

Exact Analytical Algorithm for the Solvent-Accessible Surface Area and Derivatives in Implicit Solvent Molecular Simulations on GPUs

Xin Cao, Michelle H. Hummel, Yuzhang Wang, Carlos Simmerling, and Evangelos A. Coutsias*



Cite This: <https://doi.org/10.1021/acs.jctc.3c01366>



Read Online

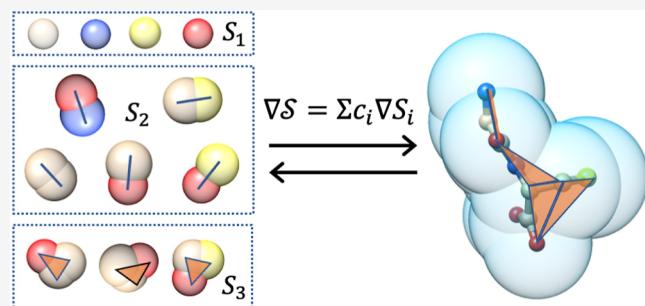
ACCESS |

Metrics & More

Article Recommendations

Supporting Information

ABSTRACT: In this paper, we present differentiable solvent-accessible surface area (dSASA), an exact geometric method to calculate SASA analytically along with atomic derivatives on GPUs. The atoms in a molecule are first assigned to tetrahedra in groups of four atoms by Delaunay tetrahedralization adapted for efficient GPU implementation, and the SASA values for atoms and molecules are calculated based on the tetrahedralization information and inclusion–exclusion method. The SASA values from the numerical icosahedral-based method can be reproduced with >98% accuracy for both proteins and RNAs. Having been implemented on GPUs and incorporated into AMBER, we can apply dSASA to implicit solvent molecular dynamics simulations with the inclusion of this nonpolar term. The current GPU version of GB/SA simulations has been accelerated up to nearly 20-fold compared to the CPU version, outperforming LCPO, a commonly used, fast algorithm for calculating SASA, as the system size increases. While we focus on the accuracy of the SASA calculations for proteins and nucleic acids, we also demonstrate stable GB/SA MD mini-protein simulations.



INTRODUCTION

An accurate description of the solvent environment is essential for biomolecular modeling, as biological machines function in an aqueous environment. The solute–solvent interactions and the rearrangement of water molecules induce the change of molecular shapes and the solvation free energy ΔG_{sol} .¹ Including water molecules in explicit solvent simulations produces more accurate simulation results at the cost of redundant calculation for the pairwise interactions between water molecules, while the friction induced by collisions with water molecules can further slow down the conformational sampling of the solute. In implicit solvent models, the solvent is treated as a continuum, and the system is simulated without explicit water molecules. The major advantage of implicit solvent models is the ability to explore the conformational space more efficiently, which can find applications in protein folding studies and structure prediction,^{2,3} solvation free-energy calculation using Poisson–Boltzmann surface area (PBSA),⁴ or generalized Born surface area (GBSA),⁵ and binding free-energy calculation using molecular mechanics Poisson–Boltzmann surface area (MM-PBSA).⁶

In implicit solvent modeling, the solvation free energy can be decomposed into polar and nonpolar contributions⁷ and the expression is

$$\Delta G_{\text{sol}} = \Delta G_{\text{pol}} + \Delta G_{\text{cav}} + \Delta G_{\text{vdW}} = \Delta G_{\text{pol}} + \Delta G_{\text{np}}$$

where the polar term ΔG_{pol} gives the difference between the work of uncharging the solute in vacuum and the work of charging the solute in solvent, ΔG_{cav} denotes the accommodation of cavity in the solvent for solute, and ΔG_{vdW} represents the van der Waals interactions between the solute and solvent. The Poisson–Boltzmann (PB) method describes the electrostatic potential in solution for a given set of boundary conditions and determines the effect of electrostatic interactions on the molecules⁸ and can be used to compute ΔG_{pol} . The ΔG_{cav} and ΔG_{vdW} are usually combined into one nonpolar term ΔG_{np} . It is considered to be proportional to the number of atoms in the solute having direct contact with solvent molecules and can be estimated in terms of the solvent-accessible surface area (SASA).⁹ Although it has been pointed out that incorporation of the volume term into the computation can provide a more complete description of the nonpolar contribution,^{10,11} the surface area-based methods can still have good performance in the prediction of the native-like conformations of proteins and the estimation of ligand-binding affinities.¹²

Received: December 13, 2023

Revised: May 6, 2024

Accepted: May 7, 2024

In implicit solvent molecular dynamics (MD) simulations, most of the effort has been devoted to the development of the dominant polar part. Due to the complexity of solving the PB equation, the generalized Born (GB) method⁵ was proposed to approximate the solution of the PB equation with simple functional forms. The recent improvement of fast GB models for proteins^{13–15} and nucleic acids¹⁶ and its implementation on graphics processing units (GPUs)¹⁷ helped GB achieve popularity in recent years. The main goal of a GB model is to achieve an agreement of electrostatic free energies with the PB approach, which can be achieved by parametrizing the GB screening parameters to provide a better match of the effective radii to those from the reference method. In particular, the pairwise decomposition version of GB¹⁸ is ideal for implementation on GPU with parallel computing. Although the performance of GB is improving, poor folding stabilities were observed in protein folding studies,² structure predictions,³ and de novo designed peptides with high helicity.¹⁹ As hypothesized in,² the neglected nonpolar solvation term might result in instability; this was supported by improved fold stability when an approximate SASA-based nonpolar term was added to the same GB model.²⁰ As shown in that paper, the inclusion of the nonpolar term in the GB/SA MD simulations can produce more stable trajectories and better simulated melting temperatures than comparable simulations with an identical setup other than the SASA term.

