Comparison of stochastic optimization methods for receptor ligand docking

H. Merlitz, W. Wenzel *

Forschungszentrum Karlsruhe GmbH, Institut fur Nanotechnologie, Postfach 3640, D 76021 Karlsruhe, Germany

Abstract

We compare the efficiency of three stochastic optimization methods, simulated annealing, parallel tempering and stochastic tunneling to locate the global minima of complex and rugged potential energy surfaces arising from atomistic models for receptor ligand docking. The stochastic tunneling method proves to be the most efficient generic approach for atomistic receptor ligand docking in the rigid ligand rigid receptor approximation.

1. Introduction

The development of methods to efficiently determine the global minima of complex and rugged energy landscapes remains a challenging problem with applications in many scientific and technological areas.

In particular for NP-hard [1,2] problems, straightforward enumerative as well as sophisticated branch-and-bound techniques become prohibitively expensive and stochastic methods offer the only acceptable compromise between the computational cost of the method and its reliability. In such techniques the global minimization is accomplished by the simulation of a fictitious dynamical process for a 'particle' on the multi-dimensional potential energy surface (PES). This particle explores the potential energy surface in a biased random walk that is designed to guide it to low-energy regions.

In one of the most widely used methods, the simulated annealing technique (SA) [3], the PES is explored in a Monte Carlo simulation with gradually decreasing temperature. The final temperature is chosen small enough to constrain the dynamics of the particle to the immediate vicinity of the nearest local minimum of the PES, while the largest temperature must be sufficiently high to allow an essentially random search of the PES. On rugged PES this strategy routinely fails, because transition states between adjacent local minima are too high to be overcome at temperatures sufficiently small to resolve the energy differences between them.

Here we investigate a family of methods that have sought to address this generic deficiency of SA by allowing the dynamical process to pass through thermodynamically inaccessible regions of

* Corresponding author. Fax: +72 47 82 6434.
E mail address: wenzel@int.fzk.de (W. Wenzel).

the PES [4]. Generalized ensemble techniques such as simulated tempering [5,6], parallel tempering (PT) [7] and multi-canonical algorithms [8,9], have been proposed to address this problem. In a previous report, we have introduced the stochastic tunneling method [10] (STUN) as an alternate approach that avoids both the numerical cost of multiple simulations as well as the estimation of auxiliary control fields. STUN was successfully applied to peptide folding [11], as were PT [12] or SA [13]. Here we compare this method with PT and SA in application to the receptor ligand docking (RLD) problem in an atomistic model. We find STUN to be an efficient and reliable docking procedure, that succeeded even in cases where neither PT nor SA were able to locate the global minimum in repeated simulations.

2. Receptor–ligand docking

In the RLD problem, suitable ligands must be selected for a given, structurally characterized receptor [14,15]. One approach is to screen large chemical databases in-silico for suitable ligands. For each ligand the best possible fit between ligand and receptor must be determined. Even in the most simple atomistic model, where both protein and ligand are treated as inflexible molecules, efficient numerical techniques to screen large databases in any reasonable timeframe are still lacking. The reason for this difficulty lies in the competition between two vastly different energy scales in the problem, where steric repulsion competes with attractive electrostatic forces and hydrogen bonding to determine the global minimum of the PES. The tight fit between receptor and ligand (key lock-principle) complicates the optimization problem significantly because it is almost impossible to reorient the ligand within the receptor, while there are few specific interactions between ligand and receptor outside the receptor pocket. For the simulations discussed below we used a scoring function:

$$S = \sum_{\text{Protein}} \sum_{\text{Ligand}} \left(\frac{R_{ij}}{r_{ij}^{12}} - \frac{A_{ij}}{r_{ij}^{6}} + \frac{q_i q_j}{\varepsilon r_{ij}} \right), \tag{1}$$

