



FI MU

Faculty of Informatics
Masaryk University Brno

Haptically Driven Travelling Through Conformational Space

by

Igor Peterlík
Aleš Křenek

FI MU Report Series

FIMU-RS-2005-02

Copyright © 2005, FI MU

January 2005

**Copyright © 2005, Faculty of Informatics, Masaryk University.
All rights reserved.**

**Reproduction of all or part of this work
is permitted for educational or research use
on condition that this copyright notice is
included in any copy.**

**Publications in the FI MU Report Series are in general accessible
via WWW:**

<http://www.fi.muni.cz/veda/reports/>

Further information can be obtained by contacting:

**Faculty of Informatics
Masaryk University
Botanická 68a
602 00 Brno
Czech Republic**

Haptically Driven Travelling Through Conformational Space

(extended version)

Igor Peterlík and Aleš Křenek

Faculty of Informatics, Masaryk University, Brno, Czech Republic

xpeterl@fi.muni.cz, ljocha@ics.muni.cz

January 7, 2005

Abstract

After reviewing our generic framework of building haptic interactive models of environments which are computationally very expensive (so that their simulation cannot be used directly) we describe in detail the deployment of the framework in modelling conformational behaviour of flexible molecules. We conclude with description of first practical results which prove correctness of our approach, and shift the whole research towards more realistic applications.

This report is an extended version of a paper accepted for presentation at the First Joint Eurohaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems, Pisa, Italy, 2005.

1 Introduction

Flexibility is a common property of virtually all molecules occurring in living organisms. A flexible molecule is able to adopt multiple spatial configurations (shapes), called *conformations*, while it does not change its chemical structure. Both experimental and theoretical results show that a molecule does not occur in a single stable shape but its population is scattered among many such shapes. On the other hand, biological activity of a molecule strongly depends on its conformation—different conformations of the same molecule may exhibit completely different biological activity. This is the main reason of studying *conformational behaviour*, the process of interchanging conformations.

One of the most challenging tasks of computational chemistry is the analysis of *molecular docking* — an interaction of a small molecule with a large one (e. g. an enzyme). However, such interaction may occur only if both the molecules adopt proper conformations. It is believed that understanding these processes is an essential step in understanding the life at its chemical level. More pragmatically, molecular docking problems are also essential in drug design.

Probably the most important quantity involved in conformational behaviour is the potential energy of the molecule. On the other hand, the human operator of an interactive application is aware of the mechanical work done with a haptic device on the virtual environment. If the work is directly transformed into changes of potential energy of the modelled virtual environment (which is expected by the user anyway due to her real-world experience), the user is able to keep notion of the energy even over fairly long time. Therefore the haptic human-computer interaction is intrinsically suitable for delivering information on potential energy.

2 Previous Work and Our Goals

Analysis of conformational behaviour has a long tradition in the field of computational chemistry. Several methods of finding a global view on behaviour of a molecule were proposed, a good overview can be found in [17]. Virtually all the methods produce a graph of conformational interconversions. Vertices of the graph are individual shapes of the molecule, either stable conformations or *transition states* — saddle points on lowest-energy paths between conformations. All vertices carry information on corresponding potential energy and other quantities. Edges represent possible interconversions between the shapes. Figure 3 shows an artificial example of such graph projected on a potential energy surface of the molecule (PES for short, a graph of the potential energy w. r. t. some set of independent variables describing the shape of the molecule). For our purposes we take the conformational graph as an input, not regarding the method with which it was computed. In practical experiments we used outputs of the CICADA family of programs [17]. At this level of view, the graph of a docking process does not differ from a conformational graph — its vertices represent geometries of both molecules, including their relative position.

Visualisation of results of chemical calculation is another area with a long tradition. Many specialised molecular visualisation packages like Insight¹ or VMD² were developed over time, molecular visualisation is also a standard subject of generic visualisation tools like OpenViz³ or

¹<http://www.accelrys.com/>

²<http://www.ks.uiuc.edu/Research/vmd/>

³<http://www.openviz.com>

IRIS Explorer⁴. Presenting molecular models is also a well-established application of immersive virtual environments, e. g. [6, 1, 2]. Therefore we do not concentrate on the visualisation aspect in our work. Currently we use only a simple ad-hoc visualisation through OpenGL, planning to deploy more sophisticated methods later.

On the contrary, only few haptic interaction applications were reported in this area. Most of the work concentrates on using haptic devices as advanced tools for manipulation with virtual objects (selection, rotation, positioning etc., e. g. the VRPN interface in the VMD package), and haptic rendering or haptic examination of molecular models treated as macro-world objects, e. g. [4]. Direct mapping of the underlying chemical forces to the haptic display is very rare. We believe that the main reason is in the bad match of the high computational complexity of chemical calculations and extremal requirements (high refresh rate) of haptic device driving. Probably the most comprehensive project with focus in haptic molecular modelling was GROPE [7]. They report a general success in adding force-feedback cues to a relatively complex specific docking task (fitting a given flexible drug molecule into a particular active site of a protein). The system built with the Argonne ARM [9] force feedback device achieved the refresh rate of about 80 Hz (reported to be sufficient for this type of device), using 3D grid of pre-computed intra-molecular forces [8].

