A Variable Neighbourhood Descent Heuristic for Conformational Search Using a Quantum Annealer

D. J. J. Marchand*, M. Noori*, A. Roberts*, G. Rosenberg*, B. Woods*, U. Yildiz*, M. Coons**, D. Devore**, and P. Margl**

*1QB Information Technologies (1QBit)
**The Dow Chemical Company
Names listed in alphabetical order

November 20, 2018

Abstract

Discovering the low-energy conformations of a molecule is of great interest to computational chemists, with applications in *in silico* materials design and drug discovery. In this paper, we propose a variable neighbourhood search heuristic for the conformational search problem. Using the structure of a molecule, neighbourhoods are chosen to allow for efficient optimization using a binary quadratic optimizer. The method is flexible with respect to the choice of molecular force field and the number of discretization levels in the search space, and can be further generalized to take advantage of higher-order binary polynomial optimizers. It is well-suited for the use of devices such as quantum annealers. After carefully defining neighbourhoods, the method easily adapts to the size and topology of these devices, allowing for seamless scaling alongside their future improvements.

1 Introduction

The study of molecular structure is foundational to attaining an understanding of chemical processes. Chemical behaviour is determined in large part by the arrangement of atoms within participating molecules as a chemical process unfolds. A widely used approach for studying aspects of molecular structure of particular interest is to limit is to limit the molecular degrees of freedom to torsions only, considering bond lengths and bond angles to be fixed at some values. By doing so, a given molecule with a specific connectivity of its constituent atoms may take on a variety of three-dimensional spatial arrangements, known as *conformational isomers*, or simply *conformations*. Whereas the internal bonds of distinct conformations are the same, the rotation angles around rotatable bonds of the molecule are what distinguish one conformation from another.

In the context of medicinal chemistry, conformational analysis often involves identifying bioactive conformations of ligand molecules, where it is important to identify conformations that are complementary to a macromolecular binding site [1, 2]. The folding of proteins is one illustration of function determined by conformation that is both very important to the fields of medicine and biochemistry and extremely difficult to compute, making it a "grand challenge" for science [3]. In a broader context, conformational analysis is a topic of utmost importance in the field of chemical and materials research. For instance, the conformational behaviour of polymers is a key factor in determining crystallinity, shape, and entanglement of individual chains [4] that in turn affect macroscopic materials' properties such as elasticity, strength, or toughness.

Importantly, the geometrical differences between conformations result in different values for the molecular potential energy [5–7], which is a key factor for molecular stability and reactivity. To this end, the problem of finding the conformations associated with the local minima of the potential energy surface (PES) of the molecule, referred to as the *conformational search problem*, has been of long-standing interest in both

academic and industrial research. The local minima of the PES are often called *conformers* [8], and the conformer with the lowest energy is referred to as the *global minimum-energy conformer*.

Over time, several practical approaches to the conformational search problem have been developed and critically reviewed [8–11]. For small molecules, one may be able to deterministically solve the conformational search problem via techniques such as branch and bound [12–14]. However, this is impractical for even moderately sized molecules, as the size of the conformational search space grows exponentially with the number of rotatable bonds. More specifically, if we consider M rotatable bonds, where each is allowed to take on d possible values, there are d^M possible conformations. Aside from the size of the exponentially growing search space, the number of local minima on the molecular PES also has a tendency to grow exponentially with the number of atoms in the molecule [15]. Such increasing difficulty in the size and shape of the PES makes the conformational search problem computationally intractable for many of the molecules that have real-world applications and an attractive target for exploration with novel computational technologies and techniques.

To address the computational complexity presented by larger molecules, many metaheuristic approaches have been studied. Examples of such approaches include genetic algorithms [16–19], conformational space annealing [20, 21], tabu search [22, 23], molecular dynamics (MD) [24, 25], and basin/funnel hopping [26, 27], amongst others. Variations of the Monte Carlo (MC) method have also been widely used [28–30] as a less computationally expensive alternative to MD [8, 31]. In addition, parallel tempering (PT), also known as replica exchange, [32–34] can be applied to both MC and MD to further improve their sampling performance of the conformational search space.

In this paper, we propose an iterative heuristic method for the conformational search problem based on *variable neighbourhood descent* (VND). In each iteration of the method, we use the molecular structure to choose specialized conformational neighbourhoods that can be minimized efficiently. More specifically, using the structural graph of a given molecule, subsets of rotatable bonds are selected at each iteration. Fixing the values of other torsion angles, the problem of minimizing the molecular energy with respect to the selected torsion angles is then formulated as a binary program with an objective function that is a polynomial of a chosen degree. This allows the method to be adapted to the specifics of the optimizer by limiting the degree of the binary program. The values of the selected torsion angles are then set to the solution of this binary program before starting each subsequent iteration during which a new subset of rotatable bonds are optimized. The process continues until some stopping criteria are met.

Although the method can be readily extended to any chosen degree, we assume in this paper that a binary quadratic program is desirable as it is well-suited for optimization using quantum annealing [35, 36]. Furthermore, by changing the parameters of the neighbourhood selection procedure, the method can be easily adapted to the size of the conformational search problem in terms of the number of rotatable bonds, as well as the size and connectivity of available quantum annealers. The flexibility of solving any conformational search problem using currently available quantum annealers, without imposing restrictions on the granularity of the conformational space, differentiates our work from a previous study on protein folding using quantum annealing [37].

The proposed method can also be used in conjunction with other conformational search techniques in devising hybrid approaches. A common strategy is to use one method to generate a series of initial low-energy conformations before applying the second method to improve their energies [9]. Our method can fit seamlessly into the second phase of this process.

We evaluate the performance of our proposed algorithm over two families of molecules: a set of transition metal molecules and a set of ortho-phenylene (oPh) oligomers that are relevant to industry. We study two methods for solving the quadratic binary programs: an algorithm that returns an exact optimal solution and the D-Wave 2000Q quantum annealer [38–40]. The latter provides an assessment using the latest available hardware at the time of writing of this work, whereas the former can be seen as a limiting ideal case. For each molecule, we compare the lowest-energy conformations found by our algorithm with those found by both parallel tempering MC (PTMC) [32, 41] and a simple local search method.

The paper is organized as follows. Section 2 provides the preliminary materials for the proposed conformational search method, as well as an overview of quantum annealing. In Section 3, the proposed

conformational search method is explained. Experiments and results for the proposed method are presented in Section 4. Finally, concluding thoughts and possible future research directions are presented in Section 5.

2 Preliminaries

The following preliminaries are needed before presenting the details of our proposed conformational search method. First, we formally define the conformational search problem we consider in this paper and describe its degrees of freedom, the representation of the molecular structure, and the molecular energy model. We also introduce the assumptions and notation used. We then provide a binary optimization formulation for the conformational search problem and explain the challenges of solving this optimization problem using a quantum annealer.

2.1 Problem Definition

As discussed above, we consider the conformational search problem as a special case of molecular structure analysis, where the structure is kept fixed except for rotations around selected bonds. Each of these torsional degrees of freedoms we hereafter refer to as a *torsion* for simplicity. We denote the i-th rotatable bond by T_i and assign its rotational angle a variable t_i, with i being the torsion index. It is convenient to identify a conformation of a molecule with M torsions by a *torsion vector* $\mathbf{t} = [t_1, \ldots, t_M]$. Without loss of generality, we assume $t_i \in [0, 2\pi)$, for all i, knowing that the method remains unchanged if each torsion has its own range chosen based on prior knowledge, experimental data, or known symmetries. For simplicity, let us assume all torsion angle values are chosen from the same set of d values $\Theta = \{\theta_1, \ldots, \theta_d\}$. The theoretical precision of this discretization scheme increases with d, while the size of the search space d^M grows exponentially with the number of torsions.

Although it is natural to describe a molecule using a molecular graph, where the atoms and their bonds are vertices and edges, respectively, we find it helpful to use the torsions to partition the molecule into $M + 1$ subsets called *rigid bodies*. The partitioning is performed such that all atoms within a rigid body are interconnected through non-torsion bonds. As a result, the relative positions of the atoms within a rigid body, denoted by R_a, remain invariant under rotation and are therefore independent of $\mathbf{t}$. This simplified representation of the molecule is now easily described by a *rigid-body graph* $G = (\mathcal{R}, \mathcal{T})$, where $\mathcal{R}$ is the set of $M+1$ vertices and $\mathcal{T}$ is the set of M edges. In G, each vertex represents a rigid body and each edge represents a rotatable bond. Two vertices are connected by an edge if their associated rigid bodies are connected by the rotatable bond that the edge represents. We will therefore use T_i to refer to both torsion i and its associated edge in the rigid-body graph. We further assume that each torsion is free to rotate independently of others, thus restricting the presence of ring systems or other cycles in the molecular graph to individual rigid bodies. Under this assumption, the rigid-body graph has no cycles and is a tree. An example of a simple molecule and its rigid-body graph is shown in Figure 1.

The search space of the conformational search problem is a hypersurface described by an energy model (function) $U : \Theta^M \to \mathbb{R}$. For a given $\mathbf{t}$, $U(\mathbf{t})$ is the molecular energy consisting of the sum of all interatomic potentials in the molecule (e.g., van der Waals, torsional, and bending), which are dependent on the relative coordinates of the atoms. Various energy models or effective *force fields* can be used for our purpose, such as the widely used "Universal force field" (UFF) [42]. The conformational search problem, with the objective of finding the global minimum-energy conformer of a given molecule, can then be formulated as

$$\begin{aligned} \min_{\mathbf{t}} \quad & U(\mathbf{t}) \\ \text{s.t.} \quad & \mathbf{t} \in \Theta^M. \end{aligned} \tag{1}$$

Upon changing the torsion angles of the molecule, some of the interatomic potential contributions will remain unchanged, while other contributions will change depending on the torsion angle vector $\mathbf{t}$. To be more specific, let us denote the values of the torsion angles on the path connecting R_a to R_b on the rigid-body graph by a vector $\mathbf{t}_{ab}$. The length of this path is represented by m_{ab}, meaning that $\mathbf{t}_{ab}$ has m_{ab} elements.