There exist several methods to calculate the molecular surface area on CPUs and GPUs. However, only a few of them can be applied to MD simulations, because most of them lack reliable derivatives for atoms. The solvent-accessible surface was proposed by Lee and Richards.²¹ A numerical implementation by Connolly²² is computationally expensive and lacks analytical derivatives. The derivatives are unreliable or unavailable in other methods, such as grid based, neighbor-counting, or machine learning.²³ The approximation methods on GPU have similar issues.^{24,25} In the CHARMM software, GBMV2/SA²⁶ can reproduce the molecular surface and has been implemented on GPUs.²⁷ However, the computation is based on grid points, so lower resolution of the grids may generate less reliable results while higher resolution results in slower simulation speed. The first available method for approximating the surface area along with derivatives in AMBER is LCPO,²⁸ which is able to estimate SASA based on the neighbors of atoms. As the MD simulations in AMBER are now primarily on GPUs,^{17,29} the data transfer between CPU and GPU deteriorates the performance because LCPO is implemented only on CPUs. Recently, another approximate pairwise method with a faster speed (up to 30 times faster), called the pairwise approximation of SASA (pwSASA), was proposed and it achieved a comparable accuracy to LCPO according to the testing of proteins.²⁰ The main purpose of pwSASA was to investigate the impact of the nonpolar term on the MD simulations of proteins with a simple and fast pairwise approximation suitable for running entirely on GPUs. The simple form of the functions provided faster simulation speed and did prove the significance of the nonpolar term in the simulations. However, the computed SASA values had various correlation coefficients from 0.6 to 0.9 compared with the numerical computation for a number of proteins, which is partly because the higher order interactions were ignored. Furthermore, pwSASA is highly empirical, and the parameters were trained only for the specific local atomic environments found in amino acids (such as hybridization and the identity of

bound atoms). Thus, in addition to relatively low accuracy on proteins, pwSASA cannot be used in simulations of more diverse systems such as those including modified amino acids, small molecule ligands, nucleic acids, or glycans, further limiting its suitability.

As every atom in a molecule can be represented by a sphere, the surface areas have also been studied by computational geometry methods. One reliable and accurate method was developed based on the Alpha Complex construction³⁰ by which the geometric descriptors of molecules can be computed through union of balls. The Alpha Complex is obtained with the atomic coordinates and radii information and can be applied to decompose the space into small segments through which the surface areas are computed by using the Gauss–Bonnet theorem³¹ or inclusion–exclusion formulas.^{32–34} Recently, the extension to other geometric measures, such as volumes, mean, and Gaussian curvatures, was proposed³⁵ in which the parallel implementation can achieve speed improvement over the serial version. However, the full implementation of such a method on GPUs is still unavailable because each step in the method requires substantial effort to utilize the parallel computing property to improve performance.

Here, we present differentiable solvent-accessible surface area (dSASA), a new method to calculate SASA and its derivatives with respect to atomic coordinates using the Alpha Complex theory and inclusion–exclusion method. This fully analytical and accurate algorithm was implemented on GPUs and incorporated into AMBER software. The basic background and the procedures employed in the method are first shown. The results for the assessment of the method are given next: the accuracy of estimation is examined by several protein systems with diverse topologies, the speed of the GPU version in GB/SA simulations is compared with the CPU version and other methods, and the performance is demonstrated by including this nonpolar term in GB/SA simulations on proteins. However, we caution readers that good practice in force-field development involves isolation and validation of the individual components, as done here. Deviations from experimental behavior can arise from inaccuracy in any of the multiple components (solute force field, polar solvation term, and nonpolar term). Likewise, a good match to the experiment can be obtained from the fortuitous cancellation of error. For this reason, we show that the dSASA term performs well in MD simulations, which verifies the efficient algorithm implementation, but our main focus remains on validating the SASA accuracy.

METHODS

Theory and Estimations for Nonpolar Solvation. In the implicit solvent model, the nonpolar solvation term is usually SASA-based.⁵ The free energy is proposed to be proportional to SASA with a surface tension parameter (γ): $\Delta G_{np} = \gamma \text{SASA}$. The surface tensions are assumed to be identical for all atoms.