We aim our research to the haptic examination of the conformational behaviour in short term and molecular docking in longer term, looking for alternate approaches in order to achieve higher quality and fidelity of the interaction, as well as extend the model to get a more global view on the underlying chemical processes. In particular, we focus on the following topics:

- Keep pace with the high refresh rate (over 1 kHz) of the current haptic devices. The high refresh rate allows to model the interaction with more detail but it also imposes extremal requirements on the speed and latency of the model evaluation.
- Drive the interaction strictly by the results of the systematic chemical calculation so that all important topological features (minima, saddle points, and valleys) of the potential energy are reflected properly.
- Allow more global view on the modelled system, i. e. map the entire conformational space, or in the case of the docking task, let the user *search for the active site* on the surface of the large molecule rather than assuming that it is known a-priori and just examining its properties. Eventually, we also want to cover flexibility of the large molecule and mutual

⁴http://www.nag.co.uk/Welcome_IEC.html

interaction of conformational changes of the two molecules (these calculations are being developed in CICADA).

Previously we developed a fast interpolation algorithm among conformations of a molecule [12, 13, 11]. Given convertible conformations of a molecule it is able to compute a continuous morphing between these shapes, heuristically obeying certain energetic constraints so that the morphing does not contradict general chemical intuition.

We attempted to drive the molecule across the conformational space with a “gear lever” paradigm—the haptic device attached to a group of atoms—but we found this approach too restrictive for an efficient interaction [10].

Because of the drawbacks of the previous model we introduced the concept of hybrid energy (see Sect. 3.2) driving an interaction on molecular surface. We demonstrated its applicability on a very simple model of flexible molecule [14].

In [15] we present a generic framework for off-line calculation of complex force fields which are then used in haptic interaction. We used a model of a vicinity of a single conformation as a prototype application of the framework. We also developed and evaluated an efficient distributed algorithm for the off-line calculation [16].

In this paper we extend this work one step further, allowing the user to “flip” the molecule among all its conformations.

3 Interaction on Molecular Surface

The model of interaction on the molecular surface described in this paper uses the generic interaction framework we introduced in [15]. In that work the mapping of the model of conformational behaviour to the interaction was outlined only. We also extended it considerably in order to be able to span over the entire conformational space eventually. Therefore, for the sake of readability of this paper, we describe the the model completely here.

As we use PHANToM⁵, a representant of *impedance control loop* (measure position, apply force) haptic device [5], all the reasoning in our work assumes this concept.

⁵<http://www.sensable.com>

3.1 Perceived Model Behaviour

Internally, the molecule is represented with a skeleton of almost stiff sticks connected with flexible joints. Conformational changes of the molecule are computed as changes of this skeleton. In vast majority the changes are composed of “twisting” several bonds.

Then the skeleton is wrapped with a *van der Waals* surface — spheres centred at atomic nuclei. The radii of the spheres are big enough to make the spheres overlap one another and to cover the skeleton completely. Union of the externally accessible parts of the spheres forms the molecular surface in turn (see Fig. 1).




Figure 1: Van der Waals surface of the alanine amino acid

The haptic device manipulated by the user is attached to a *probe*, another sphere of approximately water molecule size. When not in contact with the molecule, the probe can be moved in the three-dimensional space freely, but it is not allowed to penetrate the molecular surface. On the contrary, interaction with the surface forces the molecule to change its shape. The force required to do so is delivered back to the user via the force feedback of the haptic device.

3.2 Hybrid Energy Concept

In order to describe the behaviour of the whole model we introduce the concept of *hybrid potential energy*. The purpose of this quantity is twofold (hence the term “hybrid”):

- cover the desirable quantity that has to be delivered to the user, e. g. in the case of conformational behaviour, the potential energy of the particular molecular configuration per se,
- implement the user interaction with the virtual model in real world terms, e. g. do not allow stiff objects to penetrate each other etc.

The hybrid potential energy $H(\mathbf{u}, \mathbf{s})$ is defined as a function of two vectors

- *user input variables* \mathbf{u} describing the users input, i. e. the probe position,

- *free state variables* s that represent degrees of freedom of the model, i. e. the shape and position of the molecular model.

Currently the energy is potential only, the concept does not consider dynamic properties, e. g. kinetic energy (or temperature), velocity, etc. For our application this is not a restriction — we model also the potential energy of the molecule. A concrete definition of the hybrid energy of our model is discussed in Sect. 4.4.

3.3 Model Equations

In order to describe the interaction, we are looking for response of the model, described in terms of the state vector s , w. r. t. the user’s input (haptic device position) u . However, a straightforward mapping $u \mapsto s$ is not sufficient; there may be multiple model states corresponding to a unique device position (see an example in Fig. 2 and [15] for deeper discussion). Instead of the isolated position we have to consider the context (or the history of the user’s input) somehow.




Figure 2: When approached from opposite directions the hung ball (white) is shifted to different positions (system states), despite the final position of the probe (dashed) is the same in both cases.

We have to be aware of the computational complexity of the model (namely the hybrid energy function), and, at the same time, the required refresh rate of the haptic device, imposing a strict limit on the time available to evaluate the model response to a user action. In general, we cannot rely on being able to perform the hybrid energy computation within the haptic loop. These computationally intensive evaluations have to be done in advance instead.

