

Double variable neighbourhood search with smoothing for the molecular distance geometry problem

Leo Liberti · Carlile Lavor · Nelson Maculan ·
Fabrizio Marinelli

Received: 30 April 2007 / Accepted: 13 July 2007 / Published online: 8 August 2007
© Springer Science+Business Media, LLC 2007

Abstract We discuss the geometrical interpretation of a well-known smoothing operator applied to the Molecular Distance Geometry Problem (MDGP), and we then describe a heuristic approach based on Variable Neighbourhood Search on the smoothed and original problem. This algorithm often manages to find solutions having higher accuracy than other methods. This is important as small differences in the objective function value may point to completely different 3D molecular structures.

Keywords Molecular conformation · Distance geometry · Global optimization · Global continuation · Variable neighbourhood search · Smoothing

1 Introduction

The Molecular Distance Geometry Problem (MDGP) is the problem of finding an embedding in \mathbb{R}^3 of a weighted graph G such that all Euclidean distances between points in the embedding are the same as the corresponding edge weights in the graph. The main application is to find the three-dimensional structure of a molecule given a subset of the atomic distances (these are usually found using NMR techniques) [1, 17]. There are other applications in network

L. Liberti (✉) · F. Marinelli
École Polytechnique, LIX, Palaiseau 91128, France
e-mail: liberti@lix.polytechnique.fr

C. Lavor
Department of Applied Mathematics (IMECC-UNICAMP), State University of Campinas- CP 6065,
Campinas-SP 13081-970, Brazil
e-mail: clavor@ime.unicamp.br

N. Maculan
COPPE – Systems Engineering, Federal University of Rio de Janeiro, P.O. Box 68511,
Rio de Janeiro-RJ, 21941-972, Brazil

F. Marinelli
e-mail: marinelli@lix.polytechnique.fr

localization [5] and graph drawing [2]. It has been shown that the MDGP is NP-hard [19]; if G is a complete graph, however, the MDGP can be solved in linear time [3]. The MDGP applied to protein backbones is susceptible of a discrete formulation which leads to improved accuracy and efficiency [12]. See [20, 10] for overviews on the main solution methods used to tackle the MDGP. It should be noted straight away that we refer to the MDGP as a precisely formalized decision problem, and not as a practical chemical problem. We therefore make three assumptions that in real life may be easily challenged: (a) a subset of exact distances (as opposed to approximate) is given as part of the input; (b) no measurement errors occur; (c) the optimal 3D embedding of the graph is not influenced by a potential energy minimization term on the objective function. We refer to [7] for a more realistic formulation.

Consider an undirected graph $G = (V, E)$ with weights $d : E \rightarrow \mathbb{R}_+$ where V is the set of vertices (also called *atoms*) and E is the set of weighted edges (also called *inter-atomic distances*). Let $N = |V|$ and $d_{ij} = d(\{i, j\})$ for $\{i, j\} \in E$. A solution of the MDGP is a set of points $x_1, \dots, x_N \in \mathbb{R}^3$ satisfying

$$\forall \{i, j\} \in E \quad \|x_i - x_j\| = d_{ij}. \quad (1)$$

Notationally, each 3-vector x_i has components (x_{i1}, x_{i2}, x_{i3}) , and we indicate the vector sequence (x_1, \dots, x_N) by x . The MDGP can be naturally cast as a continuous nonconvex optimization problem $\min_x f(x)$ by minimizing the squared error on each Eq. (1):

$$f(x) = \sum_{\{i,j\} \in E} (\|x_i - x_j\|^2 - d_{ij}^2)^2. \quad (2)$$

Since each Eq. (1) must be satisfied, a candidate point x is a solution of the MDGP if and only if $f(x) = 0$. Note that (2) has a large number of local minima, so this is a practically hard global optimization problem.

One of the most promising methods for solving the MDGP is the Global Continuation Algorithm (GCA) [15, 16]. The GCA relies on a class of smoothed objective functions derived from (2) by means of a Gaussian transform. This class is indexed by a parameter λ ; for $\lambda = 0$ we recover the original function (2), and for large enough λ the smoothed function becomes convex. The GCA locally solves a sequence of smoothed problems for decreasing values of λ . To assess solution accuracy, the authors of [15, 16] rely on a different quality measure called the Largest Distance Error (LDE) defined as:

$$\text{LDE}(x) = \frac{1}{|E|} \sum_{\{i,j\} \in E} \frac{|\|x_i - x_j\| - d_{ij}|}{d_{ij}}. \quad (3)$$

It is clear that the LDE is simply a different evaluation of the error in satisfying the Eq. (2) and that x is a solution of the MDGP if and only if $\text{LDE}(x) = 0$. The GCA has been implemented in the `dgsol` code, available from <http://www.mcs.anl.gov/more/dgsol/>.

