ELSEVIER

VRDD: Applying virtual reality visualization to protein docking and design

Abraham Anderson and Zhiping Weng

Department of Biomedical Engineering, Boston University, Boston, Massachusetts, USA

We have developed an interactive docking program called VRDD. It offers various modes of displaying molecules in an immersive, three-dimensional virtual reality (VR) environment. It allows a user to interactively perform molecular docking aided by automatic docking and side chain conformational search. Binding free energies are computed in real time, and the program enables the user to explore only clash-free orientations of a ligand. VRDD also supplies visual and auditory feedback during docking and side chain search, indicating the levels of atomic overlap and interaction energy. The stunning VR graphics immerse users in the scene and can maximally stimulate their design intuition. We have tested VRDD on three cases with increasing complexity: a nine-residue-long peptide bound to a major histocompatibility complex (MHC) molecule, barstar bound to barnase, and an antibody bound to a hemagglutinin. Without prior knowledge, combinations of hand-docking and automatic refinement led to accurate complex structures for the first two complexes. The third case, for which all automatic docking algorithms failed to identify the correct complex in a previous blind test, also failed for VRDD. Our results show that the combination of VR docking and automatic docking can make unique contributions to molecular modeling. © 2000 by Elsevier Science Inc.

Keywords: virtual reality, computer graphics, CAVE, molecular docking, binding free energy calculations, interactive docking

Abbreviations: VRDD, Applying virtual reality visualization to protein docking and design, the computer program that we have developed in this article; CAVE, CAVE automatic virtual environment; Immersadesk™, a drafting table style projection area that provides a multiviewer, semiimmersive virtual reality experience; OpenGL™, a software interface to graphics hardware (GL stands for Graphics Library); GLUT™, Graphics Library Users Toolkit; SAS, solvent accessible surface representation of a macromolecule; vdW, van der Waals interactions.

Color Plates for this article are on page 217.

Corresponding author: Z. Weng, Department of Biomedical Engineering, Boston University, 44 Cummington Street, Boston, MA 02215 USA

E-mail address: zhiping@bu.edu. (Z. Weng)

INTRODUCTION

A protein performs its function by interacting with other proteins. Thus effective computational algorithms for predicting whether two proteins bind to each other and if so predicting the complex structure (protein docking), as well as for determining which sequence alterations can lead to stronger or more specific binding (protein design), have broad applications in computational biology. These algorithms employ computers to model the thermodynamic and structural nature of receptor–ligand interactions. Many ligand structures and orientations with respect to the receptor can be generated and ranked according to scoring functions and predictions can be made about the top scorers, which can then be examined experimentally.

Computer graphics has played an important role in the development of protein docking and design algorithms. Some software packages, such as Molscript,[1] Raster3D,[2] and GRASP,[3] can produce high-quality images of molecules. Some, such as RASMOL,[4] CHIME (MDL Information Systems, Inc.), and VMD,[5] allow users to display molecules using various representations and view them from different angles. Other programs designed for energy calculations, conformational search and mutations, such as INSIGHT/QUANTA© (Molecular Simulations, Inc.) and Swiss-PdbViewer[6] have graphical user interfaces.

Can user input via computer visualization help to speed up docking and design? For example, a user may be able to determine conformations or orientations that need not be explored simply by viewing the molecules. Hermans and colleagues have developed SMD, a system for interactively steering molecular dynamics calculations of proteins.[7] A number of successes have been reported in the design of small molecules bound to protein receptors.[8] By visualizing which functional

Journal of Molecular Graphics and Modelling 17, 180–186, 1999

655 Avenue of the Americas, New York, NY 10010

1093-3263/99/$–see front matter
PII S1093-3263(99)00029-7

groups are at the binding site of the receptor, experienced researchers can propose ligand designs based on their intuition. However, it is much more difficult to take advantage of user inputs in the case of protein–protein docking, since images generated by desktop computers do not give adequate depth perception, nor do they allow easy manipulation of a molecule's orientations.

Most currently available docking algorithms assume that the binding site of the receptor is known, since it is usually too time consuming to explore the entire receptor surface. Human brains are highly capable of recognizing patterns, thus a natural question arises: if we are given two plastic models of the receptor and the ligand built from their 3-D coordinates, can we find the complex structure by hand without any prior knowledge of the binding site? The advent of virtual reality (VR) and parallel processing enables us to answer this question in a general way. One major attribute of VR is immersion: it gives a user the experience of actually "being there". It can provide the in-depth perception that is lacking in conventional visualization, as well as force feedback and sound localization. To do so, a VR system contains many of the following components: surround vision, stereo cues, viewer-centered perspective, real-time interaction, tactile feedback, and directional sound. The goal of this study is to develop a VR platform that enables researchers to dock protein molecules as if handling plastic models.

