

Published in final edited form as:

J Chem Theory Comput. 2011 April 12; 7(4): 834–842. doi:10.1021/ct100670m.

The Langevin Hull: Constant pressure and temperature dynamics for non-periodic systems

Charles F. Vardeman II, Kelsey M. Stocker, and J. Daniel Gezelter*

Department of Chemistry and Biochemistry, University of Notre Dame, Notre Dame, Indiana 46556

Abstract

We have developed a new isobaric-isothermal (NPT) algorithm which applies an external pressure to the facets comprising the convex hull surrounding the system. A Langevin thermostat is also applied to the facets to mimic contact with an external heat bath. This new method, the “Langevin Hull”, can handle heterogeneous mixtures of materials with different compressibilities. These systems are problematic for traditional affine transform methods. The Langevin Hull does not suffer from the edge effects of boundary potential methods, and allows realistic treatment of both external pressure and thermal conductivity due to the presence of an implicit solvent. We apply this method to several different systems including bare metal nanoparticles, nanoparticles in an explicit solvent, as well as clusters of liquid water. The predicted mechanical properties of these systems are in good agreement with experimental data and previous simulation work.

1 Introduction

The most common molecular dynamics methods for sampling configurations from an isobaric-isothermal (NPT) ensemble maintain a target pressure in a simulation by coupling the volume of the system to a *barostat*, which is an extra degree of freedom propagated along with the particle coordinates. These methods require periodic boundary conditions, because when the instantaneous pressure in the system differs from the target pressure, the volume is reduced or expanded using *affine transforms* of the system geometry. An affine transform scales the size and shape of the periodic box as well as the particle positions within the box (but not the sizes of the particles). The most common constant pressure methods, including the Melchionna modification¹ to the Nosé-Hoover-Andersen equations of motion,^{2–4} the Berendsen pressure bath,⁵ and the Langevin Piston,^{6,7} all utilize scaled coordinate transformation to adjust the box volume. As long as the material in the simulation box has a relatively uniform compressibility, the standard affine transform approach provides an excellent way of adjusting the volume of the system and applying pressure directly via the interactions between atomic sites.

One problem with this approach appears when the system being simulated is an inhomogeneous mixture in which portions of the simulation box are incompressible relative to other portions. Examples include simulations of metallic nanoparticles in liquid environments, proteins at ice/water interfaces, as well as other heterogeneous or interfacial environments. In these cases, the affine transform of atomic coordinates will either cause numerical instability when the sites in the incompressible medium collide with each other, or will lead to inefficient sampling of system volumes if the barostat is set slow enough to avoid the instabilities in the incompressible region.

*Corresponding author. gezelter@nd.edu.

One may also wish to avoid affine transform periodic boundary methods to simulate *explicitly non-periodic systems* under constant pressure conditions. The use of periodic boxes to enforce a system volume requires either effective solute concentrations that are much higher than desirable, or unreasonable system sizes to avoid this effect. For example, calculations using typical hydration boxes solvating a protein under periodic boundary conditions are quite expensive. A 62 \AA^3 box of water solvating a moderately small protein like hen egg white lysozyme (PDB code: 1LYZ) yields an effective protein concentration of 100 mg/mL.⁸

Total protein concentrations in the cell are typically on the order of 160–310 mg/ml,⁹ and individual proteins have concentrations orders of magnitude lower than this in the cellular environment. The effective concentrations of single proteins in simulations may have significant effects on the structure and dynamics of simulated systems.

Boundary Methods

There have been a number of approaches to handle simulations of explicitly non-periodic systems that focus on constant or nearly-constant *volume* conditions while maintaining bulk-like behavior. Berkowitz and McCammon introduced a stochastic (Langevin) boundary layer inside a region of fixed molecules which effectively enforces constant temperature and volume (NVT) conditions.¹⁰ In this approach, the stochastic and fixed regions were defined relative to a central atom. Brooks and Karplus extended this method to include deformable stochastic boundaries.¹¹ The stochastic boundary approach has been used widely for protein simulations.

The electrostatic and dispersive behavior near the boundary has long been a cause for concern when performing simulations of explicitly non-periodic systems. Early work led to the surface constrained soft sphere dipole model (SCSSD)¹² in which the surface molecules are fixed in a random orientation representative of the bulk solvent structural properties. Belch *et al.*¹³ simulated clusters of TIPS2 water surrounded by a hydrophobic bounding potential. The spherical hydrophobic boundary induced dangling hydrogen bonds at the surface that propagated deep into the cluster, affecting most of the molecules in the simulation. This result echoes an earlier study which showed that an extended planar hydrophobic surface caused orientational preferences at the surface which extended relatively deep (7 Å) into the liquid simulation cell.¹⁴ The surface constrained all-atom solvent (SCAAS) model¹⁵ improved upon its SCSSD predecessor. The SCAAS model utilizes a polarization constraint which is applied to the surface molecules to maintain bulk-like structure at the cluster surface. A radial constraint is used to maintain the desired bulk density of the liquid. Both constraint forces are applied only to a pre-determined number of the outermost molecules.

Beglov and Roux have developed a boundary model in which the hard sphere boundary has a radius that varies with the instantaneous configuration of the solute (and solvent) molecules.¹⁶ This model contains a clear pressure and surface tension contribution to the free energy.