The most general approach to capture time context of the interaction is deploying a completely dynamic model. We would consider a function mapping time t to the haptic device position $u(t)$. Then we would look for the response of the modelled system, i. e. a corresponding function $t \mapsto s(t)$. Due to the second thermodynamic law the system tends to minimise its potential energy spontaneously hence its behaviour can be described in terms of the following system

of differential equations:

$$\frac{d^2\mathbf{s}(t)}{dt^2} = -A(\nabla_{\mathbf{s}}H)(\mathbf{u}(t),\mathbf{s}(t)) - B\frac{d\mathbf{s}(t)}{dt} \quad (1)$$

where $\nabla_{\mathbf{s}}$ is the standard gradient operator restricted to the \mathbf{s} vector components only (i. e. omitting partial derivatives corresponding to \mathbf{u} components), A is an acceleration constant (representing an inverse of “mass” of the system), and B represents friction or viscosity of the environment (generally required to prevent unstable oscillating behaviour). In order to get $\mathbf{s}(t)$ the equations have to be integrated over time, taking into account the actual user’s input (haptic device position). However, besides potential stability problems of numeric integration we currently can’t see any way of decoupling the actual user’s input from the model behaviour in the calculation, therefore we are not able to perform the heavy-weight calculation in advance (more discussion in [11]). Therefore we restrict the dynamic behaviour of the model according to the following assumptions:

- The model reacts fast compared to the speed of user’s action. Therefore it is always able to completely relax, i. e. reach a stable still state.

Also within the domain of conformational behaviour, we model the *potential* hypersurface of the molecule, i. e. we do not directly consider dynamic properties. Hence we are interested in the still states only in the haptic interaction too. Modelling dynamic properties of the system would be artificial and misleading from the application point of view. Therefore this restriction is completely feasible.

- The behaviour of the model does not depend on its entire history but on a limited period only.

This is a consequence of the previous restriction. As soon as the model reaches a still state, its further behaviour depends on this state and the user action; all the previous history becomes irrelevant.

Under these assumptions we can describe the behaviour in an alternate way in terms of a *conversion function* $C(\mathbf{s}_0, \mathbf{u}_0, \Delta\mathbf{u})$. The function maps an initial state \mathbf{s}_0 , corresponding user input \mathbf{u}_0 , and a change of the input (i. e. the user’s action) $\Delta\mathbf{u}$ into a new still state corresponding to the the user input $\mathbf{u}_0 + \Delta\mathbf{u}$.

A direct way of mapping the dynamic model (Eq. 1) into this framework is defining the user input time-dependent function artificially as converging from \mathbf{u}_0 to $\mathbf{u}_0 + \Delta\mathbf{u}$, i. e.:

$$\mathbf{u}(t) = \mathbf{u}_0 + (1 - e^{-t})\Delta\mathbf{u}, \quad (2)$$

and letting the system relax indefinitely, i. e.:

$$C(\mathbf{s}_0, \mathbf{u}_0, \Delta \mathbf{u}) = \lim_{t \rightarrow \infty} \mathbf{s}(t). \quad (3)$$

However, for pragmatic reasons (speed and stability of the involved numerical methods) we replace the integration of differential equations with direct minimisation of the hybrid energy function, starting from \mathbf{s}_0 w. r. t. the new user input $\mathbf{u}_0 + \Delta \mathbf{u}$. We provide more accurate reasoning in [11]; intuitively (i. e. with the experience of the 2nd thermodynamic law), the energy minimisation is natural. The system starts in a stable still state; on a user action the stability is disrupted, and the system follows a “downhill” path to reach another stable state.

3.4 State Space Discretisation Precomputation

The principal benefit of describing the model behaviour with the conversion function is its very limited dependence on the behaviour history (only a single state is required to compute the following one). Therefore we can discretise the entire state space of the model and precompute it before the user-driven interaction happens. The actual interaction interpolates among the pre-computed states only which is fast enough to be safely used within the haptic loop [15].

We sample the entire input space (haptic device position) with regular three dimensional grid⁶, and search it systematically, starting from the corner points to which we assign a-priori known “quiet” states. The conversion function is repeatedly evaluated, starting from a known state in some of the grid points, and using shifts to its direct neighbours for $\Delta \mathbf{u}$. For a single grid point multiple distinct states can be found (see Fig. 2 again)—in this approach our method is novel compared to [8] which stores only a single state in a grid point. The computation stops when all known states and all directions to neighbouring grid points were exhausted. In [15] we discuss the process in detail, including requirements on the conversion function to ensure convergence of the algorithm. Usage of the results within the haptic loop, including detailed description of the interpolation are also described there. In [16] we described an efficient distributed algorithm to perform the state space search.

⁶This is a current limitation of the method. We are restricted to three degrees of freedom of the interaction despite more may be appropriate. However, increase of this number yields an exponential growth of the state space which is infeasible due to both computational and storage requirements. We are looking for means of overcoming this restriction.




Figure 3: Potential energy surface of a fictitious molecule. The independent variables represent twisting of two bonds. Conversion paths between three conformations are shown.