One of the striking features of `dgsol` is its speed and the fact that the time taken to solve the problem seems to grow rather slowly as a function of the number of atoms in the molecule. On the other hand, `dgsol` usually finds solutions whose LDE is relatively large (in the order of 0.01 or even 0.1). Since there are many totally different 3D structures having small LDE (see for example Fig. 1, it is of paramount importance to obtain solutions whose LDEs are very close to zero.

In a previous paper [9], we tested three different general-purpose global optimization methods on several MDGP instances, concluding that Variable Neighbourhood Search (VNS)




Fig. 1 Different 3D embeddings for the same molecule with very similar objective function values for $f(x)$ (both are in the order of 10^{-11}). Notice that atomic bonds, rather than inter-atomic distances, are displayed

was the best. Notice that VNS was also used to find the 3D structure of molecules by minimizing a potential energy objective function [4]. In this paper, we combine ideas from both GCA and VNS algorithms to obtain a method which we call Double VNS with Smoothing (DVS). Although considerably slower than `dgso1`, the solutions obtained by DVS generally have a much smaller LDE. DVS is a classical two-phase global optimization algorithm with diversification in the first phase and intensification in the second phase. More precisely, DVS solves a smoothed version of the problem using VNS, then tightens the bounds around the smoothed solution, and finally applies VNS to the original problem with restricted bounds. The bounds are restricted following a geometrical interpretation of the smoothed problem, which makes it possible to relate the smoothing parameter to the bounds where the solution to the original MDGP should lie.

The rest of this paper is organized as follows. In Sect. 2 we discuss the geometrical interpretation of the MDGP smoothing proposed by [15], which leads to some indications concerning the bound restriction at each DVS iteration. In Sect. 3 we present the DVS algorithm. In Sect. 4 we report on the computational results. Section 5 concludes the paper.

2 Geometrical interpretation of smoothing

As (2) is a nonconvex function with a considerable number of local minima, it makes sense to look for a *smoothing* (i.e. a function in some sense “close” to (2) with fewer local minima). In [15], a family of smoothings for (2) is proposed which depends on a unique smoothing parameter $\bar{\lambda} \in [\bar{\lambda}^L, \bar{\lambda}^U]$: for $\bar{\lambda} = \bar{\lambda}^L$ we recover the original objective (2), and for $\bar{\lambda} = \bar{\lambda}^U$ we get a convex smoothing. The smoothing family $\bar{f}_{\bar{\lambda}}(x)$ given by [15] (and employed in the GCA) is derived as the closed analytical form of a Gaussian integral:

$$\bar{f}_{\bar{\lambda}}(x) = f(x) + \sum_{\{i,j\} \in E} (10\bar{\lambda}^2 \|x_i - x_j\|^2) + \bar{\gamma}, \tag{4}$$

where $\bar{\gamma}$ is a constant and f is given by (2). We shall call this smoothing the *Moré-Wu smoothing*. It should be clear that, since the nonconvexities in (2) arise from the negative cross product terms $-2 \sum_{\{i,j\} \in E} d_{ij}^2 \|x_i - x_j\|^2$, the smoothing $\bar{f}_{\bar{\lambda}}(x)$ tends to eliminate these nonconvexities by adding the positive term $\sum_{\{i,j\} \in E} (10\bar{\lambda}^2 \|x_i - x_j\|^2)$ to the original objective function $f(x)$. This implies that the quartic terms $\sum_{\{i,j\} \in E} \|x_i - x_j\|^4$ have a bigger effect in the minimization process. Intuitively speaking, $\bar{f}_{\bar{\lambda}}$ describes a configuration where atoms whose distance d_{ij} is small tend to be positioned even closer together, whilst the relative position of far away atoms remains roughly unchanged; in other words, close

atoms tend to be clustered together in the smoothed function. We shall now formalize this geometric intuition.