The relatively accurate SASA values can be calculated by the ICOSA numerical method³⁶ ($gbsa = 2$ in AMBER), in which starting from an icosahedron, a water probe with radius 1.4 Å is recursively rolled on the van der Waals surface of the molecule. The implementation of this method in current AMBER MD simulations is unavailable because of the absence of atomic forces. The linear combinations of pairwise overlaps (LCPO; $gbsa = 1$ in AMBER) algorithm is the first stable algorithm used in GB/SA MD simulations in AMBER. The neighbor list

of a central atom is used to subtract the pairwise overlapping from its isolated sphere areas, which results in the computation complexity $O(N^2)$, where N is the total number of solute atoms, so it will slow down quadratically as the size of the molecule increases. In the GPU version of GB/SA MD with this method, the data transfer between CPU and GPU decelerates the simulations as its implementation is CPU/GPU hybrid. The pwSASA ($gbsa = 3$ in AMBER)²⁰ is the recently developed algorithm to calculate SASA that is suitable for implementing on GPUs. The idea is similar to LCPO but only considers the interactions between two atoms, which achieves a faster speed but results in an approximate SASA. The parameters for the atom types are trained exclusively on standard amino acids, so the application to other systems, such as nucleic acids or small molecules, is unavailable.

dSASA: An Analytical Method for SASA. Here, we introduce dSASA, a geometric method ($gbsa = 4$ in AMBER) to compute SASA. As has been explored in 28 and 37, to be consistent with other methods, we assume that the heavy atoms in molecules can have a good approximation to the molecular SASA, reducing computational cost substantially. The atomic radii for four common elements (C, 1.7 Å; O, 1.5 Å; N, 1.55 Å; S, 1.8 Å) are used in ICOSA, pwSASA, and dSASA, while LCPO uses different radii (C, 1.7 Å; O, 1.6 Å; N, 1.65 Å; S, 1.9 Å). Every atom can be represented by a weighted point $p'_i = (p_i, d_i)$ in the point set $\mathbb{A} \subset \mathbb{R}^3 \times \mathbb{R}$, which contains the 3D coordinates p_i in space and one-point weight d_i . The radius of the water probe is set to 1.4 Å, and the weight $d_i = (\text{atom_radius}_i + 1.4)^2$. Here, we present the key steps, while a complete description of the analytical development of the method is given in refs 33 and 34. The method has three main steps, and the depiction of the method is shown in Figure 1.

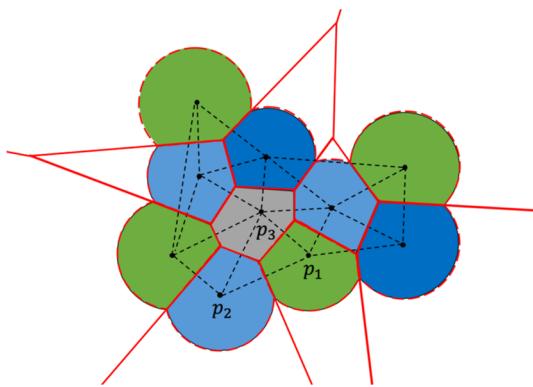


Figure 1. Depiction of the method in 2D. Every solid circle is the expanded atom including the water probe, and black points show the center of atoms. The dashed black lines denote part of the tetrahedralization of the molecule. The atoms in green and blue (such as p_1 and p_2) are exterior, and the gray atom p_3 is interior. The boundary of Laguerre cells of atoms is shown in solid red lines, and the exterior boundary of Laguerre cells of molecules in the solvent is represented by the dashed red curves which correspond to SASA.

Step 1. Calculate the 3D weighted Delaunay tetrahedralization (wDT) of all points through which the points are assigned into tetrahedra. This step has computation complexity $O(N \log N)$ and is the dominant part of the method.

Step 2. Create and classify the dual complex C which denotes the possible interactions among close atoms. The

complex C contains the lists of simplices, such as the vertices (the atoms), edges (overlap between two atoms), triangles (overlap among three atoms), and tetrahedra (overlap among four atoms) from the tetrahedralization. These simplices are then filtered and classified into interior and exterior (denoted by ∂C) based on their connection information.

Step 3. Compute the surface areas using Laguerre intersection cells and the inclusion–exclusion method based on the exterior simplices ∂C .

In the following, we show the process to compute the atomic SASA based on the exterior simplices ∂C because only the atoms having direct contact with the solvent contribute to the surface areas. The simplices have been classified into interior and exterior in Step 2. The Laguerre diagram of the point set is conjugate to the wDT. The space can be decomposed into cells by the conjugated Laguerre diagram as shown in Figure 1. The Laguerre cell for a weighted point (atom) p'_i is defined by

$$L_i = \{x \in \mathbb{R}^3 : |p'_i - x|^2 - d_i \leq |p_j - x|^2 - d_j, \forall p'_j \in \mathbb{A}\}$$

Every weighted point can also be treated as a ball $B_i = \{x \in \mathbb{R}^3 : |p'_i - x|^2 - d_i \leq 0\}$, and the union of these balls $B = \bigcup B_i$ forms the space filling model of \mathbb{A} . As shown in Figure 1, the Laguerre cells of exterior atoms can be unbounded or larger than the real size of the atoms, and then a more realistic representation for an atom is given by the Laguerre–intersection cell, the intersection part of the Laguerre cell, and the ball: $LI_i = L_i \cap B_i$.

The boundary area of the Laguerre intersection part LI_i of an exterior atom is the SASA value $S^i = \text{surf}(LI_i)$.

