybrid Best First Search. 30, 31

800 **CFN** Cost Function Network. 3, 15, 21

801 **CPD** Computational Protein Design. 30–32, 37, 38

802 **CVPR** Computer Vision and Pattern Recognition. 3, 21

803 **DFBB** Depth-First Branch and Bound. 4–7, 21, 25, 26

804 **DGVNS** Decomposition Guided VNS. 2, 3, 7–12, 14, 16, 19, 21, 22, 38

805 **EDAC** Existential Directional Arc consistency. 5

806 **FIXBB** Fixed BackBone. 3, 30, 32, 33

807 **ILDS** Iterative LDS. 5–7, 12, 15

808 **LDS** Limited Discrepancy Search. 2, 5–7, 9–14, 21, 25, 26, 38

809 **LDS<sup>r</sup>** Restricted LDS. 9–16, 18–20, 22

810 **LNS** Large Neighborhood Search. 2, 7

811 **MCS** Maximum Cardinality Search. 8

812 **min-fill** Minimum Fill-in. 8, 9, 23

813 **PIC** Probabilistic Inference Challenge. 3, 21, 23

814 **UAI** Uncertainty in Artificial Intelligence. 3, 21, 24, 31

815 **UDGVNS** Unified Decomposition Guided VNS. 2, 3, 12–18, 22, 25–28, 30–38

816 **UPDGVNS** Unified Parallel DGVNS. 3, 16, 17, 19, 20, 22, 27, 28, 30–37

817 **VAC** Virtual Arc Consistency. 21

818 **VNS** Variable Neighborhood Search. 1, 2, 7–10, 12, 15, 21, 25, 26, 30, 32, 34,  
819 37, 38

820 **VNS/LDS+CP** VNS/LDS + Constraint Programming. 2, 7, 8, 11, 30, 32–35, 37

821 **References**

- 822 [1] D. Koller, N. Friedman, Probabilistic graphical models: principles and tech-  
823 niques, The MIT Press, 2009.
- 824 [2] S. Shimony, Finding MAPs for belief networks is NP-hard, *AI* 68 (1994)  
825 399–410.
- 826 [3] R. Marinescu, R. Dechter, Memory intensive and/or search for combinatorial  
827 optimization in graphical models, *AI* 173 (16-17) (2009) 1492–1524.
- 828 [4] L. Otten, R. Dechter, Anytime AND/OR depth-first search for combinatorial  
829 optimization, *AI Communications* 25 (3) (2012) 211–227.
- 830 [5] D. Allouche, S. de Givry, G. Katsirelos, T. Schiex, M. Zytnicki, Anytime  
831 Hybrid Best-First Search with Tree Decomposition for Weighted CSP, in:  
832 *Proc. of CP*, 2015, pp. 12–28.
- 833 [6] V. Kolmogorov, Convergent Tree-Reweighted Message Passing for Energy  
834 Minimization, *IEEE T. Patt. Anal. Mach. Intell.* 28 (10) (2006) 1568–1583.
- 835 [7] M. Wainwright, M. Jordan, Graphical models, exponential families, and  
836 variational inference, *F&T. in Machine Learning* 1 (1–2) (2008) 1–305.
- 837 [8] D. Sontag, T. Meltzer, A. Globerson, Y. Weiss, T. Jaakkola, Tightening LP  
838 relaxations for MAP using message-passing, in: *Proc. of UAI*, 2008, pp.  
839 503–510.
- 840 [9] D. Sontag, D. Choe, Y. Li, Efficiently searching for frustrated cycles in MAP  
841 inference, in: *Proc. of UAI*, 2012, pp. 795–804.
- 842 [10] H. Wang, D. Koller, Subproblem-tree calibration: A unified approach to  
843 max-product message passing, in: *Proc. of ICML*, 2013, pp. 190–198.
- 844 [11] J. Park, Using weighted MAX-SAT engines to solve MPE, in: *Proc. of*  
845 *AAAI*, 2002, pp. 682–687.
- 846 [12] F. Hutter, H. Hoos, T. Stützle, Efficient stochastic local search for MPE solv-  
847 ing, in: *Proc. of IJCAI*, 2005, pp. 169–174.
- 848 [13] O. Mengshoel, D. Wilkins, D. Roth, Initialization and Restart in Stochastic  
849 Local Search: Computing a Most Probable Explanation in Bayesian Net-  
850 works, *IEEE Trans. Knowl. Data Eng.* 23 (2) (2011) 235–247.

