

Fast Search Algorithms for Computational Protein Design

Seydou Traoré,^[a,b,c] Kyle E. Roberts,^[d] David Allouche,^[e] Bruce R. Donald,^[d]
Isabelle André,^[a,b,c] Thomas Schiex,^[e] and Sophie Barbe*^[a,b,c]

One of the main challenges in computational protein design (CPD) is the huge size of the protein sequence and conformational space that has to be computationally explored. Recently, we showed that state-of-the-art combinatorial optimization technologies based on Cost Function Network (CFN) processing allow speeding up provable rigid backbone protein design methods by several orders of magnitudes. Building up on this, we improved and injected CFN technology into the well-established CPD package *Osprey* to allow all *Osprey* CPD algorithms to benefit from associated speedups. Because *Osprey* fundamentally relies on the ability of A^* to produce conformations in increasing order of energy, we defined new A^* strat-

egies combining CFN lower bounds, with new side-chain positioning-based branching scheme. Beyond the speedups obtained in the new A^* -CFN combination, this novel branching scheme enables a much faster enumeration of suboptimal sequences, far beyond what is reachable without it. Together with the immediate and important speedups provided by CFN technology, these developments directly benefit to all the algorithms that previously relied on the DEE/ A^* combination inside *Osprey** and make it possible to solve larger CPD problems with provable algorithms. © 2016 Wiley Periodicals, Inc.

DOI: 10.1002/jcc.24290

Introduction

Computational Protein Design (CPD) has become a valuable tool for creating proteins with desired biophysical and functional properties and for assessing our understanding of protein sequence–structure–function relationships. By combining physicochemical models governing relations between protein amino-acid composition and the protein three-dimensional structure with advanced computational algorithms, CPD seeks to identify one or a set of amino-acid sequences that fold into a given 3D-scaffold and that will bestow the redesigned protein with targeted properties. Such rational protein design approaches have been successfully applied to alter intrinsic properties (stability, binding affinity, ...) of existing proteins or to endow them with new functionalities, leading to the generation of novel enzymatic catalysts, therapeutic proteins, protein–protein interfaces, and self-assembling protein structures.^[1–6] The application of this technology is broad, ranging from medicine, biotechnology, and synthetic biology to nanotechnologies.^[7]

Despite notable results, substantial methodological advances are still needed to improve CPD performances and extend its effective application. The success of CPD predictions depends on several elements, which include the biologically meaningful modeling of the design problem, the accuracy of the energy and objective functions used to assess fitness of the predicted sequence-structures, and the efficiency of the search algorithms to find solutions in a timely manner. However, to face the exponential size of the search space defined by the composition of protein sequences and conformations, CPD approaches have to strike a compromise between speed and accuracy. Most of the CPD methods rely on (1) a coarse-graining of the structure as a sequence of discrete side-chain rotamers, (2) an assumption of modest backbone conforma-

tional flexibility, where often a fixed backbone or a set of possible backbones are used, and (3) an approximation of the energy model as a pairwise decomposition. Since the problem of searching for an optimal solution (GMEC: Global Minimum-Energy Conformation) over the conformational space of rotamers and possibly backbones is NP-hard,^[8] a variety of methods, both meta-heuristics (Monte Carlo simulated annealing,^[9,10] genetic algorithms,^[11] ...) and provable algorithms (Dead-End Elimination (DEE), Branch-and-Bound algorithms (BB), Integer Linear Programming (ILP), Dynamic Programming (DP),^[11–16] ...) have been proposed over the years. However, there is still a need for more efficient optimization techniques, capable of exploring vaster combinatorial spaces, representing more realistic and flexible protein models.

This article focuses on exact optimization and enumeration techniques. As provable methods know when a global

[a] S. Traoré, I. André, S. Barbe
Université De Toulouse, INSA, UPS, INP, LISBP, 135 Avenue de Rangueil,
Toulouse F-31077, France
E-mail: sophie.barbe@insa-toulouse.fr

[b] S. Traoré, I. André, S. Barbe
INRA, UMR792, Ingénierie des Systèmes Biologiques et des Procédés,
Toulouse F-31400, France

[c] S. Traoré, I. André, S. Barbe
CNRS, UMR5504, Toulouse F-31400, France

[d] K. E. Roberts, B. R. Donald
Department of Biochemistry, Department of Computer Science,
Department of Chemistry, Duke University, Durham, North Carolina

[e] D. Allouche, T. Schiex
Unité de Mathématiques et Informatique Appliquées, Toulouse, UR 875,
INRA, Castanet, Tolosan F-31320, France

Contract grant sponsor: Agence Nationale de la Recherche; Contract grant number: ANR-10-BLA-0214 and ANR-12-MONU-0015-03; Contract grant sponsor: INRA; Contract grant sponsor: Region Midi-Pyrénées; Contract grant sponsor: NIH; Contract grant number: 2-R01-GM-78031-5

© 2016 Wiley Periodicals, Inc.

optimum is reached, the search can be stopped with confidence and an exact solution obtained, sometimes in significantly less time than with meta-heuristics. It has also been observed that the accuracy of meta-heuristic approaches tends to degrade in unpredictable ways as the problem size increases.^[9] Finally, exact methods are the only methods which offer a provable basis for improving biophysical models. Indeed, they ensure that discrepancies between CPD predictions and experimental results come exclusively from modeling inadequacies and not from the algorithm. These properties are crucial to rationally tune the energy function, e.g., on the basis of a fixed-composition full redesign.^[17] Currently, the most usual provable and deterministic methods for CPD rely on the Dead-End Elimination theorem and the A^* algorithm.^[12] DEE is used as a preprocessing method. It removes rotamers which are energetically dominated by other rotamers and therefore useless to identify a global optimum. This idea has later been extended to prune pairs or higher order combinations of rotamers at different residues in order to improve the pruning power.^[18–20] However, because CPD is NP-hard, the polynomial time DEE algorithms usually cannot identify a unique sequence-conformation model. DEE preprocessing is therefore followed by an A^* search (a provable Best-First search strategy) which expands a sequence-conformation tree by tentatively assigning rotamers to residues. A^* is originally a path planning algorithm, and it has a worst-case exponential time and memory consumption. This means that it can easily choke on problems with many undominated rotamers. However it has the capacity to not only identify the GMEC but also produce an arbitrary long energy-sorted list of solutions. Thus, the same DEE/ A^* can be used to find the GMEC and to produce an energy-sorted stream of sequence-conformations.

In a recent work, we have shown that the already highly challenging usual description of the CPD problem, based on rigid backbone and discrete rotamers, could be formulated and efficiently solved as a Cost Function Network (CFN).^[21–23] CFN algorithms are able to handle complex CPD combinatorial spaces which are out of reach of a broad range of combinatorial optimization technologies including the usual DEE/ A^* , 0/1 Linear and Quadratic Programming, 0/1 Quadratic Optimization, Weighted Partial MaxSAT and Graphical Model optimization methods.^[21–23] The *toulbar2* CFN solver provides speedups of several orders of magnitude both to provably find the GMEC and to exhaustively enumerate unsorted ensembles of near-optimal solutions (within a threshold of the GMEC), offering an attractive alternative method for CPD.

The outperformance of CFN methods opens new avenues to integrate further molecular flexibility in CPD which inevitably leads to a tremendous expansion of the conformational search space to be considered. CPD methods capable of tackling different levels of molecular flexibility have raised a growing interest^[24,25] as this has been identified as being one major cause of design failures. Indeed, the lack of molecular flexibility may lead to a significant loss of sequence space accessible to properly folded and functional protein thus introducing some biases in sequence selection. However, in spite of its crucial

importance, incorporation of more realistic (macro)molecular flexibility into CPD remains a major challenge.