Restraining Potentials

Restraining *potentials* introduce repulsive potentials at the surface of a sphere or other geometry. The solute and any explicit solvent are therefore restrained inside the range defined by the external potential. Often the potentials include a weak short-range attraction to maintain the correct density at the boundary. Beglov and Roux have also introduced a restraining boundary potential which relaxes dynamically depending on the solute geometry and the force the explicit system exerts on the shell.¹⁷

Recently, Krilov *et al.* introduced a *flexible* boundary model that uses a Lennard-Jones potential between the solvent molecules and a boundary which is determined dynamically from the position of the nearest solute atom.^{18,19} This approach allows the confining potential to prevent solvent molecules from migrating too far from the solute surface, while providing a weak attractive force pulling the solvent molecules towards a fictitious bulk solvent. Although this approach is appealing and has physical motivation, nanoparticles do not deform far from their original geometries even at temperatures which vaporize the nearby solvent. For the systems like this, the flexible boundary model will be nearly identical to a fixed-volume restraining potential.

Hull methods

The approach of Kohanoff, Caro, and Finnis is the most promising of the methods for introducing both constant pressure and temperature into non-periodic simulations.^{20,21} This method is based on standard Langevin dynamics, but the Brownian or random forces are allowed to act only on peripheral atoms and exert forces in a direction that is inward-facing relative to the facets of a closed bounding surface. The statistical distribution of the random forces are uniquely tied to the pressure in the external reservoir, so the method can be shown to sample the isobaric-isothermal ensemble. Kohanoff *et al.* used a Delaunay tessellation to generate a bounding surface surrounding the outermost atoms in the simulated system. This is not the only possible triangulated outer surface, but guarantees that all of the random forces point inward towards the cluster.

In the following sections, we extend and generalize the approach of Kohanoff, Caro, and Finnis. The new method, which we are calling the “Langevin Hull” applies the external pressure, Langevin drag, and random forces on the *facets of the hull* instead of the atomic sites comprising the vertices of the hull. This allows us to decouple the external pressure contribution from the drag and random force. The methodology is introduced in section 2, tests on crystalline nanoparticles, liquid clusters, and heterogeneous mixtures are detailed in section 3. Section 4 summarizes our findings.

2 Methodology

The Langevin Hull uses an external bath at a fixed constant pressure (P) and temperature (T) with an effective solvent viscosity (η). This bath interacts only with the objects on the exterior hull of the system. Defining the hull of the atoms in a simulation is done in a manner similar to the approach of Kohanoff, Caro and Finnis.²⁰ That is, any instantaneous configuration of the atoms in the system is considered as a point cloud in three dimensional space. Delaunay triangulation is used to find all facets between coplanar neighbors.^{22,23} In highly symmetric point clouds, facets can contain many atoms, but in all but the most symmetric of cases, the facets are simple triangles in 3-space which contain exactly three atoms.

The convex hull is the set of facets that have *no concave corners* at an atomic site.^{24,25} This eliminates all facets on the interior of the point cloud, leaving only those exposed to the bath. Sites on the convex hull are dynamic; as molecules re-enter the cluster, all interactions between atoms on that molecule and the external bath are removed. Since the edge is determined dynamically as the simulation progresses, no *a priori* geometry is defined. The pressure and temperature bath interacts only with the atoms on the edge and not with atoms interior to the simulation.

Atomic sites in the interior of the simulation move under standard Newtonian dynamics,

$$m_i \dot{\mathbf{v}}_i(t) = -\nabla_i U, \quad (1)$$

where m_i is the mass of site i , $\mathbf{v}_i(t)$ is the instantaneous velocity of site i at time t , and U is the total potential energy. For atoms on the exterior of the cluster (i.e. those that occupy one of the vertices of the convex hull), the equation of motion is modified with an external force, \mathbf{F}_i^{ext} :

$$m_i \dot{\mathbf{v}}_i(t) = -\nabla_i U + \mathbf{F}_i^{ext}. \quad (2)$$

The external bath interacts indirectly with the atomic sites through the intermediary of the hull facets. Since each vertex (or atom) provides one corner of a triangular facet, the force on the facets are divided equally to each vertex. However, each vertex can participate in multiple facets, so the resultant force is a sum over all facets f containing vertex i :

$$\mathbf{F}_i^{ext} = \sum_{\substack{\text{facets } f \\ \text{containing } i}} \frac{1}{3} \mathbf{F}_f^{ext} \quad (3)$$

The external pressure bath applies a force to the facets of the convex hull in direct proportion to the area of the facet, while the thermal coupling depends on the solvent temperature, viscosity and the size and shape of each facet. The thermal interactions are expressed as a standard Langevin description of the forces,

$$\begin{aligned} \mathbf{F}_f^{ext} &= \text{external pressure} + \text{drag force} + \text{random force} \\ &= -\hat{n}_f P A_f - \Xi_f(t) \mathbf{v}_f(t) + \mathbf{R}_f(t) \end{aligned} \quad (4)$$

Here, A_f and \hat{n}_f are the area and (outward-facing) normal vectors for facet f , respectively. $\mathbf{v}_f(t)$ is the velocity of the facet centroid,

$$\mathbf{v}_f(t) = \frac{1}{3} \sum_{i=1}^3 \mathbf{v}_i, \quad (5)$$

and $\Xi_f(t)$ is an approximate (3×3) resistance tensor that depends on the geometry and surface area of facet f and the viscosity of the bath. The resistance tensor is related to the fluctuations of the random force, $\mathbf{R}(t)$, by the fluctuation-dissipation theorem,

$$\langle \mathbf{R}_f(t) \rangle = 0 \quad (6)$$

$$\langle \mathbf{R}_f(t) \mathbf{R}_f^T(t') \rangle = 2k_B T \Xi_f(t) \delta(t - t'). \quad (7)$$