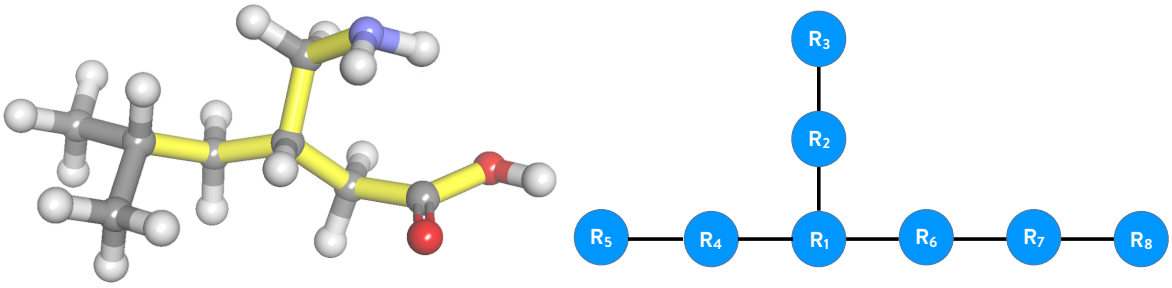

(a) (b)

Figure 1: Illustration of the rigid-body graph for a simple organic molecule (pregabalin). (a) Molecular structure, with seven rotatable bonds highlighted in yellow. The rotatable bonds connected to the two methyl groups have been discarded here, as they are less important to the system. (b) The associated rigid-body graph.

The relative positions of the atoms in R_a with respect to the atoms in R_b depend only on the torsions on this path. Now,

$$U(\mathbf{t}) = \sum_{a:R_a \in \mathcal{R}} U_a + \sum_{\substack{a,b:R_a,R_b \in \mathcal{R} \\ a \neq b}} U_{ab}(\mathbf{t}_{ab}) \,, \tag{2}$$

where U_a is the sum of the interatomic potentials of the atoms within rigid body R_a, which is invariant under rotation, while $U_{ab} : \Theta^{m_{ab}} \to \mathbb{R}$ such that for a given $\mathbf{t}_{ab}$, $U_{ab}(\mathbf{t}_{ab})$ is the sum of the interatomic potentials of all pairs of atoms where one atom is in R_a and the other is in R_b.

2.2 Binary Optimization Formulation for the Conformational Search Problem

In order to use a quantum annealer for solving the conformational search problem, one needs to reformulate the problem as a quadratic unconstrained binary optimization (QUBO) problem (see Appendix 7.1 for more details). To this end, we start by applying a one-hot encoding to the discrete values of the torsion angles, establishing a mapping between the torsion angle vector space and a binary solution space. That is, for each t_i, we assign a binary variable x_{ik}, $1 \leq k \leq d$, such that

$$x_{ik} = \begin{cases} 1 & \text{if } t_i = \theta_k \,; \\ 0 & \text{otherwise.} \end{cases} \tag{3}$$

As a result, t_i can be expressed as

$$t_i = \sum_{k=1}^{d} \theta_k x_{ik} \,, \tag{4}$$

where we add a constraint

$$\sum_{k=1}^{d} x_{ik} = 1 \tag{5}$$

to ensure t_i takes one and only one value at a time. The constraint (5) is commonly referred to as a one-hot encoding constraint. Note that after applying the one-hot encoding, any arbitrary function $f(t_i)$ can be written as

$$f(t_i) = \sum_{k=1}^{d} f(\theta_k) x_{ik} \,. \tag{6}$$

Similar to (6), a binary representation for $U(\mathbf{t}_{ab})$ can be found. For simplicity of presentation, let us assume that the torsion angles in $\mathbf{t}_{ab}$ are indexed sequentially from 1 to m, that is, $\mathbf{t}_{ab} = [t_1, t_2, \ldots, t_m]$. Now,

$$U_{ab}(\mathbf{t}_{ab}) = \sum_{k_1=1}^{d} \cdots \sum_{k_m=1}^{d} U_{ab}(t_1 = \theta_{k_1}, \ldots, t_m = \theta_{k_m}) \prod_{i=1}^{m} x_{i\,k_i}\,, \tag{7}$$

where $U_{ab}(t_1 = \theta_{k_1}, \ldots, t_m = \theta_{k_m})$ can be pre-evaluated for all possible d^m values of $\mathbf{t}_{ab}$.

Substituting $U_{ab}(\mathbf{t}_{ab})$ from (7) into (2) results in a representation of the molecular energy $U(\mathbf{t})$ in terms of the binary vector $\mathbf{x}$. We denote this representation of the energy function by $E : \{0,1\}^{Md} \to \mathbb{R}$ and write the conformational search problem (1) as

$$\begin{aligned} \min_{\mathbf{x}} \quad & E(\mathbf{x}) \\ \text{s.t.} \quad & \sum_{k=1}^{d} x_{ik} = 1, \quad \forall i \in \{1, \ldots, M\}\,, \\ & \mathbf{x} \in \{0,1\}^{Md}\,. \end{aligned} \tag{8}$$

To solve the above binary optimization problem using a quantum annealer, one faces three challenges. First, the objective function in formulation (8) is not necessarily quadratic as $U_{ab}(\mathbf{t}_{ab})$ may depend on more than two torsions, that is, $m_{ab} > 2$. Second, it is a constrained binary optimization problem. These two challenges indicate that the problem cannot be solved directly on a quantum annealer. The third challenge is that if $\max_{a,b:R_a,R_b \in \mathcal{R}} m_{ab}$ is not much smaller than M, constructing an instance of (8) becomes very computationally expensive due to the pre-evaluation of the coefficients in (7). In the extreme case where there exists an m_{ab} such that $m_{ab} = M$, formulating (8) requires a computational effort equivalent to traversing the entire conformational space. In the following section, we propose a method that addresses these challenges in order to be able to use a quantum annealer for solving the conformational search problem.

3 Variable Neighbourhood Descent for the Conformational Search Problem

Neighbourhood search, or local search, is known to be an effective heuristic algorithm for solving a large number of combinatorial optimization problems. In defining a neighbourhood relation between solutions of a problem, local search begins from an initial solution and iteratively explores the neighbourhood of the current solution for improvement. It has been shown that a solution produced by a local search algorithm will often not be globally optimal, and will often be suboptimal with respect to another neighbourhood relation [43]. When multiple neighbourhood relations are considered, the algorithm is often referred to as *variable neighbourhood search* [44]. In the context of conformational search, a solution refers to a vector of torsion angles $\mathbf{t}$.

Let N_k, for $k \in \{1, \ldots, K\}$, denote a finite set of neighbourhood structures and $N_k(\mathbf{t})$ be the set of all solutions in the k-th neighbourhood of $\mathbf{t}$. Starting from an initial solution and a neighbourhood structure, in each iteration, *variable neighbourhood descent* finds the best solution in the neighbourhood of the current solution. It then updates the current solution with the best solution found, and changes the neighbourhood structure before proceeding with the next iteration. The VND method is summarized in Procedure 1.

Procedure 1 Variable Neighbourhood Descent

Input: Initial solution $\mathbf{t}_{\text{init}}$
  $\mathbf{t}_{\text{c}} \leftarrow \mathbf{t}_{\text{init}}$ ▷ Initialize the current solution
  Select an arbitrary $N_k \in \{N_1, \ldots, N_K\}$
  while termination criteria have not been met **do**
    $\mathbf{t}' \leftarrow \arg\min_{\mathbf{t} \in N_k(\mathbf{t}_{\text{c}})} U(\mathbf{t})$ ▷ Find the best neighbour in $N_k(\mathbf{t}_{\text{c}})$
    if $U(\mathbf{t}') < U(\mathbf{t}_{\text{c}})$ **then**
      $\mathbf{t}_{\text{c}} \leftarrow \mathbf{t}'$ ▷ Update the current solution
    end if
    $N_k \leftarrow$ NeighbourhoodChange() ▷ Change neighbourhoods
  end while

A simple neighbourhood structure is obtained by considering two solutions as neighbours if and only if they differ by exactly one torsion angle value. We refer to the VND approach using this neighbourhood definition as *local search* (LS). While computationally inexpensive, the performance of LS can suffer in cases where a decrease in the molecular energy cannot be achieved by changing only a single torsion angle value in an iteration. Our proposed VND method improves upon LS by exploring more-complex neighbourhoods. In the following, we describe the components of the method.

3.1 Initial Solution

The initial solution in Procedure 1 can be selected in a variety of ways. One may simply choose a randomly generated torsion angle vector for the given molecule as the initial solution. Alternatively, one can use a greedy construction method. Another approach is to start from a known high-quality solution. This applies when using our VND method in conjunction with another conformational search method or by exploiting some prior knowledge about the molecule.

3.2 Neighbourhood Structures

We now describe a more powerful neighbourhood structure which, to our knowledge, has not been previously studied. Let $G = (\mathcal{R}, \mathcal{T})$ be a rigid-body graph, $\mathcal{T}' \subseteq \mathcal{T}$, and G' be the graph resulting from contracting all edges in $\mathcal{T} \setminus \mathcal{T}'$ (see Figure 2 for an example). If G' is a star graph (a tree graph with at most one vertex of degree > 1), then we say that $\mathcal{T}'$ has the property of *2-torsion dependency*. The motivation for using this terminology is that any two vertices of G' are connected with at most two edges (torsions). We label the maximal 2-torsion-dependent subsets of $\mathcal{T}$ as $\mathcal{T}_1, \ldots, \mathcal{T}_K$ and their associated star graphs as $G_1, \ldots, G_K$. Then, the neighbourhood structure N_k defines neighbourhoods containing all solutions which differ only in torsion angle values corresponding to edges in $\mathcal{T}_k$, for $k = 1, \ldots, K$. Solutions in neighbourhoods defined by an arbitrary N_k are also called *neighbours under* N_k.