Herein, we integrated the CFN technology in the well-established *Osprey* software (modified version of *Osprey* 2.0).^[26,27] Beyond traditional energy functions, *Osprey* includes models and provable algorithms that capture macromolecular flexibility and enable reliable biophysical modeling, making it a target of choice to disseminate CFN technology inside the CPD community. It is noteworthy that this CPD-dedicated software has been prospectively used, with experimental validation, to redesign enzymes toward noncognate substrates,^[1,28,29] design new drugs,^[30] predict drug-resistance mutations,^[2] design peptide inhibitors of protein–protein interactions,^[31] and design epitope-specific antibody probes.^[32] All provable methods inside *Osprey* intimately depend on the capacity of A^* to produce an energy-sorted stream of solutions, something that the *toulbar2* solver does not directly provide. We therefore injected the CFN bounding technology inside A^* so that all *Osprey* algorithms, including those with flexible modeling and affinity estimation, directly benefit from the enhanced bounding, leading to better efficiency.

To further improve this new CFN-based method,^[21–23] we integrated other essential components of CFN technology, such as ordering heuristics and branching schemes. Inspired by the specific nature of the CPD problem, we further developed a new branching scheme that provides further speedups. This new branching scheme can avoid the expensive enumeration of all conformations for a given sequence, allowing to directly produce an ensemble of near-optimal sequences. Thanks to this, exhaustive libraries of near-optimal sequences and energy-sorted lists of sequences can be directly produced with a much higher efficiency.

The performances of these methods have been assessed on the design of more stable proteins and cofactor-bound proteins, as well as protein–ligand and protein–protein interfaces. The results obtained for both the GMEC identification and the enumeration of near-optimal sequence-conformations were compared to those obtained using the DEE/ A^* approach implemented in *Osprey** in terms of speedup of the combinatorial optimization and enumeration step.

To make this article more self-contained, we provide hereafter a short description of the CPD underlying concepts and methods.

Background

The CPD problem formulation

The rigid backbone and discrete rotamer CPD problem is defined by: (i) a fixed backbone of a protein structure; (ii) a set of amino-acid residues to be designed, called “designable residues”; (iii) a group of allowed amino-acids for each designable residue and their respective set of discrete low energy side chain conformations, called rotamers; and (iv) pairwise atomic energy functions to evaluate the model. Rotamers correspond to cluster centers of well-represented amino-acid side-chain

*The modified *Osprey* source code consistent with this work is available at <http://www.cs.duke.edu/donaldlab/osprey.versions.php#CFNosprey>

conformations mined from a database of 3D protein structures. In the case of protein–ligand systems, the conformational flexibility of peptide ligands is also described by discrete rotamer libraries. In the case of nonpeptide ligands, the treatment of organic molecule flexibility is often left to the user to precalculate an ensemble of low-energy conformers for the ligand (used as a rotamer library).^[4]

A sequence-conformation model is defined by the choice of one specific amino-acid with one associated conformation (rotamer) for each designable residue. Its total energy (E_{total}) is defined by:

$$E_{\text{total}} = E_c + \sum_i E(i_r) + \sum_i \sum_{j < i} E(i_r, j_s) \quad (1)$$

where E_c is the constant energy contribution capturing interactions between fixed parts of the model, $E(i_r)$ depends on rotamer r at position i (and its reference energy) and $E(i_r, j_s)$ is the pairwise interaction energy between rotamer r at position i and rotamer s at position j .

The combinatorial optimization problem is to find, in the space of all complete rotamer assignments (each defined by a sequence and a conformation), one assignment that provably minimizes E_{total} .

Modeling CPD as a Cost Function Network

The problem of finding the set of rotamers that will minimize the total energy (E_{total}) can be easily formulated as a Cost Function Network problem.^[21–23]

A CFN is defined by a set of variables which are connected by a set of local cost functions.^[33] Formally, a CFN is a quadruple (X, D, C, k) where $X = \{x_1, x_2, \dots, x_n\}$ is a set of n variables indexed by $I = \{1, \dots, n\}$. Each variable $x_i \in X$ has a discrete domain $d_i \in D$ that defines the set of values that it can take. A set of local cost functions C defines a network over X . Each cost function $c_S \in C$ is defined over a subset of variables indexed by $S \subseteq I$ (called its scope), has a domain $\prod_{x_i \in S} d_i$ and takes integer values in $\{0, 1, 2, \dots, k\}$. The cost k represents a maximum (intolerable) cost. It can be infinite or set to a finite upper bound. Values or combinations of values that are forbidden by a cost function are simply mapped to k . The global cost of a complete assignment A is defined as the sum of all cost functions on this assignment (or k if this sum is larger than k). The problem of finding an assignment of all variables that minimizes this global cost is called the Weighted Constraint Satisfaction Problem defined by the CFN. It is usually assumed that C contains one constant cost function, with an empty scope, denoted as c_{\emptyset} . Since all cost functions in a CFN are non-negative, this constant cost function $c_{\emptyset} \in C$ defines a lower bound on the cost of an optimal solution. When cost functions involve at most two variables, the CFN is said to be binary and the CFN defines a graph where variables are vertices and binary scopes are edges.

As an example, we describe a simple CFN defined by a set of three variables $X = \{x_1, x_2, x_3\}$ each with the same Boolean domain $d_1 = d_2 = d_3 = \{0, 1\}$ and a set of four cost functions $C =$

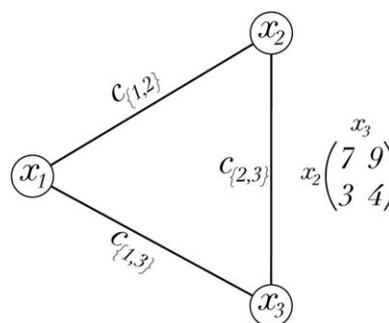


Figure 1. The graph of our CFN example together with the cost matrix of the cost function $c_{\{2,3\}}$.

$\{c_{\emptyset}, c_{\{1,2\}}, c_{\{1,3\}}, c_{\{2,3\}}\}$. The cost functions are defined as $c_{\emptyset} = 1$, $c_{\{1,2\}} = |x_1 - x_2|$, $c_{\{1,3\}} = x_1 x_3$ while $c_{\{2,3\}}$ is defined by a specific cost matrix. The upper bound $k = 9$. The graph of this CFN is represented in Figure 1, together with the cost matrix defining $c_{\{2,3\}}$. This CFN, denoted as P , defines a joint cost function $P(x_1, x_2, x_3) = \min(c_{\emptyset} + c_{\{1\}}(x_1) + c_{\{1,2\}}(x_1, x_2) + c_{\{1,3\}}(x_1, x_3) + c_{\{2,3\}}(x_2, x_3), k)$. On the triple of values $x_1 = 0, x_2 = 2, x_3 = 0$, this function has cost $1 + 1 + 0 + 3 = 5$.

The CPD optimization problem, in its pairwise-decomposed form, can be easily formulated as a binary CFN. Every designable amino-acid residue i is represented by a variable x_i and the set of rotamers available to the residue defines its domain d_i . Then, each energy term in E_{total} is represented as a cost function.^[21–23] The constant term E_c is captured as the constant cost function with empty scope (c_{\emptyset}) and terms $E(i_r)$ and $E(i_r, j_s)$ are represented by unary and binary cost functions $c_{\{i\}}$ and $c_{\{i,j\}}$, respectively, each involving the variables of the corresponding designable positions. These cost functions are positive integer cost matrices defined from precomputed energy matrices. Floating point energy terms can be mapped to positive integers through shifting and scaling according to desired precision.^[21–23] Such operations preserve the set of optimal solutions and an optimal solution of the CFN is an assignment that defines a GMEC for the CPD problem. In this model, it becomes possible to interchangeably use the notions of CFN variable and CPD designable position, the notions of value of a CFN variable and CPD rotamer, the notions of complete CFN variable assignment and of CPD sequence-conformation as well as the notions of CFN cost and CPD energy.