We introduce a function F which associates to each vector of squared NMR distances (denoted by $d^2 = (d_{ij}^2 \mid \{i, j\} \in E)$) the MDGP objective function (2):

$$F(d^2) = \sum_{\{i,j\} \in E} (\|x_i - x_j\|^4 - 2d_{ij}^2\|x_i - x_j\|^2 + d_{ij}^4).$$

Notice we obviously have $(F(d^2))(x) \equiv f(x)$. We re-state the Moré-Wu smoothing as follows, introducing the smoothing parameter $\lambda = \bar{\lambda}\sqrt{5}$ and ignoring the constant $\bar{\gamma}$:

$$f_\lambda(x) = f(x) + 2 \sum_{\{i,j\} \in E} (\lambda^2\|x_i - x_j\|^2). \tag{5}$$

Now, we define the *Moré-Wu smoothing operator* (denoted by \odot) as follows:

$$\lambda^2 \odot F(d^2) = f_\lambda(x) - \gamma, \tag{6}$$

where $\gamma = \lambda^2(\lambda^2|E| - 2 \sum_{\{i,j\} \in E} d_{ij}^2)$. It turns out that the Moré-Wu smoothing operator acts by shortening the square of the distances:

Theorem 2.1 $\lambda^2 \odot F(d^2) = F(d^2 - \lambda^2)$, where $d^2 - \lambda^2$ is the vector $(d_{ij}^2 - \lambda^2 \mid \{i, j\} \in E)$.

Proof

$$\begin{aligned} \lambda^2 \odot F(d^2) &= \sum_{\{i,j\} \in E} (\|x_i - x_j\|^4 - 2(d_{ij}^2 - \lambda^2)\|x_i - x_j\|^2 + d_{ij}^4) - \gamma \\ &= \sum_{\{i,j\} \in E} (\|x_i - x_j\|^4 - 2(d_{ij}^2 - \lambda^2)\|x_i - x_j\|^2 + (d_{ij}^2 - \lambda^2)^2) \\ &= \sum_{\{i,j\} \in E} (\|x_i - x_j\|^2 - (d_{ij}^2 - \lambda^2))^2 \\ &= F(d^2 - \lambda^2), \end{aligned}$$

where γ in the first equation was replaced by its definition in terms of λ and d_{ij} and carried into the sum. □

By Theorem 2.1, for all pairs of atoms $\{i, j\} \in E$ such that $\lambda \geq d_{ij}$, the corresponding term contributing to the objective function is $(\|x_i - x_j\|^2 + \lambda^2 - d_{ij}^2)^2$. This is a convex term with minima attained at $x_i = x_j$; in other words, atoms whose distances are shorter than λ will tend, if possible, to collapse into the same spatial positions. Pairs of atoms which are further away than λ will, if possible, be positioned at a shorter distance $\sqrt{d_{ij}^2 - \lambda^2}$. This observation is based on a term-by-term analysis of a non-separable function, so it is of course only true if λ is large enough to make \tilde{f}_λ convex (in which case we obtain the trivial solution $x_i = 0$ for all $i \in V$). Its usefulness, however, is that it formalizes an insight as to the general behaviour of a molecule in space in function of a smooth transformation of its associated objective function.

An idea that may arise naturally for generating another type of smoothing is that of considering smoothing weights proportional to the inter-atomic distances. In the rest of this section we show that this idea does not generate a smoothing at all. Let

$$\tilde{f}_\lambda(x) = f(x) + 2 \sum_{\{i,j\} \in E} (\lambda^2 d_{ij}^2 \|x_i - x_j\|^2),$$

and define the *scaling operator* (denoted by \otimes) as follows:

$$\lambda^2 \otimes F(d^2) = \tilde{f}_\lambda(x) - \tilde{\gamma}, \tag{7}$$

where $\tilde{\gamma} = (\lambda^4 - 1) \sum_{\{i,j\} \in E} d_{ij}^4$. It turns out that this type of smoothing is in fact a scaling (thus the operator name).

Proposition 2.2 $\lambda^2 \otimes F(d^2) = F(\lambda^2 d^2)$, where $\lambda^2 d^2$ is the vector $(\lambda^2 d_{ij}^2 \mid \{i, j\} \in E)$.

Proof The proof is by direct computation, similar to that of Theorem 2.1. □

The geometrical interpretation of Proposition 2.2 is simply that all inter-atomic distances are scaled by a factor λ . It shows that the weights of the smoothing terms should not be proportional to the distances, for otherwise the problem is simply scaled rather than actually smoothed.