3.5 Summary of the Interactive Model

In our haptic interactive model the molecule is represented with its van der Waals surface. The user manipulates a spherical probe with the haptic device. When the probe touches the surface the molecule changes its shape (according to its precomputed conformational behaviour), and the user feels the force required for the particular shape change.

Behaviour of the interaction is driven by the hybrid energy, a quantity encompassing both the ability of the molecule to undergo certain conformational changes, and the interaction phenomena (e. g. the probe cannot penetrate the surface).

Due to computational complexity of the hybrid energy evaluation the entire process is split into two stages: First, before the user-driven interaction starts, the entire discretised state space of the interaction is searched and precomputed. Then it is traversed with the haptic interaction, interpolating among precomputed discrete states in real time.

The interaction model described so far is quite generic — it can be used for other applications as well.

4 Geometry Parametrisation and Hybrid Energy Formula

In this section we focus on concrete definitions of the state space (i. e. the variables which form the state vector s of the system) and the hybrid energy formula.

4.1 Design Requirements

In order to implement the behaviour described in Sect. 3.1 the state vector of the system has to be defined in such a way that varying its value generates all relevant shapes of the molecule. Those are:

- All the considered conformations, i. e. vertices of the conformational graph.
- All intermediate shapes on the conformational paths between these conformations, i. e. along the edges of the conformational graph.

Due to the requirements on the hybrid energy to distinguish these paths from their vicinity (see below) the system state must distinguish them too (i. e. for a given value of s it has to be apparent whether it describes a point on a conformational path or out of it).

Finally, the generated geometry has to be continuous w. r. t. moving the state along a conformational path.

- “Out-of-path” shapes in vicinity of the conformations.

The conformations are “cross-roads”, there are multiple discrete paths originating from a single point there (C_1 in Fig. 3). As we showed in [10], it is rather difficult for the user to “hit” one of them if the model is restricted to move along these paths only. Instead, soft motion constraints have to be implemented—it has to be allowed to go “in-between” of more paths and to be attracted to one of them farther from their intersection.

On the other hand, once being farther from a conformation, the geometry can be restricted to the conformational path only. The purpose of the model is to examine these paths, not the in-between areas. For the same reason, we *do not* strictly require virtually any chemical feasibility of the “out-of-path” shapes, except of not contradicting the basic energetic constraints [12] already applied for the interpolated shapes along the paths. Those shapes play their role mainly in the energy minimisation process and almost do not appear as resulting states visible to the user.

The dimensionality of this “out-of-path” space is not, due to the complexity of the hybrid energy definition, directly mappable to the degrees of freedom of the interaction itself. As we did not investigate this relationship in deep yet, we require each of the conformational paths originating from a stable state to generate an independent dimension. It introduces certain computational overhead and may not be strictly necessary but it should be a safe approach.

There are additional requirements specific to the hybrid energy:

- Energetic properties of the conformational behaviour have to be preserved. Conformations as well as transition states have their potential energy. Changes of this energy when moving from one of these states to another have to be reflected exactly.
- The hybrid energy have to approximate the potential energy surface in its low-energy areas — namely the conformational paths have to be “valleys” on the energy graph, being separated from one another with higher-level areas.
- According to Sect. 3.2 physical constraints of the model have to be implemented, e. g. penetration of the molecular surface with the probe has to be penalised strongly.

The hybrid energy is subject to minimisation. Therefore we have to design its formula to be continuous in areas which are likely to be hit by the minimisation process, i. e. the vicinities of conformations and conformational paths. Continuity in more distant points is not strictly required. As certain terms of the hybrid energy formula are based on the geometry of the molecule (see Sect. 4.4) it is necessary that the geometry function is continuous w. r. t. the state vector in these areas too.

4.2 Interaction Restricted to a Single Conformation

In [15] we deployed a model where travelling of the molecule was restricted to a region given by one stable conformation and surrounding transition states. The shape of the molecule travelling within the region is computed from a *local context* consisting of geometries of the stable conformations and transition states and a vector of *morphing parameters*.

Given the shape of the conformation and a transition state, and a morphing parameter $m \in [0, 1]$ we are able to compute an intermediate shape such that the shape changes continuously from the conformation ($m = 0$) to the transition state ($m = 1$) [12, 11]. The morphing algorithm can be extended to consider more transition states and a corresponding vector of the morphing parameters. The resulting shape is a weighted merge of the shape changes from the central conformation to the transition state. Hence the central conformation and the several transition states can be understood as generators of a certain space of shapes of the molecule. The chemical feasibility of such intermediate shapes is very questionable, though. But according to the discussion on the “out-of-path” shapes in Sect. 4.1 they never become the final still states observed by the user (despite very close vicinity of stable conformations). Therefore they do not degrade the fidelity of the whole application.

In the following we refer to such multi-factor morphing function $morph(G_0, G_1, \dots, G_k, m_1, \dots, m_k)$ which takes geometry G_0 of some stable conformation, geometries G_i of surrounding transition states, corresponding morphing factors m_i , and computes the mixed geometry.

Within the first application [15] we used the morphing parameters directly as the components of the state vector s , applying to it the hybrid energy definition and minimisation procedure described in Sect. 4.4. We ended up with an interaction having a single stable state — the central conformation. The user could push the probe on the molecular surface, forcing it to change its shape towards one of the surrounding transition states. Once the pressure was released, the molecule relaxed to the initial stable state.