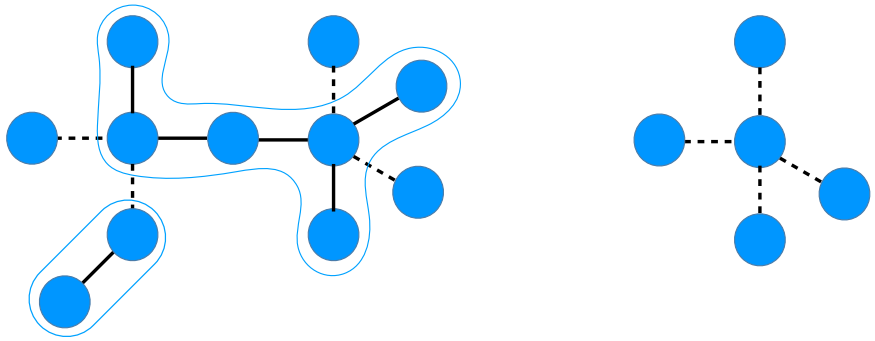

Figure 2: An example of a 2-torsion-dependent set of edges in a rigid-body graph (left). The star graph G' (right) results from selecting a 2-torsion-dependent set $\mathcal{T}'$ (dashed lines) and contracting all edges not in $\mathcal{T}'$.

Neighbourhood structures are illustrated for an example rigid-body graph $G = (\mathcal{R}, \mathcal{T})$ in Figure 3 that represents a molecule with three torsions. The maximal 2-torsion-dependent subsets of the torsions are $\mathcal{T}_1 = \{T_1, T_3\}$, $\mathcal{T}_2 = \{T_2, T_3\}$, and $\mathcal{T}_3 = \{T_1, T_2\}$. Neighbourhood structure N_1 defines neighbourhoods containing all torsion vectors that differ only in torsion angle values for T_1 and T_3. For example, $\mathbf{t} = [5°, 10°, 20°]$ and $\mathbf{t}' = [0°, 10°, 90°]$ are neighbours under N_1 while $\mathbf{t} = [5°, 10°, 20°]$ and $\mathbf{t}' = [5°, 15°, 20°]$ are not because they differ in the torsion angle value for T_2.

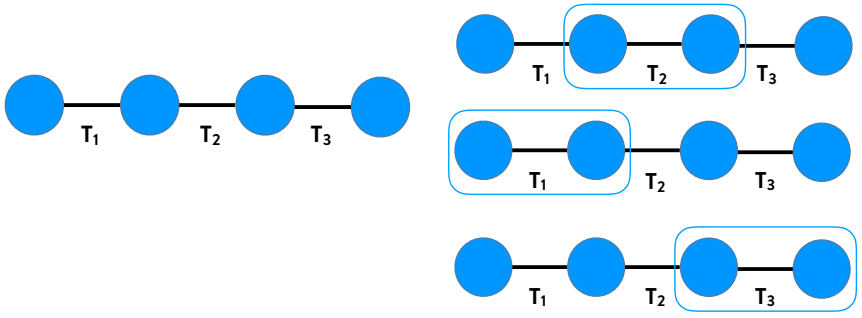

(a) (b)

Figure 3: An example of finding different neighbourhood structures of a rigid-body graph. The original rigid-body graph, $G = (\mathcal{R}, \mathcal{T})$, is shown in (a), and (b) depicts all neighbourhood structures defined by the maximal 2-torsion-dependent subsets of $\mathcal{T}$.

3.3 Neighbourhood Search

Based on the discussion in Section 2.2, the problem of finding the best solution in $N_k(\mathbf{t})$ can be formulated as a QUBO problem by restricting (8) to the binary variables corresponding to the torsions in $\mathcal{T}_k$ and moving the one-hot encoding constraints to the objective function using the quadratic penalty method. The resulting QUBO problem can be formulated as

$$\begin{aligned} \min \quad & \sum_{\substack{i,j:T_i,T_j\in\mathcal{T}_k \\ i\neq j}} \sum_{k_i=1}^{d}\sum_{k_j=1}^{d} U_{ij}(\theta_{k_i},\theta_{k_j}) x_{ik_i}x_{jk_j} + \sum_{i:T_i\in\mathcal{T}_k}\sum_{k_i=1}^{d} U_i(\theta_{k_i})x_{ik_i} + p\sum_{i:T_i\in\mathcal{T}_k}\left(\sum_{k_i=1}^{d} x_{i,k_i} - 1\right)^2, \\ \text{s.t.} \quad & x_{ik_i} \in \{0,1\}, \quad \forall i, k_i\,. \end{aligned} \tag{9}$$

In (9), the $U_{ij}(\theta_{k_i}, \theta_{k_j})$ terms represent the interaction energy of the two vertices of G_k that are connected by T_i and T_j when $t_i = \theta_{k_i}$ and $t_j = \theta_{k_j}$. The $U_i(\theta_{k_i})$ terms represent the interaction energy of the two vertices connected by T_i on G_k when $t_i = \theta_{k_i}$ and p is a sufficiently large penalty coefficient that enforces the one-hot encoding constraints. The above QUBO problem can be solved using various methods [45–54], as well as specialized hardware devices such as quantum annealers.

3.4 Neighbourhood Change

At each iteration, the neighbourhood is selected based on a random ordering, $\phi : \mathcal{T} \to \mathcal{T}$, of the torsions. The pseudocode for the neighbourhood change function is given below.

Procedure 2 Neighbourhood Change

Generate a random ordering, $\phi = (\phi(1), \phi(2), \ldots, \phi(M))$, of $\mathcal{T}$
$\mathcal{T}' \leftarrow \emptyset$
for $i = 1, \ldots, M$ **do**
  if $\mathcal{T}' \cup \{\phi(i)\}$ is 2-torsion dependent **then**
    $\mathcal{T}' \leftarrow \mathcal{T}' \cup \{\phi(i)\}$ ▷ Add $\phi(i)$ to $\mathcal{T}'$
  end if
end for
return N_k corresponding to $\mathcal{T}'$

3.5 Practical Considerations

Here, we detail the practical considerations of the proposed VND method to accommodate the use of existing and future quantum annealers.

The formulated QUBO problem in (9) is fully connected, meaning that for all i and j ($i \neq j$), the term $x_{ik_i}x_{jk_j}$ appears in the objective function. On the other hand, the connectivity of the qubits on the current D-Wave quantum annealer follows a "Chimera graph" structure [55]. Thus, the problem in (9) must be embedded onto the hardware graph using an embedding strategy [55]. There is a limit on the number of variables a a fully connected QUBO problem can have that can be embedded onto the graph. We take this limitation into consideration by imposing a limit on the number of variables in the formulated QUBO problem for each selected subset of torsion angles $\mathcal{T}_k$. We denote this parameter by s. In the following, we explain how this limit is imposed on the formulated QUBO problem.

For the selected neighbourhood structure at each iteration, we (randomly) select a total of s discrete values. That is, for each selected T_i in $\mathcal{T}_k$, we randomly choose a set of $\Theta_i \subseteq \Theta$ discrete values such that $s = \sum_{i:T_i \in \mathcal{T}_k} |\Theta_i|$.

We have already defined the $N_k(\mathbf{t})$ neighbourhood as the set of torsion angle vectors $\mathbf{t}'$ that are different from $\mathbf{t}$ only in the angle values of the torsions in $\mathcal{T}_k$. In addition, for any T_i in $\mathcal{T}_k$, t'_i (i.e., the value associated with T_i in $\mathbf{t}'$) takes on values only from Θ_i. With the above choice of neighbours, for an arbitrary $\mathbf{t}$, $N_k(\mathbf{t})$ contains

$$S_k = \prod_{i:T_i \in \mathcal{T}_k} |\Theta_i| \tag{10}$$

solutions. However, to find the best solution in $N_k(\mathbf{t})$, we need to pre-evaluate only

$$\sum_{\substack{i,j:T_i,T_j \in \mathcal{T}_k \\ i \neq j}} |\Theta_i||\Theta_j| + s \tag{11}$$

energy terms to formulate the QUBO problem. As seen above, S_k grows linearly with each $|\Theta_i|$, for all i, whereas the growth of (11) is quadratic. This means that the number of energy pre-evaluations grows more slowly than the size of the neighbourhood as $\mathcal{T}_k$ increases.

Another practical consideration for our proposed VND method is the stopping criteria. The first stopping criterion sets a limit on the computational effort of the method by introducing a maximum number of iterations, denoted by B. The second stopping criterion aims to terminate the method early if it becomes stuck at a local minimum or finds the global minimum of the problem. For this purpose, we introduce a parameter called A that represents the maximum number of consecutive iterations to can be performed without decreasing the energy.

With the above-mentioned practical considerations, the implemented VND method is summarized in Procedure 3.

Procedure 3 Conformational Search VND Method

Input: A random initial conformation of the molecule, denoted by a torsion vector $\mathbf{t}_{\text{init}}$

$a, b \leftarrow 0$ ▷ Initialize the counters
Select an arbitrary $N_k \in \{N_1, \ldots, N_K\}$ ▷ Initialize the neighbourhood structure
$\mathbf{t}_{\text{c}} \leftarrow \mathbf{t}_{\text{init}}$ ▷ Initialize the current solution $\mathbf{t}_{\text{c}}$
while $a < A$ and $b < B$ **do**
  Formulate the QUBO problem to find the best solution in $N_k(\mathbf{t}_{\text{c}})$
  $\mathbf{t}' \leftarrow \text{argmin}_{\mathbf{t} \in N_k(\mathbf{t}_{\text{c}})} U(\mathbf{t})$ ▷ Find the best solution in $N_k(\mathbf{t}_{\text{c}})$
  if $U(\mathbf{t}') < U(\mathbf{t}_{\text{c}})$ **then**
    $\mathbf{t}_{\text{c}} \leftarrow \mathbf{t}'$ ▷ Update the current solution
    $b \leftarrow 0$ ▷ Reset the no-improvement counter
  else
    $b \leftarrow b + 1$ ▷ Update the no-improvement counter
  end if
  $a \leftarrow a + 1$ ▷ Update the counter
  $N_k \leftarrow$ NeighbourhoodChange() ▷ Change neighbourhoods
end while

Two important features of the VND method is its scalability and ease of adaptation to the size and connectivity of the available quantum annealer. To be more specific, one can simply increase s in the described method to take advantage of improvements in the number of qubits and their connectivity. The rest of the method remains intact. It is worth mentioning that although the focus of this section has been on quantum annealers, other QUBO problem solvers could be used.