2.1 Bounds tightening

The aim of this section is to show how to use the geometrical interpretation of smoothing to derive some indications as to how to tighten the bounds at each iteration of the DVS algorithm (i.e. after the solution of each smoothed problem and before the solution of the original problem).

Suppose \bar{x} is a global optimum of the smoothing $\lambda^2 \odot F(d^2)$ for some $\lambda > 0$. According to Theorem 2.1, this is likely to represent a molecule where the closer atoms are clustered together (possibly even at the same position), whereas the more distant atoms are at a smaller distance. With this information, we can try to enforce meaningful lower and upper bounds x^L, x^U on the atomic position variables x to restrict the search in a subsequent iteration. We recall that as $\lambda \rightarrow 0, \tilde{f}_\lambda \rightarrow f$; in other words, the smoothings become more and more nonconvex and more difficult to minimize. In order for intensification to occur in the second phase, the desirable behaviour is that the bounds should get tighter as λ decreases.

For a given $\lambda > 0$, the smoothed function $\lambda^2 \odot F(d^2) = \tilde{f}_\lambda$ roughly describes a configuration where the inter-atomic distance between atoms i and j is $\tilde{d}_{ij} = \sqrt{d_{ij}^2 - \lambda^2}$.

Let $\xi_{ij} = d_{ij} - \tilde{d}_{ij}$ be the difference between the smoothed and the original distance, let \bar{x} be the global minimizer of $\tilde{f}_\lambda(x)$ and x^* the globally optimal solution of the original objective function $f(x)$. For any atom pairs $\{i, j\} \in E$, if we assume that $x_j^* = \bar{x}_j$ (i.e. atom j is fixed after the solution of the smoothed problem), then x_i^* is on the surface of a sphere S_{ij} centered at x_j^* with radius d_{ij} . Similarly, for any atom $i \in V$, if we assume that $x_j^* = \bar{x}_j$ for all $j \in V : \{i, j\} \in E$, x_i^* is on the intersection of the surfaces of the spheres S_{ij} . Naturally, since the smoothing may not actually describe any feasible molecule, this intersection may be empty (see Fig. 2; the figure shows the pairwise intersections of the spheres, which give a meaningful boundary for the position of the atom once the neighbouring atoms are perturbed).

If we now allow the previously fixed atoms $j \in V : \{i, j\} \in E$ to be perturbed, the position of the i -th atom is likely to be in $S_i = \bigcap_j S_{ij}$. For algorithmic reasons (see Sect. 3.1) we are interested in a box-shaped neighbourhood of S_i . More precisely, for each $i \in V$ we wish to find the tightest interval $[x_{ik}^L, x_{ik}^U]$ containing x_{ik} (for $1 \leq k \leq 3$) such that the box $[x_{i1}^L, x_{i1}^U] \times [x_{i2}^L, x_{i2}^U] \times [x_{i3}^L, x_{i3}^U]$ contains S_i . For each $i \in V$ we can therefore solve $6N$ convex problems defined for different $1 \leq k \leq 3$ and $b \in \{0, 1\}$:

$$(-1)^b \min_{h \leq 3} (-1)^b x_{ik} \sum_{h \leq 3} (x_{ih} - \bar{x}_{jh})^2 \leq d_{ij}^2 \quad \forall j \in V : \{i, j\} \in E \tag{8}$$




Fig. 2 Likely positions of atom i when the adjacent atoms are fixed (viewed in 2D)

where the optimum yields x_{ik}^L for $b = 0$ and x_{ik}^U for $b = 1$. Because of the (relative) computational cost of solving $6N$ convex problems (8) at each step of the DVS algorithm (see Sect. 3), we content ourselves with an approximation.

Since each point $\tilde{x}_i \in S_i$ must fall in each sphere S_{ij} , one has $\bar{x}_{jk} - d_{ij} \leq \tilde{x}_{ik} \leq \bar{x}_{jk} + d_{ij}$ for any k, j . Motivated by this fact, for all $i \in V, k \leq 3$ we define:

$$x_{ik}^L = \max\{-d_{ij} + \bar{x}_{jk} \mid j \in V : \{i, j\} \in E\} \tag{9}$$

$$x_{ik}^U = \min\{d_{ij} + \bar{x}_{jk} \mid j \in V : \{i, j\} \in E\}. \tag{10}$$

It is easy to show that the cartesian product across all $k \leq 3$ of $[x_{ik}^L, x_{ik}^U]$ contains the set S_i (for all $i \in V$), although it may not in general be the tightest such box.

