

A System for Interactive Molecular Dynamics Simulation

John E. Stone
Theoretical Biophysics Group
Beckman Institute
405 N. Mathews Ave.
Urbana, IL 61801
johns@ks.uiuc.edu

Justin Gullingsrud
Beckman Institute
University of Illinois
405 N. Mathews Ave.
Urbana, IL 61801
justin@ks.uiuc.edu

Klaus Schulten
Beckman Institute
University of Illinois
405 N. Mathews Ave.
Urbana, IL 61801
kschulte@ks.uiuc.edu

ABSTRACT

We have implemented a system termed Interactive Molecular Dynamics (IMD), which permits manipulation of molecules in molecular dynamics simulations with real-time force feedback and graphical display. Communication is achieved through an efficient socket connection between the visualization program (VMD) and a molecular dynamics program (NAMD) running on single or multiple machines. A natural force feedback interface for molecular steering is provided by a haptic device. We model the effect of simulation speed on the haptic feedback, and discuss results of an IMD study of a 4,000 atom system, the gramicidin A channel.

Categories and Subject Descriptors

H.5.2 [Information Interfaces]: User Interfaces—*Haptic I/O*; I.3.6 [Computer Graphics]: Methodology and Techniques—*Interaction Techniques*; I.6.8 [Simulation and Modeling]: Types of Simulation—*Parallel*

General Terms

Algorithms, Design

Keywords

Molecular Dynamics, Haptic Feedback, Steered Molecular Dynamics

1. INTRODUCTION

The binding properties of biomolecules and their response to mechanical forces can now be studied directly using single-molecule micromanipulation experiments. The techniques involved in such experiments include atomic force microscopy (AFM), optical tweezers, and the surface force apparatus; see [2] for a review. The insight gained by these investigations has inspired us and others [4, 15, 12] to adopt a similar approach for the study of biomolecules by means of computer simulations.

One such simulation approach, termed Steered Molecular Dynamics (SMD) [7], has already provided important qualitative insights into biologically relevant problems, including identification of binding pathways and the explanation of elastic properties of proteins. In SMD, selected atoms are steered by a moving restraint point attached via a linear spring. A limitation of SMD is that the pulling direction and spring constant of the restraint must be set before the simulation is begun. This restriction has been a consequence of the high computational cost of molecular dynamics simulations and the need to run in batch mode at supercomputing facilities. With modern high-performance workstations, however, SMD can be implemented as an interactive system, rather than in batch mode. Such an arrangement permits the user to make adjustments to the applied forces based on the progress of the simulation. We term this approach Interactive Molecular Dynamics (IMD).

Our goal for IMD is to provide real-time feedback and steering for molecular dynamics simulations of biomolecules. Force feedback in particular will enhance the already substantial insight that can be gained from SMD. Incorporating force feedback poses a significant challenge to an IMD system; therefore, we began with software designed to handle large, complex biomolecular systems. In this article, we describe our design and implementation of a communication scheme connecting a force feedback device, visualization program, and simulation engine. We also analyze the performance of these components, both in terms of an abstract model and in the context of a representative biological application.

2. SOFTWARE COMPONENTS

IMD consists of three primary components: a haptic device controlled by a server which generates the force environment, a molecular dynamics simulation for determining the effects of force application, and a visualization program for display of the results; see Fig. 1. We shall refer to the latter two elements by the names of the programs used in our implementation: NAMD and VMD. All three components communicate through TCP/IP network connections and therefore can reside on different machines, though this is not required.

VMD [5] was a nearly ideal program to serve as a front-end for IMD. Since its earliest versions, VMD was designed to work with 3D spatial trackers and other non-traditional input devices. Also, unlike many other programs which are geared towards the display of static molecular structures,

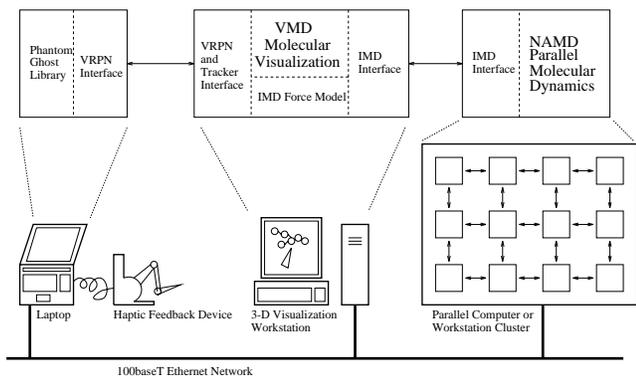


Figure 1: Decomposition of IMD components into asynchronous communicating processes.