3.6 The Effect of the Molecular Structure

As previously discussed, due to the limitations of current quantum annealers, we are able to jointly optimize only a subset of torsions that are 2-torsion dependent at each iteration of VND. The potential speedup of the quantum annealer over a naïve exhaustive QUBO solver depends on the cardinality of the selected subset of torsions.

Since $\sum_{i:T_i \in \mathcal{T}_k} |\Theta_i|$ is fixed, the number of solutions in $N_k(\mathbf{t})$, S_k, in (10) is maximized when the cardinality of each Θ_i is the same. That is, when $|\Theta_i| = s/|\mathcal{T}_k|$ for all $i \in \{1, \ldots, |\mathcal{T}_k|\}$, resulting in the maximum number of solutions

$$S_k = \left(\frac{s}{|\mathcal{T}_k|}\right)^{|\mathcal{T}_k|}. \tag{12}$$

The maximum value of S_k in (12) is an increasing function of $\mathcal{T}_k$ if $s/|\mathcal{T}_k| \geq e$, where e is the base of the natural logarithm.

As the potential speedup of the quantum annealer over an exhaustive QUBO solver is dependent on S_k, it is favourable to select each $\mathcal{T}_k$ with a large cardinality. The speedup diminishes when $|\mathcal{T}_k| = 2$, which is the case for molecules with linear rigid-body graphs. The above discussion suggests that star-like molecules that have rigid-body graphs with high-degree nodes stand to benefit more than those that do not from using a quantum annealer to solve the QUBO problem at each VND iteration.

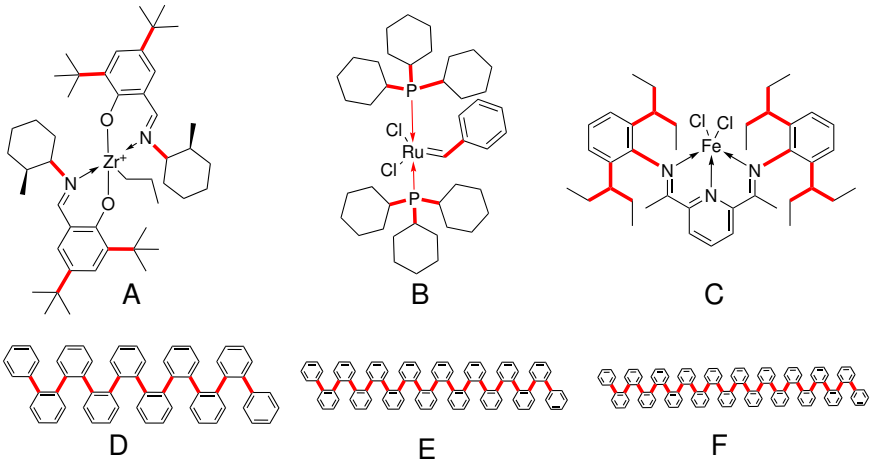

Figure 4: A graphical representation of the six model systems studied: three organometallic compounds (A–C) and three ortho-phenylenes (D–F). The thick (red) lines represent the torsion bonds. See Table 5 in the Appendix for additional details about each model system.

4 Experimental Results

In this section, we evaluate the performance of the proposed VND method. We first provide details about the molecules used in our experiments and then present the experimental results.

4.1 Trial Molecules

The performance of VND was evaluated using a testbed containing the six molecules depicted in Figure 4. These model systems include three organometallic molecules useful for catalyzing reactions relevant to industry (labelled "A" [56–58], "B" [59–61], and "C" [62–64]) and a set of three ortho-phenylene oligomers that are of interest as electronic materials and nanomaterials (labelled "D", "E", and "F") [65, 66].

The choice of model systems A–F was motivated by several considerations. First, it is important to show that our method can be applied to a wide variety of conformational search problems of relevance to industry. Second, the model systems are representative of a diversity of molecular graphs: systems A–C have star-like graphs, whereas D–F have linear graphs. Third, the model systems represent a significant variety of active torsions, ranging from very modest (e.g., A) to very substantial (e.g., F). Another motivation for choosing these systems was the existence of experimental data pertaining to their three-dimensional structure (see Appendix for details).

4.2 Energy Model

In our experiments, we use the Lennard-Jones 6-12 potential that models the interaction energy of two atoms α and β as

$$V(\alpha,\beta) = \epsilon_{\alpha\beta}\left[\left(\frac{\sigma_{\alpha\beta}}{r_{\alpha\beta}}\right)^{12} - 2\left(\frac{\sigma_{\alpha\beta}}{r_{\alpha\beta}}\right)^{6}\right], \tag{13}$$

where $\epsilon_{\alpha\beta}$ is the depth of the potential well, $\sigma_{\alpha\beta}$ is the van der Waals bond length, and $r_{\alpha\beta}$ is the distance between the two atoms. An expression for (2) is then obtained by summing over all pair of atoms. Values for the parameters $\epsilon_{\alpha\beta}$ and $\sigma_{\alpha\beta}$ are taken from the Universal force field [42] without modification.

4.3 Results

In this section, we present the results of solving the conformational search problem for the selected molecules using the proposed VND method. Since the D-Wave 2000Q quantum annealer is not guaranteed to find an optimal solution, we first present results for VND, where an exact QUBO problem solver is used for optimizing the selected torsions at each iteration. These results focus solely on the performance of the proposed VND method as a conformational search approach by preventing any deterioration of the results attributable to the use of a device that has imperfections. For comparison purposes, we also present the results achieved by performing a local search heuristic (LS), and a hybrid of the two methods (LS–VND) in which a random conformation is first optimized using LS and then passed to VND for further optimization. The comparison with LS is helpful in understanding the improvement that can be achieved through the use of a more complex neighbourhood than what is used for LS. We also present the VND results with the quantum annealer used as the underlying QUBO problem solver. Finally, we compare the results with those found by our implementation of the PTMC algorithm for the conformational search problem, a state-of-the-art metaheuristic for conformational sampling. Details of the parameters used for VND, the quantum annealer, and PTMC are presented in the Appendix.

4.3.1 Reference Conformations

We compare our results against *reference* conformations found using a PTMC conformational search method. More specifically, we use an *initial* conformation for each of the selected molecules as input to a PTMC algorithm and let it run with a sufficiently large number of sweeps, such that the resulting *reference* geometry can be assumed to represent the absolute minimum conformation with a high degree of confidence. Details of this procedure are given in the Appendix. To ensure the fairness and accuracy of this comparison, the reference conformations are generated from the initial ones with the same potential energy model as the one we used for VND.

4.3.2 Success Metrics

The following metrics have been used for performance evaluation:

1. Success rate: the fraction of runs that found a conformation with an energy within 1 kcal/mol (roughly chemical accuracy) of the *reference* conformation's energy.

2. Residual: the energy difference between the best conformation found in a run and the *reference* conformation.

3. Time to solution (TTS): the time it took to find the best solution (conformation) found in a run.

Each run was terminated once it found a conformation within 0.1 kcal/mol of the *reference* conformation, or if it reached some other stopping criterion.

4.3.3 VND vs. LS vs. LS–VND

As a baseline for comparison, we used an exact QUBO problem solver for optimizing the selected set of torsions at each VND iteration in both VND and LS–VND runs. The residual and TTS for each of the three methods, for all six model systems, are presented in Figure 5. In the residual plots in Figure 5a, the three methods show significant overlap, so it is difficult to draw concrete conclusions about their relative performance. However, LS–VND appears to have a small advantage for the catalysts (top row), and for the ortho-phenylenes (bottom row) VND and LS–VND are comparable, whereas LS performs worse.

The TTS plots in Figure 5b show that LS is significantly faster than VND and LS–VND for the catalysts and ortho-phenylenes. This is presumably due to the conformational search problem being optimally solved by LS, leaving no room for further improvement by VND. These results suggest that for VND and LS–VND to have a time advantage over LS, a fast QUBO problem solver is needed. Note that TTS is dependent on

the shape of the PES, the size of the search space (a function of the number of torsions and the number of allowed discrete torsion angles), and the cost of energy evaluations (dependent on the number of atoms). We

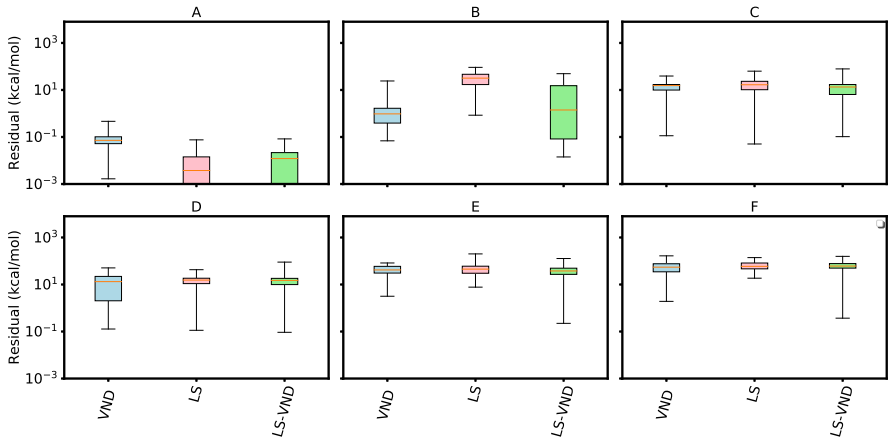

(a)

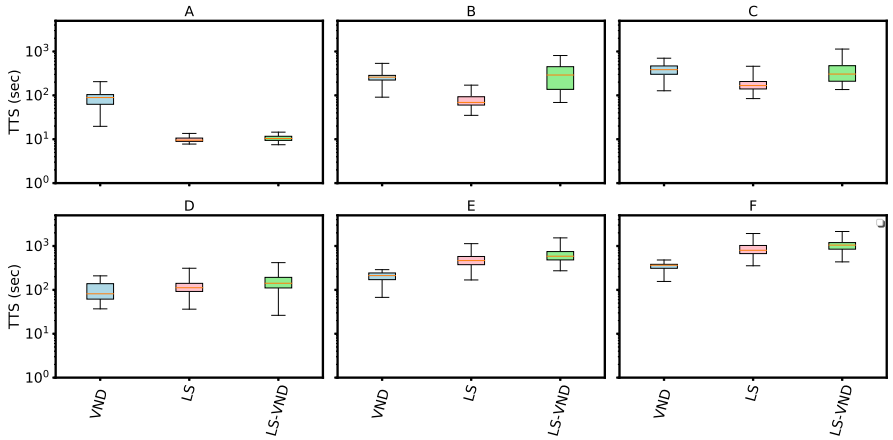

(b)

Figure 5: Residual (a) and time to solution (b) for VND, LS, and LS–VND, when using an exact solver as the underlying solver for the VND iterations. Each box represents the spread from the 25th to the 75th percentile, the red line represents the median, and the whiskers represent the minimum and maximum, including outliers.

present the results for VND, LS, and LS-VND using an exact solver subroutine in Table 1. Since LS–VND combines both the LS and VND methods, it is not surprising that the success probability of LS–VND is generally at least as high as that of LS or VND.