A more accurate approximation can be derived by iteratively determining a sequence of box-shaped domains the first of which corresponds to the tightest box containing a given sphere S_{i1} (in an arbitrary sequence), the last of which is a box-shaped neighbourhood of S_i , and the h -th of which is the tightest box containing the exact intersection of the h -th sphere with the $(h - 1)$ -th box-shaped domain. At each iteration, the box-shaped domain can be determined by finding the points on the frontier of (11) that minimize/maximize the coordinate axes. More precisely, let $N_i = |\delta(i)|$ for all $i \in V, B_0 = \mathbb{R}^3$ and B_h be the tightest box $[L_{h1}, U_{h1}] \times [L_{h2}, U_{h2}] \times [L_{h3}, U_{h3}]$ containing $S_{ih} \cap B_{h-1}, 1 \leq h \leq |\delta(i)|$. It is easy to see that, for $1 \leq k \leq 3, [L_{1k}, U_{1k}] = [\bar{x}_{1k} - d_{i1}, \bar{x}_{1k} + d_{i1}]$ and $[x_{ik}^L, x_{ik}^U] = [L_{N_i k}, U_{N_i k}]$.

The approximation of the box-shaped neighbourhood of S_i can be therefore obtained starting from B_0 and iteratively determining B_h , until $h = N_i$. By definition B_h is the smallest box containing all the points y satisfying

$$\begin{cases} \sum_{k=1}^3 (y_k - \bar{x}_{hk})^2 \leq d_{hk}^2 \\ L_{h-1,k} \leq y_k \leq U_{h-1,k} \end{cases} \quad 1 \leq k \leq 3 \tag{11}$$

Without loss of generality suppose $B_h \neq \emptyset$. Let $C^l, 1 \leq l \leq q \leq 6$, be the intersection arc between S_{ih} and the l -th facet of B_{h-1} , and $C = \bigcup_{l \leq q, k \leq 3} \{\bar{y} \in C^l \mid \bar{y}_k \leq y_k, \forall y \in C^l\}$. Moreover, let Y be the set of intersection points between S_{ih} and the edges of B_{h-1} . Therefore, for $1 \leq k \leq 3$,

$$L_{hk} = \min\left\{ \min_{y \in C \cup Y} \{y_k\}, \Gamma_k^L \right\} \quad \text{and} \quad U_{hk} = \max\left\{ \max_{y \in C \cup Y} \{y_k\}, \Gamma_k^U \right\},$$

where

$$\Gamma_k^L = \begin{cases} \bar{x}_{hk} - d_{ih} & \text{if } \bar{x}_{hk} - d_{ih} \text{ satisfies (11)} \\ +\infty & \text{otherwise} \end{cases}$$

and

$$\Gamma_k^U = \begin{cases} \bar{x}_{hk} + d_{ih} & \text{if } \bar{x}_{hk} + d_{ih} \text{ satisfies (11)} \\ -\infty & \text{otherwise.} \end{cases}$$

3 The algorithm

The DVS algorithm is basically a two-level VNS search within a GCA-type framework. The GCA works by locally solving a sequence of Moré-Wu smoothings $\min \tilde{f}_{\lambda_k}(x)$. At step k , the GCA obtains the next iterate x_{k+1} by locally solving the current smoothing $\min \tilde{f}_{\lambda_k}(x)$ from the initial starting point x_k . The problem of this approach is that its search scope is not global enough. There is no guarantee that the global optimum of the initial smoothing will be traced homotopically to the global optimum of the original problem. Moreover, the local solution algorithm may fail to solve even the initial “mildly nonconvex” smoothing to global optimality.

In the DVS, the local search phase is replaced by a two-level VNS. The global solution \bar{x} to the smoothed problem $\min f_{\lambda_k}(x)$ (see Eq. (5)) is found by means of a formulation-based VNS for NLPs [9, 11]. Next, an attempt to find the global solution x^* of $\min f(x)$ is performed by restricting the original problem to a neighbourhood of \bar{x} (as detailed in Sect. 2) and solving this restriction with the same formulation-based VNS.