Once the resistance tensor is known for a given facet, a stochastic vector that has the properties in Eq. (7) can be calculated efficiently by carrying out a Cholesky decomposition to obtain the square root matrix of the resistance tensor,

$$\Xi_f = \mathbf{S}\mathbf{S}^T, \quad (8)$$

where \mathbf{S} is a lower triangular matrix.²⁶ A vector with the statistics required for the random force can then be obtained by multiplying \mathbf{S} onto a random 3-vector \mathbf{Z} which has elements chosen from a Gaussian distribution, such that:

$$\langle \mathbf{Z}_i \rangle = 0, \quad \langle \mathbf{Z}_i \cdot \mathbf{Z}_j \rangle = \frac{2k_B T}{\delta t} \delta_{ij}, \quad (9)$$

where δt is the timestep in use during the simulation. The random force, $\mathbf{R}_f = \mathbf{S}\mathbf{Z}$, can be shown to have the correct properties required by Eq. (7).

Our treatment of the resistance tensor is approximate. Ξ_f for a rigid triangular plate would normally be treated as a 6×6 tensor that includes translational and rotational drag as well as translational-rotational coupling. The computation of resistance tensors for rigid bodies has been detailed elsewhere,²⁷⁻³⁰ but the standard approach involving bead approximations would be prohibitively expensive if it were recomputed at each step in a molecular dynamics simulation.

Instead, we are utilizing an approximate resistance tensor obtained by first constructing the Oseen tensor for the interaction of the centroid of the facet (f) with each of the subfacets $\ell = 1, 2, 3$,

$$T_{\ell f} = \frac{A_\ell}{8\pi\eta R_{\ell f}} \left(I + \frac{\mathbf{R}_{\ell f} \mathbf{R}_{\ell f}^T}{R_{\ell f}^2} \right) \quad (10)$$

Here, A_ℓ is the area of subfacet ℓ which is a triangle containing two of the vertices of the facet along with the centroid. $\mathbf{R}_{\ell f}$ is the vector between the centroid of facet f and the centroid of sub-facet ℓ , and I is the (3×3) identity matrix. η is the viscosity of the external bath.

The tensors for each of the sub-facets are added together, and the resulting matrix is inverted to give a 3×3 resistance tensor for translations of the triangular facet,

$$\Xi_f(t) = \left[\sum_{i=1}^3 T_{if} \right]^{-1}. \quad (11)$$

Note that this treatment ignores rotations (and translational-rotational coupling) of the facet. In compact systems, the facets stay relatively fixed in orientation between configurations, so this appears to be a reasonably good approximation.

We have implemented this method by extending the Langevin dynamics integrator in our code, OpenMD.^{31,32} At each molecular dynamics time step, the following process is carried out:

1. The standard inter-atomic forces ($\nabla_i U$) are computed.
2. Delaunay triangulation is carried out using the current atomic configuration.
3. The convex hull is computed and facets are identified.
4. For each facet:
 - a. The force from the pressure bath ($-\hat{n}_f P A_f$) is computed.
 - b. The resistance tensor ($\Xi_f(t)$) is computed using the viscosity (η) of the bath.
 - c. Facet drag ($-\Xi_f(t)\mathbf{v}_f(t)$) forces are computed.
 - d. Random forces ($\mathbf{R}_f(t)$) are computed using the resistance tensor and the temperature (T) of the bath.
5. The facet forces are divided equally among the vertex atoms.
6. Atomic positions and velocities are propagated.

The Delaunay triangulation and computation of the convex hull are done using calls to the qhull library.³³ There is a minimal penalty for computing the convex hull and resistance tensors at each step in the molecular dynamics simulation (roughly $0.02 \times$ cost of a single force evaluation), and the convex hull is remarkably easy to parallelize on distributed memory machines (see Appendix A).

3 Tests & Applications

To test the new method, we have carried out simulations using the Langevin Hull on: 1) a crystalline system (gold nanoparticles), 2) a liquid droplet (SPC/E water),³⁴ and 3) a heterogeneous mixture (gold nanoparticles in an SPC/E water droplet). In each case, we have computed properties that depend on the external applied pressure. Of particular interest for the single-phase systems is the isothermal compressibility,

$$\kappa_T = -\frac{1}{V} \left(\frac{\partial V}{\partial P} \right)_T. \quad (12)$$

One problem with eliminating periodic boundary conditions and simulation boxes is that the volume of a three-dimensional point cloud is not well-defined. In order to compute the compressibility of a bulk material, we make an assumption that the number density, $\rho = \frac{N}{V}$, is uniform within some region of the point cloud. The compressibility can then be expressed in terms of the average number of particles in that region,

$$\kappa_T = -\frac{1}{N} \left(\frac{\partial N}{\partial P} \right)_T. \quad (13)$$

The region we used is a spherical volume of 20 Å radius centered in the middle of the cluster with a roughly 25 Å radius. N is the average number of molecules found within this

region throughout a given simulation. The geometry of the region is arbitrary, and any bulk-like portion of the cluster can be used to compute the compressibility.

One might assume that the volume of the convex hull could simply be taken as the system volume V in the compressibility expression (Eq. 12), but this has implications at lower pressures (which are explored in detail in the section on water droplets).