4.3.4 VND Using a Quantum Annealer

The results using VND with a quantum annealer as the underlying solver are presented in Table 2. The larger catalysts and ortho-phenylenes were solved to within 14–51 kcal/mol. A comparison of these results with the results of VND when using an exact QUBO problem solver (see Table 1) shows that the results are of lower quality. It is expected, however, that spending more effort on tuning the the quantum annealer's parameters would improve these results. Another observation from Table 2 is that VND employing a quantum annealer to solve QUBO problems has a larger TTS than VND using an exact QUBO problem solver. The quantum annealer solves a QUBO problem much faster than the exact solver, so one might think that its TTS should

Method	Model System	Success Rate	Residual (kcal/mol) Min	50th	75th	TTS (seconds) Min	50th	75th
VND	A	1.00	0.0	0.1	0.1	19.7	90.2	110.2
	B	0.60	0.1	0.9	1.6	90.9	271.6	320.3
	C	0.04	0.1	14.1	17.1	127.3	418.7	498.3
	D	0.16	0.1	14.4	21.4	37.0	101.1	153.5
	E	0.00	3.2	41.1	51.1	67.5	211.6	243.4
	F	0.00	1.9	60.7	76.1	155.7	374.1	391.4
LS	A	1.00	0.0	0.0	0.0	7.8	9.6	10.6
	B	0.03	0.8	33.1	46.5	35.0	69.7	91.6
	C	0.02	0.1	16.9	22.4	84.4	162.4	198.9
	D	0.03	0.1	15.4	18.5	36.3	112.9	148.2
	E	0.00	7.7	45.5	59.8	168.5	468.1	580.7
	F	0.00	18.6	59.2	78.5	354.6	784.2	1078.2
LS–VND	A	1.00	0.0	0.0	0.1	7.5	10.3	11.4
	B	0.44	0.0	1.4	15.2	68.9	294.5	421.7
	C	0.04	0.1	13.0	16.8	135.6	305.8	474.3
	D	0.17	0.1	15.0	18.4	26.3	142.0	199.0
	E	0.03	0.2	37.7	49.5	272.6	576.6	742.7
	F	0.02	0.4	60.9	78.1	435.1	1048.3	1284.8

Table 1: Results for VND, LS, and LS–VND, when using an exact solver as the underlying solver for the VND iterations. For each model system we report the minimum, median, and 75th percentile of the residual and TTS, over 100 runs.

also be lower. To explain this observation, one should note that it is not sufficient to merely solve QUBO problems fast, because if they are not solved optimally by the quantum annealer, the VND method may take longer to converge. Further, a significant portion of time is spent on transforming the QUBO problems into Ising problems (see Appendix), communicating with the quantum annealer, and mapping the results from the quantum annealer back to the logical bits. It is worth noting that elapsed real time is a fair measure of the time required to solve actual problems using a quantum annealer. This is in contrast to the customary approach of reporting only the annealing time, which is very small in comparison (on the order of a few microseconds per annealing cycle).

There are several known factors that make these QUBO problems challenging for the current generation of quantum annealers. First, the problem graphs are fully connected, whereas the quantum annealer's connectivity graph is extremely sparse, resulting in each logical bit being embedded onto chains of 17 qubits on the quantum annealer. It is difficult to maintain identical states for those qubits, resulting in a higher error rate. Second, these problems have a very large range of coefficients, due to the $1/r^{12}$ and $1/r^6$ terms in the Lennard-Jones 6-12 energy model. On the other hand, the quantum annealer's couplers have a limited bit precision and a fixed range. The large range of coefficients results in a loss of precision, which manifests itself in a lower success probability. Third, the current generation of quantum annealers has a high level of noise, referred to as intrinsic control error (ICE), causing a significant loss in precision. Future quantum annealers are expected to mitigate these factors, with more-dense hardware graphs, higher bit precision, and lower ICE levels.

Model System	Success Rate	Residual (kcal/mol)			TTS (sec)		
		Min	50th	75th	Min	50th	75th
A	0.00	1.8	5.8	7.5	156.1	330.2	366.4
B	0.00	14.6	56.4	90.7	136.0	288.9	395.5
C	0.00	15.2	40.7	65.4	300.9	507.3	726.9
D	0.00	8.3	24.2	34.6	385.8	788.3	990.1
E	0.00	31.6	69.4	103.8	445.7	982.2	1107.3
F	0.00	51.1	104.2	151.0	511.1	1221.2	1648.8

Table 2: Results for VND using a quantum annealer as the underlying QUBO solver. For each model system we report the minimum, median, and 75th percentile of the residual and TTS, over 25 runs.

4.3.5 A Comparison with Parallel Tempering

The results for PTMC are presented in Table 3. PTMC had a high success rate for all model systems except for system F. It is worth noting that PTMC's time to solution is generally significantly higher than that of VND, even for the model systems for which both algorithms had a 100% success rate.

Model System	Success Rate	Residual (kcal/mol)			TTS (seconds)		
		Min	50th	75th	Min	50th	75th
A	1.00	0.0	0.1	0.1	40.1	107.5	143.3
B	1.00	0.0	0.1	0.1	301.0	787.5	1092.8
C	1.00	0.0	0.1	0.1	585.5	2068.3	3025.0
D	1.00	0.0	0.1	0.1	154.8	710.4	1085.0
E	0.60	0.0	0.1	7.8	1051.4	4711.6	5975.8
F	0.12	0.0	14.2	18.0	3025.9	5758.3	6631.8

Table 3: Results for PTMC. For each model system, we report the minimum, median, and 75th percentile of the residual and TTS, over 100 runs.

4.3.6 Effect of the Neighbourhood Size on VND's Performance

In the VND experiments discussed above, we restricted the neighbourhood size, s, in each VND iteration to 63, the largest size of QUBO problem, whose underlying graph is complete, that can be solved using an equal-length embedding on the D-Wave 2000Q. In what follows, we study the effect of s on the performance of VND. This is useful in predicting the performance improvement achievable by increasing the number of qubits or their connectivity.

Table 4 presents the results for molecules B and C for different neighbourhood sizes when an exact solver is used to solve the QUBO problem at each iteration. The reason for choosing these two molecules is that they have a star-like structure and, as discussed in Section 3.6, the advantage of using a quantum annealer over an exact solver is expected to be more pronounced for these molecules.

As shown, increasing the neighbourhood size from 30 to 60 and then from 60 to 90 noticeably improves the results for both molecules. However, the improvement exhibits diminishing returns when the neighbourhood size is increased beyond 90. We expect similar behaviour to occur at different neighbourhood sizes for different families of molecules.

Neighbourhood Size (s)	Model System	Success Rate	Residual (kcal/mol) Min	50th	75th
30	B	0.07	0.2	16.7	26.0
	C	0.00	1.1	19.4	25.6
60	B	0.39	0.0	1.3	1.9
	C	0.04	0.1	13.4	17.1
90	B	0.55	0.0	0.9	1.5
	C	0.09	0.0	8.7	14.8
120	B	0.58	0.0	0.9	1.5
	C	0.10	0.1	8.7	14.1

Table 4: Effect of the neighbourhood size on the performance of VND. For each model system, we report the minimum, median, and 75th percentile of the residual over 500 runs.

5 Conclusion

In this paper, we present a variable neighbourhood descent (VND) method for conformational search. We introduce the rigid-body graph and use this simplified molecular structure to carefully define a neighbourhood structure to allow for efficient optimization using a binary quadratic optimizer. Based on current quantum annealing hardware, we select a 2-torsion-dependent neighbourhood at each iteration such that finding the best solution in the selected neighbourhood can then be formulated as a QUBO problem. The size of the neighbourhood can be chosen such that the method can be adapted to the number of available qubits as well as to their connectivity on the quantum annealer. As a result, the proposed method is not only well-suited for current hardware, but can be easily adapted to take advantage of hardware improvements. Whereas the proposed method can be used as a standalone conformational search approach, it could be combined with existing conformational search methods for potentially improved performance.

Beyond a simple presentation of the method, we also conduct a preliminary case study based on an implementation of the VND method using the D-Wave 2000Q quantum annealer to perform a conformational search for two families of molecules. In this exploration, we compare the results of our method with those of PTMC, a state-of-the-art solver for conformational search. To understand how much of the gap between the results of PTMC and those of VND used with a quantum annealer can be attributed to the imperfections of the quantum annealer, we replaced it with an exact QUBO problem solver. VND used along with the exact solver finds noticeably better conformations than those found with VND and the quantum annealer used together. This observation points to the potential improvement achievable in the short term through more-advanced tuning of the existing quantum annealer, and in the long term using improved hardware.

This work suggests a number of possible future research directions. For example, investigating refined neighbourhood change functions rather than using a random ordering of the torsions to choose the neighbourhood could lead to significant improvements. This could involve further exploitation of the molecular graph or the solutions from previous VND iterations. In addition, an improved selection of the s discrete points at each VND iteration to account for hardware limitations could further improve the results. We leave these improvements for future work.