3.1 VNS solver

We employed the VNS solver described in [11]. The search space is defined as the hyper-parallelepiped given by the set of variable ranges $x^L \leq x \leq x^U$. At first we pick a random point \tilde{x} in the search space, we start one (or optionally, more) local searches and we store the local optimum x^* . Then, until k does not exceed a pre-set k_{\max} , we iteratively select new starting points \tilde{x} in an increasingly larger neighbourhood $N_k(x^*)$ and start new local searches from \tilde{x} leading to local optima x' . As soon as we find a local optimum x' better than x^* , we update $x^* = x'$, re-set $k = 1$ and repeat. Otherwise the algorithm terminates.

For each $k \leq k_{\max}$ we consider hyper-parallelepipeds $H_k(x^*)$ proportional to $x^L \leq x \leq x^U$, centered at x^* , whose sides have been scaled by $\frac{k}{k_{\max}}$. More formally, let $H_k(x^*)$ be the hyper-parallelepiped $y^L \leq x \leq y^U$ where, for all $i \leq n$,

$$y_i^L = x_i^* - \frac{k}{k_{\max}}(x_i^* - x_i^L)$$

$$y_i^U = x_i^* + \frac{k}{k_{\max}}(x_i^U - x_i^*).$$

This construction forms a set of hyper-parallelepiped-shaped shells centered at x^* and proportional to $x^L \leq x \leq x^U$. As has been mentioned above, we define each neighbourhood $N_k(x^*)$ as $H_k(x^*) \setminus H_{k-1}(x^*)$.

3.2 The implementation

The conceptual simplicity of the proposed algorithm does not reflect a corresponding simplicity in the implementation. Our algorithm uses an VNS global optimization solver [11] for continuous nonconvex programming problems as a black-box. In turn, the VNS solver calls a local NLP solver (SNOPT [6]) as a black-box. Careful software architecture and code re-entrancy is required. Our implementation uses the *ooOPS* optimization framework library [14], which allows to easily formulate nonlinear programming problems and solve them with a variety of global and local optimization algorithms. The main adjustable parameters of our algorithm are: the k_{\max} terminating parameter for VNS and the number s of local searches in each VNS neighbourhood.

4 Computational results

In this section we report on the computational results, obtained on an Intel 2.66 GHz Pentium IV CPU with 1 GB RAM running Linux. The algorithmic parameters have been set to the following default values: $k_{\max} = 10$, $s = 1$. The initial value λ_0 for the smoothing parameter is chosen so that a reasonable number of close atoms are clustered together:

$$\lambda_0 = \frac{1}{|V|} \sum_{i \in V} \frac{1}{|\delta(i)|} \sum_{j \in \delta(i)} d_{ji}, \quad (12)$$

where $\delta(i)$ is the neighbourhood of vertices adjacent to vertex i in the graph G .

4.1 Instances

Our computational tests refer to two sets of instances: the “Moré-Wu” instances [15] and the “Lavor” instances [8].

The “Moré-Wu” instances are simply a cubic lattice with s^3 atoms ($s = 1, 2, 3, \dots$) defined by

$$\{(i_1, i_2, i_3) \in \mathbb{R}^3 : 0 \leq i_k \leq s - 1, k = 1, 2, 3\}. \quad (13)$$

See Fig. 3 for an example with $s = 3$. An order is defined for the atoms of the lattice by letting atom i be the atom at position (i_1, i_2, i_3) , where

$$i = 1 + i_1 + si_2 + s^2i_3, \quad (14)$$

and the set E is defined by

$$E = \{\{i, j\} : |i_3 - j_3| \leq 1\} \quad (15)$$

i.e. all distances to atoms on the same or adjacent i_3 -coordinate plane are considered. For example, for a molecule with 8 atoms ($s = 2$), the sequence of atoms is

$$\begin{aligned} x_1 &= (0, 0, 0), x_2 = (1, 0, 0), x_3 = (0, 1, 0), x_4 = (1, 1, 0), \\ x_5 &= (0, 0, 1), x_6 = (1, 0, 1), x_7 = (0, 1, 1), x_8 = (1, 1, 1), \end{aligned}$$

and the set E is the set of edges of the complete graph on eight vertices.

The “Lavor” instances, described in [8], are based on the model proposed by [18], whereby a molecule is represented as a linear chain of atoms. Bond lengths and angles are kept fixed, and a set of likely torsion angles is generated randomly based on an energy function including




Fig. 3 The $s = 3$ Moré-Wu instance with 27 atoms. Not all distances are drawn




Fig. 4 The `lavor11_7` instance

some Lennard-Jones potential terms. Depending on the initial choice of bond lengths and angles, the Lavor instances give rather more realistic models of proteins than the Moré-Wu instances do. We generated 10 different Lavor instances for each size $N = 10, \dots, 70$. These are called `lavorN-m`, where N is the number of atoms in the molecule and m is an instance ID (since there is a random element of choice in the generation of the Lavor instances, many different instances can be generated having the same atomic size). See Fig. 4 for an example.

