

GRAMM-X public web server for protein–protein docking

Andrey Tovchigrechko^{1,*} and Ilya A. Vakser^{1,2,*}

¹Center for Bioinformatics and and ²Department of Molecular Biosciences, The University of Kansas, 2030 Becker Drive, Lawrence, KS 66047, USA

Received February 15, 2006; Revised and Accepted March 22, 2006

ABSTRACT

Protein docking software GRAMM-X and its web interface (<http://vakser.bioinformatics.ku.edu/resources/gramm/grammx>) extend the original GRAMM Fast Fourier Transformation methodology by employing smoothed potentials, refinement stage, and knowledge-based scoring. The web server frees users from complex installation of database-dependent parallel software and maintaining large hardware resources needed for protein docking simulations. Docking problems submitted to GRAMM-X server are processed by a 320 processor Linux cluster. The server was extensively tested by benchmarking, several months of public use, and participation in the CAPRI server track.

INTRODUCTION

The growing needs of experimental and computational biology require reliable computational procedures for modeling of protein interactions. Recent progress in docking algorithms and computer hardware makes it possible to implement such procedures as automated web servers, which greatly improves the utility of the docking approaches in the biological community. Such servers also allow an objective test of the underlying docking methodologies, unbiased by expert human intervention. In 2005, we launched our docking web server GRAMM-X (<http://vakser.bioinformatics.ku.edu/resources/gramm/grammx>). GRAMM-X grew out of the original Fast Fourier Transformation (FFT) GRAMM methodology (1–3). It represents a new implementation that uses a smoothed Lennard-Jones potential on a fine grid during the global search FFT stage, followed by the refinement optimization in continuous coordinates and rescoring with several knowledge-based potential terms.

The field of protein–protein docking is currently in the state of rapid development and expansion (4). A number of existing docking methods rely on large database scans, such

as PSI-BLAST searches for evolutionary conserved residues. The minimization procedures used in protein docking are computationally demanding, requiring parallel execution on large supercomputers/Linux clusters. These factors make the standard model for research software distribution inconvenient for an average biologist (the software has to be downloaded, installed and configured by the user). The paradigm of a web server interface acting as the front end for the developers' own computational cluster solves the problems of installation complexity, frequent updates, and the availability of large uniformly configured computational resources.

In a few months since its launch, the GRAMM-X server has processed >1000 jobs submitted by >250 users. The new features were extensively evaluated on the benchmark of unbound protein pairs with known complexed structures. The server is also regularly subjected to peer review by ~30 professional groups working in the field of protein docking through our participation in the CAPRI blind prediction experiment (<http://capri.ebi.ac.uk>) (5).

ORIGINAL GRAMM SOFTWARE AND METHOD

The original GRAMM docking methodology has been available to the public for a number of years as downloadable software compiled for different platforms, including Linux and Windows. The best surface match between molecules is determined by correlation technique using FFT. An important feature of GRAMM is the ability to smooth the protein surface representation to account for possible conformational change upon binding within the rigid body docking approach. The simplicity of the interface and installation, as well as its availability on the Windows platform are other strong points contributing to the popularity of GRAMM in the biological community. GRAMM success provided important guidelines for the interface and distribution model of GRAMM-X. The last release of the original GRAMM version is available for download at <http://vakser.bioinformatics.ku.edu/resources/gramm/gramm1>. GRAMM has been installed in > 6000 sites worldwide.