We believe that the proposed method, based on careful hardware-aware neighbourhood selection, holds the potential to provide promising solutions to important optimization problems. Our method has opened a scalable path forward to leveraging emerging quantum technologies for one such problem of critical importance in the field of chemical and materials science.

6 Acknowledgements

The authors would like to thank Jamie Cohen, Alejandro Garza Gonzalez, Steve Arturo, Maritza Hernandez, Helmut G. Katzgraber, Kausar N. Samli, Takeshi Yamazaki, and Arman Zaribafiyan for their insightful

comments, Varinia Bernales for her assistance in preparing molecular geometries, and Marko Bucyk for reviewing and editing the manuscript.

References

[1] Emanuele Perola and Paul S Charifson. Conformational analysis of drug-like molecules bound to proteins: an extensive study of ligand reorganization upon binding. *Journal of medicinal chemistry*, 47(10):2499–2510, 2004.

[2] Johannes Kirchmair, Christian Laggner, Gerhard Wolber, and Thierry Langer. Comparative analysis of protein-bound ligand conformations with respect to Catalyst's conformational space subsampling algorithms. *Journal of chemical information and modeling*, 45(2):422–430, 2005.

[3] Ken A Dill and Justin L MacCallum. The protein-folding problem, 50 years on. *Science*, 338(6110):1042–1046, 2012.

[4] Richard H Boyd and Paul J Phillips. *The science of polymer molecules*. Cambridge University Press, 1996.

[5] Yinglong Miao and J Andrew McCammon. Unconstrained enhanced sampling for free energy calculations of biomolecules: A review. *Molecular simulation*, 42(13):1046–1055, 2016.

[6] Robert A Copeland. Conformational adaptation in drug–target interactions and residence time. *Future medicinal chemistry*, 3(12):1491–1501, 2011.

[7] Ekaterina I Izgorodina, Ching Yeh Lin, and Michelle L Coote. Energy-directed tree search: an efficient systematic algorithm for finding the lowest energy conformation of molecules. *Physical Chemistry Chemical Physics*, 9(20):2507–2516, 2007.

[8] Paul CD Hawkins. Conformation generation: The state of the art. *Journal of chemical information and modeling*, 57(8):1747–1756, 2017.

[9] Andrew R Leach. A survey of methods for searching the conformational space of small and medium-sized molecules. *Reviews in computational chemistry*, pages 1–55, 1991.

[10] Marcus PD Hatfield and Sándor Lovas. Conformational sampling techniques. *Current pharmaceutical design*, 20(20):3303–3313, 2014.

[11] Markus J Loferer, István Kolossváry, and András Aszódi. Analyzing the performance of conformational search programs on compound databases. *Journal of Molecular Graphics and Modelling*, 25(5):700–710, 2007.

[12] Robert E Bruccoleri and Martin Karplus. Prediction of the folding of short polypeptide segments by uniform conformational sampling. *Biopolymers: Original Research on Biomolecules*, 26(1):137–168, 1987.

[13] Jonathan M Goodman and W Clark Still. An unbounded systematic search of conformational space. *Journal of computational chemistry*, 12(9):1110–1117, 1991.

[14] Markus Christen and Wilfred F Van Gunsteren. On searching in, sampling of, and dynamically moving through conformational space of biomolecular systems: A review. *Journal of computational chemistry*, 29(2):157–166, 2008.

[15] Stefan Goedecker. Minima hopping: An efficient search method for the global minimum of the potential energy surface of complex molecular systems. *The Journal of chemical physics*, 120(21):9911–9917, 2004.

[16] Mikko J Vainio and Mark S Johnson. Generating conformer ensembles using a multiobjective genetic algorithm. *Journal of chemical information and modeling*, 47(6):2462–2474, 2007.

[17] Alex Strizhev, Edmond J Abrahamian, Sun Choi, Joseph M Leonard, Philippa RN Wolohan, and Robert D Clark. The effects of biasing torsional mutations in a conformational GA. *Journal of chemical information and modeling*, 46(4):1862–1870, 2006.

[18] Fang Bai, Xiaofeng Liu, Jiabo Li, Haoyun Zhang, Hualiang Jiang, Xicheng Wang, and Honglin Li. Bioactive conformational generation of small molecules: a comparative analysis between force-field and multiple empirical criteria based methods. *BMC bioinformatics*, 11(1):545, 2010.

[19] Adriana Supady, Volker Blum, and Carsten Baldauf. First-principles molecular structure search with a genetic algorithm. *Journal of chemical information and modeling*, 55(11):2338–2348, 2015.

[20] Jooyoung Lee, Harold A Scheraga, and Shelly Rackovsky. New optimization method for conformational energy calculations on polypeptides: conformational space annealing. *Journal of computational chemistry*, 18(9):1222–1232, 1997.

[21] Seungryong Heo, Juyong Lee, Keehyoung Joo, Hang-Cheol Shin, and Jooyoung Lee. Protein loop structure prediction using conformational space annealing. *Journal of chemical information and modeling*, 57(5):1068–1078, 2017.

[22] Luis B Morales, Ramón Garduño-Juárez, JM Aguilar-Alvarado, and FJ Riveros-Castro. A parallel tabu search for conformational energy optimization of oligopeptides. *Journal of computational chemistry*, 21(2):147–156, 2000.

[23] Svetlana Stepanenko and Bernd Engels. Tabu search based strategies for conformational search. *The Journal of Physical Chemistry A*, 113(43):11699–11705, 2009.

[24] Zhenqin Li, Keith E Laidig, and Valerie Daggett. Conformational search using a molecular dynamics–minimization procedure: Applications to clusters of coulombic charges, lennard–jones particles, and waters. *Journal of computational chemistry*, 19(1):60–70, 1998.

[25] Urmi Doshi and Donald Hamelberg. Towards fast, rigorous and efficient conformational sampling of biomolecules: Advances in accelerated molecular dynamics. *Biochimica et Biophysica Acta (BBA)-General Subjects*, 1850(5):878–888, 2015.

[26] Ralf Gehrke and Karsten Reuter. Assessing the efficiency of first-principles basin-hopping sampling. *Physical Review B*, 79(8):085412, 2009.

[27] Longjiu Cheng, Yan Feng, Jie Yang, and Jinlong Yang. Funnel hopping: Searching the cluster potential energy surface over the funnels. *The Journal of chemical physics*, 130(21):214112, 2009.

[28] Stephen R Wilson, Weili Cui, Jules W Moskowitz, and Kevin E Schmidt. Applications of simulated annealing to the conformational analysis of flexible molecules. *Journal of computational chemistry*, 12(3):342–349, 1991.

[29] Dominik Gront, Andrzej Kolinski, and Jeffrey Skolnick. Comparison of three monte carlo conformational search strategies for a proteinlike homopolymer model: Folding thermodynamics and identification of low-energy structures. *The Journal of Chemical Physics*, 113(12):5065–5071, 2000.

[30] Olivier Sperandio, Marc Souaille, François Delfaud, Maria A Miteva, and Bruno O Villoutreix. Med-3dmc: A new tool to generate 3D conformation ensembles of small molecules with a monte carlo sampling of the conformational space. *European journal of medicinal chemistry*, 44(4):1405–1409, 2009.

[31] William L Jorgensen and Julian Tirado-Rives. Monte carlo vs molecular dynamics for conformational sampling. *The Journal of Physical Chemistry*, 100(34):14508–14513, 1996.

[32] Chris Thachuk, Alena Shmygelska, and Holger H Hoos. A replica exchange monte carlo algorithm for protein folding in the HP model. *BMC bioinformatics*, 8(1):342, 2007.

[33] Sarah Rauscher, Chris Neale, and Régis Pomes. Simulated tempering distributed replica sampling, virtual replica exchange, and other generalized-ensemble methods for conformational sampling. *Journal of chemical theory and computation*, 5(10):2640–2662, 2009.

[34] Young Min Rhee and Vijay S Pande. Multiplexed-replica exchange molecular dynamics method for protein folding simulation. *Biophysical journal*, 84(2):775–786, 2003.

[35] Giuseppe E Santoro, Roman Martoňák, Erio Tosatti, and Roberto Car. Theory of quantum annealing of an ising spin glass. *Science*, 295(5564):2427–2430, 2002.

[36] Satoshi Morita and Hidetoshi Nishimori. Mathematical foundation of quantum annealing. *Journal of Mathematical Physics*, 49(12):125210, 2008.

[37] Alejandro Perdomo-Ortiz, Neil Dickson, Marshall Drew-Brook, Geordie Rose, and Alán Aspuru-Guzik. Finding low-energy conformations of lattice protein models by quantum annealing. *Scientific reports*, 2:571, 2012.

[38] The D-Wave 2000Q Quantum Computer: Technology Overview. https://www.dwavesys.com/sites/default/files/D-Wave%202000Q%20Tech%20Collateral_0117F.pdf.

[39] SHW Van der Ploeg, A Izmalkov, M Grajcar, U Hubner, S Linzen, S Uchaikin, Th Wagner, A Yu Smirnov, A Maasen Van Den Brink, MHS Amin, et al. Adiabatic quantum computation with flux qubits, first experimental results. *IEEE Transactions on Applied Superconductivity*, 17(2):113–119, 2007.

[40] Mark W Johnson, Mohammad HS Amin, Suzanne Gildert, Trevor Lanting, Firas Hamze, Neil Dickson, R Harris, Andrew J Berkley, Jan Johansson, Paul Bunyk, et al. Quantum annealing with manufactured spins. *Nature*, 473(7346):194, 2011.

[41] Robert H Swendsen and Jian-Sheng Wang. Replica monte carlo simulation of spin-glasses. *Physical review letters*, 57(21):2607, 1986.

[42] A. K. Rappé, C. J. Casewit, K. S. Colwell, W. A. Goddard, and W. M. Skiff. UFF, a Full Periodic Table Force Field for Molecular Mechanics and Molecular Dynamics Simulations. *Journal of the American Chemical Society*, 114(25):10024–10035, 1992.