4.2 Comparative results

In Table 1, we report on the comparative computational results (seconds of user CPU time and Largest Distance Errors) obtained by running the GCA and DVS algorithms on a selection of Moré-Wu and Lavor instances of molecular sizes in the range 8–70 atoms. For molecules with fewer than 50 atoms, the DVS algorithm improved on the straight VNS in only around half of the instances. For molecules with more than 40 atoms, the DVS improves on a straight VNS for the large majority of the instances.

In Table 2, we give arithmetic average values for user CPU time and Largest Distance Error for each sample of 10 Lavor instances. The average LDE is calculated considering only

Table 1 Computational results for a sample of Moré-Wu, and Labor instances

Instance			GCA		DVS	
Name	N	$ E $	CPU	LDE	CPU	LDE
mmorewu-2	8	28	0.02	2.63E+5	1.09	2.51E-8
mmorewu-3	27	331	0.23	6.99	27.05	2.25E-9
mmorewu-4	64	1882	0.67	7.79E-6	642.97	2.75E-10
lavor10_0	10	33	0.02	1.57E-5	2.76	1.55E-9
lavor15_0	15	57	0.10	4.04E-5	10.01	3.77E-9
lavor20_0	20	105	0.14	2.77E-5	18.14	2.68E-9
lavor25_0	25	131	0.84	1.18E-4	60.41	2.24E-9*
lavor30_0	30	169	0.40	1.75E-5	231.02	6.51E-9*
lavor35_0	35	171	0.81	9.33E-5	624.72	4.54E-3*
lavor40_0	40	295	2.84	0.096	770.78	2.46E-5
lavor45_0	45	239	3.33	0.170	538.25	3.13E-4
lavor50_0	50	271	3.45	0.696	971.79	8.66E-6
lavor55_0	55	551	5.80	0.257	870.50	1.13E-8*
lavor60_0	60	377	5.15	0.049	1800.35	2.85E-4
lavor65_0	65	267	2.61	0.065	1119.82	3.94E-3
lavor70_0	70	431	8.73	0.107	2165.81	4.97E-4

LDE values marked with * have been found by the pre-processing VNS run on the unrestricted original problem (i.e. the smoothing run did not improve the value)

Table 2 Average statistics for Labor instances (taken over 10 instances for each molecular size)

Instance	GCA/avg.		DVS/avg.		
	CPU	LDE	CPU	LDE-st	Unsuitable
10	0.03	4.40E-01	2.81	3.04E-9	0
15	0.08	1.96E-02	10.07	3.34E-9	0
20	0.23	3.20E-03	22.01	3.59E-9	1
25	0.56	1.58E-02	46.19	3.64E-5	1
30	0.65	1.03E-02	276.40	2.94E-5	2
35	1.10	5.43E-02	465.72	2.37E-3	1
40	1.41	2.61E-02	486.64	4.03E-4	1
45	2.13	5.80E-02	752.93	1.69E-3	1
50	2.54	1.65E-01	863.14	1.06E-3	0
55	4.10	7.29E-02	762.78	4.03E-4	0
60	4.47	1.59E-01	2172.23	8.06E-4	0
65	4.64	1.16E-01	1404.57	1.97E-3	0
70	7.63	9.28E-02	1912.97	4.12E-4	0

those instances in the sample for which DVS calculated a suitable solution (we consider a solution clearly unsuitable if its associated LDE exceeds 0.01, as such a large LDE usually indicates the wrong 3D structure). We therefore included two columns in Table 2, labelled “LDE-st” and “Unsuitable” that report: the arithmetic average of the LDE for instances with $LDE < 0.01$, and the number of unsuitable instances in each 10-sample.

As can be easily seen from the results, DVS outperforms the GCA in accuracy, whilst the GCA is superior to the DVS in terms of computation times. Since usually looking for the 3D structure of a molecule is a task where accuracy is more important than short CPU times, we feel these computational results validate the soundness of the proposed approach.