Levine et al. have developed a docking algorithm[9] that can be carried out in a projection-based VR system called CAVE (CAVE Automatic Virtual Environment).[10,11] The system allows a user to be visually immersed in the docking arena and watch in VR the computer docking two molecules. The user can also interrupt the computer at any time and translate the receptor or ligand and then let the computer restart docking calculations from that point. However, the authors have encountered a number of difficulties: (1) a significant delay in redrawing the molecule when using the hand tracking system, because of the large number of spheres to be redrawn and a slow sampling rate of the hand tracker, leading to "oversteering"; (2) when immersed within the molecule, difficulty in determining what part of the molecule they were viewing, especially when attempting to steer the ligand. The awkward steering led to the conclusion that they did not take advantage of the system's immersive capabilities, but instead "tried to hand-dock the ligand with the wand from afar".[9] Attempts were made to simplify the system by removing side chain atoms from the display, but this led to the inability to determine overlapping atoms at each new conformation. In the end they found the "hand-docked" solutions were always worse than the best solution generated by a computer search algorithm.

In this study, we have applied virtual reality visualization to protein docking and design (and the program thus developed is called VRDD). We have implemented VRDD on an Immersadesk™, a drafting table-style projection system that provides a multiviewer, semiimmersive VR experience. Two proteins can be displayed and manipulated independently of each other with the wand. A floating menu activated by pressing the right wand button enables the user to choose from different modes, such as viewing, docking, energy calculation, side chain conformational search, etc. The user can also use a keyboard, which allows precise selection of residues or individual atoms. Many attempts have been made to improve the speed of steering. Surface representations of molecules have been implemented to simplify the scene, as well as to increase speed. A Monte Carlo algorithm can be evoked at any stage to sample the vicinity of any ligand orientation, which can be a hand-docked one. Ligand–receptor interaction is quantified using a previously calibrated binding free energy function. Throughout the course of docking, a list of top ligand orientations (the ones with the lowest binding free energies) are updated and the user can revert to any of them, as well as print the list out for further analysis.

We have tested VRDD on three cases with increasing complexity: a nine-residue-long peptide bound to an MHC molecule (PDB code 1HHI[12]), barstar bound to barnase (PDB code 1BRS[13]), and an antibody bound to hemagglutinin (PDB code 1QFU[14]). Without prior knowledge of the binding site, combinations of hand docking and Monte Carlo refinement led to accurate complex structures (less than 3 Å root mean square deviation from the X-ray structure orientations) for the first two cases. The antibody–hemagglutinin case, which is a difficult (and almost pathologic) one,[15] interactive docking failed to identify the complex structure. For debugging purpose, we also have a version of VRDD that runs on desktop Silicon Graphics (SGI) workstations. Using this version, hand docking is practically impossible since it is difficult to see where to steer the ligand. Our results show that VR computer graphics can make unique contributions to molecular modeling.

PROGRAM DESIGN

Hardware and Software Platforms

The hardware used in this study is a Pyramid Systems Immersadesk. The scene is displayed on a 4-ft by 5-ft rear-projected screen, which receives input from a 4-processor SGI Power Onyx equipped with 2GB of memory and InfiniteReality2 Graphics. CrystalEyes emitters and stereo glasses are used for viewing. An Ascension SpacePad system driven by a PC is used to track the position and orientation of the user's head and wand. Stereo sounds are generated by a separate SGI Indy and a Kurzweil 2500R sampling synthesizer, and delivered to a multichannel speaker array via an MIDI-controlled Yamaha 03D digital mixing console.

VRDD is written in C and OpenGL to facilitate portability. Two versions are available. The VR version running on Immersadesk also uses the CAVE library, which is a standard VR library implemented on most VR apparatus. The other version uses purely OpenGL, and runs on a desktop SGI workstation. It was developed to facilitate the debugging of the VR version, as well as for comparing the performances of the two versions. Except for the VR aspect, both versions have similar features. Inputs to the program are atomic coordinates of the receptor and the ligand in the Protein Data Bank (PDB; http://www.rcsb.org/pdb/) format. Unless otherwise noted, all discussions hereafter refer to the VR version of VRDD.

The source and executable codes of VRDD are available to academic users free of charge. They can be obtained via contacting the corresponding author (zhiping@bu.edu). The manual can also be accessed on-line from http://engpub1.bu.edu/zhiping/VRDD.html.

Display Models Supported by VRDD

VRDD can be instructed to display proteins in several commonly used styles, including C_α backbone traces, ball-and-stick

models, van der Waals (vdW) space-filling models, and solvent-accessible surface (SAS) representation. Among them, vdW space-filling models are the most visually appealing; however, they are the slowest to render. On a 4-processor SGI Power Onyx with InfiniteReality2 graphics, a midsized protein (<100 residues) in vdW space-filling models can be interactively manipulated without noticeable delay. However, for hand docking with energy calculations, vdW space-filling models are feasible only for short peptides (<20 residues). The colors of an atom and its bonds depend on the atom type; the solvent-accessible surface can also be colored by atom type or residue type. All surfaces are rendered with shadows, which enhance the user's depth perception. The user can choose the desired model from the floating menu.