[43] Fred W Glover and Gary A Kochenberger. *Handbook of metaheuristics*, volume 57. Springer Science & Business Media, 2006.

[44] Nenad Mladenović and Pierre Hansen. Variable neighborhood search. *Computers & operations research*, 24(11):1097–1100, 1997.

[45] Kengo Katayama and Hiroyuki Narihisa. Performance of simulated annealing-based heuristic for the unconstrained binary quadratic programming problem. *European Journal of Operational Research*, 134(1):103–119, 2001.

[46] Peter Merz and Bernd Freisleben. Greedy and local search heuristics for unconstrained binary quadratic programming. *Journal of heuristics*, 8(2):197–213, 2002.

[47] Peter Merz and Kengo Katayama. Memetic algorithms for the unconstrained binary quadratic programming problem. *BioSystems*, 78(1-3):99–118, 2004.

[48] Gintaras Palubeckis. Iterated tabu search for the unconstrained binary quadratic optimization problem. *Informatica*, 17(2):279–296, 2006.

[49] Zhipeng Lü, Fred Glover, and Jin-Kao Hao. A hybrid metaheuristic approach to solving the ubqp problem. *European Journal of Operational Research*, 207(3):1254–1262, 2010.

[50] Jialong Shi, Qingfu Zhang, Bilel Derbel, and Arnaud Liefooghe. A parallel tabu search for the unconstrained binary quadratic programming problem. In *IEEE Congress on Evolutionary Computation (CEC 2017)*, pages 557–564, 2017.

[51] Yang Wang, Zhipeng Lü, Fred Glover, and Jin-Kao Hao. Backbone guided tabu search for solving the ubqp problem. *Journal of Heuristics*, 19(4):679–695, 2013.

[52] Fubito Toyama, Kenji Shoji, Hiroshi Mori, and Juichi Miyamichi. An iterated greedy algorithm for the binary quadratic programming problem. In *Soft Computing and Intelligent Systems (SCIS) and 13th International Symposium on Advanced Intelligent Systems (ISIS), 2012 Joint 6th International Conference on*, pages 2183–2188. IEEE, 2012.

[53] Yang Wang, Zhipeng Lü, Fred Glover, and Jin-Kao Hao. Path relinking for unconstrained binary quadratic programming. *European Journal of Operational Research*, 223(3):595–604, 2012.

[54] Gary Kochenberger, Jin-Kao Hao, Fred Glover, Mark Lewis, Zhipeng Lü, Haibo Wang, and Yang Wang. The unconstrained binary quadratic programming problem: a survey. *Journal of Combinatorial Optimization*, 28(1):58–81, 2014.

[55] Rupak Biswas, Zhang Jiang, Kostya Kechezhi, Sergey Knysh, Salvatore Mandra, Bryan O'Gorman, Alejandro Perdomo-Ortiz, Andre Petukhov, John Realpe-Gómez, Eleanor Rieffel, et al. A nasa perspective on quantum computing: Opportunities and challenges. *Parallel Computing*, 64:81–98, 2017.

[56] Daniel J. Arriola, Edmund M. Carnahan, Phillip D. Hustad, Roger L. Kuhlman, and Timothy T. Wenzel. Catalytic production of olefin block copolymers via chain shuttling polymerization. *Science*, 312(5774):714–719, 2006.

[57] Haruyuki Makio, Norio Kashiwa, and Terunori Fujita. FI Catalysts: A New Family of High Performance Catalysts for Olefin Polymerization. *Advanced Synthesis and Catalysis*, 344(5):477–493, 2002.

[58] Terunori Fujita, Yasusushi Tohi, Makoto Mitani, Shigekazu Matsui, Junji Saito, Masatoshi Nitabaru, Kyoaki Sugi, Haruyuki Makio, and Toshiyuki Tsutsui. Olefin Polymierzation catalysts, transition metal compounds, processes for olefin polymerization,and Alpha-olefin/conjugated diene copolymers, EP0874005B1, 2003.

[59] Tina M Trnka. *Catalysts for olefin metathesis: ruthenium alkylidene complexes with phosphine and N-heterocyclic carbene ligands*. PhD thesis, California Institute of Technology, 2003.

[60] Sebastian Torker, Andre Müller, Raphael Sigrist, and Peter Chen. Tuning the Steric Properties of a Metathesis Catalyst for Copolymerization of Norbornene and Cyclooctene toward Complete Alternation. *Organometallics*, 29(12):2735–2751, 2010.

[61] Peter Schwab, Marcia B France, Joseph W Ziller, and Robert H Grubbs. A series of well-defined metathesis catalysts- synthesis of $[RuCl_2(=CHR')(PR_3)_2]$ and its reactions. *Angewandte Chemie International Edition*, 34:2039–2041, 1995.

[62] Takafumi Kawakami, Shingo Ito, and Kyoko Nozaki. Iron-catalysed homo- and copolymerisation of propylene: Steric influence of bis(imino)pyridine ligands. *Dalton Transactions*, 44(47):20745–20752, 2015.

[63] Highly Active Iron and Cobalt Catalysts for the Polymerization of Ethylene. *Journal of the American Chemical Society*, 120(16):4049–4050, 1998.

[64] George J.P. Britovsek, Michael Bruce, Vernon C. Gibson, Brian S. Kimberley, Peter J. Maddox, Sergio Mastroianni, Stuart J. McTavish, Carl Redshaw, Gregory A. Solan, Staffan Strömberg, Andrew J.P. White, and David J. Williams. Iron and cobalt ethylene polymerization catalysts bearing 2,6- bis(imino)pyridyl ligands: Synthesis, structures, and polymerization studies. *Journal of the American Chemical Society*, 121(38):8728–8740, 1999.

[65] Jian He, Jason L. Crase, Shriya H. Wadumethrige, Khushabu Thakur, Lin Dai, Shouzhong Zou, Rajendra Rathore, and C. Scott Hartley. Ortho-Phenylenes: Unusual conjugated oligomers with a surprisingly long effective conjugation length. *Journal of the American Chemical Society*, 132(39):13848–13857, 2010.

[66] Alexander J. Berresheim, Markus Müller, and Klaus Müllen. Polyphenylene Nanostructures. *Chem. Rev.*, 99(7):1747–1785, 1999.

[67] C. R. Groom, I. J. Bruno, M. P. Lightfoot, and S. C. Ward. The Cambridge Structural Database . *Acta Crystallographica B*, pages 171–179, 2016.

[68] Takafumi Kawakami, Shingo Ito, and Kyoko Nozaki. CCDC 1416774: Experimental crystal structure determination, 2015.

[69] Hiroshi Terao, Sei Ichi Ishii, Junji Saito, Sadahiko Matsuura, Makoto Mitani, Naoshi Nagai, Hidetsugu Tanaka, and Terunori Fujita. Phenoxycycloalkylimine ligated zirconium complexes for ethylene polymerization: Formation of vinyl-terminated low molecular weight polyethylenes with high efficiency. *Macromolecules*, 39(25):8584–8593, 2006.

[70] Sanyo Mathew, Laura A. Crandall, Christopher J. Ziegler, and C. Scott Hartley. Enhanced helical folding of ortho-phenylenes through the control of aromatic stacking interactions. *Journal of the American Chemical Society*, 136(47):16666–16675, 2014.

[71] L.R Gomes, l.M.N.B.F. Santos, and J.N. Low. CSD communication, 2012.

[72] L.R Gomes, l.M.N.B.F. Santos, and J.N. Low. CCDC 808554: Experimental crystal structure determination, 2014.

[73] H. Terao, Ishii S., J. Saito, S. Matsuura, M. Mitani, N. Nagai, H. Tanaka, and T. Fujita. CCDC 635962: Experimental crystal structure determination, 2014.

[74] H. Terao, Ishii S., J. Saito, S. Matsuura, M. Mitani, N. Nagai, H. Tanaka, and T. Fujita. CCDC 635963: Experimental crystal structure determination, 2014.

[75] S. Torker, A. Muller, R. Sigrist, and P. Chen. CCDC 811503: Experimental crystal structure determination, 2014.

[76] Frank Neese. The ORCA program system. *Wiley Interdisciplinary Reviews: Computational Molecular Science*, 2(1):73–78.

[77] FAccTs GmbH. ORCA 4.0.1, 2017.

[78] Yan Zhao and Donald G. Truhlar. The M06 suite of density functionals for main group thermochemistry, thermochemical kinetics, noncovalent interactions, excited states, and transition elements: Two new functionals and systematic testing of four M06-class functionals and 12 other function. *Theoretical Chemistry Accounts*, 120(1-3):215–241, 2008.

[79] Florian Weigend and Reinhart Ahlrichs. Balanced basis sets of split valence, triple zeta valence and quadruple zeta valence quality for h to rn: Design and assessment of accuracy. *Phys. Chem. Chem. Phys.*, (18):3297–3305.

[80] Florian Weigend. Accurate Coulomb-fitting basis sets for H to Rn. *Physical Chemistry Chemical Physics*, 8(9):1057–1065, 2006.

[81] D. Andrae, U. Häußermann, M. Dolg, H. Stoll, and H. Preuß. Energy-adjustedab initio pseudopotentials for the second and third row transition elements. *Theoretica chimica acta*, 77(2):123–141, Mar 1990.

[82] Florian Weigend. A fully direct RI-HF algorithm: Implementation, optimised auxiliary basis sets, demonstration of accuracy and efficiency. *Physical Chemistry Chemical Physics*, 4(18):4285–4291, 2002.

[83] Frank Neese, Frank Wennmohs, Andreas Hansen, and Ute Becker. Efficient, approximate and parallel Hartree-Fock and hybrid DFT calculations. A 'chain-of-spheres' algorithm for the Hartree-Fock exchange. *Chemical Physics*, 356(1-3):98–109, 2009.

[84] Dassault Systemes BIOVIA. Materials studio 2016, 2017.

[85] Hiroyasu Sato, Shinji Ando, Atsuko Kosaka, Takanori Fukushima, Daisuke Hashizume, Mikio Yamasaki, Kimiko Hasegawa, Azusa Muraoka, Hiroshi Ushiyama, Koichi Yamashita, and Takuzo Aida. Redox-responsive molecular helices with highly condensed π-clouds. *Nature Chemistry*, 3:68, 2011.