Solving the Optimization and Enumeration of CPD Problems

Existing approaches to provable CPD rely on the combination of several techniques. In the most usual CPD approach, DEE is first used to prune rotamers and if needed, an A^* algorithm is used to find the GMEC or enumerate sequence-conformations. The A^* algorithm is a famous path-planning search algorithm which belongs to the larger family of Branch and Bound algorithms. The idea of Branch & Bound is to split the space of all complete assignments in disjoint subspaces in some way (branching). Then a specific lower bounding mechanism is

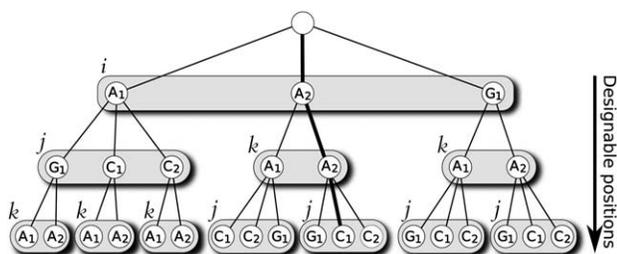


Figure 2. Tree Search using n -ary branching. We assume a trivial CPD problem with three designable positions denoted i , j , and k with respectively 3 (A_1, A_2, G_1), 3 (G_1, C_1, C_2), and 2 (A_1, A_2) rotamers. To expand a node in children nodes (branching), a yet unfixed designable position is chosen (indicated in the top left of each grey box). The children nodes (inside the grey box) correspond each to a subspace, where the position is fixed to an available rotamer. A path in the tree, as shown in bold, corresponds to a complete rotamer assignment. Note that with a dynamic variable ordering, the same designable position needs not be chosen at a given level.

used to quickly compute an underestimation of the energy of the best assignment in each subspace. By comparing this lower bound to the energy of any complete assignment (defining an upper bound on the optimum), it is possible to prove that some subspaces do not need to be explored, because the best assignment they contain has a worst cost than the available upper bound. This way, it is possible to avoid the exhaustive exploration of all subspaces.

Because of the recursive nature of splitting spaces in subspaces through branching, all the possible subspaces can be organized in a rooted tree where the root node is the space of all possible sequence-conformations and the children of a node which is not reduced to a singleton can be obtained by applying the branching scheme. In the usual A^* algorithm used in CPD, the branching scheme is a standard n -ary branching. A yet unfixed designable position x_i is chosen. Then, for each possible rotamer at this position in the considered node, one child is created where the position x_i is set to this rotamer. Notice that this branching scheme alone does not define a unique tree as the branching scheme must choose a designable position each time a subspace is split. In CPD, it is usual to choose the same variable at each level of the tree (a static variable ordering), in their protein sequence order. A leaf node defines a complete assignment of amino-acid rotamers. An internal node defines a partial protein conformation assignment. In Figure 2, we show the structure of such a n -ary branching tree for a very simple CPD problem (see caption), using an arbitrary dynamic variable ordering. But other more powerful branching schemes exist. Algorithms that aim at producing a provable GMEC or enumerating complete near-optimal assignments must traverse the defined tree in some order. A simple strategy is to use *Depth-First Search* and always explore the deepest node, leading to a *Depth-First Branch and Bound* (DFBB) algorithm. Because there are typically many nodes at the same depth, a secondary criteria, called *the value ordering heuristics*, can be used to choose the next explored node. The lower bound of an explored node is immediately computed. If this lower bound is larger than or equal to the cost of the best-known solution (the current upper bound), the branch below this node can be safely pruned and search

proceeds in the rest of the tree. Thanks to the very ordered exploration of the tree defined by depth first search, remembering the current position of the search requires only polynomial space. Ultimately, the order in which the search will proceed in the tree, whose shape is defined by the branching scheme and the variable ordering heuristics, is conditioned only by the value ordering heuristics, the lower bounding mechanism, and the initial value of the upper bound k . Indeed it may be easy to produce an initial non-naïve complete assignment that will provide a hopefully tight initial upper bound k , leading to more abundant pruning. Then, when only the GMEC is sought, each time a new complete assignment is explored (a sequence-conformation), its energy is used to tighten the upper bound k and improve pruning. In this case, the last encountered leaf will be optimal. If instead one wishes to enumerate all solutions within a Δ threshold of the GMEC, a two-step process is followed: first the GMEC, with energy e^* , is sought using the above algorithm. Then, the upper bound k is set to $(e^* + \Delta)$ and the search restarted but without updating k when a solution is found. This is guaranteed to enumerate all sequence-conformations within Δ of the GMEC.

In CPD, A^* Best-First Search is used instead. As the name indicates, the next node to explore is always the known unexplored node with the best lower bound. This usually uniquely defines the node and no secondary value ordering criteria is needed. Because this requires to remember which nodes have been explored, Best-First search requires exponential space in the worst-case: upon exploration of a node, new children nodes are pushed in a priority list of "open" nodes (originally containing only the root) from which the smallest lower bound node will be successively extracted until a leaf is extracted. In this case, it is known that all other open nodes cannot lead to a better solution than this one: the first extracted leaf is the GMEC. Without any change in the algorithm, if the extraction of best nodes continues, the stream of leaves explored will be an energy-sorted list of sequence-conformations. The order of the exploration of the tree, whose shape is defined by the branching scheme and the variable ordering heuristics, depends on the lower bounding mechanism used. DEE/ A^* relies on a simple bound^[34] which we denote in the rest of the article as the "vanilla" bound. In A^* , the lower bound is usually denoted as "the admissible heuristics."

Cost Function Network processing

In contrast with the most usual exact CPD method, where dominance analysis through the DEE theorem is widely used, the fundamental idea in CFNs relies on so-called *Local Consistencies*. From a given CFN, it is always possible to define a family of small (local) subproblems. A Local Consistency property requires that each subproblem must be sufficiently explicit about local optimal costs. As an example, the node consistency of a variable x_i with associated cost function c_i requires that d_i contains at least one value v such that $c_i(v) = 0$ and no value w such that that $c_{\emptyset} + c_i(w) \geq k$ (the forbidden cost). Equivalently, this means that there is at least one value

that does not locally increase cost and no value that is locally infeasible. If a problem does not satisfy a given Local Consistency property, it is possible to transform it in an equivalent problem (defining the same joint cost/energy distribution) satisfying the property thanks to *Equivalence Preserving Transformations* (EPTs). An EPT is a local transformation of the CFN which can shift cost (or energy) between cost functions of intersecting scopes without changing the global energy distribution. These EPTs are iteratively applied in so-called Local Consistency enforcing algorithms that iterate EPTs until the CFN satisfies the Local Consistency property. For example, if a variable violates node consistency then by deleting infeasible values and by shifting costs to c_{\emptyset} (i.e., by subtracting the minimum of all $c_i(v)$ from each $c_i(v)$ and adding it to c_{\emptyset}), the variable can be made node consistent. Beyond node consistency, many other Local Consistency properties and associated polynomial time enforcing algorithms have been defined.^[35–37] Depending on the locality of the property, which may apply to one variable, one cost function or more, they are called Node, Arc or higher order consistencies. In this paper, we only consider Node and Arc consistencies. The most desirable effects of enforcing Local Consistency is that infeasible values, which can only lead to solutions of cost larger than k , may be pruned and that the constant cost function c_{\emptyset} , which remains a lower bound on the optimum, may be increased. The amount of pruning and the strength of the lower bound may increase as the upper bound k decreases. For a given upper bound, these effects are obtained without changing the global energy distribution. Compared to Local Consistency enforcing, DEE, which has also been studied in CFN under the name of substitutability^[23,38,39] does not preserve the global energy distribution as it may remove feasible suboptimal solutions. In CFN, the pruning and non-naïve lower bounding based on c_{\emptyset} is done incrementally at each visited node during DFBB, using various branching schemes, variable and value ordering schemes as well as specialized upper bounding algorithms. Because the amount of pruning depends on the value of the upper bound k , it is important to have a reasonable upper bound at the beginning of the tree search. This can be obtained by doing bounded tree search using Limited Discrepancy Search.^[40]

Our recent studies^[21–23] highlighted the power of these CFN-based method to efficiently solve various CPD problems and hence, spurred new developments to tackle more complex and challenging CPD problems. Novel CFN-based methodological advances are provided herein. In particular, a variant of the A^* search algorithm which uses the CFN lower bound as its admissible heuristics instead of the usual CPD lower bound^[41] is defined. It is named “ A^* -CFN.” The performances of this new search method were compared to those of the A^* search commonly used in CPD, denoted as “ A^* -vanilla.” In addition, we enhanced CPD algorithms by injecting more CFN technology such as sophisticated variable ordering heuristics (for both A^* and DFBB), value ordering heuristics, upper bounding and branching schemes. The specific nature of the CPD problems was specifically used to design a new family of branching strategies based on amino-acid types and Local Consistency enforcement. The open source CPD framework

Osprey 2.0^[27] in which the CFN-based methods have been implemented, allows for discrete and continuous modeling of the protein conformation at the side-chain and backbone level.^[25,42] All those models (including the DEEPer flexibility models) enabling the consideration of more flexibility will benefit from CFN-based developments as they are ultimately based on the optimization of a pairwise energy matrix.