Solvent-Accessible Surface Representation

Surface representation is widely used in molecular graphics and modeling. Compared with line-trace or ball-and-stick models, surface representation omits interior details of a molecule and thus reduces the visual complexity of the scene tremendously. Surface rendering is also highly efficient, since triangulated surfaces can be used. In comparison, the sphere rendering for vdW space-filling models is much slower, and the performance deteriorates quickly for large molecules. A large number of algorithms have been developed to compute the solvent-accessible surface (SAS) and its variants.[16–23] We have adopted the MSMS program by Sanner et al.,[21] which generates a triangulated surface from the analytic solution. VRDD interprets the two output files from MSMS and creates a shaded SAS in VR. The SAS can be colored according to the atom or residue beneath each triangulated surface patch.

We found that the SAS representation substantially improves the rendering and navigation speed. We do not experience the delay in redrawing molecules encountered by Levine et al.,[9] who used vdW space-filling models. VRDD can render the SASs of a complex with as many as 500 residues without noticeable delay. The SAS representation also simplifies the geometric feature of a molecule, so that ridges and valleys stand out from the scene. This is especially important for hand docking, since the user can most effectively match the surface shapes and color patterns of the receptor and the ligand. In the uncolored SAS representation of an MHC molecule (Color Plate 1), the B and F pockets of the peptide-binding site are clearly visible as two deep grooves. The user can "get in" and "slide down" the binding cleft, explore the two pockets that contribute to the majority of the binding free energy, and look outside from the cleft.

Docking

Docking is one of the main features of VRDD. A user can enter the docking mode by selecting "docking" from the floating menu. Then the user can translate and rotate the ligand by using the left and middle wand buttons. For each ligand orientation, the receptor–ligand binding free energy, which is the quantitative measurement of the binding strength, can be computed using a previously developed function [Eq. (1)][24–27]:

$$\Delta G = E_{\text{elec}} + E_{\text{vdW}} + \Delta G_{\text{solv}} \qquad (1)$$

where E_{elec} and E_{vdW} are the electrostatic and van der Waals interaction energies between the receptor and the ligand calculated according to the CHARMm potential,[28] and ΔG_{solv} is the desolvation free energy on binding, calculated according to Zhang et al.[26,27] The above binding free energy function has been shown to perform well in evaluating the energetics of protein–protein binding, protein-folding/unfolding transitions, and the docking of flexible peptides to protein peptides.[26,27] It has also been used as the target function to rapidly predict the conformations of several interacting side chains.[29]

Although evaluating the binding free energy function is reasonably rapid (it takes less than 1 s for one evaluation of the barnase–barstar complex), it can still slow the program down considerably if performed in real time. A space-compartmentalization scheme has been developed to reduce the number of atom pairs whose interactions need to be included. The entire space is divided into small cubes, and an occupancy matrix containing the receptor atoms in each cube is created during the program's initialization. This needs to be done only once, since the receptor is fixed throughout docking. The default size of the cubes is set to 8.5 Å, compared with the cutoff distances for the electrostatic energy (17 Å), for the vdW energy (8.5 Å), and for the solvation energy (6.5 Å). For each ligand atom, we can determine which cube it is located in by using its Cartesian coordinates. Only those receptor atoms that occupy neighboring cubes are considered for energy calculation with this ligand atom. For electrostatic energy, two layers of neighboring cubes need to be considered while only one layer is sufficient for the other two energy terms. The scheme can speed up energy calculation by approximately threefold without sacrificing accuracy.

A floating text line at the top-center of the VR scene is updated with the current binding energy value calculated according to Eq. (1). To avoid unnecessary energy calculations, the user can explore only the orientations of a ligand that are free of vdW clashes with the receptor, by choosing the "solid" mode from the floating menu. In other words, the two molecules are treated as solid and can slide along each other's surface. To do so, VRDD saves the rotational and translational matrices of the ligand for every orientation. If the ligand is steered to an orientation that clashes with the receptor, it is reverted to the previous orientation. In the meantime the user is alarmed with the metal-hitting sound. Testing if two molecules clash is extremely rapid and can be done in real time. Energy calculations need to be carried out only for clash-free ligand orientations. We have found this feature essential for hand docking, since it substantially decreases the search space. The correct ligand orientation must be one that is clash free and yet has good surface complementarity to the receptor. By sliding two "solid" molecules along each other's surface, the user can explore only orientations that are likely to be correct.

During the entire course of docking, some number (currently set to 10) of clash-free ligand orientations with the most favorable binding free energies is kept in a list, along with their corresponding rotational and translational matrices. This list is updated whenever the user has sampled an orientation that has lower energy than any orientations in the list. To avoid keeping similar orientations, we require any two orientations in the list to have more than 1 Å root mean square deviation. The user can revert to any of these orientations by choosing from the floating menu. We find this useful for hand docking. Many times the user would want to inspect and compare top orientations to decide on the best one. The user can also print the list out for further analysis.