*To whom correspondence should be addressed. Tel: 785 864 1057; Fax: 785 864 5558; Email: andrey@ku.edu

*Correspondence may also be addressed to Ilya A. Vakser. Tel: 785 864 1057; Fax: 785 864 5558; Email: vakser@ku.edu

GRAMM-X METHOD

In the original GRAMM method, the intermolecular energy potential is a step function approximating Lennard-Jones potential, based on the grid representation of the molecules (6). The smoothing of the intermolecular energy landscape is achieved by increasing potential range and lowering the value of the repulsion part. Since the range is the step of the grid, the step becomes larger and the structural representation of the molecules is reduced to lower resolution. Such approach has an important advantage of implementation simplicity and computational speed. It also allows the study of fundamental molecular recognition characteristics focusing on underlying simplicity of the basic principles (7–9). However, in practical docking applications aimed at maximizing the chances of the correct prediction, the association of the potential range with the grid step often becomes a disadvantage. The grid step association with the potential range is not suitable for more sophisticated forms of the potential. At lower resolution, it is also sensitive to the positioning of the molecules for the grid digitization, introducing a significant degree of random noise. Disassociation of the interatomic energy potential from the grid, implemented in GRAMM-X, provides a possibility to alleviate these negative factors.

The procedure uses a fine-grid projection of a softened Lennard-Jones potential function (10) calculated for a probe atom:

$$V_{ij}(r) = \frac{1}{\alpha\sigma_{ij}^6 + r^6} \left(\frac{4\epsilon_{ij}\sigma_{ij}^{12}}{\alpha\sigma_{i,j}^6 + r^6} - 4\epsilon_{ij}\sigma_{ij}^6 \right).$$

The benchmarking docking showed that the optimal values of the parameters for a typical protein in unbound conformation are $\alpha = 0.4$, $\sigma = 0.33$ nm and $\epsilon = 0.5$. These uniform values of σ and ϵ applied to all non-hydrogen atoms yielded better results than the values taken from the AMBER atom types. The docking runs also showed that the results do not improve for translation grid steps <1.5 Å and rotation steps $<10^\circ$, which can be explained by the subsequent minimization of the grid predictions in continuous space.

The top 4000 grid-based predictions are subjected to a conjugate gradient minimization in continuous 6D rigid body space with the same soft potential. The minimization accumulates many points, initially located on the grid, in a few local minima. One representative prediction for each minimum is stored and the number of initial predictions falling into this minimum is marked as the volume of the minimum. The average radius of such minima on our smoothed landscape is 5 Å. The local minimization of a smoothed landscape can be viewed as clustering on the original rugged Lennard-Jones landscape, and helps locate the protein binding funnel. For each minimized prediction the following terms are calculated: soft Lennard-Jones potential, evolutionary conservation of predicted interface, statistical residue–residue preference, volume of the minimum, empirical binding free energy and atomic contact energy. To eliminate predictions that are likely to be located far from the correct binding site, we apply the Support Vector Machine

filter trained on a subset of the benchmark set using the above mentioned set of potential terms. The remaining predictions are re-scored by a weighted sum of the potential terms. A detailed description of the algorithm was reported earlier (5). New GRAMM-X features are regularly made available to the web server back-end after extensive benchmarking (5).

The limited number of available test cases [Protein Data Bank (PDB) complexes for which unbound structures also exist in PDB] does not yet allow one to obtain a quantitative estimate for the expected accuracy of docking versus the properties of the input structures. The main factor that affects the quality of the docking prediction is the degree of conformational change of the input structures. Especially important is the degree of such change at the interface area. There is no reliable method at this time to estimate this without the actual knowledge of the bound conformations. An attempt to address that problem was presented in our earlier study (5) where we were substituting bound conformations of the most flexible interface side chains into the unbound structures from the benchmark complexes. It demonstrated that the knowledge of the conformational change upon binding of only three critical interface side chains per complex would provide a 40% improvement of the benchmarking results, and beyond that other factors such as backbone changes and force field accuracy would dominate. However, the critical question is how well the existing PDB benchmarks represent the real population of all interacting proteins. The Dockground (<http://dockground.bioinformatics.ku.edu>) project currently under development in our group aims at increasing the statistical significance of benchmarking results by generating simulated unbound structures for known complexes. That should increase the number of test cases by about an order of magnitude. Such expanded benchmark will also improve our understanding of other properties that influence docking: the type of interaction (antibody–antigen, enzyme–inhibitor), the size of the proteins, etc. Currently, only some qualitative considerations can be provided: the antibody–antigen complexes are more difficult to predict than enzyme–inhibitor complexes, the limited conformational change upon binding can be tolerated with a reasonable chance of success [interface should have <3 Å root mean square deviation (r.m.s.d.) of conformational change], and complexes with a significant backbone movement at the interface area are typically out of scope. The users of the docking server should treat the output structures as potential candidates and critically evaluate them using the available biological information such as interacting residue data from mutational studies and general knowledge about the interaction being studied.