Suppose there are k points in a subset $T \subset \mathbb{A}$. The number of points in T is $|T| = k$ and the centers of points are denoted by T' . The convex hull (*conv*) of points in T' is written as σ_T , and the term S_T is the surface area of the intersection of the balls in T . By applying the inclusion–exclusion method shown in Figure 2, the SASA of the molecule is

$$S = \sum_{\sigma_T \in \partial C} (-1)^{k+1} c_T S_T, \quad k = 1, 2, 3$$

The coefficients c_T are given by the corresponding simplices. When $|T| = 1$, i.e., a vertex v_i : $c_T = \Omega_T$ is the fraction of the ball i outside the tetrahedra in C . Ω_T is the normalized outer solid angle subtended by the union of tetrahedra in C containing v_i . When $|T| = 2$, i.e., an edge e_{ij} : $c_T = \Phi_T$ is the normalized outer dihedral angle of the union of tetrahedra in C which contain the edge e_{ij} . When $|T| = 3$, i.e., a triangle t_{ijk} : $c_T = 1$ or 0.5 is the fraction of S_T that is outside the union of tetrahedra in C . Here, $v_i = p_i$, $e_{ij} = \text{conv}(\{p_i, p_j\})$, and $t_{ijk} = \text{conv}(\{p_i, p_j, p_k\})$.

The contribution of an individual atom p'_i to the molecular surface area is

$$S^i = \Omega_i S_i^{(i)} - \sum_{e_{ij} \in \partial C} \Phi_{ij} S_{ij}^{(i)} + \sum_{t_{ijk} \in \partial C} c_{ijk} S_{ijk}^{(i)} \quad (1)$$

where $S_T^{(i)}$ is the contribution of S_T to S^i , and $\sum_{i=1}^n S^i = S$.

Then, the corresponding derivative for the points can be obtained on the basis of the atomic surface area with respect to the atomic coordinates in the molecule as

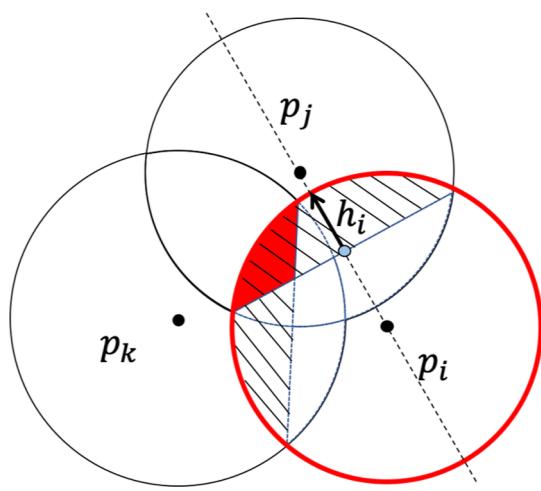


Figure 2. Inclusion–exclusion process of p_i . The red circle denotes the standalone ball S_i of p_i ; the shaded region indicates the intersection S_{ij} and S_{ik} with p_j and p_k ; and the small red region is the intersection of three atoms S_{ijk} .

$$\begin{aligned} \nabla S^i = \nabla \Omega_i S_i^{(i)} - \sum_{e_{ij} \in \partial C} (\nabla \Phi_{ij} S_{ij}^{(i)} + \Phi_{ij} \nabla S_{ij}^{(i)}) \\ + \sum_{t_{ijk} \in \partial C} c_{ijk} \nabla S_{ijk}^{(i)} \end{aligned} \quad (2)$$

The terms in the equations are calculated as follows. When $|T| = 1$, $S_T^{(i)} = 4\pi d_i$. Let \mathcal{I}_1 be the set of tetrahedra in C to which the point p_i' belongs to. For $\zeta_I \in \mathcal{I}_1$, define ω^I as the normalized inner solid angle subtended by the tetrahedron ζ_I from the point p_i' . Then

$$\Omega_T = 1 - \sum_{\zeta_I \in \mathcal{I}_1} \omega^I, \text{ with the derivative } \nabla \Omega_T = - \sum_{\zeta_I \in \mathcal{I}_1} \nabla \omega^I_T \quad (3)$$

The normalized inner solid angle, ω^I , of a tetrahedron $p_i p_j p_k p_l$ subtended by the vectors $\mathbf{a} = p_j - p_i$, $\mathbf{b} = p_k - p_i$, and $\mathbf{c} = p_l - p_i$ is shown in the left of **Figure 3** and given by the equation

$$\omega^I = \frac{1}{2\pi} \arctan \left(\frac{|\mathbf{a} \cdot (\mathbf{b} \times \mathbf{c})|}{abc + (\mathbf{a} \cdot \mathbf{b})c + (\mathbf{a} \cdot \mathbf{c})b + (\mathbf{b} \cdot \mathbf{c})a} \right)$$

where $a = |\mathbf{a}|$ is the length of \mathbf{a} and likewise for b and c .

When $|T| = 2$, the intersection of two atoms p_i' and p_j' is shown in **Figure 2**. $S_T^{(i)} = 2\pi\sqrt{d_i}h_i$. Let \mathcal{I}_2 be the set of

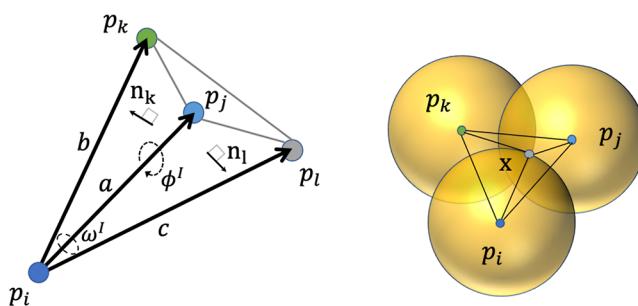


Figure 3. Left: Tetrahedron formed by points p_i , p_j , p_k , and p_l . Right: Center x of the intersection of three balls and the resulting tetrahedron.