Monte Carlo Local Search

The main idea behind hand docking is that visual inspection may be able to narrow down the search space. Once the user has identify a likely good fit, the refinement should be carried out by the computer, which is much faster and more systematic. VRDD has a Monte Carlo local search mode, which explores around current ligand orientation for a user-defined number of steps (currently set to 100) according to the Metropolis Monte Carlo algorithm.[30] For each of the steps, a random translational and/or rotational perturbation is applied to the ligand orientation. If the new orientation has a lower free energy, it is accepted. Otherwise it is accepted with a probability that is an exponential function of the negative of the energy increase. All proteins undergo conformational changes on binding, some more than others. Therefore we allow some clashes under the "solid" docking mode. However, both electrostatic and vdW energy terms in Eq. (1) are sensitive to clashes, thus hand docked orientations usually have high energies even when they are very similar to the correct orientation. The Monte Carlo local search is highly effective for removing such "slight clash" and generating meaningful energy values.

Amino Acid Selection and Side Chain Conformational Search

When the bound and the free structures of the same protein are compared, many surface side chains are found to have different conformations, especially those side chains at the binding site. Even for protease–inhibitor complexes that are commonly considered as rigid-body docking examples, key residues at the binding site loop of the inhibitors can adopt quite different conformations on binding.[25] Side chain conformational search has been shown to improve docking results.[25] VRDD allows the user to select any side chain dihedral angle and interactively search the conformational space. Currently, side chain selections can be entered only from the keyboard, since we have not implemented any efficient algorithm for 3-D picking using the wand or for voice activation. At each conformation, the binding free energy between the side chain and its environment can be computed according to Eq. (1) and displayed at the top-center VR scene. The volume of atomic overlap (in $Å^3$) is also calculated and this is indicated by the change in the residue's color. Red represents high overlap, white indicates no overlap, and the colors in between are based on the natural spectrum. In this manner the user gets immediate visual feedback about the residue's interaction with its environment and can find the residue's lowest energy state. To speed up the side chain search, atoms that are further than a predefined cutoff (currently set to 17 Å) away from the selected residue are excluded from the residue's neighbor list and not used in the interaction energy calculations.

This feature is useful when the user has some residues in mind, for example residues that have been experimentally shown to affect binding. Sometimes adjusting the conformation of a single side chain can lead to a much better fit. A well-known example is the P1 residue (lysine 15) of BPTI, which is an inhibitor of trypsin. The conformation of this side chain in the free state (PDB code 4PTI) is more than 2.5 Å (root mean square deviation) away from that of the bound state (PDB code 2PTC), and without adjustment no clash-free ligand orientation can be found. With the help of side chain conformational search, the correct complex structure can be easily found using VRDD.

Analysis of Experimentally Determined Complex Structures

VRDD stores the starting orientation of the ligand. At any point of VR docking, the user can reset the ligand orientation to its starting value by selecting "reset" from the floating menu. If the user selects "open", the ligand will be pulled away from its starting orientation with respect to the receptor by 4 Å along the line connecting their centers of mass. Then the user can traverse the "binding cleft" formed by the contact surfaces of the two molecules. When the molecules are represented as SASs colored according to atom types, the user can easily see the charge complementarity and hydrophobic packing by looking toward the left and the right of the "binding cleft". We find this useful for analyzing experimentally determined complex structures.

TEST CASES AND RESULTS

We have tested VRDD on a representative set of protein complexes: a nine-residue-long peptide bound to an MHC molecule (PDB code 1HHI[12]), barstar bound to barnase (PDB code 1BRS[13]), and an antibody bound to hemagglutinin (PDB code 1QFU[14]). The general docking strategy is as follows: (1) first display the two protein molecules in SAS models. Rotate them around to study the shape and charge patterns; (2) identify concave regions on the larger molecule that are likely to be binding sites; (3) hand dock the ligand to each of the identified receptor regions based on shape and charge complementarity. After each hand docking, the Monte Carlo local search is carried out; and (4) compare top ligand orientations VRDD kept in a list. Carry out side chain conformational searches for a small number of residues if necessary.

The MHC–peptide complex (PDB code 1HHI[12]) is the easiest case, if we do not consider the backbone conformational change of the peptide on binding. The peptide has nine residues, and is in an extended conformation. The MHC molecule has 180 residues (only the α_1 and α_2 domains were used). The binding site on MHC is apparent as soon as the molecule is displayed: a long and deep cleft, within which there are two smaller but even deeper pockets (Color Plate 1). There are only two ways the peptide can align with the binding cleft; however, the decision is relatively easy to make since the N terminus of the peptide is positively charged and the C terminus is negatively charged. Furthermore, the two smaller pockets look ideal for fitting the side chains of peptide positions 2 and 9. After studying the charge and hydrophobic patterns of the cleft, the correct orientation can be quickly identified. A few rounds of Monte Carlo local search lead to a peptide orientation that is within 2 Å of the one determined by X-ray crystallography.