The metallic force field in use for the gold nanoparticles is the quantum Sutton-Chen (QSC) model.³⁵ In all simulations involving point charges, we utilized damped shifted-force (DSF) electrostatics³⁶ which is a variant of the Wolf summation³⁷ that has been shown to provide good forces and torques on molecular models for water in a computationally efficient manner.³⁶ The damping parameter (α) was set to 0.18 \AA^{-1} , and the cutoff radius was set to 12 \AA . The Spohr potential was adopted in depicting the interaction between metal atoms and the SPC/E water molecules.³⁸

3.1 Bulk Modulus of gold nanoparticles

The compressibility (and its inverse, the bulk modulus) is well-known for gold, and is captured well by the embedded atom method (EAM)³⁹ potential and related multi-body force fields. In particular, the quantum Sutton-Chen potential gets nearly quantitative agreement with the experimental bulk modulus values, and makes a good first test of how the Langevin Hull will perform at large applied pressures.

The Sutton-Chen (SC) potentials are based on a model of a metal which treats the nuclei and core electrons as pseudo-atoms embedded in the electron density due to the valence electrons on all of the other atoms in the system.⁴⁰ The SC potential has a simple form that closely resembles the Lennard Jones potential,

$$U_{tot} = \sum_i \left[\frac{1}{2} \sum_{j \neq i} D_{ij} V_{ij}^{pair}(r_{ij}) - c_i D_{ii} \sqrt{\rho_i} \right], \quad (14)$$

where V_{ij}^{pair} and ρ_i are given by

$$V_{ij}^{pair}(r) = \left(\frac{\alpha_{ij}}{r_{ij}} \right)^{n_{ij}}, \quad \rho_i = \sum_{j \neq i} \left(\frac{\alpha_{ij}}{r_{ij}} \right)^{m_{ij}}. \quad (15)$$

V_{ij}^{pair} is a repulsive pairwise potential that accounts for interactions between the pseudoatom cores. The $\sqrt{\rho_i}$ term in Eq. (14) is an attractive many-body potential that models the interactions between the valence electrons and the cores of the pseudo-atoms. D_{ij} and D_{ii} set the appropriate overall energy scale, c_i scales the attractive portion of the potential relative to the repulsive interaction and α_{ij} is a length parameter that assures a dimensionless form for ρ . These parameters are tuned to various experimental properties such as the density, cohesive energy, and elastic moduli for FCC transition metals. The quantum Sutton-Chen (QSC) formulation matches these properties while including zero-point quantum corrections for different transition metals.^{35,41}

In bulk gold, the experimentally-measured value for the bulk modulus is 180.32 GPa, while previous calculations on the QSC potential in periodic-boundary simulations of the bulk

crystal have yielded values of 175.53 GPa.⁴¹ Using the same force field, we have performed a series of 1 ns simulations on gold nanoparticles of three different radii: 20 Å (1985 atoms), 30 Å (6699 atoms), and 40 Å (15707 atoms) utilizing the Langevin Hull at a variety of applied pressures ranging from 0 – 10 GPa. For the 40 Å radius nanoparticle we obtain a value of 177.55 GPa for the bulk modulus of gold, in close agreement with both previous simulations and the experimental bulk modulus reported for gold single crystals.⁴² The smaller gold nanoparticles (30 and 20 Å radii) have calculated bulk moduli of 215.58 and 208.86 GPa, respectively, indicating that smaller nanoparticles are somewhat stiffer (less compressible) than the larger nanoparticles. This stiffening of the small nanoparticles may be related to their high degree of surface curvature, resulting in a lower coordination number of surface atoms relative to the the surface atoms in the 40 Å radius particle.

We obtain a gold lattice constant of 4.051 Å using the Langevin Hull at 1 atm, close to the experimentally-determined value for bulk gold and the value for gold simulated using the QSC potential and periodic boundary conditions (4.079 Å and 4.088Å, respectively).⁴¹ The slightly smaller calculated lattice constant is most likely due to the presence of surface tension in the non-periodic Langevin Hull cluster, an effect absent from a bulk simulation. The specific heat of a 40 Å gold nanoparticle under the Langevin Hull at 1 atm is $24.914 \frac{\text{J}}{\text{mol K}}$, which compares very well with the experimental value of $25.42 \frac{\text{J}}{\text{mol K}}$.

We note that the Langevin Hull produces rapidly-converging behavior for structures that are started far from equilibrium. In Fig. 4 we show how the pressure and temperature respond to the Langevin Hull for nanoparticles that were initialized far from the target pressure and temperature. As expected, the rate at which thermal equilibrium is achieved depends on the total surface area of the cluster exposed to the bath as well as the bath viscosity. Pressure that is applied suddenly to a cluster can excite breathing vibrations, but these rapidly damp out (on time scales of 30 – 50 ps).

3.2 Compressibility of SPC/E water clusters

Prior molecular dynamics simulations on SPC/E water (both in NVT⁴³ and NPT^{44,45} ensembles) have yielded values for the isothermal compressibility that agree well with experiment.⁴⁶ The results of two different approaches for computing the isothermal compressibility from Langevin Hull simulations for pressures between 1 and 3000 atm are shown in Fig. 5 along with compressibility values obtained from both other SPC/E simulations and experiment.