tetrahedra in C containing the edge σ_T . For $\zeta_I \in \mathcal{I}_2$, define ϕ^I as the normalized inner dihedral angle of ζ_I along σ_T . Then

$$\Phi_T = 1 - \sum_{\zeta_I \in \mathcal{I}_2} \phi^I_T$$

The normalized inner dihedral angle between planes $p_i p_j p_k$ and $p_i p_j p_l$ with normals \mathbf{n}_k and \mathbf{n}_l is shown in the left of **Figure 3** and given by

$$\begin{aligned} \phi^I = \frac{\arccos(\mathbf{n}_k \cdot \mathbf{n}_l)}{2\pi}, \\ \text{with } \nabla \phi^I = \frac{-\nabla(\mathbf{n}_k \cdot \mathbf{n}_l)}{2\pi\sqrt{1 - (\mathbf{n}_k \cdot \mathbf{n}_l)^2}} \end{aligned} \quad (4)$$

Then, the derivative terms for **eq 2** are

$$\nabla S_T^{(i)} = 2\pi\sqrt{d_i} \nabla h_i, \quad \nabla \Phi_T = - \sum_{\zeta_I \in \mathcal{I}_2} \nabla \phi^I_T \quad (5)$$

When $|T| = 3$, consider $T = \{p_i', p_j', p_k'\}$. The surface area of the common intersection of the three balls can be written as a weighted sum of the surface areas of the single and double intersections. If p_i' , p_j' , and p_k' have a nonempty intersection then there are two points in common with the surfaces of all three balls, shown in the right of **Figure 3**. Denote one of the two points x , define $p_x' = (x, 0) \in \mathbb{R}^3 \times \mathbb{R}$, and let $T_x = \{p_i', p_j', p_k', p_x'\}$. Let S_2^x be the set of edges defined by σ_{T_x} and S_1^x be the set of vertices in σ_{T_x} . The atomic surface area of p_i' from the intersection of p_i', p_j' , and p_k' is given by

$$\frac{1}{2} S_T^{(i)} = \Phi_{ij}^x S_{ij}^{(i)_x} + \Phi_{ik}^x S_{ik}^{(i)_x} - \Omega_i^x S_i^x \quad (6)$$

where Φ_{ij}^x is the normalized dihedral angle of σ_{T_x} along the edge σ_{ij} , Ω_i^x is the normalized solid angle of σ_{T_x} subtended from p_i and similarly for other combinations i , j , and k . The derivative of this term is

$$\frac{1}{2} \nabla S_T^{(i)} = \nabla \Phi_{ij}^x S_{ij}^{(i)_x} + \Phi_{ij}^x \nabla S_{ij}^{(i)_x} + \nabla \Phi_{ik}^x S_{ik}^{(i)_x} + \Phi_{ik}^x \nabla S_{ik}^{(i)_x} - \nabla \Omega_i^x S_i^x \quad (7)$$

By substituting **eqs 3, 4, 5, 6**, and **7** into **eqs 1** and **2**, we can obtain the surface area and the derivative for every atom. More details of the description and the equations can be found in refs 33 and 34. However, the atomic derivatives may become discontinuous when atoms are approaching. Through detailed analysis in ref 33, it can be seen that the most common type of discontinuities occur when two atoms become externally tangent. We examine here the simplest case where only two atoms with extended radii $r_i = 2.9$ and $r_j = 3.1 \text{ \AA}$ are considered to show the possible singularity in the calculation. The changing process of the SASA values and the atomic derivatives is given in **Figure 4**. In the left, the summation of SASA values is continuous as two atoms are approaching from a long distance, and the minimal value occurs when two atoms are internally tangent. In the right, the magnitude of the atomic derivatives has a leap when two atoms are externally tangential, at which the singularity occurs. As investigated in ref 33, the singularity rarely happens when the solvent radius is 1.4 \AA . In the implementation, we need one threshold value ϵ to determine the external tangency if $|p_i - p_j| - (r_i + r_j) | < \epsilon$. When $\epsilon = 10^{-6}$, the external tangency was not detected after simulating a few systems for thousands of steps. When $\epsilon =$

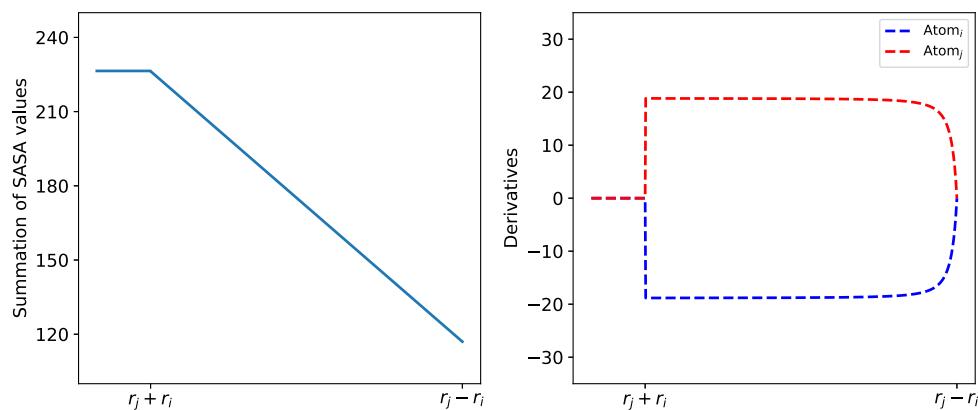


Figure 4. Trend of SASA and the magnitude of atomic derivatives as the distance of two atomic centers is changing. The *x*-axis shows the distance between the atomic centers, decreasing to the right. The summation of SASA values are on the left, while the atomic derivatives are on the right.