The barnase–barstar complex (PDB code 1BRS[13]) is a typical example of enzyme–inhibitor complexes. Barnase has 108 residues and barstar has 87 residues. In Color Plate 2, barnase is drawn in colored SAS reflecting charge patterns and barstar in uncolored SAS. The binding site of barnase, a deep cleft, is easily identifiable. However, barstar is a relatively globular molecule and it is not obvious which region binds to barnase. Hand docking and Monte Carlo local search resulted in six distinct orientations for barstar. The one with the lowest bind-

ing free energy (more than 5 kcal/mol lower than the second lowest energy) corresponded to a structure less than 3 Å away from the X-ray structure.

The hemagglutinin–antibody complex (PDB code 1BRS[13] is an extremely difficult case for a number of reasons. First, both proteins are large molecules. Monomeric hemagglutinin has two noncovalently bound chains (A chain has 328 residues and B chain has 175 residues). The Fab fragment of the antibody has 205 residues (108 for the heavy chain and 97 for the light chain), excluding constant domains. Second, the binding site of an antibody is known to experience large conformational changes on binding. Third, and perhaps most importantly, the binding mode this antibody adopts is counterintuitive. The user would think that all six CDR (complement determining region) loops of the antibody molecule, or at least both CDR3 loops, tend to participate in the binding. Therefore, most of the hand-docked hemagglutinin orientations have some flat region of the hemagglutinin docked into the large surface of the antibody that is made up by most of the CDR loops. However, the experimentally determined complex structure reveals that the antibody light chain hardly participates in the binding. The lowest binding energy produced by the combination of hand docking and Monte Carlo local search (around −4 kcal/mol) was much less favorable than the value calculated for the correct complex (−18 kcal/mol), which indicates that hand docking did not sample any orientations close to the correct answer.

DISCUSSION

We have developed a program, VRDD, that enables molecular visualization and interactive docking in a VR environment. Compared with other VR programs, the unique contribution of VRDD is that it engages the user's visual and motor skills in an intuitive setting that is optimized for interactive docking. The ease with which the ligand can be navigated is key to a successful VR docking program. Levine et al.[9] expressed concern over the significant delay they experienced when they tried to steer the ligand with the wand, which contributed to poor hand-docked structures. In this study we tackled this problem by the following approaches: (1) we implemented SAS models. They may be rendered much faster than spheres, without sacrificing visual effects. In addition, interior atoms do not contribute to SAS models, which leads to further improvement in rendering speed; (2) the compartmentalization scheme increases the speed of energy calculation by severalfold; and (3) the "solid" docking mode eliminates the need to search the part of conformational space that cannot be correct.

Another key point for a successful VR docking algorithm is to prevent the user from being disoriented when he or she is immersed in the VR scene. This was also a difficulty experienced by Levine et al.[9] Our SAS representation (especially when it is colored according to residue types) substantially decreases visual complexity. After studying a molecule in its SAS representation, an experienced user can even memorize the molecule in terms of general shapes and surface coloring patterns. Since human brains are good at recognizing spatial patterns, the user can readily pinpoint where to steer the ligand for a likely good fit.

Other useful VR docking features that VRDD offers include (1) real-time calculation of binding free energies according to a function that has been validated previously, (2) Monte Carlo local refinement after hand docking, (3) side chain conformational search for key residues, and (4) sound and color feedback during docking.

Many automated docking algorithms are available.[31–45] Compared with them, VRDD actively involves human intuition in a VR setting. Visual inspections can frequently speed up automated searches by restricting the search space to regions that are likely to be the binding site. This is especially important when no information regarding the binding site is available. Moreover, visually inspecting ligand orientations generated by automated docking algorithm can facilitate the identification of the correct orientation. For example, in the first blind test of protein-docking algorithms, a number of the participating groups used visual inspection to rank solutions generated by automated docking algorithms.[46]

The result of applying VRDD to the three test cases indicates that the scheme of hand docking and automated local refinement can lead to accurate (<3 Å) ligand orientations for reasonably difficult cases (MHC–peptide and barnase–barstar complexes). VRDD failed to identify the correct ligand orientation for the most difficult case, the antibody–hemagglutinin complex. However, this complex is almost a pathologic example since the antibody light chain hardly participates in the binding, which is atypical for antibody–antigen complexes. It was presented as a blind test challenge at the second meeting of critical assessment of structure prediction (CASP2), and none of the participating groups predicted the correct complex structure.[15]

Written in C, OpenGL, and CAVE with an open architecture, VRDD readily allows for future enhancements. Possible algorithmic improvements include the incorporation of more sophisticated automated docking algorithms, backbone flexibility, and bonded energy. A number of VR features can also improve the performance of VRDD: (1) voice activation, especially for residue selection, (2) localized sound feedback, and (3) force feedback reflecting binding energies.