VMD was designed to animate the progress of molecular dynamics simulations. VMD’s code structure facilitated the extension of the program’s tracker support to include a haptic interface and the handling of unbuffered communication between VMD and the simulation (see Sec. 3). VMD’s support for hardware 3D acceleration enabled coordinate sets from the molecular dynamics simulation to be rendered in real time.

The molecular dynamics platform for IMD has been NAMD [11, 8], a fast, scalable [3], program which implements the popular Charmm force field for molecular dynamics. The absolute speed of NAMD is an essential ingredient of a responsive IMD system. NAMD’s scalability enables us to use IMD to study large, biologically important systems as well as small ones, given enough computational power. Our use of an established molecular dynamics force field allows us to compare results from IMD simulations to other molecular dynamics investigations.

3. IMPLEMENTATION

IMD requires efficient network communication between the visualization front-end and the molecular dynamics back-end. While the network bandwidth requirements for performing IMD are quite low relative to the computational requirements, latency is a major concern. Our earlier efforts in the area of interactive steering of molecular dynamics simulations used MDComm [10]. We found that MDComm’s coordinate buffering mechanism would have imposed unacceptably high communication latency when coupled with a haptic device. We replaced the MDComm modules in VMD and NAMD with custom sockets code to allow VMD and NAMD to communicate efficiently. The communications protocol consists of a small 16 byte header, followed by a data segment whose content and length is specified by the header. For maximum efficiency, the header and the data are sent together in one socket write, ideally in a single network packet. The entire IMD interface consists of a small set of C-callable functions which can be adapted to any molecular dynamics and/or visualization programs. An IMD interface was recently written for the ProtoMol [9] molecular dynamics framework.

Each time through its event loop, VMD checks for a coordinate set from the MD program, updates the representation

geometry if a new coordinate set was received, redraws the screen, and updates the restraint position of the haptic device. If the user has applied a force through the haptic device, VMD routes this force to NAMD, which then integrates the force into the equations of motion for the molecule. VMD communicates with a server provided by the VRPN library [14]; the server controls the haptic environment experienced by the user handling the device. Updates to the haptic restraint point are made when VMD receives a new coordinate set from NAMD; while awaiting an update, the haptic server applies smooth feedback forces based on the most recent restraint point position. This scheme of splitting the haptic, visualization, and simulation components into three communicating, asynchronous processes has been employed successfully elsewhere [6].

4. FORCE MODEL

In IMD, physical forces generated by the user’s hand are translated into forces on simulated atoms. Three scaling parameters define the relationship between the virtual haptic environment and the simulation environment. The first, which we call α , is the ratio of wall clock time required for a molecular dynamics timestep to the physical simulation time, and is presently at least 10^{12} . The second parameter, β , is the amount by which forces applied by the user are scaled before being applied to atoms in the simulation. Unlike α , β can be set to any value the user desires; however, as we shall see, the value of β has important implications for the sensitivity and responsiveness of the IMD system. Finally, atomic coordinates are scaled up by a factor γ in order to be perceptible to humans. This factor is ultimately limited by the workspace of the haptic device, although one could in practice focus on a small part of the simulation region in order to increase the effective spatial resolution of the haptic interface.

With these parameters in mind, we now derive a model of the forces experienced by a user in an IMD session. An important result is the quantification of the role of simulation speed in the response of the haptic device to applied forces.

In the user-haptic coordinate system, let F_{haptic} be the force exerted by user’s hand on the haptic device. The particle experiences a force βF_{haptic} , taking into account the scaling of applied forces. Using the scaling parameters described above, we can relate the dynamics of the particle as perceived in the haptic environment to the dynamics of the particle in the simulation. Let x and t be the space and time coordinate, respectively, of the particle in the haptic environment, and χ and τ be the corresponding simulation variables. The dynamics of an otherwise free particle in the simulation are given by

$$m \frac{d^2 \chi}{d\tau^2} = \beta F_{haptic} = m \frac{\alpha^2}{\gamma} \frac{d^2 x}{dt^2} . \quad (1)$$

The user interprets the motion of the particle in terms of F_{haptic} and an effective mass m^* , defined by

$$F_{haptic} = m^* \frac{d^2 x}{dt^2} . \quad (2)$$

Combining equations (1) and (2), we arrive at an expression for the effective mass in terms of the actual mass of the

particle and the scaling parameters:

$$m^* = \frac{\alpha^2}{\beta\gamma} m. \quad (3)$$

Consider now the special case in which the applied force is proportional to the distance of the particle in the virtual environment from a restraint point that moves with the haptic controller. In this case $F_{haptic} = -kx$, where k is the spring constant defining the strength of the restraint. The atom experiences an external force $\beta kx = k\beta\gamma\chi$. If the particle is also subject to a potential U due to the presence of other atoms in the system, the dynamics of the particle will be given by