10^{-3} , one such case may occur after dozens of simulation steps. On the other hand, the magnitude of the leap in the atomic derivatives is less than 20 \AA^2 and the surface tension is usually less than $0.01 \text{ kcal}/(\text{mol}\cdot\text{\AA}^2)$, so the impact of such singularity is insignificant compared to the overall forces for exterior atoms whose magnitude can be more than 10 kcal/mol . Given the small chance of happening and the tiny impact of the singularity on the overall forces of the atoms, we can expect that the discontinuity in SASA calculation will have a negligible impact on the trajectories in GB/SA simulations.

Implementation of dSASA on GPUs. As the implementation of MD in AMBER has been done on GPUs,^{17,29} our main concern here is to implement the method on GPUs to speed up the calculation, making it applicable to longer time MD simulations, so we discuss more details about the necessary treatments on GPUs below.

Weighted Delaunay Tetrahedralization. In the wDT, all possible connections among local points can be detected. Based on the coordinates and weights, four local points are assigned to a tetrahedron having the property that no other points will be inside the circumsphere of this tetrahedron. The wDT can be sequentially computed on CPUs with the complexity $O(N \log N)$. One parallel algorithm on GPUs to generate exact wDT for a large data set remains a challenge. The typical GPU implementation of wDT often obtains a near-Delaunay tetrahedralization followed by transformation to CPU to generate the valid wDT. However, an algorithm gReg3D,³⁸ which is able to compute the exact wDT on GPUs, was implemented a few years ago. In the algorithm, all points are first included in a cube with an appropriate size, and next the points will be assigned to smaller cubes. The local tetrahedralization for these small cubes is calculated in parallel, followed by checking the consistency and applying necessary modifications to achieve the final wDT. The largest size of the initial cube in this algorithm is 512, so the point set with a greater size will be rescaled to fit into the cube, which may introduce a few errors. As the SARS-CoV-2 spike protein containing nearly 4000 residues typically occupies a cube with a size of 256, this algorithm can compute an exact wDT for proteins with thousands of residues on GPUs. As shown in Figure 1, the black points are the center of atoms and dashed lines connecting black points denote the edges in the tetrahedralization. The purpose of this step is to obtain the dual complex C and the conjugated Laguerre diagram. In this diagram, every point is represented by a cell enclosed by the

red solid lines and dashed curves; then, the surface area can be calculated with such information.

Extraction of Exterior Simplices and Calculation of Surface Areas. The information on close atoms in the tetrahedra has been achieved in the wDT above. However, some unrealistic connections and the interior contacts will not have a contribution to surface areas. Then, we need to extract the exterior simplices ∂C containing the feasible vertices, edges, and triangles along with the connections among them, which will finally be used to calculate SASA. The extraction of simplices on CPUs is serial and cannot be implemented on GPUs directly. Here, we provide more details of the extraction of simplices on GPUs.

We first need to create lists of unique triangles, edges, and vertices with connection information. Take one tetrahedron as an example. It contains 4 triangles as faces, 6 edges, and 4 vertices. One element in the list of triangles contains the information on three vertices and the tetrahedra it belongs to, one edge element contains the information on two vertices and the triangles and tetrahedra it is in, while one vertex element has the information on the associated edges and tetrahedra.

The exported information from gReg3D provides the vertex indices of all tetrahedra, so we can pull out all information on edges and triangles in tetrahedra by taking advantage of the parallelism on GPUs. One list of all possible triangles is created, and every element has three vertices and the associated tetrahedra. The list of all possible edges is constructed with each element containing two vertices and the associated tetrahedra. Suppose there are N tetrahedra, then $4N$ possible triangle elements and $6N$ possible edge elements are in the lists, respectively. However, some duplicate elements need to be filtered out to create lists of unique elements. Take the list of triangles as an example shown in Figure 5. We first perform lexical sorting for the list of triangles by the first three indices, bringing the same triangles next to each other. Next, with parallel computing, we can identify the number of associated tetrahedra for a triangle. One triangle can belong to at most two tetrahedra. In the following assignment, every element will check the three vertices of the next element in the list, if three vertices are identical, then we will assign “1” to the newly created array “idxTri”, which indicates that the triangle is duplicate, otherwise we assign “0”. We then can identify the unique triangle elements in the list and record the indices in one array “Tri”: a “1” in “idxTri” is a new element, and the second “0” in two consecutive “0’s indicates a new element as