4.3 Global Geometry Parametrisation

We aim at extending the interaction from the local vicinity of a stable conformation, beyond the surrounding transition states. As long as we consider the edges of the conformational graph only, the extension is quite easy — we would keep the context and switch it to the vicinity of the other conformation once we reached the transition state. However, considering the high dimensionality of the “out-of-path” space, the process becomes more ambiguous when these states are involved, and the explicit context bookkeeping would require non-trivial overhead.

Further on we describe parametrisation of the geometry which satisfies the requirements given in Sect. 4.1 but which avoids an explicit context bookkeeping by including it in the parametrisation formulae. Instead of the morphing factors from the previous section we use “weights” of all stable conformations simultaneously. Informally, a weight w_i represents the contribution of the geometry of the associated stable conformation to the final shape of the molecule.

4.3.1 Formal definition of the weights

Let $C = \{C_1, \dots, C_n\}$ be a set of stable conformations of the examined molecule, $T = \{T_{ij}\}$ a set of transition states. For the sake of simplicity we consider at most one transition state T_{ij} for each pair (C_i, C_j) . This is not a principal restriction, all the following reasoning can be easily extended to allow multiple transition states between a given pair of conformations. Whenever we talk about conversion from C_i to C_j , we assume that the transition state T_{ij} exists in T .

Let $\mathbf{w} = (w_1, \dots, w_n)$, where n is the number of considered stable conformations, be the *vector of weights*. Domain of valid weights is defined by the following conditions:

$$(w1) \quad w_i \in [0, 1] \text{ for } i = 1, \dots, n$$

$$(w2) \sum_{i=1}^n w_i = 1$$

(w3) If there are 2 or more non-zero elements in \mathbf{w} let's M denote the index of the maximal one. Then for all $w_j > 0, j \neq M$ there exists the transition state $T_{Mj} \in \mathbf{T}$.

Informally, we don't allow two conformations to contribute to the state simultaneously if they are not directly connected.

(w4) Any other \mathbf{w} value is not valid.

4.3.2 Interpretation of the weights

In previous section, we defined the domain of the weight vector $\mathbf{w} = (w_1, \dots, w_n)$. In the following list we discuss reasons of the restrictions and provide certain interpretation to characteristic classes of the \mathbf{w} values. We suggest the reader to watch Fig. 3 throughout the following sections in order to understand the discussions on various positions w. r. t. the conformational graph more easily.

(a) Let $w_s = 1$ for some $s \in \{1, \dots, n\}$. Then, due to (w2), this is the only non-zero element of \mathbf{w} , and \mathbf{w} represents the stable conformation C_s exactly.

(b) Let $s, t \in \{1, \dots, n\}, s \neq t$ such that $w_s \geq w_t > 0$ and $w_j = 0$ for all other j 's. Then, due to (w3), there is a transition state $T_{st} \in \mathbf{T}$ between the stable conformations C_s and C_t , and \mathbf{w} represents the state on the path $C_s - T_{st}$, at the distance proportional to the ratio $w_s : w_t$. In particular, when $w_s = w_t = \frac{1}{2}$ (due to (w2)), \mathbf{w} represents the transition state T_{st} exactly.

(c) Let $w_B, w_{B_1}, \dots, w_{B_k}, k > 1$ be all the non-zero elements of \mathbf{w} , and $w_B > w_{B_j}$ for all $j = 1, \dots, k$. Due to (w3) there exist all the transition states T_{BB_j} .

We treat such \mathbf{w} value as an intermediate state within the vicinity of C_B , shifted from it in the direction $(w_{B_1}, \dots, w_{B_k})$. The conformational analysis does not give any direct cues on these states, and we do not interpret them to have a particular chemical meaning. The only purpose of introducing them is the user driven interaction which requires a continuous coverage of the local vicinities of "cross-roads" (see Sect. 4.1).

The three enumerated cases above cover the requirements given in Sect. 4.1. Hence we can proceed to the exact definition of the mapping of a valid \mathbf{w} value to a geometry of the molecule.

In all the cases we also identify a maximal weight and derive the matching state to be in the vicinity of the conformation corresponding to the maximal weight, considering only the

surrounding transition states. Hence the interpretation of the weight vector implicitly covers the local context switching from the previous section.

4.3.3 Mapping to the morphing factors

In Sect. 4.2 we briefly described the function $morph(G_0, G_1, \dots, G_k, m_1, \dots, m_k)$ which computes an intermediate geometry from geometries of a stable conformation, several transition states, and corresponding morphing factors. On the other hand, we defined the weight vector giving it an interpretation such that the state of the molecule described by any valid value of \mathbf{w} falls into the vicinity of a uniquely identifiable stable conformation. Consequently, it makes sense to look for a mapping of the \mathbf{w} values to a certain set of indices B, B_1, \dots, B_k and matching morphing factors m_1, \dots, m_k . Once those are given, we are able to use $morph$ to assign a concrete geometry to any valid value of \mathbf{w} .