$$m \frac{d^2\chi}{dt^2} = -k\beta\gamma\chi - \frac{dU}{d\chi}. \quad (4)$$

Equations (3) and (4) together describe the interplay of the scaling parameters in determining the environment experienced by the user in IMD. In order for the system to respond to applied forces, m^* must be reasonably small, i.e. a few kilograms, since the maximum force that can be exerted by modern haptic devices is $\sim 1 - 10$ N. According to Eq. (3) m^* depends strongly on the speed of the simulation. The effective mass can be made smaller by scaling up the force applied by the user. However, this cannot be done without limit since, as Eq. (4) suggests, the applied forces will overwhelm the forces due to the simulation potential. The user would not be able to feel the potential, only the response of the particle restrained by the haptic device. For purposes of quantitative analysis of the system’s response to applied forces, it has been shown that stiff restraints are necessary for accurate sampling of the interatomic potential [1].

5. RESULTS

Table 1: Dependence of VMD frame rate on the rate of coordinate updates for two different representation styles. Solid reps corresponds to the representation depicted in Fig. 2; line reps was a lines-only representation. Frame rates were gathered from timings of five sets of 200 screen redraws.

coordinate updates/sec	solid reps (fps)	line reps (fps)
(disabled)	42.9 \pm .09	124.5 \pm .05
1.22	36.7 \pm 1.8	99.4 \pm 5.7
2.45	37.9 \pm .24	99.4 \pm .60
4.90	35.0 \pm .44	85.6 \pm .88
9.80	29.4 \pm .25	57.6 \pm .35
19.60	18.2 \pm .45	26.0 \pm 2.6

The ion channel gramicidin A was used as a testbed both for developing the VMD interface for IMD as well as evaluating the performance of the implementation. Researchers can use SMD to study the process of ion permeation through the channel, but its modest size (4385 atoms) also make it amenable to study with IMD; see Fig. 2. As described in Sec. 4, the scaling of the haptic workspace relative to the simulation, the scaling of the forces sent to the simulation, and the stiffness of the restraint all play a role in the sensitivity of the haptic device to events in the simulation, and



Figure 2: Example of using IMD to study the permeation of ions through a gramicidin A channel. Depicted are a sodium ion (large sphere), a gramicidin A dimer in ribbon representation, four of the lipid molecules comprising the membrane in which the channel is embedded (for clarity, not all are shown), and the surrounding water. Researchers using a haptic device or 3D tracker can interactively steer the ion through the channel and gather applied force data for analysis.

vice versa. In VMD, all three of these parameters are user-configurable.

Performance of our IMD system may be assessed quantitatively by means of Table 1. Simulations were performed on two dual-processor 500 MHz Alpha workstations connected via 100baseT Ethernet to each other and to the visualization program; network latency was less than 1 ms. The visualization program ran on a single processor of a 450 MHz Sun Ultra 80 with Expert 3D graphics. Redraws were decoupled from the vertical refresh rate of the monitor. Viewpoint-dependent rendering of pre-computed geometry was done for each frame; molecular geometry was recalculated only when new coordinate sets were received from the MD program.

At the highest coordinate transfer rates, VMD’s performance declined somewhat due to the need to retrieve all available data from the IMD connection before proceeding with graphics computations. However, for simpler representations (column 3 of Table 1), VMD’s frame rate is always higher than NAMD’s coordinate update rate. Since the position of the haptic restraint is updated only as fast as new coordinate frames are received, the best use of our system was to set the coordinate transfer rate as high as possible and use a molecular representation just simple enough to keep the rendering rate higher than the coordinate transfer rate.

6. CONCLUSIONS

Our implementation of IMD has demonstrated that advances in available computational power make possible a new, interactive mode of biophysical investigation. We have described quantitatively the effect of simulation speed on the sensitivity and responsiveness of the haptic system for any

IMD implementation employing our paradigm; these results should prove to be a useful guide for future designs.

In a related work [13], molecular dynamics simulations were visualized in real time in a virtual environment, with 3D trackers used for steering. Our system goes further in delivering force as well as visual feedback to the user. As we have seen, this is a nontrivial extension, as the force experienced by the user depends strongly on the speed of the simulation. A production-quality parallel molecular dynamics program such as NAMD as well as high-speed computers were key to IMD's success.

Our current implementation of IMD has shown that it is important to minimize sources of latency and unnecessary synchronization between the components of the system. The existing decomposition of the IMD components into three concurrent communicating processes was sufficient for our early investigations into IMD, however a finer-grained decomposition utilizing multithreading could provide further reductions in communications latency and entirely eliminate the remaining synchronization of haptic constraint updates with visual display. With these improvements in place, only the coordinate update rate of a given molecular dynamics simulation would limit the sensitivity of the haptic device and the overall performance of the IMD system.