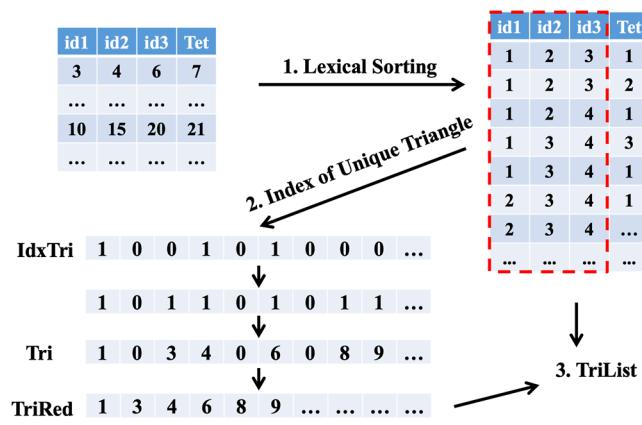


Figure 5. Procedures to create the list for unique triangles.

well. Now we have the indices for all unique triangles and the number of associated tetrahedra, and then we can create the data structure “TriList” for these triangles. When creating triangles, we extract a list of edges and assign the associated triangle indices to this new list, which will provide partial information to edges. Following the similar process above, we can create lists of unique edges and unique vertices, respectively.

With the lists for vertices, edges, and triangles, we next can classify them into interior and exterior and record the indices of exterior simplices. The basic rule is the triangle with at most one associated tetrahedron is exterior, and one edge is exterior only if its two vertices are both exterior. As shown in Figure 1, the atoms in green and blue are exterior while the gray atom p_3 is interior whose SASA value will be 0; the edge p_1p_2 is exterior and the edges p_1p_3 and p_2p_3 are interior. For a molecule with 1700 atoms, the numbers of tetrahedra, triangles, and edges are more than 11,200, 22,600, and 13,000, respectively. These

simplices can be independently implemented in parallel on a GPU to speed up the classification. Given the lists of exterior vertices, edges, and triangles, we then can calculate atomic surface areas and gradients by eqs 1 and 2 in parallel and assign the values to individual atoms.

Simulated Protein Systems and RNA Systems. Trp-cage variant Tc5b (PDB code, 1L2Y³⁹) has 20 residues with the sequence NLYIQWLKDGGPSSGRPPPS, and the burial of the hydrophobic tryptophan side chain provides a driving force for its folding. It contains an α -helix, a 3_{10} -helix, and a C-terminal PP_{II} helix and the Trp indole ring is encapsulated in a cluster of Pro rings. The structure was determined via the NMR experiment.

Homeodomain variant (PDB code, 2P6J⁴⁰) contains 52 residues with sequence MKQWSENVEEKLKEFVKRH δ QRT QEEELH δ QYAQLGLNEEAIRQFFEEFEQRK. It is a variant of *Drosophila melanogaster* engrailed homeodomain and was solved by NMR. It consists of three α -helices connected by loop regions.

The 14-mer cUUCGg tetraloop hairpin RNA (PDB code, 2KOC⁴¹) is an NMR-solved model and contains 14 bases with sequence GGCACUUCGGUGGCC, with the common and highly stable UUCG loop.

The stem loop C 5'AUA3' triloop of brom mosaic virus RNA (PDB code, 1ESH⁴²) is composed of 13 bases with sequence GGUGCAUAGCACC. It is designed to contain the triloop AUA in the middle and is an NMR-solved model.

The CD experiments provide the melting curves for Trp-cage³⁹ and the melting temperature for the homeodomain variant.⁴⁰ These two protein systems were studied in ref 20 and the ab initio folding experiments in ref 2 using the same force field and solvent model, providing a good reference to quantify the possible improvement by addition of a nonpolar solvation term. As mentioned above, it is important to compare results in which only the component of interest is varied.