GRAMM-X is implemented in Python and C++, thus combining fast prototyping power of Python and numerical performance of C++ modules. The message passing interface (MPI) library is used for parallelization. On the benchmark of unbound complexes (11) the full docking protocol (FFT + refinement) for a single complex, on average, completes in 2 min, running on 16 2.0 GHz Opteron processors. When the simulation request is submitted from the web server to our cluster of 320 CPUs, it will be queued, and the wait time will vary depending on the current load of the cluster.

WEB INTERFACE AND BACK-END LINUX CLUSTER

The rapid development of protein docking methodology presented a challenge for the traditional research software distribution model in which the end user was supposed to download, install, configure and regularly update the package. Some docking algorithms rely on searches in databases of known structures or sequences. For example, the evolutionary residue conservation term requires BLAST search for the NCBI 'nr' database. The application either has to maintain the local mirror of that large and frequently updated database, or rely on remote calls to NCBI or some other server. The very computationally demanding nature of protein docking calculations (basically the global minimization on the energy landscapes of large molecules) dictates the parallel model of computations. Installing and configuring MPI and possibly maintaining a cluster of workstations are usually beyond the commitment of a biologist. Thus, the growing complexity of docking software presents a high entry barrier that can limit the spread of docking methods among its target audience—the experimental lab scientists. When we embarked on the major extension of GRAMM algorithms and its reimplementations, we realized the need to change the distribution model.

GRAMM-X installation is maintained, along with all related databases and parallel libraries, on our computational cluster (currently 320 Opteron processors). A simple web interface accepts two PDB protein structures from the user, forms a job request and submits it to the execution queue on the cluster. The queue manager ensures that the cluster is used in full capacity without degrading its performance by too many concurrent processes. The web server creates a temporary page for the future simulation results. The user can periodically refresh the page until the docking is done and the results are posted, or wait for an Email notification from the web server (contains the URL of the web page with the results). The Email address is provided by the user during job submission, and is used only once to send the resulting URL. The output PDB file contains 10 models ranked as the most probable prediction candidates according to our scoring function. The models are separated by the MODEL keyword, similar to the PDB files with multiple NMR structures. Output in the PDB format ensures that the user can freely process or view the results with any of the standard structure analysis and visualization tools, such as RasMol, VMD or Swiss PDB Viewer. The output generator tries to maintain the original chain identifiers from the input files as long as they do not conflict when combined in the output file. Currently, the server does not report to the user the gradual progress of the simulation in real time. Instead, the user receives a high level log of the simulation stages once the job has finished. Because the final set of 10 predictions is selected from the thousands of candidate structures after rescoring has been performed during the last seconds of the simulation, there are no meaningful intermediate prediction results to report before the process is done completely.

The experience with the original GRAMM software showed that very few users are willing to learn non-obvious combinations of parameters. Many users would often stop

using the software altogether unless it can quickly offer satisfactory results by intelligently choosing initial parameters. Thus the GRAMM-X input form was designed to be as simple as possible. The docking server back-end analyzes the input structures and selects the best course of action automatically.