5 Conclusion

In this paper we presented an algorithm called Double VNS with Smoothing (DVS) used to solve the Molecular Distance Geometry Problem. The DVS is based on VNS for global optimization problems and a smoothed version of the problem. We tested this algorithm on two classes of problems from the literature. It turns out that the DVS algorithm finds solutions with high accuracy, compared to the Global Continuation Algorithm. This is important insofar as a small error in the objective function may lead to a completely different molecular structure.

Acknowledgements Some of the authors (CL, NM, LL) would like to thank FAPESP and CNPq for their financial support.

References

1. Crippen, G.M., Havel, T.F.: Distance Geometry and Molecular Conformation. Wiley, New York (1988)
2. Cruz, I.F., Twarog, J.P.: 3D graph drawing with simulated annealing. In: Brandenburg, F.-J. (ed.) 3D Graph Drawing – GD95 Proceedings, LNCS, vol. 1027, pp. 162–165. Springer, Berlin (1996)
3. Dong, Q., Wu, Z.: A linear-time algorithm for solving the molecular distance geometry problem with exact inter-atomic distances. *J. Global Optim.* **22**, 365–375 (2002)
4. Dražić, M., Lavor, C., Maculan, N., Mladenović, N.: A continuous variable neighbourhood search heuristic for finding the tridimensional structure of a molecule. *Eur. J. Oper. Res.* (to appear)
5. Eren, T., Goldenberg, D.K., Whiteley, W., Yang, Y.R., Morse, A.S., Anderson, B.D.O., Belhumeur P.N.: Rigidity, computation, and randomization in network localization. *IEEE Infocom Proc.* pp. 2673–2684 (2004)
6. Gill, P.E.: User’s Guide for SNOPT 5.3. Systems Optimization Laboratory, Department of EESOR, Stanford University, California, February 1999
7. Klepeis, J.L., Floudas, C.A., Morikis, D., Lambris, J.D.: Predicting peptide structures using NMR data and deterministic global optimization. *J. Comput. Chem.* **20**(13), 1354–1370 (1999)
8. Lavor, C.: On generating instances for the molecular distance geometry problem. In: Liberti, and Maculan [13], pp. 405–414
9. Lavor, C., Liberti, L., Maculan, N.: Computational experience with the molecular distance geometry problem. In: Pintér, J. (ed.) *Global Optimization: Scientific and Engineering Case Studies.*, pp. 213–225. Springer, Berlin (2006)
10. Lavor, C., Liberti, L., Maculan, N.: An overview of distinct approaches for the molecular distance geometry problem. In: Floudas, C.A., Pardalos, P. (eds.), *Encyclopedia of Optimization*, 2nd edn, Springer (to appear)
11. Liberti, L., Dražić, M.: Variable neighbourhood search for the global optimization of constrained NLPs. In *Proceedings of GO Workshop, Almeria, Spain (2005)*
12. Liberti, L., Lavor, C., Maculan, N.: The discretizable molecular distance geometry problem. arXiv: q-bio/0608012, 2006.
13. Liberti, L., Maculan, N. (eds.): *Global Optimization: From Theory to Implementation.* Springer, Berlin (2006)
14. Liberti, L., Tsiakis, P., Keeping, B., Pantelides, C.C.: *ooOPS*. Centre for Process Systems Engineering Chemical Engineering Department, Imperial College, London, UK (2001)
15. Moré, J.J., Wu, Z.: Global continuation for distance geometry problems. *Siam J. Optim.* **7**(3), 814–846 (1997)
16. Moré, J.J., Wu, Z.: Distance geometry optimization for protein structures. *J. Global Optim.* **15**, 219–234 (1999)
17. Neumaier, A.: Molecular modeling of proteins and mathematical prediction of protein structure. *SIAM Rev.* **39**, 407–460 (1997)
18. Phillips A.T., Rosen J.B., Walke V.H.: Molecular structure determination by convex underestimation of local energy minima. In: Pardalos P.M., Shalloway D., Xue G. (eds) *Global Minimization of Nonconvex Energy Functions: Molecular Conformation and Protein Folding*, vol 23, pp. 181–198 American Mathematical Society, Providence (1996)

19. Saxe, J.B.: Embeddability of weighted graphs in k -space is strongly NP-hard. Proceedings of 17th Allerton Conference in Communications, Control and Computing, vol. 23; pp. 480–489 (1979)
20. Yoon, J.-M., Gad, Y., Wu, Z.: Mathematical modeling of protein structure using distance geometry. Technical Report TR00-24, Dept. Comput. Applied Maths, Rice University, Houston (2000)