which contains the empirical Pauli repulsion (first term), the Van de Waals attraction (second term) and the electrostatic Coulomb potential (third term). Neither solvation effects nor dielectric screening were used in the simulations because such terms alter the specifics of the affinity of a given ligand to the receptor, but not the nature of the optimization problem. We note, however, that the inclusion of such terms generally leads to nondifferentiable potentials, which are not amenable to optimization methods such as basin hopping [16]. The ligands are simulated as rigid bodies, there are six degrees of freedom in the simulations. In cases where rotatable bonds exist, the X-ray crystallographic structures of the docked ligands were taken from the PDB database. The force field parameters R_{ij} and A_{ij} are taken from the OPLSAA force field [17] and the scoring function is pre-calculated on grids. The atomic affinity grids are interpolated using a logarithmic interpolation technique [18].

3. Methods

In SA the PES is explored in a stochastic process, where a given configuration with energy E_1 is modified and the new configuration with energy E_2 is accepted with probability:

$$p = \exp(-\beta(E_1 - E_2)), \tag{2}$$

where $\beta = 1/k_{\mathrm{b}}T$ is an inverse fictious temperature. For fixed β the simulation samples the configuration space with a thermodynamic equilibrium distribution corresponding to this temperature. In SA, the temperature is gradually reduced to zero according to a cooling schedule, so that the system should end in principle in its ground state, or global energy minimum at the end of the simulation. For rugged or glassy PES, however, ergodicity is routinely lost in the cooling process, which 'freezes' in some metastable state.

For such PES, the freezing problem may be circumvented by allowing a trapped particle to escape from a local minimum by increasing the temperature of its simulation. Following this idea the parallel tempering method replaces the unidirectional cooling of SA by a set of concurrent

simulations at different temperatures $\{T_i | i = 1 \ldots n\}$, which occasionally exchange configurations with probability:

$$p = \exp(-(\beta_1 - \beta_2)(E_1 - E_2)), \tag{3}$$

where β_i and E_i $(i = 1,2)$ are the inverse temperatures and energies of the two simulations/configurations, respectively. This mechanism permits each particle to alternate between low-temperature simulations where only the closest local minimum is explored and high-temperature simulations where it diffuses freely across potential barriers. The specific choice in Eq. (3) allows all simulations to remain in thermal equilibrium so that thermal averages can be computed at a variety of temperatures simultaneously (detailed balance). Compared to straightforward SA, PT incurs an n-fold increase in cost for a given total simulation length. On rugged or glassy PES, however, where the escape time from a given local minium can be exponentially long, this overhead may be more than compensated for. Recently a number of methods have been proposed that provide similar mechanisms by generalizing the Monte Carlo method [5,6,12] to simulate ensembles other than the canonical. However, the efficiency of at least some of these techniques has been questioned for glassy PES [19].

The stochastic tunneling method [10] incorporates this ability by letting the particle in the minimization process 'tunnel' forbidden regions of the PES. As in MC and SA, we retain the idea of a biased random walk, but apply a non-linear transformation to the potential energy surface:

$$E_{\mathrm{STUN}}(x) = 1 - \exp[-\gamma(E(x) - E_0)], \tag{4}$$

where E_0 is the lowest minimum encountered by the dynamical process so far. Alternately a suitable upper bound for the global minimum can be used for E_0. This effective potential preserves the locations of all minima, but maps the entire energy space from E_0 to the maximum of the potential onto the interval $[0, 1]$. At a given finite temperature of O(1), the dynamical process can therefore pass through energy barriers of arbitrary height, while the low-energy region is resolved even better than in the original potential. The degree of steepness of the cutoff is controlled by the tunneling parameter γ. Figs. 1b and c illustrate the STUN potential energy surface for a 1D model potential at a hypothetical point in the simulation where the minima indicated by the arrows have been found as the present best estimate for the ground state.