Isothermal compressibility values calculated using the number density (Eq. 13) expression are in good agreement with experimental and previous simulation work throughout the 1 – 1000 atm pressure regime. Compressibilities computed using the Hull volume, however, deviate dramatically from the experimental values at low applied pressures. The reason for this deviation is quite simple: at low applied pressures, the liquid is in equilibrium with a vapor phase, and it is entirely possible for one (or a few) molecules to drift away from the liquid cluster (see Fig. 6). At low pressures, the restoring forces on the facets are very gentle, and this means that the hulls often take on relatively distorted geometries which include large volumes of empty space.

At higher pressures, the equilibrium strongly favors the liquid phase, and the hull geometries are much more compact. Because of the liquid-vapor effect on the convex hull, the regional number density approach (Eq. 13) provides more reliable estimates of the compressibility.

In both the traditional compressibility formula (Eq. 12) and the number density version (Eq. 13), multiple simulations at different pressures must be done to compute the first derivatives.

It is also possible to compute the compressibility using the fluctuation dissipation theorem using either fluctuations in the volume,⁴⁷

$$\kappa_T = \frac{\langle V^2 \rangle - \langle V \rangle^2}{Vk_B T}, \quad (16)$$

or, equivalently, fluctuations in the number of molecules within the fixed region,

$$\kappa_T = \frac{\langle N^2 \rangle - \langle N \rangle^2}{Nk_B T}. \quad (17)$$

Thus, the compressibility of each simulation can be calculated entirely independently from other trajectories. Compressibility calculations that rely on the hull volume will still suffer the effects of the empty space due to the vapor phase; for this reason, we recommend using the number density (Eq. 13) or number density fluctuations (Eq. 17) for computing compressibilities. We obtained the results in Fig. 5 using a sampling radius that was approximately 80% of the mean distance between the center of mass of the cluster and the hull atoms. This ratio of sampling radius to average hull radius excludes the problematic vapor phase on the outside of the cluster while including enough of the liquid phase to avoid poor statistics due to fluctuating local densities.

A comparison of the oxygen-oxygen radial distribution functions for SPC/E water simulated using both the Langevin Hull and more traditional periodic boundary methods – both at 1 atm and 300K – reveals an understructuring of water in the Langevin Hull that manifests as a slight broadening of the solvation shells. This effect may be due to the introduction of a liquid-vapor interface in the Langevin Hull simulations (an interface which is missing in most periodic simulations of bulk water). Vapor-phase molecules contribute a small but nearly flat portion of the radial distribution function.

3.3 Molecular orientation distribution at cluster boundary

In order for a non-periodic boundary method to be widely applicable, it must be constructed in such a way that they allow a finite system to replicate the properties of the bulk. Early non-periodic simulation methods (e.g. hydrophobic boundary potentials) induced spurious orientational correlations deep within the simulated system.^{13,14} This behavior spawned many methods for fixing and characterizing the effects of artificial boundaries including methods which fix the orientations of a set of edge molecules.^{12,15}

As described above, the Langevin Hull does not require that the orientation of molecules be fixed, nor does it utilize an explicitly hydrophobic boundary, or orientational or radial constraints. Therefore, the orientational correlations of the molecules in water clusters are of particular interest in testing this method. Ideally, the water molecules on the surfaces of the clusters will have enough mobility into and out of the center of the cluster to maintain bulk-like orientational distribution in the absence of orientational and radial constraints. However, since the number of hydrogen bonding partners available to molecules on the exterior are limited, it is likely that there will be an effective hydrophobicity of the hull.

To determine the extent of these effects, we examined the orientations exhibited by SPC/E water in a cluster of 1372 molecules at 300 K and at pressures ranging from 1 – 1000 atm. The orientational angle of a water molecule is described by

$$\cos\theta = \frac{\vec{r}_i \cdot \vec{\mu}_i}{|\vec{r}_i| |\vec{\mu}_i|} \quad (18)$$

where \vec{r}_i is the vector between molecule i 's center of mass and the cluster center of mass, and $\vec{\mu}_i$ is the vector bisecting the H-O-H angle of molecule i . Bulk-like distributions will result in $\langle \cos\theta \rangle$ values close to zero. If the hull exhibits an overabundance of externally-oriented oxygen sites, the average orientation will be negative, while dangling hydrogen sites will result in positive average orientations.

Fig. 7 shows the distribution of $\cos\theta$ values for molecules in the interior of the cluster (squares) and for molecules included in the convex hull (circles).

As expected, interior molecules (those not included in the convex hull) maintain a bulk-like structure with a uniform distribution of orientations. Molecules included in the convex hull show a slight preference for values of $\cos\theta < 0$. These values correspond to molecules with oxygen directed toward the exterior of the cluster, forming dangling hydrogen bond acceptor sites.

The orientational preference exhibited by water molecules on the hull is significantly weaker than the preference caused by an explicit hydrophobic bounding potential. Additionally, the Langevin Hull does not require that the orientation of any molecules be fixed in order to maintain bulk-like structure, even near the cluster surface.

Previous molecular dynamics simulations of SPC/E liquid/vapor interfaces using periodic boundary conditions have shown that molecules on the liquid side of interface favor a similar orientation where oxygen is directed away from the bulk.⁴⁸ These simulations had well-defined liquid and vapor phase regions equilibrium and it was observed that *vapor* molecules generally had one hydrogen protruding from the surface, forming a dangling hydrogen bond donor. Our water clusters do not have a true vapor region, but rather a few transient molecules that leave the liquid droplet (and which return to the droplet relatively quickly). Although we cannot obtain an orientational preference of vapor phase molecules in a Langevin Hull simulation, but we do agree with previous estimates of the orientation of *liquid phase* molecules at the interface.