Given \mathbf{w} , we choose B to be the index of the maximal component of \mathbf{w} (if there are more, the choice is arbitrary but fixed for the sake of uniqueness). We define the set of indices $\{B_1, \dots, B_k\}$ to be $\{j : T_{Bj} \in \mathbf{T}\}$, i. e. we consider all existing transition states connected to C_B . Then the mapping to the morphing factors is defined by the formula

$$m_i = \frac{2w_{B_i}}{w_B + w_{B_i}} \quad \text{for all } i = 1, \dots, k. \quad (4)$$

Due to the domain condition (w1) m_i is always within $[0, 1]$.

In the following list we review the classes of \mathbf{w} values as discussed in Sect. 4.1, focusing on the properties of the resulting geometry:

- (a) For a single non-zero element $w_B = 1$ all the morphing factors are set to zero, and the resulting geometry is exactly G_B due to the properties of $morph$.
- (b) For exactly two non-zero elements $w_s \geq w_t$ we get $B = s, B_i = t$ for some i . The only non-zero morphing factor is m_i , therefore the resulting geometry remains on the path $C_s - T_{st}$. In particular, when $w_s = w_t = \frac{1}{2}$, Eq. 4 yields $m_i = 1$, hence the resulting geometry is exactly the geometry of the transition state T_{st} .
- (c) Three or more non-zero elements of \mathbf{w} yield more than one non-zero m_i hence the resulting geometry falls between the conformational paths. The morphing factors are proportional to the corresponding weights, therefore, due to the properties of $morph$, the higher a particular weight w_{B_i} is, the more favoured is the geometry of the transition state T_{B_i} .

The function *morph* is continuous, as well as Eq. 4. Unfortunately, the continuity is not preserved across the implicit context switches — values of \mathbf{w} such that in an arbitrary ε -vicinity of \mathbf{w} there exist values yielding distinct index B according to the rules above. Except of the transition states (values of \mathbf{w} having $w_s = w_t = \frac{1}{2}$) which are guaranteed to yield the same geometry when approached from both sides, the geometry is generated from completely different set of base geometries, and there is no guarantee of continuity.

Currently we are investigating the possible impact of these discontinuities on the minimisation process. If it is proven to cause problems we foresee a solution in certain compression transformation of the resulting morphing parameters which would restrict the geometries closer to the stable conformations. Consequently, the discontinuities would be reduced to isolated points in high energy areas (hence not reached by the minimisation) only.

4.4 Hybrid Energy and Minimisation Process

According to Sect. 3.2 we denote hybrid energy with $H(\mathbf{u}, \mathbf{s})$. The vector $\mathbf{u} = (u_x, u_y, u_z)$ corresponds to a position of the probe P .

In the model with a single stable conformation the state vector \mathbf{s} includes the morphing factors $\mathbf{m} = (m_1, \dots, m_k)$ (see Sect. 4.2). For the purpose of adaptive fixing of the molecule in 3D space (Sect. 4.4.3) we also add other six degrees of freedom — rotation and translation. They are represented by variables $\mathbf{r} = (r_x, r_y, r_z)$ and $\mathbf{t} = (t_x, t_y, t_z)$ respectively and these variables are included in the state vector as well.

In the model with conformational interconversions the state vector \mathbf{s} consists of weight vector $\mathbf{w} = (w_1, \dots, w_n)$. As the adaptive fixing does not work in this case there are no additional degrees of the state vector.

In the following sections we describe computation of individual terms which contribute to the overall hybrid energy: approximation of the potential energy surface, probe interaction constraints, stabilisation in 3D space, and constraints (w1–w3, Sect. 4.3.1) on the domain of the weight vector. All the terms are associated their specific constants expressing the strength of contribution of these terms. We determine these values empirically, they differ among different test cases.

4.4.1 Paths between stable conformations

Pairs of stable conformations which can be converted to each other are connected by paths containing a corresponding transition state.




Figure 4: Function $f(x) = 3x^2 - 2x^3$ modelling conversion paths

According to Sect. 4.3.2 let w_B be the maximal weight component determining the local context. If there is exactly one morphing factor $m_{B_i} > 0$, then molecule is being converted from initial conformation C_B towards stable conformation C_{B_i} ; $m_i = 0$ yields exactly the geometry of C_B , $m_i = 1$ yields the transition state T_{BB_i} . Since the energy profile of this conversion has to be modelled smoothly, it is computed using the function $3x^2 + 2x^3$ (Fig. 4). It satisfies the following conditions:

- on $[0, 1]$ its local minimum is in 0 and local maximum in 1,
- the first derivative is defined and continuous,
- the first derivative in both 0 and 1 is 0, hence such functions can be “sewed” together smoothly in both conformations and transition states to form continuous valleys on approximated PES.

Although the profile of the paths modelled by the function given above differs from the real paths between stable conformations on PES, they represent the sufficient approximation for our purposes.

The hybrid energy of the molecule on the path between two stable conformations C_i and C_j — H_{val} is given by

$$H_{val}^i = c_{val}(3m_i^2 - 2m_i^3) \quad (5)$$

where m_{ij} is the only non-zero morphing factor and c_{val} is constant for this type of term in the hybrid energy formula.

4.4.2 Barriers among the paths between stable conformations

The paths connecting the stable conformations are separated by energy barriers which fill the space of intermediate states among the stable conformations and the paths.