While the transformation in Eq. (4) is not the only possible functional, we believe that there are a number of features that constrain its construction. (i) The transformation must be strongly nonlinear in the high-energy regime, as only such a transformation will lead to a nearly constant effective PES for high energies and true 'tunneling'. (ii) There must be a parameter that modulates the degree of compression (γ), since the ratio of the energy differences of adjacent local minima to the transition state energy separating them varies from problem to problem. (iii) Requiring an essentially flat PES at high energy (for typically unbounded PES) requires a transformation that maps the interval $[E_0, \infty]$ onto some finite interval, which can be chosen as $[0, 1]$ without loss of generality. (iv) It is possible to use a fixed inverse temperature β as a second parameter and to quench the configuration whenever a configuration with an energy lower than E_0 is encountered. The optimization of this additional parameter can be avoided, when one adopts the self-adjusting cooling schedule introduced previously [10].

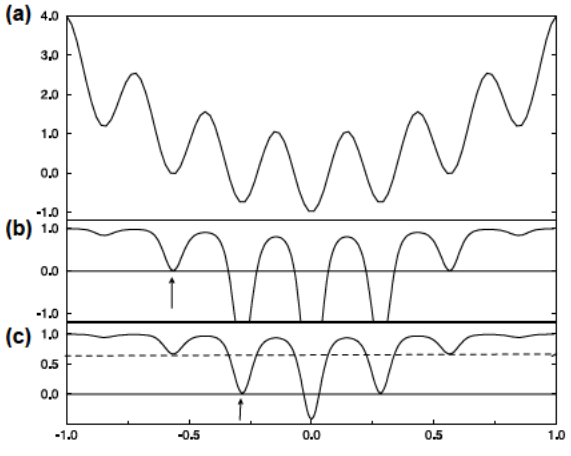

Fig. 1. Sample rugged one dimennsional PES (a) and STUN effective potentials ((b) and (c)) for two hypothetical snapshots where the local minima indicated by the arrows have been lo cated as the best estimate for the global minimum.

4. Results

RLD simulations were carried out for three representative natural ligands, benzamidine, methotrexate and retinol in their natural receptors, β-trypsin (3 ptb), dihydrofolate reductase (4 dfr), and retinol binding protein (1 rbp), respectively. In general RLD simulations are complicated by the fact that (a) not much is known a priori about the center of mass location of the ligand with respect to the receptor pocket and (b) a large fraction of the search space is sterically forbidden.

To succeed in this scenario we optimized the dynamics of the underlying random walks as follows: the ligand is embedded into a symmetrical ellipsoid, where analytical expressions for the friction coefficients f_t (translation), f_a (rotation around long axis a) and f_b (rotation around short axis b) exist [20]. The Brownian dynamics technique in the diffusive regime is based on the update scheme [21]

$$\mathbf{r}(t+\delta t) = \mathbf{r}(t) + \mathbf{F}(t)\delta t/f_t + \mathbf{X}, \tag{5}$$

where $\mathbf{F}$ is a drift term and $\mathbf{X}$ is the random displacement sampled from a Gaussian with zero mean and width $\langle X^2\rangle = 2k_B T\delta t/f_t$. Analogous equations exist for the rotational updates $\delta\omega_a$ and $\delta\omega_b$ around axes a and b, respectively. Taking retinol, the ratio $a/b \approx 4.5$ is reflected in the different friction coefficients f_i which determine the values for the angular shifts, leading to a ratio $\delta\omega_a/\delta\omega_b \approx 2.4$. Similarly, the translational update is physically related to the rotational updates and adapted to the specific shape of the ligand. In contrast to Brownian dynamics, the time step δt has no physical counterpart but can be used as a free parameter to adjust the acceptance rate.