3.4 Heterogeneous nanoparticle/water mixtures

To further test the method, we simulated gold nanoparticles ($r = 18 \text{ \AA}$, 1433 atoms) solvated by explicit SPC/E water clusters (5000 molecules) using a model for the gold/water interactions that has been used by Dou *et al.* for investigating the separation of water films near hot metal surfaces.³⁸ The Langevin Hull was used to sample pressures of 1, 2, 5, 10, 20, 50, 100 and 200 atm, while all simulations were done at a temperature of 300 K. At these temperatures and pressures, there is no observed separation of the water film from the surface.

In Fig. 8 we show the density of water and gold as a function of the distance from the center of the nanoparticle. Higher applied pressures appear to destroy structural correlations in the outermost monolayer of the gold nanoparticle as well as in the water at the near the metal/water interface. Simulations at increased pressures exhibit significant overlap of the gold and water densities, indicating a less well-defined interfacial surface.

At even higher pressures (500 atm and above), problems with the metal - water interaction potential became quite clear. The model we are using appears to have been parameterized

for relatively low pressures; it utilizes both shifted Morse and repulsive Morse potentials to model the Au/O and Au/H interactions, respectively. The repulsive wall of the Morse potential does not diverge quickly enough at short distances to prevent water from diffusing into the center of the gold nanoparticles. This behavior is likely not a realistic description of the real physics of the situation. A better model of the gold-water adsorption behavior would require harder repulsive walls to prevent this behavior.

4 Discussion

The Langevin Hull samples the isobaric-isothermal ensemble for non-periodic systems by coupling the system to a bath characterized by pressure, temperature, and solvent viscosity. This enables the simulation of heterogeneous systems composed of materials with significantly different compressibilities. Because the boundary is dynamically determined during the simulation and the molecules interacting with the boundary can change, the method inflicts minimal perturbations on the behavior of molecules at the edges of the simulation. Further work on this method will involve implicit electrostatics at the boundary (which is missing in the current implementation) as well as more sophisticated treatments of the surface geometry (alpha shapes^{25,49} and Tight Cocone⁵⁰). The non-convex hull geometries are significantly more expensive ($\mathcal{O}(N^2)$) than the convex hull ($\mathcal{O}(N \log N)$), but would enable the use of hull volumes directly in computing the compressibility of the sample.

Acknowledgments

Support for this project was provided by the National Science Foundation under grant CHE-0848243. Computational time was provided by the Center for Research Computing (CRC) at the University of Notre Dame.

Molecular graphics images were produced using the UCSF Chimera package from the Resource for Biocomputing, Visualization, and Informatics at the University of California, San Francisco (supported by NIH P41 RR001081).

References

1. Melchionna S, Ciccotti G, Holian BL. *Mol Phys.* 1993; 78:533–544.
2. Hoover WG. *Phys Rev A.* 1985; 31:1695. [PubMed: 9895674]
3. Andersen HC. *J Chem Phys.* 1980; 72:2384–2393.
4. Sturgeon J, Laird B. *J Chem Phys.* 2000; 112:3474–3482.
5. Berendsen HJC, Postma JPM, VanGunsteren WF, Dinola A, Haak JR. *J Chem Phys.* 1984; 81:3684–3690.
6. Feller SE, Zhang YH, Pastor RW, Brooks BR. *J Chem Phys.* 1995; 103:4613–4621.
7. Jakobsen A. *J Chem Phys.* 2005:122.
8. Asthagiri D, Paliwal A, Abras D, Lenhoff A, Paulaitis M. *Biophys J.* 2005; 88:3300–3309. [PubMed: 15792969]
9. Brown GC. *J Theor Biol.* 1991; 153:195–203. [PubMed: 1787736]
10. Berkowitz M, McCammon JA. *Chem Phys Lett.* 1982; 90:215–217.
11. Brooks CL III, Karplus M. *J Chem Phys.* 1983; 79:6312–6325.
12. Warshel A. *Chem Phys Lett.* 1978; 55:454–458.
13. Belch A, Berkowitz M. *Chem Phys Lett.* 1985; 113:278–282.
14. Lee C, McCammon J, Rossky P. *J Chem Phys.* 1984; 80:4448–4455.
15. King G, Warshel A. *J Chem Phys.* 1989; 91:3647–3661.
16. Beglov D, Roux B. *J Chem Phys.* 1994; 100:9050–9063.
17. Beglov D, Roux B. *Biopolymers.* 1995; 35:171–178.
18. Li Y, Krilov G, Berne B. *J Phys Chem B.* 2005; 109:463–470. [PubMed: 16851037]
19. Zhu W, Krilov G. *J Mol Struct - THEOCHEM.* 2008; 864:31–41.