The molecule gets into such a state if there are more than one non-zero morphing factors. We model the energy of the barrier as:

$$H_{bar} = c_{bar} \sum_{i,j|m_i>0,m_j>0} m_i m_j \quad (6)$$

where c_{bar} is a constant of this term of the hybrid energy.

Since the terms H_{bar} and H_{val} from 4.4.1 are added in the overall formula (Eq. 10), and $H_{bar} = 0$ if there is just one non-zero morphing factor (i. e. a point on a conformational path), the paths from the previous section become valleys on the hybrid energy graph, approximating the real PES.

Once again, let's emphasise that such a modelling of barriers and valleys is completely artificial, and it is connected directly with the real shape of the PES. However, for our purposes it is a sufficiently approximation.

4.4.3 Probe penetration and transformations of the entire structure

Using the function *morph* (see Sect. 4.2) and the vector of morphing factors, we compute a geometry corresponding to given morphing factors. This geometry is transformed eventually using rotation \mathbf{r} and translation \mathbf{t} .

Since the interaction is based on not allowing the probe to penetrate the molecule, a configuration of the system when the probe penetrates any of the atoms has to be penalised. The appropriate energy term is

$$H_{pen} = c_{pen} \sum_{A \in G | r_A + r_P - |AP| > 0} r_A + r_P - |AP| \quad (7)$$

where r_A and r_P are radii of the atom A and probe P respectively, $|AP|$ is distance between an atom A and the probe P and c_{pen} is a penalisation constant of the surface penetration.

In case of model with single stable conformation, we added other six degrees of freedom—rotations and translations along the coordinate axes. Since the molecule cannot float in space unrestricted (this would not allow any interaction), we make each atom be attracted towards its original position. This attraction is implemented with a virtual string between the original and the transformed position of the atom. Strength of the string depends on the distance between the atom and the probe—consequently atoms which are farther from the probe are attracted more strongly than those in direct contact with the probe.

The corresponding energy contribution is

$$H_{stab} = c_{stab} \sum_{A \in G} |AA_0|^2 (|AP| - r_A - r_P) \quad (8)$$

where r_A , r_P and $|AP|$ are given above, and $|AA_0|$ is distance from the original position.

Unfortunately, this adaptive stabilisation does not work for the other model. If the atoms were always attracted to the positions in the initial conformations, the attraction would interfere with the shape change too strongly. Hence, for the time being, we omit the six degrees of freedom and fix the molecule unconditionally on a set of atoms specified by the user. We are looking for re-introducing more sophisticated adaptive stabilisation in the same time.

4.4.4 Constrains of the weight vector

In Sect. 4.3.1 we defined the constrained domain of weight vector \mathbf{w} . We keep the vector in the defined domain by adding *penalty terms* [3, Chapter 4] to the hybrid energy whenever a \mathbf{w} value violates a domain constraint. This may result in slight constraint violation but we did not observe any impact in our practical results. On the other hand, we benefit from avoiding more complex constrained optimisation numerical methods. The energy term corresponding (w1) ($w_i \in [0, 1]$) and (w3) (contribution of directly connected conformations only) is

$$H_{dom} = c_1 \sum_{i|w_i \notin [0,1]} |w_i| + c_3 \sum_{i|T_{B_i} \notin T} w_i w_j \quad (9)$$

The remaining constraint (w2) is implemented by omitting one of the \mathbf{w} components from the state vector, and calculating its value according to (w2). This also improves behaviour of numeric optimisation methods which tend to misbehave when the free variables are dependent on one another.

4.4.5 Total hybrid energy

The hybrid energy of the entire system is computed using terms defined above as:

$$H_v = H_{bar} + H_{val} + H_{pen} + H_{stab} + H_{dom} \quad (10)$$

This energy is being minimised for particular values of the parameter $\mathbf{u} = (u_x, u_y, u_z)$ w. r. t. variables \mathbf{m} , \mathbf{r} and \mathbf{t} in case of model with one stable state and variables \mathbf{w} in case of the model with conformational interconversions.

5 Implementation and Results

For both the models (single and multiple conformations), the prototype implementation consists of precomputation (computations of the state vector at 3D grid) and interaction (haptic rendering and visualisation).

Due to the computational complexity the precomputation was implemented in parallel environment using MPI (we provide quantitative assessment in [16]). The hybrid energy function was minimised using the Powell optimisation routine taken from [18].

The two models differ only slightly in the hybrid energy function (precomputation phase), and interpretation of the state vector (interaction phase).

The interaction was implemented using several VR devices. The most important part—the haptic rendering—is done with PHANToM. It allows the user to manipulate the virtual probe as well as to feel the force required for particular shape changes.

The molecule with the probe are visualised in OpenGL 3D display. For better depth perception, namely the mutual position of the probe and the molecule where no other spatial cues are available, the scene is displayed on a large stereoscopic projection.

Moreover, we enriched the application with several interactively controlled transformations of the entire scene—rotations, translations, and zooming which allow the user to focus on any part of the molecule. For this purpose, we use other VR devices—motion tracking system *Nest of Birds*⁷ and *Pinch glove*⁸. The user, wearing the pinch glove with a mounted motion sensor on her non-dominant hand, is able to perform the desired transformations using natural gestures any time during the interaction while holding PHANToM in her dominant hand all the time.