New CFN-Based Algorithms for CPD

The strategies offered in the CFN-based CPD framework developed in this work are described in Figure 3. They are defined through a customizable tree search algorithm that can rely on different branching strategies (n -ary, dichotomic, binary, and a new SCP-based branching scheme), search strategies (Best First A^* and DFBB), and variable/value orderings. This algorithm always relies on Local Consistency combined with a simple DEE pruning^[23] to prune rotamers and compute a lower bound at preprocessing but also at each node of the tree search. Through the use of different combinations of branching schemes and rules for updating k in DFBB, it is possible to offer a larger variety of services than the usual DEE/ A^* combination.

Tree search: strategies, heuristics, and branching schemes

Tree search-based exact optimization and enumeration of near-optimal solutions can be achieved by different combinations of basic ingredients. In the rest of the article, we examine specific combinations for CPD, some of which include new ingredients. There are however some crucial features that will remain unchanged in all the considered combinations. First of all, we will always use Local Consistency at preprocessing and at each node of the search. By increasingly simplifying the problem and strengthening the lower bound c_{\emptyset} , Local Consistency provides information that can be used to prune and heuristically guide the search in a tree defined by branching schemes and variable and value ordering. Because the pruning strength of Local Consistency depends on the tightness of the upper bound k , we also always perform an initial upper-bounding using Limited Discrepancy Search.^[40]

Best-First Search with a CFN Lower Bound: The A^* -CFN. The completeness of the A^* search method (A^* -vanilla) used in the DEE/ A^* approach^[41] relies on the use of a so-called “admissible heuristic function”: this function must provide an optimistic estimate of the energy of all complete conformations below the current node i.e., must be a lower bound on the optimum of the current problem. Thus, the lower bound defined by c_{\emptyset} and Local Consistency enforcing can be directly plugged in the A^* algorithm and replace the existing heuristics.

The lower bound used in current DEE/ A^* implementations is equivalent to the CFN “Directed Arc Consistency Counts” bound.^[43] This lower bound has been obsolete by several Arc Consistency properties.^[44] Hence, we expect that replacing A^* lower bound with stronger Local Consistency based bounds will improve the pruning efficiency of A^* . We use Existential Directed Arc Consistency (EDAC),^[44] an Arc Consistency property with a fast incremental enforcing algorithm. When Local

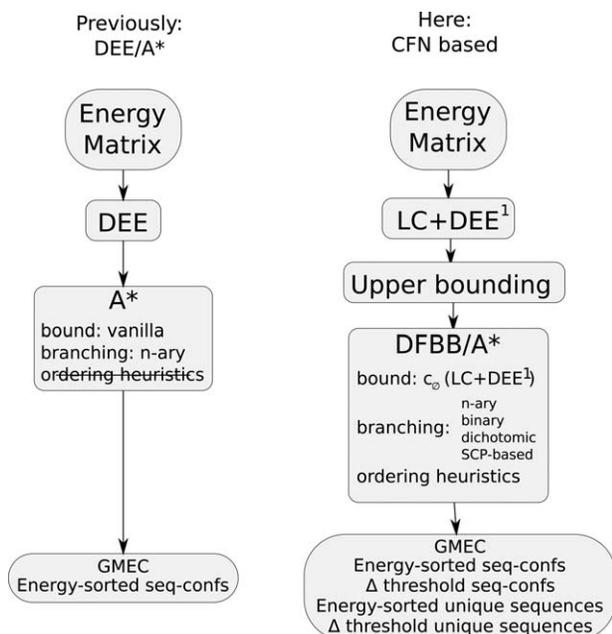


Figure 3. CPD frameworks. By upgrading components of the usual DEE/A* combination (left) using local consistency (LC), new branching schemes, and ordering heuristics, it is possible to increase efficiency for provably finding the GMEC and enumerating suboptimal sequence-conformations within a Δ threshold of the GMEC. It also becomes possible to directly enumerate unique near-optimal sequences (right).

Consistency is enforced on an open node, domains may get pruned and the updated lower bound c_{\emptyset} is used as a priority. Since upper bounding is always done as preprocessing, any developed node with a lower bound larger than the current upper bound k needs not be inserted in the queue, which saves space. This defines the A*-CFN algorithm. The systematic enforcement of Local Consistency at each node contrasts with DEE/A* where DEE is only used as a preprocessing pruning technique. Hence, A*-CFN has a permanent opportunity to reduce its search space and uses a stronger lower bound than the A* lower bound.^[44]

Variable and value ordering heuristics. Besides a search strategy, tree search methods also need to choose the next variable that will be used to branch. This is achieved using a variable ordering heuristics. Our algorithm supports dynamic variable ordering where a choice of variable is done at each node (and not just at each level of the tree). Variable ordering heuristics may have a tremendous effect on the efficiency of the search algorithm. They are all based on the “fail-first” principle^[45]: “To succeed, try first where you are most likely to fail.” Several measures have been used to try to evaluate the likelihood of failing by fixing a variable. One simple measure is the current size of the domain (dom).^[45] Under this measure, the variable which has the smallest domain should be assigned next. To take into account the number of cost functions that involve a variable (the so-called degree of the variable), more sophisticated heuristics select the variable that has the minimum ratio of the domain size over the current degree ($dom/ddeg$)^[46], over the degree weighted by the number of failures observed in the past for each cost function ($dom/wdeg$)^[21,47]

or by the sum of the median cost of cost functions involving the variable ($dom/cmed$)^[22,23]. Additionally, the last conflict heuristics^[48] simply try to select the last variable that led to inconsistency during search (if any).

Once a variable is chosen, Depth-First search and different branching scheme need to choose the next value to consider. The effect of value ordering heuristics is often less dramatic than for variables. However, a good value ordering heuristics may help to quickly find a good solution (upper bound). The most usual value ordering heuristics in CFNs is to choose first a value a that has a unary cost $c_i(a)=0$. Such a value always exists thanks to EDAC.^[44]

Branching schemes. We already described n -ary branching, where branching consists in choosing a variable x_i and all its available values to fix the variable value. Another branching scheme consists in selecting a variable x_i and a value/rotamer a and branch on the fact that the variable x_i either takes the value a , or not (and it must take one of its remaining values). This is called binary branching. By exploiting results in proof theory, binary branching has been previously shown to be more powerful than n -ary branching (with a given Local Consistency being maintained at each node, binary branching may explore an exponentially smaller number of nodes than n -ary branching, and the converse is impossible.^[49] n -ary branching can however be polynomially better on some problems). In Figure 4, a search tree of the same toy CPD problem is shown using binary branching.

An even more general branching method is dichotomic branching where the domain of a chosen variable i is split in two chosen sets. The split can be done in the middle of the domain range or in the middle of the unary energy range (similar to Ref. [50]), where low-energy rotamers are part of the first set and high-energy rotamers form the second set.