ACKNOWLEDGMENTS

We are grateful to Dr. Charles DeLisi for valuable advice and continuing support. We also thank Roy Kimura for some initial programming work, and Drs. Eric Brisson, Kathleen Curry, Robert Putnam, and Roscoe Giles for the use of Immersadesk and for stimulating discussions. We are especially grateful to Eric Brisson and Kathleen Curry for help on OpenGL and CAVE programming. The work has been funded by an NSF grant (DBI-9806002).

REFERENCES

1 Kraulis, P.J. Molscript: A program to produce both detailed and schematic plots of protein structures. *J. Appl. Crystallogr*. 1991, **24,** 946–950
2 Merritt, E.A., and Murphy, M.E.P. Raster3D, version 2.0, a program for photorealistic molecular graphics. *Acta Crystallogr. D* 1994, **50,** 869–873
3 Nicholls, A., Bharadwaj, R., and Honig, B. GRASP—graphical representation and analysis of surface properties. *Biophys. J.* 1993, **64,** A166–A166
4 Sayle, R.A., andMilner-White, E.J. RasMol: Biomolecular graphics for all. *Trends Biochem. Sci.* 1995, **20,** 374–376

5 Humphrey, W., Dalke, A., and Schulten, K. VMD: Visual molecular dynamics. *J. Mol. Graphics* 1996, **14,** 27–28, 33–38
6 Guex, N. Swiss-Pdb Viewer: A new fast and easy to use PDB viewer for the Macintosh. *Experientia* 1996, **52,** A26
7 Leech, J., Prins, J.F., and Hermans, J. SMD: Visual steering of molecular dynamics for protein design. *IEEE Comput. Sci. Eng.* 1996, **3,** 38–45
8 Aronov, A.M., Verlinde, C.L., Hol, W.G., and Gelb, M.H. Selective tight binding inhibitors of trypanosomal glyceraldehyde-3-phosphate dehydrogenase via structure-based drug design. *J. Med. Chem.* 1998, **41,** 4790–4799
9 Levine, D., Facello, M., Hallstrom, P., Reeder, G., Walenz, B., and Stevens, F. Stalk: An interactive system for virtual molecular docking. *IEEE Comput. Sci. Eng.* 1997, **4,** 55–65
10 Akkiraju, N., Edelsbrunner, H., Fu, P., and Qian, J. Viewing geometric protein structures from inside a CAVE. *IEEE Comput. Graphics Appl.* 1996, **16,** 58–61
11 Czernuszenko, M., Pape, D., Sandin, D.J., DeFanti, T.A., Dawe, G.L., and Brown. M. The ImmersaDesk and Infinity Wall projection-based virtual reality displays. *Comput. Graphics* 1997, **31,** 46–49
12 Madden, D.R., Garboczi, D.N., and Wiley, D.C. The antigenic identity of peptide–MHC complexes: A comparison of the conformations of five viral peptides presented by HLA-A2. *Cell* 1993, **75,** 693–708. [Published erratum appears in *Cell* 1994;76:410]
13 Buckle, A.M., Schreiber, G., and Fersht, A.R. Protein–protein recognition: Crystal structural analysis of a barnase–barstar complex at 2.0-A˚ resolution. *Biochemistry* 1994, **33,** 8878–8889
14 Fleury, D., Barrère, B., Bizebard, T., Daniels, R.S., Skehel, J.J., and Knossow, M. A complex of influenza hemagglutinin with a neutralizing antibody that binds outside the virus receptor binding site. *Nature Struct. Biol.* 1999, **6,** 530–534
15 Sternberg, M.J., Gabb, H.A., and Jackson, R.M. Predictive docking of protein–protein and protein–DNA complexes. *Curr. Opin. Struct. Biol.* 1998, **8,** 250–256
16 Lee, B., and Richards, F.M. The interpretation of protein structures: Estimation of static accessibility. *J. Mol. Biol.* 1971, **55,** 379–400
17 Connolly, M.L. Solvent-accessible surfaces of proteins and nucleic acids. *Science* 1983, **221,** 709–713
18 Varshney, A., Brooks, F.P., and Wright, W.V. Computing smooth molecular surfaces. *IEEE Comput. Graphics Appl.* 1994, **14,** 19–25
19 Duncan, B.S., and Olson, A.J. Texture mapping parametric molecular surfaces. *J. Mol. Graphics* 1995, **13,** 258–264