- 851 [14] R. Marinescu, K. Kask, R. Dechter, Systematic vs. non-systematic algo-  
852 rithms for solving the MPE task, in: Proc. of UAI, 2003, pp. 394–402.
- 853 [15] P. Shaw, Using constraint programming and local search methods to solve  
854 vehicle routing problems, in: Proc. of CP, 1998, pp. 417–431.
- 855 [16] L. Perron, P. Shaw, V. Furnon, Propagation guided large neighborhood  
856 search, in: Proc. of CP, 2004, pp. 468–481.
- 857 [17] M. Lombardi, P. Schaus, Cost impact guided lns, in: Proc. of Integration of  
858 AI and OR Techniques in Constraint Programming, 2014, pp. 293–300.
- 859 [18] J. J. Dekker, M. G. de la Banda, A. Schutt, P. J. Stuckey, G. Tack, Solver-  
860 independent large neighbourhood search, in: Proc. of CP, 2018, pp. 81–98.
- 861 [19] E. Demirovic, G. Chu, P. J. Stuckey, Solution-based phase saving for CP:  
862 A value-selection heuristic to simulate local search behavior in complete  
863 solvers, in: Proc. of CP, 2018, pp. 99–108.
- 864 [20] N. Mladenović, P. Hansen, Variable Neighborhood Search, *Comput. Oper.*  
865 *Res.* 24 (11) (1997) 1097–1100.
- 866 [21] S. Loudni, P. Boizumault, Solving constraint optimization problems in any-  
867 time contexts, in: Proc. of IJCAI, 2003, pp. 251–256.
- 868 [22] W. Harvey, M. Ginsberg, Limited discrepancy search, in: Proc. of IJCAI,  
869 1995, pp. 607–615.
- 870 [23] M. Fontaine, S. Loudni, P. Boizumault, Exploiting tree decomposition for  
871 guiding neighborhoods exploration for VNS, *RAIRO OR* 47 (2) (2013) 91–  
872 123.
- 873 [24] T. Davidovic, T. Crainic, MPI parallelization of variable neighborhood  
874 search, *Electronic Notes in Discrete Mathematics* 39 (2012) 241–248.
- 875 [25] A. Ouali, S. Loudni, L. Loukil, P. Boizumault, Y. Lebbah, Replicated parallel  
876 strategies for decomposition guided VNS, *ENDM* 47 (2015) 93–100.
- 877 [26] D. Simoncini, D. Allouche, S. de Givry, C. Delmas, S. Barbe, T. Schiex,  
878 Guaranteed discrete energy optimization on large protein design problems,  
879 *J. of Chemical Theo. and Comput.* 11(12) (2015) 5980–5989.

- 880 [27] P. Meseguer, F. Rossi, T. Schiex, Soft constraints processing, in: Handbook  
881 of Constraint Programming, Elsevier, 2006, Ch. 9, pp. 279–326.
- 882 [28] B. Hurley, B. O’Sullivan, D. Allouche, G. Katsirelos, T. Schiex, M. Zyt-  
883 nicki, S. de Givry, Multi-Language Evaluation of Exact Solvers in Graphical  
884 Model Discrete Optimization, *Constraints* 21 (3) (2016) 413–434.
- 885 [29] R. Dechter, I. Rish, Mini-buckets: A general scheme for bounded inference,  
886 *Journal of the ACM (JACM)* 50 (2) (2003) 107–153.
- 887 [30] J. Larrosa, S. de Givry, F. Heras, M. Zytnicki, Existential arc consistency:  
888 getting closer to full arc consistency in weighted CSPs, in: *Proc. of IJCAI*,  
889 2005, pp. 84–89.
- 890 [31] M. Cooper, S. de Givry, M. Sanchez, T. Schiex, M. Zytnicki, T. Werner, Soft  
891 arc consistency revisited, *AI* 174 (2010) 449–478.
- 892 [32] W. Karoui, M.-J. Huguet, P. Lopez, W. Naanaa, Yields: A yet improved lim-  
893 ited discrepancy search for csps, in: *Proc. of Integration of AI and OR Tech-  
894 niques in Constraint Programming for Combinatorial Optimization Prob-  
895 lems, 2007*, pp. 99–111.
- 896 [33] P. Prosser, C. Unsworth, Limited discrepancy search revisited, *ACM Journal*  
897 *of Experimental Algorithmics* 16 (2011) 1.6:1.1–1.6:1.18.
- 898 [34] H. Bodlaender, A. Koster, Treewidth computations i. upper bounds, *Inf.*  
899 *Comput.* 208 (3) (2010) 259–275.
- 900 [35] H. Bodlaender, A. Koster, F. Van den Eijkhof, Preprocessing rules for trian-  
901 gulation of probabilistic networks, *Computational Intelligence* 21 (3) (2005)  
902 286–305.
- 903 [36] S. Arnborg, et al., Complexity of finding embeddings in a  $k$ -tree, *SIAM*  
904 *Journal on Algebraic and Discrete Methods* 8 (1987) 277–284.
- 905 [37] U. Kjrulff, Triangulation of graphs – algorithms giving small total state  
906 space, Tech. rep., Aalborg University (1990).
- 907 [38] R. E. Tarjan, et al., Simple linear-time algorithms to test chordality of graphs,  
908 test acyclicity of hypergraphs, and selectively reduce acyclic hypergraphs,  
909 *SIAM Journal on Computing* 13 (3) (1984) 566–579.