- A new SCP-based branching scheme

Compared to CFN, where all values are considered as uniform objects, the rotamers available at a designable position can be classified according to the chemical nature of the amino-acid

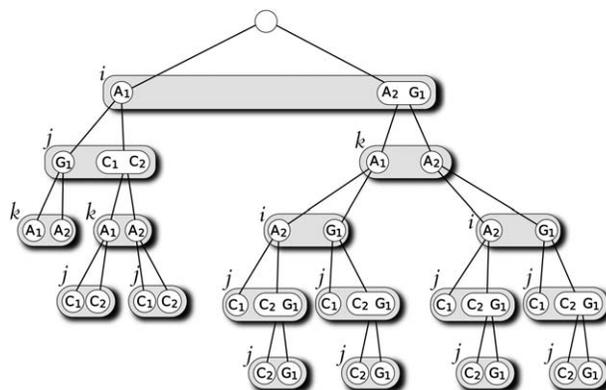


Figure 4. Binary branching scheme applied to the same naïve CPD problem as in Figure 1. To properly define the tree shape, ordering heuristics must be defined.

they define. In CPD, ultimately, it is the nature of the amino-acid that matters more than its precise conformation (as far as the optimal energy is known). Using the amino-acid nature of rotamers, it is possible to define a new branching scheme (Figure 5) where a variable is chosen first among variables which have more than one possible amino-acid type in their available rotamers. In this case, a dichotomic split will be done between all rotamers corresponding to a *chosen* amino-acid and all other rotamers. If all variables already have rotamers with a single amino-acid identity (the protein sequence is fixed in the current node), then a simple binary branching strategy is followed. The n -ary branching counterpart of this search mechanism is achieved by simply creating a new SCP-subproblem for each amino acid and branch onto them using an amino-acid selection heuristic. As for binary branching, once all the identities are fixed, the SCP problem is solved by n -ary branching.

We propose two heuristics to select the amino-acid type (AA) that will be used to split the domain in SCP branching. The idea behind these heuristics is to try to select an amino-acid that leads to a very good solution quickly, thus allowing to tighten the upper bound k . Inspired by the value ordering heuristics, the first heuristic selects the amino-acid associated to a rotamer v with a cost $c_i(v)=0$ (Zero-Cost – AA) which always exists because of EDAC enforcing. The second one selects the wild-type amino-acid if it is still available after Local Consistency enforcing (Wt – Zero-Cost – AA) and otherwise uses the Zero-Cost – AA choice. This branching incrementally reduces CPD to SCP: each time a branching is done, the domain of a variable is restricted to a single amino-acid in the first branch. Once there is no mutable position left in a given node, either a binary or n -ary branching scheme is applied to select rotamers. This branching strategy is denoted as SCP-based branching in what follows.

Interestingly, thanks to the separation that SCP-branching creates between a top tree that sets sequences and a large number of subtrees that solve SCP problems, it becomes possible to address the problem of the enumeration of all unique sequences within an energy threshold of the GMEC together with their optimal conformation and associated energy. As the ultimate goal of CPD is to provide a set of unique sequences to experimentalists, the restriction of the enumeration to a single conformation per sequence makes sense^[22] and should allow to drastically reduce the running time because of the exponential reduction in the number of enumerated solutions. In GMEC search, the upper bound k is updated each time a new solution is found. For enumeration, the upper bound is never updated. Since SCP-based branching defines the amino acid identity at all positions prior to conformation search, it becomes possible to update the upper bound locally only (in the current SCP subproblem) in order to just identify the best conformation for the current sequence, leading to additional pruning capabilities during enumeration. This defines two new search strategies: within a DFBB search strategy, the upper bound is locally updated once there is no mutable residue left in a given SCP branch. Similarly, for A^* -CFN with SCP-branching scheme, each time there is no mutable residue left, we simply use a DFBB search to get the best solution for the current sequence before adding the node to the list of open nodes (with its optimal energy).

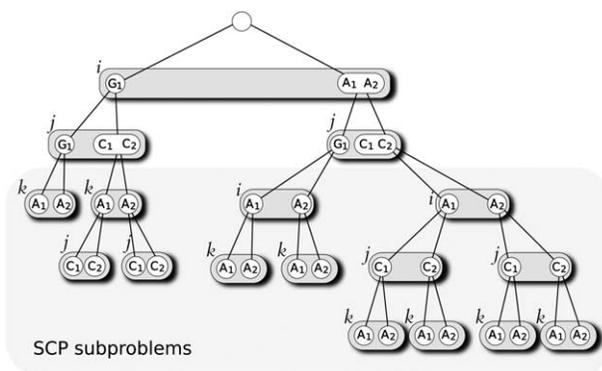


Figure 5. SCP-branching scheme illustrated on our toy example. Rotamer labels indicate the nature of the amino acid and the number of the rotamer. Designable position with more than one amino acid (i, j) are selected first. For each such choice, two nodes are added to the search tree: the node on the left restricts available rotamers to a preferred amino acid and the node on the right corresponds to the remaining rotamers. Once all designable positions have their amino-acid type set, search reduces to an SCP problem.

This hybrid A^* /DFBB strategy enables A^* -CFN to automatically enumerate unique sequences in an energy sorted order (the SCP part being explored using depth first search).

Benchmark Set for Method Performance Assessment

The benchmark set of 30 CPD instances (including structural models and associated precomputed energy matrices using *Osprey*) previously prepared^[22] (Supporting Information, Table S1) was used to assess the performances of the CFN-based methods in terms of runtime to provably identify the GMEC or enumerate suboptimal solutions for each protein design case. This benchmark set includes a variety of protein structures, alone or in complex with a protein or a ligand, derived from high-resolution structures deposited in the Protein Data Bank (PDB).^[51] These benchmarks vary by the number of designable residues (mutable or flexible) ranging from 23 to 97, leading to a wide range of sequence-conformation search spaces going from 4.36×10^{26} to 2.17×10^{94} . The penultimate rotamer library was used.^[52] All pairwise energy terms were precomputed and stored using *Osprey* 2.0.^[27] The used energy functions consist in the sum of the Amber electrostatic terms (with a distance-dependent dielectric constant), van der Waals, and the EEF1 implicit solvation energy term.^[53] No cutoff was used for nonbonded interactions. A more detailed description of the benchmark set can be found in Ref. [22].

All computations were performed on one core of an AMD Opteron™ Processor 6176@2.3 GHz. We used 128 GB of RAM and a 9,000 s timeout. CPU times reported hereafter correspond to runtime from the preprocessing to the identification of the optimal or suboptimal solutions. Problems not solved within the time limit were considered as using 9,000 s (a lower bound on the real time they would require). Unless specified otherwise, the EDAC Arc Consistency with incremental DEE (DEE¹ option) pruning^[21–23] was activated in all CFN-based experiments (both during preprocessing and during search), the median cost variable ordering heuristic

Table 1. Solving the GMEC identification problem using A*-vanilla and A*-CFN.

PDB accession code	Sequence-conformation space size	A*-vanilla			A*-CFN		
		nd	Time (s)	nd/min	nd	Time(s)	nd/min
1MJC	4.36E+26	0	5.1	-	0	4.8	-
1CSP	5.02E+30	127	94.6	80	10	38.7	15
1BK2	1.18E+32	5,234	285.3	1,100	21	25.6	49
1SHG	2.13E+32	558	77.5	432	28	76.8	21
1CSK	4.09E+32	113	7.8	869	16	3.7	259
1SHF	1.05E+34	21	13.9	90	8	14.5	33
1FYN	5.04E+36	58,380	3,355.0	1,044	81	160.7	30
1PIN	5.32E+39	24	1,480.4	0	11	1,515.8	0
1NXB	2.61E+41	45	3.8	710	9	2.0	270
1TEN	6.17E+43	243	37.1	392	9	37.9	14
1POH	8.02E+43	169	29.8	340	14	29.1	28
2DRI	1.16E+47	-	-	-	-	-	-
1FNA	3.02E+47	109,523	638.6	10,290	298	571.8	31
1UBI	2.43E+49	193,077	6,895.0	1,680	829	1,913.0	26
1C9O	3.77E+49	63,974	960.6	3,995	20	245.0	4
1CTF	3.95E+51	22,658	2,363.0	575	417	1,949.0	12
2PCY	2.34E+52	14,929	299.1	2,994	119	50.4	141
1DKT	3.94E+58	202,611	1,078.6	11,270	111	855.9	7
2TRX	9.02E+59	29,845	123.5	14,499	245	63.0	233
1CM1	3.73E+63	-	-	-	-	-	-
1BRS	1.67E+64	-	-	-	-	-	-
1CDL	5.68E+65	-	-	-	-	-	-
1LZ1	1.04E+72	-	-	-	4,899	1,481.2	198
1GVP	1.51E+78	-	-	-	-	-	-
1RIS	1.23E+80	-	-	-	-	-	-
2RN2	3.68E+80	15,845	1,276.7	744	104	1,146.3	5
1CSE	8.35E+82	26	79.0	19	8	187.6	2
1HNG	3.70E+88	16	1,739.0	0	4	1,901.5	0
3CHY	2.36E+92	-	-	-	-	-	-
1L63	2.17E+94	19,567	1,409.9	832	74	1,358.7	3
Number of cases solved in 900 s			14			16	
Number of cases solved in 9,000 s			22			23	
For each method are reported the number of nodes expanded (nd), the runtime in second (time), and the speed (nd/min). Dash line corresponds to cases where no GMEC solution was found in the timeout.							