For the drift term we introduce a point $\mathbf{p}$ somewhere inside the cavity of the receptor. The drift force $\delta\mathbf{f}_d \sim -k_d T(\mathbf{r}-\mathbf{p})$ defines an additional, systematic contribution to the dynamics. The strength of the drift is proportional to the distance $|\mathbf{r}-\mathbf{p}|$ as it were in a harmonic oscillator field. If there were no further external forces, this drift would lead to a Gaussian localization of the center of mass coordinates of the ligand [22]. The advantage of this approach over the introduction of a penalty function is that the structure of the potential surface remains unchanged. No more than a bias in the sampling procedure is added to the random displacements. The ligand is localized in a way that its probability distribution fills the cavity of the receptor and is significantly reduced far outside the cavity. The localization volume does not depend on the step size δt and the temperature T.

In all simulations ligands were placed in a random position *outside the cavity* and we averaged the results of 50 runs of predescribed step number. A ligand was defined as 'docked' if the average RMS deviation of the atoms from the global minimum was less than 0.1 nm. The potential values were pre-calculated on cubic grids with a grid constant of 0.04 nm and a dimension of $3 \times 3 \times 3$ nm^3. Generally, a grid whose size is well adapted to the cavity dimension can improve the efficiency of the simulation, but in the presence of the drift term the actual size and shape of the grid becomes secondary. As a drift center we have chosen the actual center of mass of the docked ligand, but due to the diffuseness of the localization volume there is no need for an accurate estimate for that. Any drift term which is centered near the cavity and whose delocalization allows the ligand to probe the entire cavity allows for the desired docking.

The optimum cooling schedule for SA was to start with an initial temperature of 3000 K and then cooling down to 300 K toward the end of the run. In PT, we face the problem that n simultaneous simulations produce an n-fold increase of the number of steps for a given simulation length. A certain minimum simulation length, however, is mandatory for the ligands to find their way into the cavity. This imposes the upper limit of three configurations for PT to be competitive with STUN and SA, a number too low to exploit the full potential of this technique. The temperatures were found to be best at 1000, 1500 and 2250 K for the three configurations.

4.1. Benzamidine docking

Benzamidine is an oblate shaped ligand which, in the ellipsoidal approximation, has an axial ratio of $a/b = 11$, i.e., it is a flat discus. The docking to

β-trypsin is straightforward, since the cavity is easily accessible and there is no potential barrier to be tunneled. It has been observed that even a simple random dynamics reliably docked after 10^6 steps [18]. The potential minimum is quite flat with respect to rotations of the docked ligand around its short axis, which may be an artefact of the inaccurate treatment of hydrogen bonds. We therefore relaxed the docking criteria in this particular case and regarded configurations with energies not more than 5% higher than the global minimum as docked.

The performance of STUN (tunnel parameter $\gamma = 0.05$) and PT compared to SA is displayed in Fig. 2 which shows the success rate as a function of the number of energy evaluations. The data are averaged over 50 independent runs and indicate that all methods perform almost equally well. The success of SA appears surprising, but is entirely due to the drift term which was absent in the simulations reported by Diller [18]. The present simulations confirm the observation made by other groups that this system is too simple to differentiate competing global optimization techniques with respect to their efficiency. Nevertheless, it has been pointed out that in a database of small molecules to be docked, many compounds fall into this category, and therefore any algorithm to be applied to the RLD problem should pass this test efficiently.

4.2. Methotrexate docking

Methotrexate is a prolate shaped ligand with a axial ratio of $a/b = 2.5$ in the ellipsoidal approximation, i.e., a fat cigar. This system presents several problems: the ligand features plenty of polarized endgroups and tends to dock wherever there are residues with partial electric charges. This leads to a rugged potential surface with a large number of local minima. To make things worse, there exists a very deep local minimum just at the entrance of the cavity. The global minimum, the energy of which is only a few percent lower, is separated from the metastable state by a barrier of hundreds kJ/mol. The ligand has to tunnel through the barrier, which requires a high temperature, and then to localize the minimum to an accuracy of a few percent in order to distinguish it from the local minimum in front of the barrier.