- 910 [39] M. Luby, A. Sinclair, D. Zuckerman, Optimal speedup of Las Vegas algo-  
911 rithms, in: Proc. of TCS, 1993, pp. 128–133.
- 912 [40] F. Boussemart, F. Hemery, C. Lecoutre, L. Sais, Boosting systematic search  
913 by weighting constraints, in: Proc. of ECAI, 2004, pp. 146–150.
- 914 [41] J. Kappes, B. Andres, F. Hamprecht, C. Schnörr, S. Nowozin, D. Batra,  
915 S. Kim, B. Kausler, T. Kröger, J. Lellmann, N. Komodakis, B. Savchyn-  
916 skyy, C. Rother, A comparative study of modern inference techniques for  
917 structured discrete energy minimization problems, *Int. J. of Computer Vi-*  
918 *sion* 115 (2) (2015) 155–184.
- 919 [42] A. Favier, S. Givry, A. Legarra, T. Schiex, Pairwise decomposition for com-  
920 binatorial optim. in graphical models, in: Proc. of IJCAI, 2011, pp. 2126–  
921 2132.
- 922 [43] S. de Givry, S. Prestwich, B. O’Sullivan, Dead-end elimination for weighted  
923 CSP, in: Proc. of CP, 2013, pp. 263–272.
- 924 [44] A. Ihler, N. Flerova, R. Dechter, L. Otten, Join-graph based cost-shifting  
925 schemes, in: Proc. of UAI, 2012, pp. 397–406.
- 926 [45] L. Otten, A. Ihler, K. Kask, R. Dechter, Winning the PASCAL 2011 MAP  
927 challenge with enhanced AND/OR branch-and-bound, in: Proc. of NIPS  
928 Workshop DISCML, 2012.
- 929 [46] B. Neveu, G. Trombettoni, F. Glover, ID Walk: A Candidate List Strategy  
930 with a Simple Diversification Device, in: Proc. of CP, 2004, pp. 423–437.
- 931 [47] J. M. Mooij, libDAI: A free and open source C++ library for discrete ap-  
932 proximate inference in graphical models, *JMLR* 11 (2010) 2169–2173.
- 933 [48] L. Otten, R. Dechter, And/or branch-and-bound on a computational grid,  
934 *JAIR* 59 (2017) 351–435.
- 935 [49] D. Allouche, J. Davies, S. de Givry, G. Katsirelos, T. Schiex, S. Traoré,  
936 I. André, S. Barbe, S. Prestwich, B. O’Sullivan, Computational Protein De-  
937 sign as an Optimization Problem, *AI* 212 (2014) 59–79.
- 938 [50] P.-S. Huang, S. E. Boyken, D. Baker, The coming of age of de novo protein  
939 design, *Nature* 537 (2016) 320–327.

- 940 [51] S. M. Lippow, B. Tidor, Progress in computational protein design, *Current*  
941 *Opinion in Biotechnology* 18 (4) (2007) 305 – 311, protein technologies /  
942 *Systems biology*. doi:[https://doi.org/10.1016/j.copbio.](https://doi.org/10.1016/j.copbio.2007.04.009)  
943 2007.04.009.  
944 URL [http://www.sciencedirect.com/science/article/](http://www.sciencedirect.com/science/article/pii/S0958166907000778)  
945 [pii/S0958166907000778](http://www.sciencedirect.com/science/article/pii/S0958166907000778)
- 946 [52] D. L. Trudeau, D. S. Tawfik, Protein engineers turned evolution-  
947 iststhe quest for the optimal starting point, *Current Opinion in*  
948 *Biotechnology* 60 (2019) 46 – 52, pharmaceutical Biotechnology.  
949 doi:<https://doi.org/10.1016/j.copbio.2018.12.002>.  
950 URL [http://www.sciencedirect.com/science/article/](http://www.sciencedirect.com/science/article/pii/S095816691830209X)  
951 [pii/S095816691830209X](http://www.sciencedirect.com/science/article/pii/S095816691830209X)
- 952 [53] M. Abseher, N. Musliu, S. Woltran, Improving the efficiency of dynamic  
953 programming on tree decompositions via machine learning, *JAIR* 58 (2017)  
954 829–858.