(dom/cmed) was used in all the experiments (with the "last conflict" heuristics^[48]) as well as the zero unary cost value ordering heuristic and initial upper bounding by Limited Discrepancy Search.^[40] The statistic metrics reported hereafter ignore the few instances for which the difference in the time needed to solve the targeted problem was essentially similar (difference of <1 s).

Results and Discussion

Effects of the CFN lower bound

Performances of the new A*-CFN method (using the CFN lower bound) were assessed against those of the A*-vanilla algorithm (usually used in CPD after a DEE preprocessing) for the provable GMEC identification problem. In order to ensure that the difference of performances between both methods is only due to the quality of the lower bound, a static variable ordering was used. The usual *n*-ary branching was used and no upper bounding performed to stick to usual DEE/A* conditions (favoring this combination). The criteria used for comparison are the runtime, the number of visited nodes, and the number of solved problems. To fairly compare the search methods, we

performed a full DEE preprocessing (*Osprey algOption*=3) and a Local Consistency preprocessing step on each instance and then applied either A*-vanilla or A*-CFN. Thence, both strategies start with an identical search space for each of the 30 CPD instances.

A*-vanilla and A*-CFN solved, respectively, 22 and 23 CPD instances before timeout (Table 1). A*-CFN outperforms A*-vanilla in terms of runtime for 15 out of 18 discriminating cases (with more than 1 s difference). If we assign an optimistic 9,000 s to unsolved cases, A*-CFN allowed to save 978.9 s in average. Except for 1MJC problem which was directly solved at the end of the preprocessing step, A*-CFN expanded fewer nodes than A*-vanilla in all cases solved by both methods. For example, A*-vanilla expanded 63,974 nodes to solve 1C9O case while A*-CFN expanded only 20 nodes (corresponding to 3,199-fold decrease in the number of expanded nodes). As expected, the inexpensive A*-vanilla lower bound led to a higher number of nodes expanded per minute than the A*-CFN lower bound, which is however not enough to compensate for the much higher number of nodes that A*-vanilla needs to explore for solving the CPD problem.

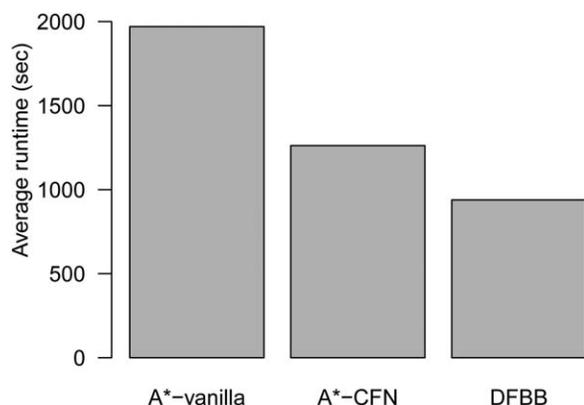


Figure 6. Comparison of search strategies for solving the GMEC identification problem: the average runtime (in second) for respectively A*-vanilla; A*-CFN, and DFBB. Cases unsolved by any method are excluded from the comparison. The average runtime is 1970.13, 1262.12, and 938.70 s for respectively A*-vanilla, A*-CFN, and DFBB.

We then compared the performances of the DFBB we implemented in *Osprey* against those of A*-CFN (Figure 6 and Supporting Information, Table S2), both using the same lower bounding mechanism. Since both DFBB and A*-CFN allow for it, default value and variable ordering heuristics were activated along with an initial upper bounding using Limited Discrepancy Search. *n*-ary branching was used.

Out of the 30 design cases, A*-CFN and DFBB managed to find 24 GMEC within the timeout. DFBB was faster than A*-CFN in 15 cases and saved in average 323.42 s. It was also faster than A*-vanilla in 16 cases, saving in average 1,031.43 s. Therefore, DFBB search with incremental Local Consistency enforcing algorithm and ordering heuristics remains overall more efficient than traditional A* search methods.

While A*, as a Best-First search method, is known to always explore less nodes than Depth-First search when the same tree is explored with the same lower bound, DFBB can quickly improve its upper bound *k* and therefore pruning, whereas A* cannot. This probably explains the performance gain of the Depth-First search strategy combined with the CFN lower bound. The polynomial space use of DFBB is also likely to help in the context of CPU processors using multiple levels of increasingly slow caches.

Branching schemes

Binary, *n*-ary, dichotomic, and SCP-based branching strategies were assessed against each other (Figure 7). To select the amino-acid in SCP-based branching, the Zero-Cost – AA amino-acid heuristic was used. Supporting Information, Table S3 reports the results for all 30 CPD instances.

As expected, binary branching was faster than *n*-ary branching (by an average of 84.35 s) and solved more CPD problems (11 over 19 discriminating problems). Compared to the (unary cost-based) dichotomic branching, it also showed a favorable average runtime gain (400.63 s). However, the number of problems solved faster does not increase (only 5 cases out of 18). In contrast, the SCP-based branching scheme is clearly better than

all other branching schemes. Out of 20 cases and all other contenders, a minimum of 18/20 cases were solved more quickly by SCP-based branching with a minimum average runtime gain of 328.25 s per instance. Hence, these results show that SCP-based branching outperforms all the other branching schemes, both in terms of runtime and number of problems solved faster. The poor performance of dichotomic branching is somehow unexpected since it splits apart rotamers at a given position into two groups depending on their energies.

Enumeration of near-optimal sequence-conformation ensemble

The problem of enumerating suboptimal sequence-conformation models within a given energy range above the GMEC (arbitrarily set here to 2 kcal mol⁻¹) was considered for all 30 CPD instances. First, we applied the A*-CFN and DFBB schemes using the binary branching scheme and with the SCP-based branching activated (default ordering heuristics and upper bounding activated in all cases).

For the A*-based strategies, the average runtime is 5959.4 and 4590.8 s for, respectively, A* and A*-SCP (Figure 8). For the DFBB-based methods, the average runtime is 4432.3 and 3958.4 s for, respectively DFBB and DFBB-SCP. As expected, SCP brings significant improvement to both A*-CFN and DFBB search strategies also in the context of enumeration. It is important to remind that A*-based algorithms preserve their natural advantage of being able to produce an energy-sorted stream of sequence-conformations. Instead, DFBB-based algorithms produce an unsorted list.

Enumeration of near-optimal unique sequences

Finally, the performances of the CFN-based algorithms were also compared to enumerate ensembles of suboptimal unique sequences using SCP-branching to separate sequence and conformation search. A comparison of runtime performances of all methods is shown in Figure 8 and more detailed results are included in Supporting Information, Tables S4–S6. Ensembles of unique sequences were generated using both A*-SCP

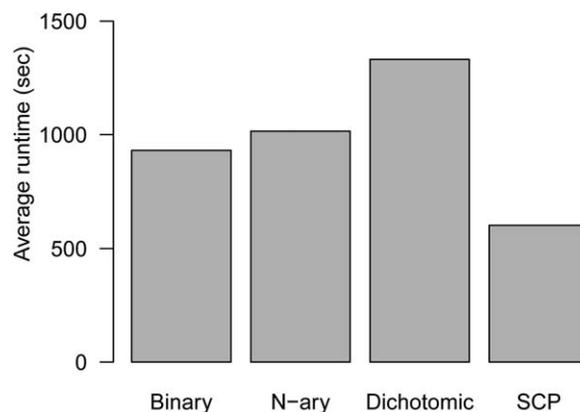


Figure 7. Comparison of branching schemes for solving the GMEC identification problem. The average runtime is 930.83, 1015.18, 1331.46, and 602.58 s for respectively Binary, N-ary, Dichotomic, and SCP branching schemes.