Fig. 3 shows the success distribution, where STUN reached a reliability of 0.5 after 50.000 steps, PT required 150.000 while SA required 200.000 steps. In previous work [23], SA was reported to fail completely for this system in the absence of a drift term. During the simulations we observed frequently that the ligand reached the inner region of the docking site in early stages of

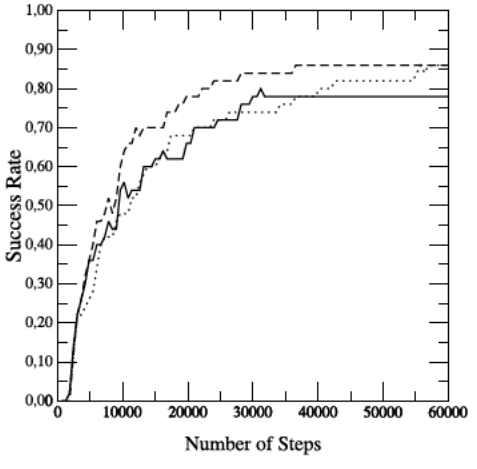

Fig. 2. Sucess rate of SA (full line), PT (dotted line) and STUN (dashed line) in docking benzamidine versus the number of steps.

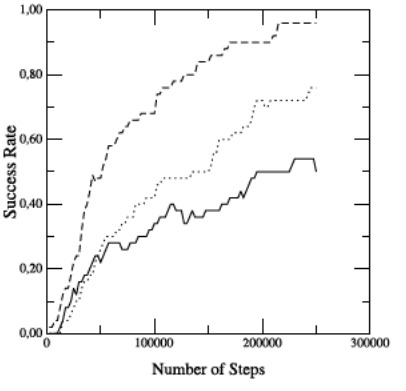

Fig. 3. Sucess rate of SA (full line), PT (dotted line) and STUN (dashed line) in docking methotrexate versus the number of steps.

the simulation, when the temperature was still too high to probe the potential minimum so that shortly afterwards a better score was found again in the low-energy region in front of the cavity. This effect was less pronounced for STUN (tunnel parameter $\gamma = 0.05$), where the temperature is regulated by an automatic mechanism. The energy difference between the minima is resolved much earlier in the simulation.

4.3. Retinol docking

As mentioned above, retinol is a prolate shaped ligand with an axial ratio of $a/b = 4.5$, i.e., a slim cigar. There exists only one polarized endgroup in retinol. In the rigid receptor approximation the binding site is almost completely enclosed and the molecule has to tunnel through a barrier of several thousands kJ/mol to reach the global minimum. On the other hand, this system presents no low-energy secondary minima. Since the cavity is quite elongated (length 1.65 nm) only a weak drift term was applied.

In agreement with previous studies [18], SA completely failed to pass the ligand through the barrier, the same held true for PT. Among 50 runs there was no successful docking. The success distribution for STUN ($\gamma = 0.002$) (Fig. 4) demonstrates that this technique is capable for a fast and reliable docking of retinol, reaching a success rate of 0.5 after about 40 000 energy evaluations.

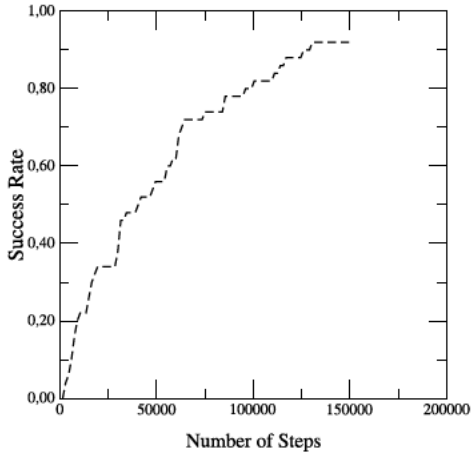

Fig. 4. Success rate of docking retinol with STUN.