20 Duncan, B.S., and Olson, A.J. Approximation and visualization of large-scale motion of protein surfaces. *J. Mol. Graphics* 1995, **13,** 250–257
21 Sanner, M.F., Olson, A.J., and Spehner, J.C. Reduced Surface: An efficient way to compute molecule surfaces. *Biopolymers* 1996, **38,** 305–320
22 Sanner, M.F., and Olson, A.J. Real time surface reconstruction for moving molecular fragments. *Pacific Symp. Biocomput.* 1997, **71,** 385–396
23 Olson, A.J., and Pique, M.E. Visualizing the future of molecular graphics. *SAR QSAR Environ. Res.* 1998, **8,** 233–247
24 Vajda, S., Weng, Z., Rosenfeld, R., and DeLisi, C. Effect of conformational flexibility and solvation on receptor–ligand binding free energies. *Biochemistry* 1994, **33,** 13977–13988
25 Weng, Z., Vajda, S., and Delisi, C. Prediction of protein complexes using empirical free energy functions. *Protein Sci.* 1996, **5,** 614–626
26 Zhang, C., Cornette, J.L., and Delisi, C. Consistency in structural energetics of protein folding and peptide recognition. *Protein Sci.* 1997, **6,** 1057–1064
27 Zhang, C., Vasmatzis, G., Cornette, J.L., and DeLisi. C. Determination of atomic desolvation energies from the structures of crystalized proteins. *J. Mol. Biol.* 1997, **267,** 707–726
28 Brooks, B.R., Bruccoleri, R.E., Olafson, B.D., States, D.J., Swaminathan, S., and Karplus, M. CHARMm: A program for macromolecular energy, minimization, and dynamics calculations. *J. Comput. Chem.* 1983, **4,** 187–217
29 Vasmatzis, G., Zhang, C., Cornette, J.L., and DeLisi, C. Computational determination of side chain specificity for pockets in class I MHC molecules. *Mol. Immunol.* 1996, **33,** 1231–1239
30 Binder, K., and Heermann, D.W. *Monte Carlo Simulations in Statistical Physics: An Introduction*, 2nd Ed. Springer-Verlag, Berlin, 1992
31 Bacon, D.J., and Moult, J. Docking by least-squares fitting of molecular surface patterns. *J. Mol. Biol.* 1992, **225,** 849–858
32 Caflisch, A., Niederer, P., and Anliker, M. Monte-Carlo docking of oligopeptides to proteins. *Proteins Struct. Funct. Genet.* 1992, **13,** 223–230
33 Cherfils, J., Duquerroy, S., and Janin, J. Protein–protein recognition analyzed by docking simulation. *Proteins Struct. Funct. Genet.* 1991, **11,** 271–280
34 Cherfils, J., Bizebard, T., Knossow, M., and Janin, J. Rigid-body docking with mutant constraints of influenza hemagglutinin with antibody HC19. *Proteins* 1994, **18,** 8–18
35 Fischer, D., Lin, S.L., Wolfson, H.L., and Nussinov, R. A geometry-based suite of molecular docking processes. *J. Mol. Biol.* 1995, **248,** 459–477
36 Goodsell, D.S., and Olson, A.J. Automated docking of substrates to proteins by simulated annealing. *Proteins Struct. Funct. Genet.* 1990, **8,** 195–202
37 Hart, T.N., and Read, R.J. A multiple-start Monte-Carlo docking method. *Proteins Struct. Funct. Genet.* 1992, **13,** 206–222
38 Helmer-Citterich, M., and Tramontano, A. PUZZLE: A new method for automated protein docking based on surface shape complementarity. *J. Mol. Biol.* 1994, **235,** 1021–1031
39 Janin, J. Principles of protein–protein recognition from structure to thermodynamics. *Biochimie* 1995, **77,** 497–505
40 Jiang, F., and Kim, S. Soft docking: Matching of molecular surface cubes. *J. Mol. Biol.* 1991, **219,** 79–102
41 Judson, R.S., Tan, Y.T., Mori, E., Melius, C., Jaeger, E.P., Treasurywala, A.M., and Mathiowetz, A. Docking flexible molecules: A case study of three proteins. *J. Comput. Chem.* 1995, **16,** 1405–1419