20. Kohanoff J, Caro A, Finnis M. *Chem Phys Chem*. 2005; 6:1848–1852. [PubMed: 16144013]
21. Baltazar SE, Romero AH, Rodriguez-Lopez JL, Terrones H, Martonak R. *Comp Mat Sci*. 2006; 37:526–536.
22. Delaunay B. *Bull Acad Science USSR VII: Class Sci Mat Nat*. 1934:793–800.
23. Lee DT, Schachter BJ. *Int J Parallel Programming*. 1980; 9:219–242.
24. Barber CB, Dobkin DP, Huhdanpaa HT. *ACM Trans Math Software*. 1996; 22:469–483.
25. Edelsbrunner H, Mücke EP. *ACM Trans Graphics*. 1994; 13:43–72.
26. Schlick, T. *An Interdisciplinary Guide*. Springer-Verlag New York, Inc; Secaucus, NJ, USA: 2002. *Molecular Modeling and Simulation*.
27. García de la Torre J, Huertas ML, Carrasco B. *Biophys J*. 2000; 78:719–730. [PubMed: 10653785]
28. García de la Torre J. *Biophys Chem*. 2001; 94:265–274. [PubMed: 11804736]
29. García de la Torre J, Carrasco B. *Biopolymers*. 2002; 63:163–167. [PubMed: 11787004]
30. Sun X, Lin T, Gezelter JD. *J Chem Phys*. 2008; 128:234107. [PubMed: 18570491]
31. Meineke MA, Vardeman CF, Lin T, Fennell CJ, Gezelter JD. *J Comp Chem*. 2005; 26:252–271. [PubMed: 15614795]
32. Gezelter, JD.; Kuang, S.; Marr, J.; Stocker, KM.; Li, C.; Vardeman, CF.; Lin, T.; Fennell, CJ.; Sun, X.; Daily, K.; Zheng, Y.; Meineke, MA. OpenMD, an open source engine for molecular dynamics. <http://openmd.net>
33. Barber, CB.; Huhdanpaa, H. Qhull. <http://www.qhull.org>
34. Berendsen H, Grigera J, Straatsma T. *J Phys Chem*. 1987; 91:6269–6271.
35. Qi Y, Çağın T, Kimura Y, Goddard WA III. *Phys Rev B*. 1999; 59:3527–3533.
36. Fennell CJ, Gezelter JD. *J Chem Phys*. 2006; 124(12):234104. [PubMed: 16821904]
37. Wolf D, Keblinski P, Phillpot SR, Eggebrecht J. *J Chem Phys*. 1999; 110:8254–8282.
38. Dou Y, Zhigilei L, Winograd N, Garrison B. *J Phys Chem A*. 2001; 105:2748–2755.
39. Foiles SM, Baskes MI, Daw MS. *Phys Rev B*. 1986; 33:7983–7991.
40. Sutton AP, Chen J. *Phil Mag Lett*. 1990; 61:139–146.
41. Kimura, Y.; Qi, Y.; Cagin, T.; Goddard, W, III. Technical Report 003. Caltech ASCI; 1998. The Quantum Sutton-Chen Many Body Potential for Properties of fcc Metals.
42. Collard S, McLellan R. *Acta metall mater*. 1991; 39:3143–3151.
43. Glättli A, Daura X, van Gunsteren W. *J Chem Phys*. 2002; 116:9811–9828.
44. Motakabbir K, Berkowitz M. *J Phys Chem*. 1990; 94:8359–8362.
45. Pi HL, Aragoes JL, Vega C, Noya EG, Abascal JL, Gonzalez MA, McBride C. *Mol Phys*. 2009; 107:365–374.
46. Fine RA, Millero FJ. *J Chem Phys*. 1973; 59:5529–5536.
47. Debenedetti P. *J Chem Phys*. 1986; 84:1778–1787.
48. Taylor RS, Dang LX, Garrett BC. *J Phys Chem*. 1996; 100:11720–11725.
49. Edelsbrunner H. *Discrete Comput Geom*. 1995; 13:415–440.
50. Dey T, Giesen J, Goswami S, Zhao W. *Discrete Comput Geom*. 2003; 29:419–434.

Appendix A: Computing Convex Hulls on Parallel Computers

In order to use the Langevin Hull for simulations on parallel computers, one of the more difficult tasks is to compute the bounding surface, facets, and resistance tensors when the individual processors have incomplete information about the entire system's topology. Most parallel decomposition methods assign primary responsibility for the motion of an atomic site to a single processor, and we can exploit this to efficiently compute the convex hull for the entire system.

The basic idea involves splitting the point cloud into spatially-overlapping subsets and computing the convex hulls for each of the subsets. The points on the convex hull of the

entire system are all present on at least one of the subset hulls. The algorithm works as follows:

1. Each processor computes the convex hull for its own atomic sites (left panel in Fig. 9).
2. The Hull vertices from each processor are communicated to all of the processors, and each processor assembles a complete list of hull sites (this is much smaller than the original number of points in the point cloud).
3. Each processor computes the global convex hull (right panel in Fig. 9) using only those points that are the union of sites gathered from all of the subset hulls. Delaunay triangulation is then done to obtain the facets of the global hull.

The individual hull operations scale with $\mathcal{O}(\frac{n}{p} \log \frac{n}{p})$ where n is the total number of sites, and p is the number of processors. These local hull operations create a set of p hulls, each with approximately $\frac{n}{3pr}$ sites for a cluster of radius r . The worst-case communication cost for using a “gather” operation to distribute this information to all processors is $\mathcal{O}(\alpha(p-1) + \frac{n\beta(p-1)}{3rp^2})$, while the final computation of the system hull scales as $\mathcal{O}(\frac{n}{3p} \log \frac{n}{3r})$.

For a large number of atoms on a moderately parallel machine, the total costs are dominated by the computations of the individual hulls, and communication of these hulls to create the Langevin Hull sees roughly linear speed-up with increasing processor counts.