PARTICIPATION IN CAPRI SERVER TRACK

Analysis of user submissions and input data irregularities helped improve the stability of results. An important continuing test is participation in CAPRI. GRAMM-X was first used in a fully unsupervised mode in CAPRI Round 5, and after that has been participating in a special track for unattended public servers (an example of CAPRI prediction is shown in Figure 1). In this track, a server must return the results within 72 h after receiving the input structures, and human intervention in the docking process is prohibited. The CAPRI competition exposes the server to the scrutiny of other groups active in protein docking and thus provides important feedback from the professional community. The full results of CAPRI are made available on its website (<http://capri.ebi.ac.uk>). After we began participating with the new GRAMM-X in unsupervised mode in Round 5, we



Figure 1. Prediction for CAPRI Target 18 (TAXI xylanase inhibitor and *Aspergillus niger* xylanase; 1.8 Å r.m.s.d. prediction accuracy for the ligand interface area). The correct and the predicted structures are shown in different colors.

predicted conformations of the complex with r.m.s.d. of the ligand interface atoms 0.68 Å for Target 14 and 1.88 Å for Target 18. Thus, our predictions for Targets 14 and 18 were evaluated in CAPRI as ‘acceptable’ according to the number of predicted native residue contacts. For Target 19, the interface areas on both proteins were correctly predicted, for Target 22—only the interface area of the receptor. Overall, CAPRI targets proved to be difficult for the servers—from the official start of the server track to this moment (before Round 9 results), only for Target 22 there were server predictions (from ClusPro and PatchDock) ranked as ‘acceptable’. The following other public web servers took part in a server track (as of CAPRI Round 9): ClusPro (<http://nrc.bu.edu/cluster/>) (12), PatchDock/SymmDock (<http://bioinfo3d.cs.tau.ac.il/>) (13), Proteus (<http://graylab.jhu.edu/proteus>) (14), SKE-DOCK (http://www.pharm.kitasato-u.ac.jp/biomoleculardesign/files/SKE_DOCK.html) (15) and SmoothDock (<http://structure.pitt.edu/servers/smoothdock/>) (16). The servers are based on different algorithms, thus allowing a user to obtain an ‘orthogonal’ set of predictions. ClusPro and SmoothDock start from a few thousand predictions obtained from FFT-based global search programs such as ZDOCK (17), DOT (18) or GRAMM (1–3), and then employ rigid body minimization with a combination of empirical and standard force field energy terms and clustering. In general outline, ClusPro and SmoothDock are similar to GRAMM-X, and all use FFT-based initial global search, but the refinement/rescoring protocols and potential terms differ in each case. PatchDock and its sister server SymmDock (for symmetrical multimeric docking) employ the geometric hashing technique as the search procedure. Proteus (using RosettaDock engine from the same group) does a Monte Carlo search: rigid body minimization stage followed by the simultaneous optimization of backbone displacement and side-chain conformations. SKE-DOCK server finds the possible binding sites by building benzene clusters around the receptor molecule, rescoring with a shape complementarity on a grid, and then repacking interface side chains with a homology modeling program. A detailed comparison of the methods is beyond the scope of this paper. The advantages of the server concept is that it is easy to try each server for a specific biological target without committing substantial resources.

FUTURE DIRECTIONS

GRAMM-X algorithms and implementation will be further developed as part of our group’s structural genomics and bioinformatics effort. GRAMM-X docking engine is also being incorporated as a major data generation component in our other public services: DOCKGROUND, the integrated system of databases for protein recognition studies, and GWIDD, the resource for 3D structure prediction of protein–protein complexes on genome scale (<http://vakser.bioinformatics.ku.edu/resources>). The development of GRAMM-X web interface will include the Advanced Options page for users preferring direct control over the docking parameters, and a graphical output and basic analysis of the results in addition to the PDB file output. The features of the graphical output will be modeled after the PDB site (bitmaps of ribbon diagrams,

with highlighted predicted interface areas and Rasmol scripts). The analysis will include the list of predicted pairs of contact residues. For the users planning their own further refinement or rescoring of multiple predictions, a list of first 1000 predictions will be provided in a compact coordinates format, accompanied by a short Python script to generate corresponding PDB files.

ACKNOWLEDGEMENTS

This work was supported by NIH grant R01 GM61889. Funding to pay the Open Access publication charges for this article was provided by NIH.

Conflict of interest statement. None declared.