42 Meng, E.C., Shoichet, B.K., and Kuntz, I.D. Automated docking with grid-based energy evaluation. *J. Comput. Chem.* 1992, **13,** 505–524
43 Meng, E.C., Gschwend, D.A., Blaney, J.M., and Kuntz, I.D. Orientational sampling and rigid-body minimization in molecular docking. *Proteins Struct. Funct. Genet.* 1993, 17, 266–278
44 Shoichet, B.K., and Kuntz, I.D. Protein docking and complementarity. *J. Mol. Biol.* 1991, **221,** 327–346
45 Shoichet, B.K., Bodian, D.L., and Kuntz, I.D. Molecular docking using shape descriptors. *J. Comput. Chem.* 1992, **13,** 380–397
46 Strynadka, N.C., Eisenstein, M., Katchalski-Katzir, E., Shoichet, B.K., Kuntz, I.D., Abagyan, R., Totrov, M., Janin, J., Cherfils, J., Zimmerman, F., Olson, A., Duncan, B., Rao, M., Jackson, R., Sternberg, M., and James, M.N. Molecular docking programs successfully predict the binding of a beta-lactamase inhibitory protein to TEM-1 beta-lactamase. *Nature Struct. Biol.* 1996, **3,** 233–239. [See comments]

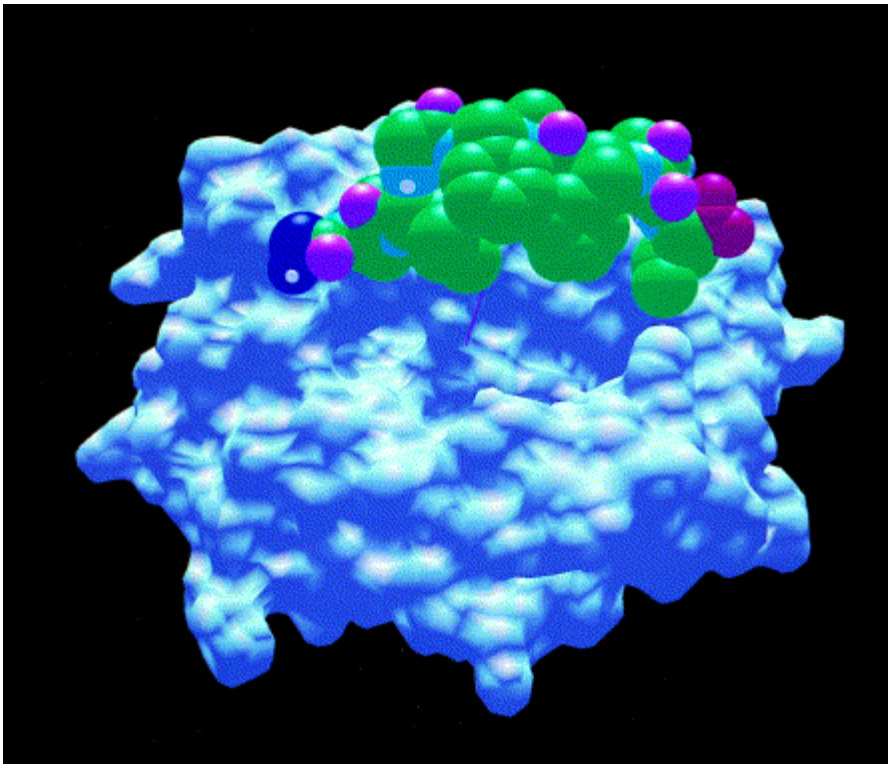

Color Plate 1. The MHC–peptide complex. This is a snapshot of the VRDD display (the pure OpenGL version). The receptor, a 180-residue HLA-A2 MHC molecule, is shown in the noncolored solvent-accessible surface (SAS) representation. The deep groove in the center of the molecule is the binding site. The ligand is a nine-residue-long peptide shown as a van der Waals filling model. The coloring scheme is as follows: blue, charged nitrogen atoms; cyan, polar nitrogen atoms; magenta, polar oxygen atoms; red, charged oxygen atoms; green, carbon atoms. A line connecting the centers of the two molecules is also shown (in red).

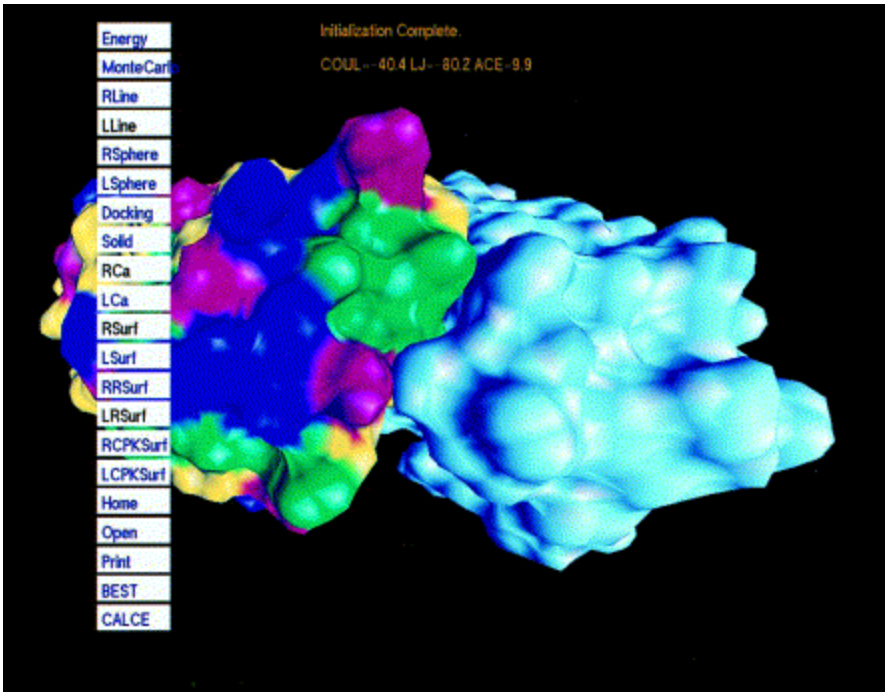

Color Plate 2. The barnase–barstar complex. This is a snapshot of the VRDD display (the CAVE version). The receptor barnase is shown in the noncolored solvent-accessible surface (SAS) representation. The ligand barstar is shown in the SAS representation colored according to residue types. Its orientation with respect to the receptor is being explored. The white menu that controls different modes of the program can be activated by pressing the right wand button at any time.