Figure 1. Affine scaling methods use box-length scaling to adjust the volume to adjust to under-or over-pressure conditions. In a system with a uniform compressibility (e.g. bulk fluids) these methods can work well. In systems containing heterogeneous mixtures, the affine scaling moves required to adjust the pressure in the high-compressibility regions can cause molecules in low compressibility regions to collide.




Figure 2. The external temperature and pressure bath interacts only with those atoms on the convex hull (grey surface). The hull is computed dynamically at each time step, and molecules can move between the interior (Newtonian) region and the Langevin Hull.




Figure 3. The resistance tensor Ξ for a facet comprising sites i , j , and k is constructed using Oseen tensor contributions between the centroid of the facet f and each of the sub-facets (i, f, j) , (j, f, k) , and (k, f, i) . The centroids of the sub-facets are located at 1, 2, and 3, and the area of each sub-facet is easily computed using half the cross product of two of the edges.




Figure 4.

The response of the internal pressure and temperature of gold nanoparticles when first placed in the Langevin Hull ($T_{\text{bath}} = 300\text{K}$, $P_{\text{bath}} = 4 \text{ GPa}$), starting from initial conditions that were far from the bath pressure and temperature. The pressure response is rapid (after the breathing mode oscillations in the nanoparticle die out), and the rate of thermal equilibration depends on both exposed surface area (top panel) and the viscosity of the bath (middle panel).




Figure 5.
Compressibility of SPC/E water




Figure 6.

At low pressures, the liquid is in equilibrium with the vapor phase, and isolated molecules can detach from the liquid droplet. This is expected behavior, but the volume of the convex hull includes large regions of empty space. For this reason, compressibilities are computed using local number densities rather than hull volumes.

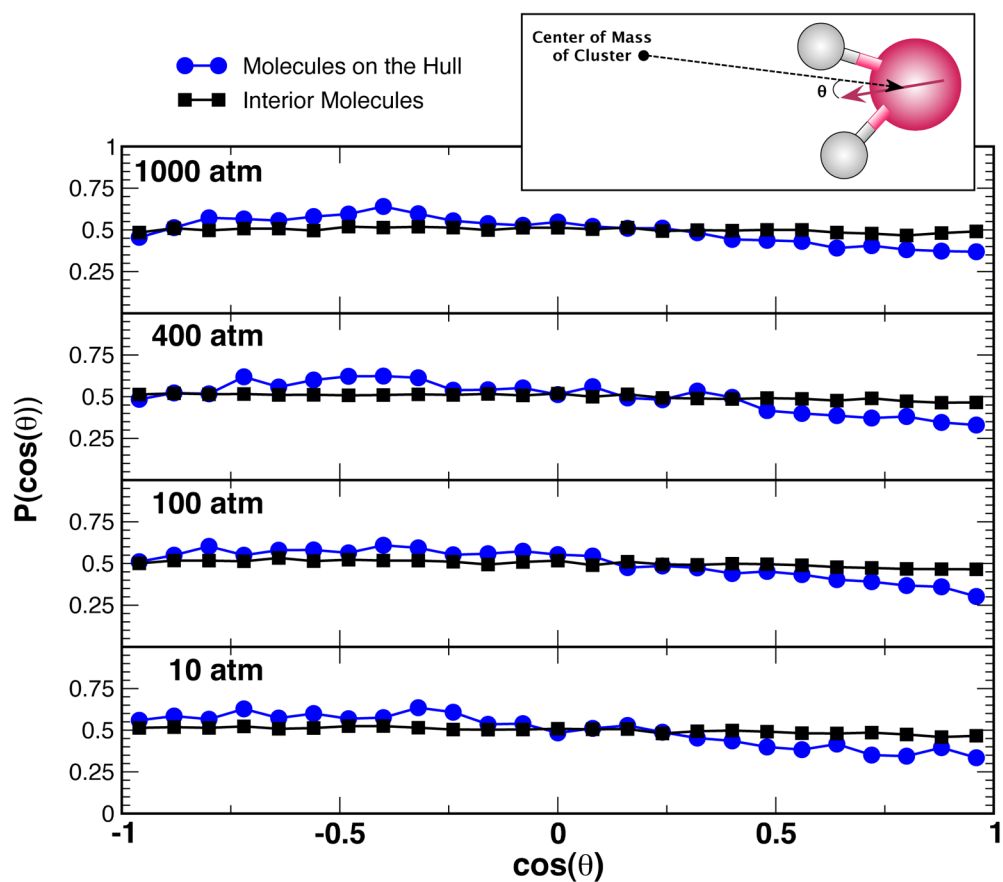


Figure 7. Distribution of $\cos \theta$ values for molecules on the interior of the cluster (squares) and for those participating in the convex hull (circles) at a variety of pressures. The Langevin Hull exhibits minor dewetting behavior with exposed oxygen sites on the hull water molecules. The orientational preference for exposed oxygen appears to be independent of applied pressure.




Figure 8. Density profiles of gold and water at the nanoparticle surface. Each curve has been normalized by the average density in the bulk-like region available to the corresponding material. Higher applied pressures de-structure both the gold nanoparticle surface and water at the metal/water interface.




Figure 9.

When the sites are distributed among many nodes for parallel computation, the processors first compute the convex hulls for their own sites (dashed lines in left panel). The positions of the sites that make up the subset hulls are then communicated to all processors (middle panel). The convex hull of the system (solid line in right panel) is the convex hull of the points on the union of the subset hulls.