7. ACKNOWLEDGMENTS

The authors wish to acknowledge support from the NIH Resource for Macromolecular Modeling and Bioinformatics (NIH PHS 5 P41 RR05969, <http://www.ks.uiuc.edu>), Simulations of Supramolecular Biological Systems (NRAC MCA93S028), Computational - Experimental Collaborations on Biological Nanostructures (Roy J. Carver Charitable Trust), and the NIH National Research Resource in Molecular Graphics and Microscopy at the University of North Carolina at Chapel Hill.

8. ADDITIONAL AUTHORS

Paul Grayson, Massachusetts Institute of Technology, email: pdg@mit.edu).

9. REFERENCES

- [1] M. Balsera, S. Stepaniants, S. Izrailev, Y. Oono, and K. Schulten. Reconstructing potential energy functions from simulated force-induced unbinding processes. *Biophys. J.*, 73:1281–1287, 1997.
- [2] S. Block and K. Svoboda. Biological applications of optical forces. *Ann. Rev. Biophys. Biomol. Struct.*, 23:247–285, 1994.
- [3] R. K. Brunner, J. C. Phillips, and L. V. Kalé. Scalable molecular dynamics for large biomolecular systems. In *Proceedings of the 2000 ACM/IEEE SC2000 Conference*, 2000.
- [4] H. Grubmüller, B. Heymann, and P. Tavan. Ligand binding and molecular mechanics calculation of the streptavidin-biotin rupture force. *Science*, 271:997–999, 1996.
- [5] W. F. Humphrey, A. Dalke, and K. Schulten. VMD – Visual Molecular Dynamics. *J. Mol. Graphics*, 14:33–38, 1996.
- [6] F. D. IX, H. Qin, A. Kaufman, and J. El-Sana. Haptic sculpting of dynamic surfaces. In *1999 Symposium on Interactive 3D Graphics*, pages 103–110, 1999.
- [7] S. Izrailev, S. Stepaniants, B. Isralewitz, D. Kosztin, H. Lu, F. Molnar, W. Wriggers, and K. Schulten. Steered molecular dynamics. In P. Deuffhard, J. Hermans, B. Leimkuhler, A. E. Mark, S. Reich, and R. D. Skeel, editors, *Computational Molecular Dynamics: Challenges, Methods, Ideas*, volume 4 of *Lecture Notes in Computational Science and Engineering*, pages 39–65. Springer-Verlag, Berlin, 1998.
- [8] L. Kalé, R. Skeel, M. Bhandarkar, R. Brunner, A. Gursoy, N. Krawetz, J. Phillips, A. Shinozaki, K. Varadarajan, and K. Schulten. NAMD2: Greater scalability for parallel molecular dynamics. *J. Comp. Phys.*, 151:283–312, 1999.
- [9] T. Matthey and J. A. Izaguirre. Protomol: A molecular dynamics framework with incremental parallelization. In *Proceedings of 10th SIAM Conference on Parallel Processing for Scientific Computing*, 2001.
- [10] M. Nelson, W. Humphrey, A. Gursoy, A. Dalke, L. Kalé, R. Skeel, K. Schulten, and R. Kufryn. MDScope – A visual computing environment for structural biology. In S. Atluri, G. Yagawa, and T. Cruse, editors, *Computational Mechanics 95*, volume 1, pages 476–481, 1995.
- [11] M. Nelson, W. Humphrey, A. Gursoy, A. Dalke, L. Kalé, R. D. Skeel, and K. Schulten. NAMD – A parallel, object-oriented molecular dynamics program. *Int. J. Supercomput. Appl. High Perform. Comput.*, 10:251–268, 1996.
- [12] E. Paci and M. Karplus. Forced unfolding of fibronectin type 3 modules: An analysis by biased molecular dynamics simulations. *J. Mol. Biol.*, 288:441–459, 1999.
- [13] J. F. Prins, J. Hermans, G. Mann, L. S. Nyland, and M. Simons. A virtual environment for steered molecular dynamics. *Future Generation Computer Systems*, 15:485–495, 1999.
- [14] R. M. Taylor II. VRPN: Virtual Reality Peripheral Network, 1998. <http://www.cs.unc.edu/Research/vrpn/>.
- [15] R. C. Wade, R. R. Gabdouliline, S. K. Ludemann, and V. Lounnas. Electrostatic steering and ionic tethering in enzyme-ligand binding: Insights from simulations. *Proc. Natl. Acad. Sci. USA*, 95:5942–5949, 1998.