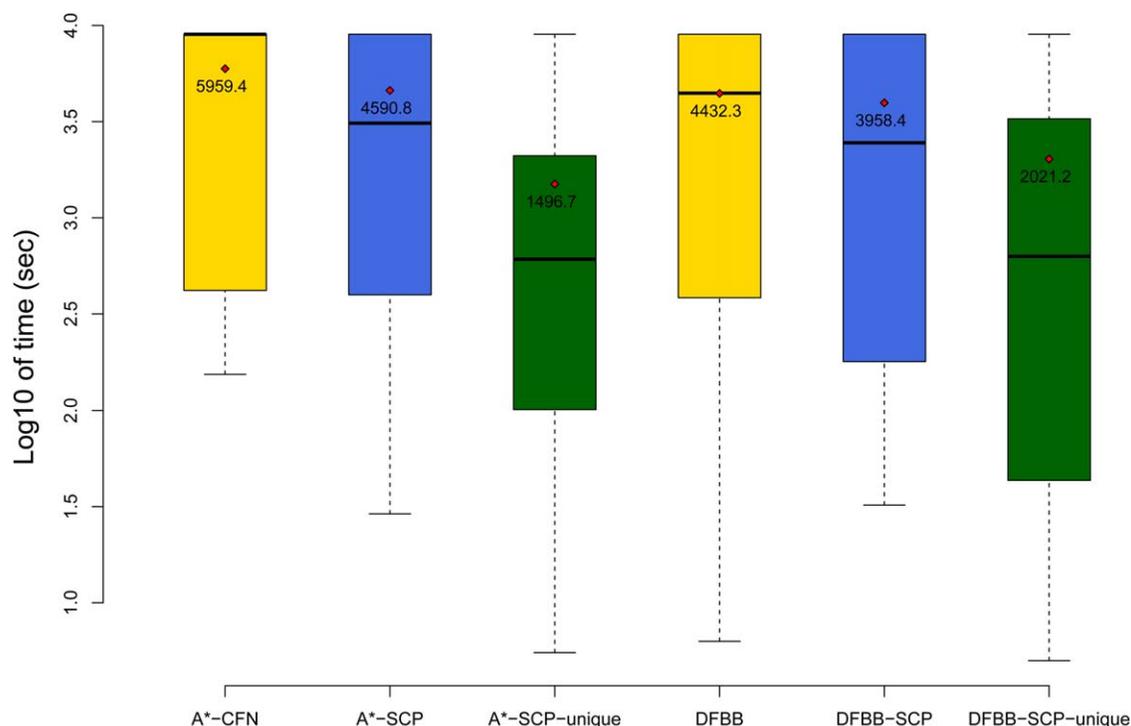


Figure 8. Boxplot of the runtime (in second) for respectively A*-CFN; A*-SCP; A*-SCP-unique; DFBB; DFBB-SCP; and DFBB-SCP-unique. The bold horizontal lines indicate the median (as usual for boxplots). The points in red indicate the average runtimes. The SCP branching schemes significantly improve both A* and DFBB search strategies. This improvement is even more drastic when the unique sequence constraint is activated.

and DFBB-SCP methods. More cases were solved and a significant gain in runtime was always obtained when unique sequence SCP is activated, leading to 1496.7 s average runtime for A*-SCP-unique and 2021.2 s for DFBB-SCP-unique (Supporting Information, Table S6). Consequently, SCP-based branching schemes enable enumerating unique sequences over a much larger energy window above the GMEC. Indeed, taking as an illustration the 1MJC instance and an interval of 2 kcal mol⁻¹, the A*-CFN required 232.6 s to enumerate 2,110,737 sequence-conformations (4,408,541 nodes). DFBB, A*-SCP and DFBB-SCP required, respectively, 140.7 s (4,537,935 nodes), 143.7 s (4,415,342 nodes), and 178.7 s (4,555,451 nodes) to handle the same instance. However, this ensemble corresponds to only 91 unique sequences. For a larger energy interval of 20 kcal mol⁻¹, using DFBB-SCP method restricted to the enumeration of unique sequences, 1,034 sequences were generated in 295.85 s (with a number of nodes lower than 50,489). Because of the sheer number of conformations per sequence, this energy interval is out of reach when all conformations have to be enumerated. SCP-branching is thus an effective solution to handle CPD problems and gives access to the enumeration of unique sequences within much larger energy windows than what can be achieved using standard branching schemes.

Conclusion

This work reports the introduction of Cost Function Network optimization techniques to solve Computational Protein Design

problems. Compared to prior work,^[21–23] it introduces novel search strategies in order to speed up search methods, which have been implemented in the CPD-dedicated software *Osprey*. The presented search algorithms have shown to be more efficient than optimization methods based on the DEE/A* framework for both GMEC search and suboptimal solutions enumeration (which is often performed in order to account for inaccuracies and approximations made in the modeling of the design problem and what is of interest to guide construction of experimentally manageable sizes of mutant libraries). In addition, the new SCP-based branching scheme makes it possible to generate unique sequences within a larger energy threshold than what was possible with other methods up to now. The implementation in *Osprey* of existing CFN algorithm together with these new methods enables novel opportunities. Indeed, these algorithms can be used in combination with all the other functionalities in this software, in particular for molecular flexibility modeling. Any macromolecular flexibility method that can be represented as an optimization problem with a single matrix can be performed with the presented CFN-based methods, even if they intrinsically rely on an energy-sorted stream of solutions.

Acknowledgments

We thank the Computing Center of Region Midi-Pyrénées (CALMIP, Toulouse, France) and the GenoToul Bioinformatics Platform of INRA-Toulouse for providing computing resources and support. We thank Pablo Gainza for discussion on search algorithms in *Osprey*.

Keywords: computational protein design · computer-aided protein design · exact combinatorial optimization · cost function networks · deterministic search methods · search heuristics · global minimum energy conformation · near-optimal solutions

How to cite this article: S. Traoré, K. E. Roberts, D. Allouche, B. R. Donald, I. André, T. Schiex, S. Barbe. *J. Comput. Chem.* **2016**, *37*, 1048–1058. DOI: 10.1002/jcc.24290

 Additional Supporting Information may be found in the online version of this article.