[86] Alexander J. Blake, Paul A. Cooke, Kevin J. Doyle, Susan Gair, and Nigel S. Simpkins. Poly-orthophenylenes: Synthesis by Suzuki coupling and solid state helical structures. *Tetrahedron Letters*, 39(49):9093–9096, 1998.

[87] Sanyo M. Mathew, James T. Engle, Christopher J. Ziegler, and C. Scott Hartley. The role of arene-arene interactions in the folding of ortho-phenylenes. *Journal of the American Chemical Society*, 135(17):6714–6722, 2013.

[88] Sanyo M. Mathew, James T. Engle, Christopher J. Ziegler, and C. Scott Hartley. CCDC 954742: Experimental crystal structure determination, 2014.

[89] Tomas Boothby, Andrew D King, and Aidan Roy. Fast clique minor generation in chimera qubit connectivity graphs. *Quantum Information Processing*, 15(1):495–508, 2016.

7 Appendix

7.1 A Brief Overview of Quantum Annealing

A quantum annealer solves problems in the form of an Ising Hamiltonian defined by

$$\begin{aligned} \arg\min_{\mathbf{s}} \quad & \mathbf{s}^T\mathbf{J}\mathbf{s} + \mathbf{h}^T\mathbf{s}, \\ \text{s.t.} \quad & \mathbf{s} \in \{-1, 1\}^N, \end{aligned}$$

where N is the number of qubits and $\mathbf{h}$ and $\mathbf{J}$ represent the local fields and couplers, respectively. The above Ising problem can be transformed into a quadratic unconstrained binary optimization (QUBO) problem

$$\begin{aligned} \min_{\mathbf{x}} \quad & \mathbf{x}^T\mathbf{Q}\mathbf{x}, \\ \text{s.t.} \quad & \mathbf{x} \in \{0, 1\}^N \end{aligned}$$

by substituting $\mathbf{s} = 2\mathbf{x} - \mathbf{1}$. To utilize the quantum annealer, the local fields and the pairwise couplings of the Ising formulation ($\mathbf{h}$ and $\mathbf{J}$) must be specified.

Quantum annealing is inspired by the adiabatic principle: a system at equilibrium, in its ground state, will remain in the ground state provided the system is evolved sufficiently slowly. During each annealing cycle, the Hamiltonian of the system is continuously deformed from a known, simple Hamiltonian H_0 to the problem Hamiltonian H_{Ising}, which is typically a non-trivial task. This evolution is represented by the Hamiltonian

$$H(\tau) = A(\tau)H_0 + B(\tau)H_{\text{Ising}}, \quad \tau \in [0, 1], \tag{14}$$

where $\tau = t/t_a$, with t_a representing the annealing time. The functions $A(\tau)$ and $B(\tau)$ specify the annealing schedule. Typically, $A(\tau)$ and $B(\tau)$ are monotonically decreasing and increasing functions, respectively.

Given that this device operates at a very low but finite temperature, the quantum annealer is a heuristic solver. For this reason, typical use of the quantum annealer involves obtaining an ensemble that consists of the results of hundreds, if not thousands, of reads from the quantum annealer. Additional specifications of the quantum annealer include the annealing time, the scaling of the $\mathbf{h}$ and $\mathbf{J}$ values, and using gauges to compensate for the errors due to noise.

7.2 Choice of Molecules

7.2.1 *Source* (Experimental) Structures

For systems A–C, atomic coordinates from single-crystal X-ray structures are available, either for the compounds themselves (B and C) [60, 67, 68] or for close analogues (A) [69]. Systems A and C represent the active, cationic (A) and inactive, initial form (C) of two olefin polymerization catalysts. System B embodies a catalyst for olefin metathesis. X-ray experimental data regarding the 3D structures of ortho-phenylenes (oPh) are currently confined to substituted analogues up to chain lengths of ten phenylene units [70]; the longest-chain unsubstituted oPh^x instance is for $x = 5$ [67, 71, 72]. Complementary computational studies exist for $x = 12$ and support the experimental 3D structures [65]. Experiments and computation agree that the preferred morphology for oPh^x oligomers is the tightly helical, "closed-helix" conformation. In the following, we refer to the geometries that correspond directly to experiments (or high-level computation for system A) as *source* geometries.

7.2.2 *Initial* Conformations

In the present work, we have focused exclusively on the torsional degrees of freedom, while other molecular degrees of freedom such as the bond lengths and angles are fixed at the values they assumed in the *initial* structure and excluded from the optimization. We employed two strategies to create *initial* structures. The details of these methods are given below; the geometries resulting from them are referred to as *initial* geometries in the following.

- **Method a:** Guesses pertaining to the structure of metal compounds were created using off-the-shelf molecular modelling software based on the structure of experimentally characterized analogues (compound A) [69, 73, 74] or directly obtained from the X-ray diffraction structures (compounds B and C) [67, 68, 75]. These structures were locally optimized using the ORCA 4 program [76, 77] in order to remove minor inconsistencies in the atomic positions. For this, we used the popular M06 density functional [78] and a def2-SVP basis set [79], along with the corresponding fit basis set on non-metal atoms and a def2-TZVP basis set [80] on metal atoms, together with the appropriate ECPs [81]. The RI-JK [82] and RIJCOSX [83] approximations were employed for acceleration. An unrestricted wave function with four unpaired spins was used for system C.

- **Method b:** Ortho-phenylene (oPh) geometries were generated as closed helices [65] using Materials Studio [84] by optimizing them with the Universal force field [42] implemented in Materials Studio. There is evidence that the open and closed helical states of oPh^n are close in energy [65] to the majority of experimental evidence gathered in the solid state pointing to a general preference for the closed helix [70, 85–88]. Using UFF to optimize all bond lengths, angles, and torsions, we found that the closed and open helix conformers of oPh^{20} differ by 1.5 kcal/mol per phenylene unit, favouring the closed helix.

7.2.3 From *Initial* Conformations to *Reference* Conformations

Reference conformations for all nine molecules, embodying a best-effort guess of the global minima available to the *initial* conformations in the torsional space, were generated using an implementation of the parallel tempering Monte Carlo (PTMC) method. All other experiments described in the rest of this paper were stopped once a conformation with an energy within 0.1 kcal/mol of the energy of the *reference* conformation was found.

The agreement between *source* structures, which closely mirror or are identical to experimental structures, and *reference* structures found by PTMC is significant for several reasons. *Source* structures embody the best available information regarding the global minimum geometry of the system. In order to be useful, a conformational search algorithm should be able to locate this geometry starting without prior knowledge of the optimal set of torsion angles. The *reference* structure, therefore, should closely approximate the *source* structure. If they do not agree, several causes may be contributing to the disagreement:

- **Force-field failure:** The force field guiding the search of the torsional energy surface may be insufficiently realistic to capture the true global minimum.
- **Structural bias:** The initial set of bond lengths and bond angles supplied with the *initial* structure may bias the conformational search in the torsional space toward regions inconsistent with the *source* geometry.
- **Experimental bias:** The *source* geometry itself may be dominated by effects not present in an isolated, molecular system, as might be the case if the *source* geometry is taken from a crystal with significant intermolecular interactions.

Table 5 summarizes the differences between *source* and *reference* geometries. Agreement between the *reference* structures and the *source* structures is indicated by the contents of the "*Reference* Matches *Source*" column. In general, *Reference* structures correspond closely to the *source* structures; except for the case of system B, *source* and *reference* structures occupy the same energy basin. For system B, the PTMC search identified a basin with a minimum energy lower than the *source* structure by ≈ 1.5 kcal/mol. This deviation is attributable either to force-field failure (due to the use of a simplified Universal force field during PTMC) or experimental bias (see above) due to subtle effects of crystal packing on the *source* geometry.

Model System	Type	Torsions	Source	Method	*Reference* Matches *Source*
A	metal comp'd	6	TW	a	Y
B	metal comp'd	9	[60, 67]	a	N
C	metal comp'd	14	[67, 68]	a	Y
D	o-phenylene	9	TW	b	Y
E	o-phenylene	15	TW	b	Y
F	o-phenylene	19	TW	b	Y

Table 5: Geometric properties of model systems A–F. "Model System" gives the label assigned to each model, "Type" is the type of the model system, "Torsions" is the number of the free rotating torsion bonds, "Source" is the source for the initial geometry (where "TW" stands for "this work"), "Method" is the method for optimizing the source geometry to yield the initial geometry, "*Reference* Matches *Source*" indicates if *source* and *reference* geometries occupy the same energy basin.

7.3 Further Details of the Experiments

Parameters of different methods used in Section 4 are explained in the following.

- **PTMC**: Our implementation of PTMC performs 10 parallel MC experiments (replicas) at different temperatures. Each sweep of PTMC consists of proposing M moves, one for each torsion, for each of the replicas (i.e., a total of $10M$ proposals to sweep over all replicas). The proposed moves are accepted according to the Metropolis criterion. Initial temperatures of the replicas are distributed geometrically, given the temperatures of the first and last replicas. The temperatures are automatically updated during the experiment to maintain reasonable Monte Carlo and replica exchange acceptance probabilities.
- **D-Wave 2000Q**: To solve each of the QUBO problems created for optimizing a subset of torsions, a total of 5000 solutions (reads) from the the quantum annealer are obtained and the best solution from among them is used. To improve the results obtained from the annealer, 10 gauges are applied, meaning that the 5000 solutions are obtained by making 10 calls, each with a different gauge, to the quantum annealer, where each call returns 500 solutions. The annealing time is 5 microseconds for each of the reads. Using a deterministic embedding solver that finds complete embeddings with chains of

equal length [89], the largest complete graph that can be embedded on the quantum annealer has a size of 63. Majority voting is used to settle the disagreement between the qubits of the chain assigned to a logical binary variable.

- **VND**: The VND results are obtained for $b_{\max} = 200$ and $b_{\text{nc}} = 10$. To respect the size of the largest complete graph embeddable onto the quantum annealer's hardware graph, we set $s = 63$.