5. Discussion

In this investigation we have compared three stochastic optimization strategies for rugged low-dimensional potential energy surfaces that arise in RLD. For this problem STUN proved to be a generic reliable approach superior to both PT and SA. SAS, a close cousin of the DEM, was investigated in a previous study [18] and located the minimum with 10 000 and 8000 energy/gradient evaluations for the Benzamidine and retinol, respectively. The success of the SAS docking procedure depends strongly on the definition of the binding pocket. In applications of computational high throughput screening little can be said a priori about the location of the center of mass of the ligand with respect to the receptor. If one allows for the possibility that the CM of the ligand remains outside the cavity, it is clear that the global minimum at large diffusion times must always lie outside the protein. In this scenario the DEM-like methods must rely on finding the branching point of the potential when the weak attractive potential inside the binding area splits from the large minimal area outside the protein.

We want to point out that the three optimization methods are not necessarily exclusive. For example both, PT and STUN, yield lower energies if a short SA (optimization) run with low temperature is appended to the ordinary (search) run. Other combinations like using PT to sample the STUN energy landscape are worth to be investigated.

Acknowledgements

We thank the Fond der Chemischen Industrie, the BMBF, the Deutsche Forschungsgemeinschaft (grant WE) and the Kurt Eberhard Bode Stiftung for financial support.

References

[1] Christos, H. Papadimitriou, Computational Complexity, Addison Wesley, Reading, MA, 1994.

[2] Michael, R. Garey, David, S. Johnson, Computers and Intractability A Guide to the Theory of NP Completeness, Freeman and Company, New York, 1979.
[3] S. Kirkpatrick, C.D. Gelatt, M.P. Vecchi, Science 220 (1983) 671.
[4] S. Schelstrate, W. Shepens, H. Verschelde, in: Molecular Dynamics: From Classical to Quantum Methods, Elsevier, Amsterdam, 1999, p. 129 (chapter: Energy Minimization by Smoothing Techniques).
[5] A.P. Lyubartsev, A.A. Martinovski, S.V. Shevkunov, P.N. Vorontsov Velyaminov, J. Chem. Phys. 96 (1992) 1776.
[6] E. Marinari, G. Parisi, EPL 451 (1992) 19.
[7] U.H.E. Hansmann, CPL 281 (1997) 140.
[8] B.A. Berg, T. Neuhaus, Phys. Letters B 267 (1991) 249.
[9] B.A. Berg, T. Neuhaus, Phys. Rev. Lett. 68 (1992) 9.
[10] W. Wenzel, K. Hamacher, Phys. Rev. Lett. 82 (1999) 3003.
[11] U.H.E. Hansmann, Eur. Phys. J. B 12 (1999) 607.
[12] Ulrich, H.E. Hansmann, Yuko Okamoto, J. Comput. Chem. 18 (1997) 920.
[13] Ulrich, H.E. Hansmann, Masato Masuya, Yuko Okamoto, Proc. Natl. Acad. Sci. USA 94 (1997) 10652.
[14] K. Gubernator (Ed.), Structure Based Ligand Design, Wiley, New York, 1998.
[15] H.J. Bohm, G. Schneider (Eds.), Virtual Screening for Bioactice Molecules, Wiley, New York, 2001.
[16] J.P.K. Doyle, D. Wales, J. Chem. Phys. 105 (1996) 8428.
[17] W.L. Jorgensen, N.A. McDonald, J. Mol. Struct. 424 (1998) 145.
[18] David, D. Diller, Christophe, L.M.J. Verlinde, J. Comp. Chem. 20 (1999) 1740.
[19] Kamal, K. Bhattacharya, James, P. Sethna, Phys. Rev. E 57 (1998) 2553.
[20] C.R. Cantor, P.R. Schimmel, in: Biophysical Chemistry, vol. 2, Freeman, New York, 1980.
[21] W.F. van Gunsteren, H.J.C. Berendsen, Molecular Phys. 3 (1982) 637.
[22] S. Chandrasekhar, Rev. Mod. Phys. 15 (1943) 1.
[23] G.M. Morris, J. Comp. Chem. 19 (1998) 1639.