- [1] C. Y. Chen, I. Georgiev, A. C. Anderson, B. R. Donald, *Proc. Natl. Acad. Sci. USA* **2009**, *106*, 3764.
- [2] K. M. Frey, I. Georgiev, B. R. Donald, A. C. Anderson, *Proc. Natl. Acad. Sci. USA* **2010**, *107*, 13707.
- [3] N. P. King, J. B. Bale, W. Sheffler, D. E. McNamara, S. Gonen, T. Gonen, T. O. Yeates, D. Baker, *Nature* **2014**, *510*, 103.
- [4] A. E. Miklos, C. Kluwe, B. S. Der, S. Pai, A. Sircar, R. A. Hughes, M. Berrondo, J. Xu, V. Codrea, P. E. Buckley, A. M. Calm, H. S. Welsh, C. R. Warner, M. A. Zacharko, J. P. Carney, J. J. Gray, G. Georgiou, B. Kuhlman, A. D. Ellington, *Chem. Biol.* **2012**, *19*, 449.
- [5] J. Karanicolas, B. Kuhlman, *Curr. Opin. Struct. Biol.* **2009**, *19*, 458.
- [6] J. B. Siegel, A. Zanghellini, H. M. Lovick, G. Kiss, A. R. Lambert, J. L. St Clair, J. L. Gallaher, D. Hilvert, M. H. Gelb, B. L. Stoddard, K. N. Houk, F. E. Michael, D. Baker, *Science* **2010**, *329*, 309.
- [7] J. M. Palomo, M. Filice, *Biotechnol. Adv.* **2015**, *33*, 605.
- [8] N. A. Pierce, E. Winfree, *Protein Eng.* **2002**, *15*, 779.
- [9] C. A. Voigt, D. B. Gordon, S. L. Mayo, *J. Mol. Biol.* **2000**, *299*, 789.
- [10] B. Kuhlman, D. Baker, *Proc. Natl. Acad. Sci. USA* **2000**, *97*, 10383.
- [11] L. Wernisch, S. Hery, S. J. Wodak, *J. Mol. Biol.* **2000**, *301*, 713.
- [12] J. Desmet, M. De Maeyer, B. Hazes, I. Lasters, *Nature* **1992**, *356*, 539.
- [13] D. B. Gordon, S. L. Mayo, *Structure* **1999**, *7*, 1089.
- [14] E. Althaus, O. Kohlbacher, H. P. Lenhof, P. Müller, *J. Comput. Biol.* **2002**, *9*, 597.
- [15] A. Leaver-Fay, B. Kuhlman, J. Snoeyink, *Pac. Symp. Biocomput.* **2005**, *10*, 16.
- [16] C. L. Kingsford, B. Chazelle, M. Singh, *Bioinformatics* **2005**, *21*, 1028.
- [17] O. Alvizo, S. L. Mayo, *Proc. Natl. Acad. Sci. USA* **2008**, *105*, 12242.
- [18] R. F. Goldstein, *Biophys. J.* **1994**, *66*, 1335.
- [19] N. A. Pierce, J. A. Spriet, J. Desmet, S. L. Mayo, *J. Comput. Chem.* **2000**, *21*, 999.
- [20] D. B. Gordon, G. K. Hom, S. L. Mayo, N. Pierce, *J. Comput. Chem.* **2003**, *24*, 232.
- [21] D. Allouche, S. Traoré, I. André, S. de Givry, G. Katsirelos, S. Barbe, T. Schiex, In Proceedings of the 18th International Conference on Principles and Practice of Constraint Programming, CP 2012, Quebec, Canada, **2012**, 840.
- [22] S. Traoré, D. Allouche, I. André, S. De Givry, G. Katsirelos, T. Schiex, S. Barbe, *Bioinformatics* **2013**, *29*, 2129.
- [23] D. Allouche, I. André, S. Barbe, J. Davies, S. de Givry, G. Katsirelos, B. O'Sullivan, S. Prestwich, T. Schiex, S. Traoré, *Artif. Intell.* **2014**, *212*, 59.
- [24] E. L. Jackson, N. Ollikainen, A. W. Covert, T. Kortemme, C. O. Wilke, *PeerJ* **2013**, *1*, e211.
- [25] P. Gainza, K. E. Roberts, B. R. Donald, *PLoS Comput. Biol.* **2012**, *8*, e1002335.
- [26] B. R. Donald, Cambridge, MA: MIT Press, 2011.
- [27] P. Gainza, K. E. Roberts, I. Georgiev, R. H. Lilien, D. Keedy, C. Y. Chen, F. Reza, A. C. Anderson, D. C. Richardson, J. S. Richardson, B. R. Donald, *Methods Enzymol.* **2013**, *523*, 87.
- [28] R. H. Lilien, B. W. Stevens, A. C. Anderson, B. R. Donald, *J. Comput. Biol.* **2005**, *12*, 740.
- [29] B. W. Stevens, R. H. Lilien, I. Georgiev, B. R. Donald, A. C. Anderson, *Biochemistry* **2006**, *45*, 15495.
- [30] M. J. Grczynski, J. Grembecka, Y. Zhou, Y. Kong, L. Roudaia, M. G. Douvas, M. Newman, I. Bielnicka, G. Baber, T. Corpora, J. Shi, M. Sridharan, R. Lilien, B. R. Donald, N. A. Speck, M. L. Brown, J. H. Bushweller, *Chem. Biol.* **2007**, *14*, 1186.
- [31] K. E. Roberts, B. R. Donald, P. Boisguerin, P. R. Cushing, D. R. Madden, P. Boisguerin, D. R. Madden, B. R. Donald, *PLoS Comput. Biol.* **2012**, *8*, e1002477.
- [32] I. Georgiev, P. Acharya, S. Schmidt, Y. Li, D. Wycuff, G. Ofek, N. Doria-Rose, T. Luongo, Y. Yang, T. Zhou, B. Donald, J. Mascola, P. Kwong, *Retrovirology* **2012**, *9*, 50.
- [33] T. Schiex, H. Fargier, G. Verfaillie, *Int. Jt. Conf. Artif. Intell.* **1995**, *14*, 631.
- [34] I. Georgiev, R. H. Lilien, B. R. Donald, *J. Comput. Chem.* **2008**, *29*, 1527.
- [35] M. Cooper, T. Schiex, *Artif. Intell.* **2004**, *154*, 199.
- [36] J. Larrosa, T. Schiex, *Artif. Intell.* **2004**, *159*, 1.
- [37] M. Cooper, S. de Givry, T. Schiex, In Proceedings of the 8th International on Principles and Practice of Constraint Programming CP 2006 Workshop, on Preferences and Soft Constraints, Nantes, France, **2006**, p. 14.
- [38] C. Lecoutre, O. Roussel, D. E. Dehani, In Proceedings of the 18th International Conference on Principles and Practice of Constraint Programming, CP 2012, Quebec, Canada, **2012**, *7514*, 406.
- [39] S. de Givry, S. Prestwich, B. O'Sullivan, In Proceedings of International Conference on Principles and Practice of Constraint Programming, CP 2013, Uppsala, Sweden, **2013**, 263.
- [40] W. D. Harvey, M. L. Ginsberg, In Proceedings of the 14th International Joint Conference on Artificial Intelligence IJCAI-95, Montreal, Canada, **1995**, *1*, 607.
- [41] A. R. Leach, A. P. Lemon, *Proteins* **1998**, *33*, 227.
- [42] M. A. Hallen, D. A. Keedy, B. R. Donald, *Proteins* **2013**, *81*, 18.
- [43] R. Wallace, In: M. Meyer (Ed.), Selected papers from the ECAI-94 Workshop on Constraint Processing, Lecture Notes in Computer Science, Springer, Berlin, **1995**, *923*, 121.
- [44] S. D. Givry, M. Zytznicki, F. Heras, J. Larrosa, *IJCAI* **2005**, *19*, 84.
- [45] R. M. Haralick, G. L. Elliot, *Artif. Intell.* **1980**, *14*, 263.
- [46] C. Bessière, J. C. Régin, In Proceedings of the 2nd International Conference on Principles and Practice of Constraint Programming, CP 1996, Heidelberg, Germany, **1996**, *1118*, 61.
- [47] C. Lecoutre, L. Saïs, S. Tabary, V. Vidal, In proceedings of the 17th European Conference on Artificial Intelligence, ECAI'06, Trento, Italy, **2006**, 133.
- [48] C. Lecoutre, L. Saïs, S. Tabary, V. Vidal, *Artif. Intell.* **2009**, *173*, 1592.
- [49] D. Mitchell, In Proceedings of the 9th International Conference on Principles and Practice of Constraint Programming, CP 2003, Kinsale, Ireland, **2003**, *2833*, 363.
- [50] E. Hong, S. M. Lippow, B. Tidor, T. Lozano-pérez, *J. Comput. Chem.* **2009**, *30*, 1923.
- [51] F. C. Bernstein, T. F. Koetzle, G. J. Williams, E. F. Meyer, M. D. Brice, J. R. Rodgers, O. Kennard, T. Shimanouchi, M. Tasumi, *Eur. J. Biochem.* **1977**, *80*, 319.
- [52] S. C. Lovell, J. M. Word, J. S. Richardson, D. C. Richardson, *Proteins* **2000**, *40*, 389.
- [53] T. Lazaridis, M. Karplus, *Proteins* **1999**, *35*, 133.

Received: 14 April 2015
Revised: 23 September 2015
Accepted: 27 November 2015
Published online on 2 February 2016