REFERENCES

- Katchalski-Katzir, E., Shariv, I., Eisenstein, M., Friesem, A.A., Aflalo, C. and Vakser, I.A. (1992) Molecular surface recognition: determination of geometric fit between proteins and their ligands by correlation techniques. *Proc. Natl Acad. Sci. USA*, **89**, 2195–2199.
- Vakser, I.A. and Aflalo, C. (1994) Hydrophobic docking: a proposed enhancement to molecular recognition techniques. *Proteins*, **20**, 320–329.
- Vakser, I.A. (1995) Protein docking for low-resolution structures. *Protein Eng.*, **8**, 371–377.
- Marshall, G.R. and Vakser, I.A. (2005) Protein–protein docking methods. In Waksman, G. (ed.), *Proteomics and Protein–Protein Interaction: Biology, Chemistry, Bioinformatics, and Drug Design*. Springer, NY, pp. 115–146.
- Tovchigrechko, A. and Vakser, I.A. (2005) Development and testing of an automated approach to protein docking. *Proteins*, **60**, 296–301.
- Vakser, I.A. (1996) Long-distance potentials: an approach to the multiple-minima problem in ligand–receptor interaction. *Protein Eng.*, **9**, 37–41.
- Vakser, I., A., Matar, O.G. and Lam, C.F. (1999) A systematic study of low-resolution recognition in protein–protein complexes. *Proc. Natl Acad. Sci. USA*, **96**, 8477–8482.
- Tovchigrechko, A. and Vakser, I.A. (2001) How common is the funnel-like energy landscape in protein–protein interactions? *Protein Sci.*, **10**, 1572–1583.
- Jiang, S., Tovchigrechko, A. and Vakser, I.A. (2003) The role of geometric complementarity in secondary structure packing: a systematic docking study. *Protein Sci.*, **12**, 1646–1651.
- Schafer, H., Van Gunsteren, W.F. and Mark, A.E. (1999) Estimating relative free energies from a single ensemble: hydration free energies. *J. Comput. Chem.*, **20**, 1604–1617.
- Chen, R., Mintseris, J., Janin, J. and Weng, Z. (2003) A protein–protein docking benchmark. *Proteins*, **52**, 88–91.
- Comeau, S.R., Gatchell, D.W., Vajda, S. and Camacho, C.J. (2004) ClusPro: an automated docking and discrimination method for the prediction of protein complexes. *Bioinformatics*, **20**, 45–50.
- Schneidman-Duhovny, D., Inbar, Y., Nussinov, R. and Wolfson, H.J. (2005) PatchDock and SymmDock: servers for rigid and symmetric docking. *Nucleic Acids Res.*, **33**, W363–W367.
- Daily, M.D., Masica, D., Sivasubramanian, A., Somarouthu, S. and Gray, J.J. (2005) CAPRI rounds 3–5 reveal promising successes and future challenges for RosettaDock. *Proteins*, **60**, 181–186.
- Terashi, G., Takeda-Shitaka, M., Takaya, D., Komatsu, K. and Umeyama, H. (2005) Searching for protein–protein interaction sites and docking by the methods of molecular dynamics, grid scoring, and the pairwise interaction potential of amino acid residues. *Proteins*, **60**, 289–295.

16. Camacho,C.J. (2005) Modeling side-chains using molecular dynamics improve recognition of binding region in CAPRI targets. *Proteins*, **60**, 245–251.
17. Wiehe,K., Pierce,B., Mintseris,J., Tong,W., Anderson,R., Chen,R. and Weng,Z. (2005) ZDOCK and RDOCK performance in CAPRI rounds 3, 4, and 5. *Proteins*, **60**, 207–221.
18. Mandell,J.G., Roberts,V.A., Pique,M.E., Kotlovyyi,V., Mitchell,J.C. and Nelson,E. (2001) Tsigelny I and Ten Eyck LF, Protein docking using continuum electrostatics and geometric fit. *Protein Eng.*, **14**, 105–113.