The first test case of single-conformation application was one stable conformation and two transition states of alanine amino acid (22 atoms). During the interaction, the user could manipulate the probe and to force the molecule to change its shape towards both the transition states. According to the model, the molecule tends to get back to the one central conformation. We achieved similar results with a more complex molecule of a 77-atom peptide.

For testing the multi-conformation setup we used alanine again. We computed the state space for distinct pairs of three stable conformations. In all cases the user was able to flip the molecule from one stable state to another. Unfortunately, due to certain numerical stability problems we are currently not able to configure the application to involve all the conformations simultaneously. Anyway, the practical success in flipping the molecule between stable states proves the basic correctness of our approach.

⁷<http://www.ascension-tech.com/products/nestofbirds.php>

⁸<http://www.vrealities.com/pinch.html>

6 Conclusion and Future Work

In this paper we described in detail the application of our generic haptic interaction framework [15] in the area of modelling conformational behaviour of flexible molecules. In this way we are able to create highly responsive model of behaviour which otherwise takes rather long time (incomparable to the speed of interaction) to simulate numerically. Therefore this work validates the generic framework.

From the point of view of the application, this work represents a significant qualitative change—instead of examining a local vicinity of a single conformation (as it was done in our previous work) the new model lets the user “flip” the molecule over an energetic barrier to another stable state. In principle this allows to examine the entire conformational space of the molecule. A practical success of doing so with a simple test molecule proves the feasibility of the approach.

In short term we are going to focus on improving numerical stability of the state space calculation which is necessary for performing tests on more data. In longer term we will look for extensions of the model in order to apply it to molecular docking.

Besides the chemical applications we also foresee using the generic framework in the area of modelling soft tissues.

Acknowledgement

This work was supported by the research programme of the Ministry of Education of the Czech Republic no. MSM-143300003. Special thanks to the project MetaCentre (within the research programme no. MSM-000000001) for providing the necessary computing resources. We also thank prof. Luděk Matyska for his valuable comments to a draft of this manuscript.

References

- [1] Z. Ai and T. Frölich. Molecular dynamics simulation in virtual environments. *J. Eurographics Assoc.*, 17(3):267–273, 1998.
- [2] A. Anderson and Z. Weng. VRDD: applying virtual reality visualization to protein docking and design. *J. Mol. Graph.*, 17, 2000.
- [3] A. C. Bajpai, L. R. Mustoe, and D. Walker. *Advanced Engineering Mathematics*. John Willey and Sons, 1977.

- [4] A. Breckenridge and B. Hamlet. Using haptics interaction in bioinformatics application. In *Proc. PHANTOM Users Group*, 2003.
- [5] C. R. Carignan and K. R. Cleary. Closed-loop force control for haptic simulation of virtual environments. *Haptics-e*, 1(2), 2000. <http://www.haptics-e.org/>.
- [6] D. Levine et al. STALK: an interactive virtual molecular docking system. *IEEE Computational Science & Engineering*, 4, 1997.
- [7] F. P. Brooks et al. Project GROPE – haptic displays for scientific visualization. *Computer Graphics*, 24(4), 1990.
- [8] N. Pattabiraman et al. Computer graphics and drug design: Real time docking, energy calculation and minimization. *J. Comp. Chem*, 6:432–436, 1985.
- [9] R. C. Goertz et al. The ANL model 3 master slave manipulators—its design and use in a cave. In *Proc. 9th Conference on Hot Laboratories and Equipment*, 1961.
- [10] Z. Kabeláč and A. Křenek. Studying conformational behaviour with PHANTOM device(s). In *Proc. PHANTOM Users Research Symposium*, Heidelberg, Germany, 1999.
- [11] A. Křenek. *Towards Interactive Molecular Models*. PhD thesis, Faculty of Informatics, Masaryk University in Brno, Czech Republic. In preparation.
- [12] A. Křenek. An algorithm on interpolating between two shapes of molecule. In W. Straßer, editor, *Proc. SCCG '97*, pages 77–84. Comenius University, Bratislava, 1997.
- [13] A. Křenek. Handling cycles in conformational behaviour visualization. In *Proc. WSCG'00*, 2000.
- [14] A. Křenek. Haptic rendering of molecular flexibility. In M. Harders and S. Huber, editors, *Proc. PURS 2000*, pages 19–26, 2000. ISBN: 3-89649-579-8.
- [15] A. Křenek. Haptic rendering of complex force fields. In *Proc. 9th Eurographics Workshop on Virtual Environments*, pages 231–240. Zurich, Switzerland, 22–23 May 2003. ISBN 3-905673-00-2.
- [16] A. Křenek, I. Peterlík, and L. Matyska. Building 3D state spaces of virtual environments with a TDS-based algorithm. In *Proc. Recent Advances in Parallel Virtual Machine and Message Passing Interface*, pages 529–536. Springer-Verlag, 2003.

- [17] J. Koča. Traveling through conformational space: an approach for analyzing the conformational behaviour of flexible molecules. *Progress in Biophys. and Mol. Biol.*, 70:137–173, 1998.
- [18] W. H. Press et al. *Numerical Recipes in C*. Cambridge University Press, 2002.