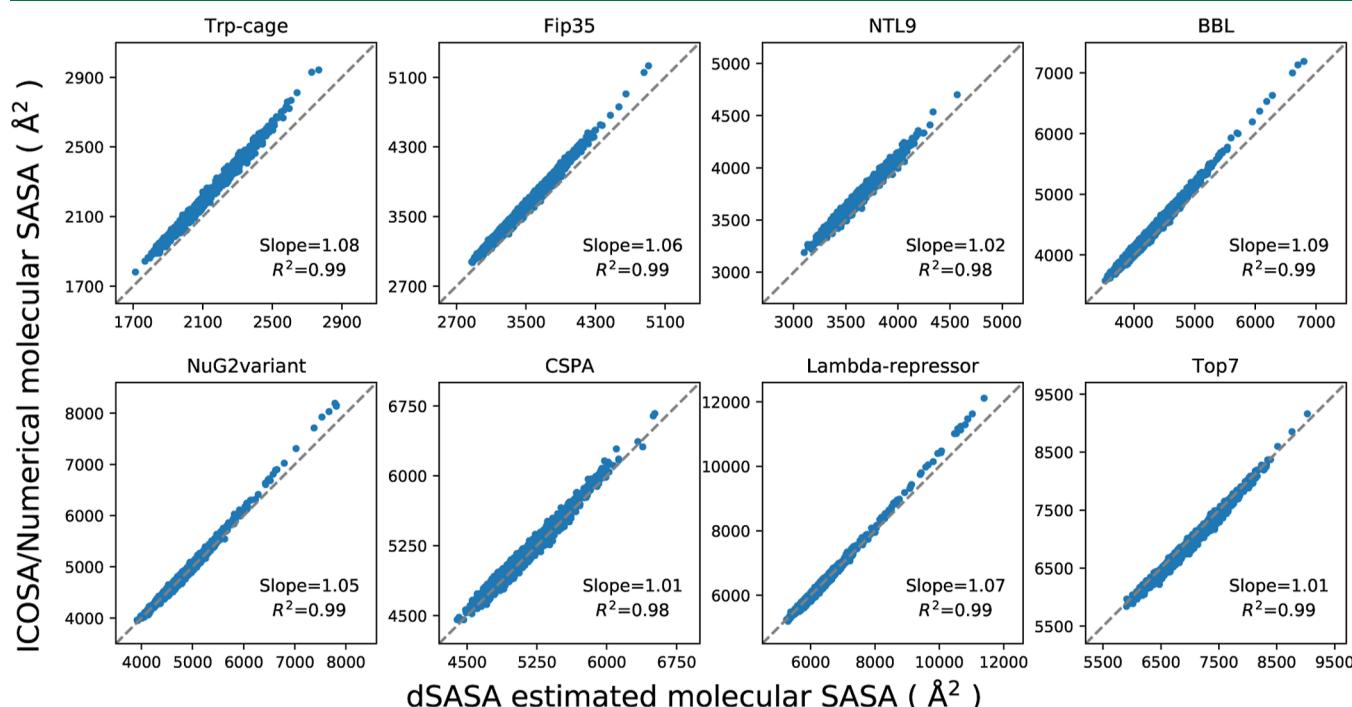


Figure 6. Comparison of molecular SASA between the dSASA and ICOSA numerical method for the proteins. Each point represents one conformation of the protein. The diagonal dashed lines indicate perfect agreement.

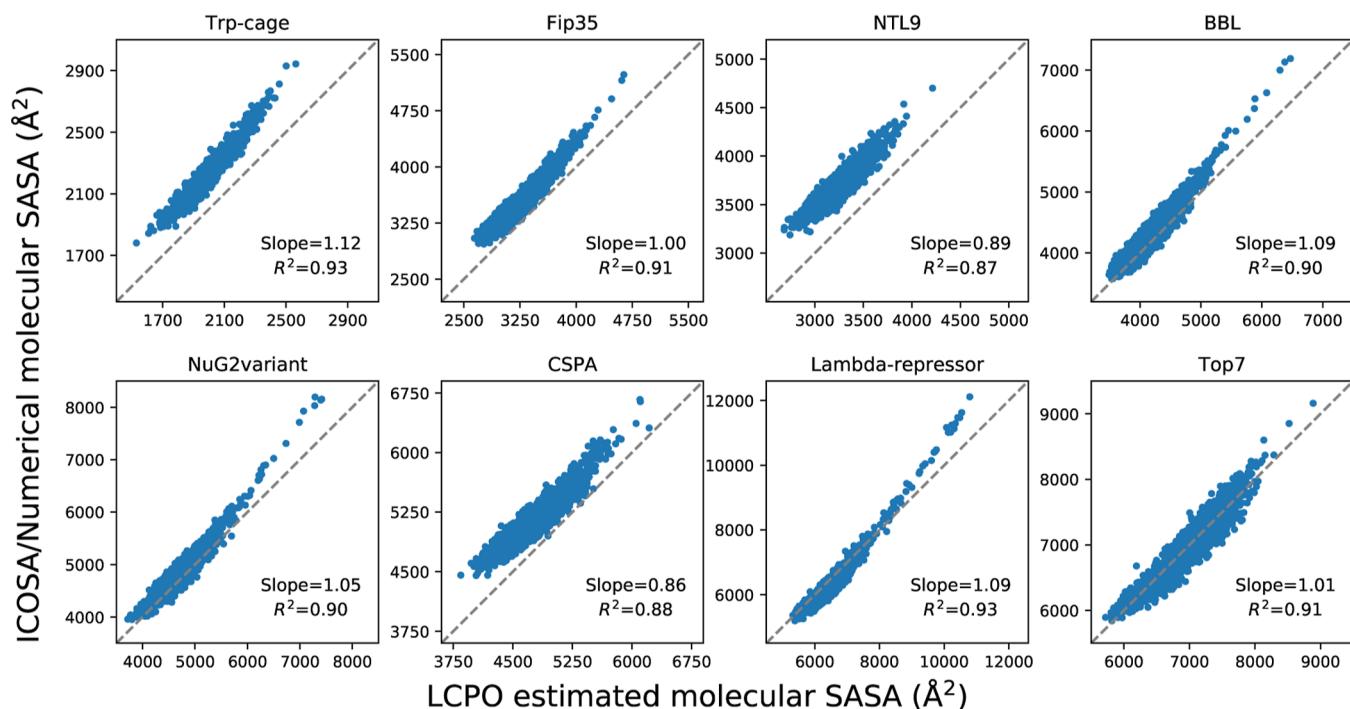


Figure 7. Comparison of molecular SASA between the LCPO and ICOSA numerical method for the proteins. Each point represents one conformation of the protein. The diagonal dashed lines indicate perfect agreement.

Details of MD Simulation. In the GB and GB/SA MD simulations, replica exchange molecular dynamics (REMD) was applied to enhance the efficiency of sampling. The setting for the system is as follows: all bonds involving hydrogen were added SHAKE constraints; Langevin dynamics ($ntt = 3$) with 1 ps^{-1} collision frequency was used; the time step was 4 fs by following the protocol in ref 43 through which the masses of hydrogen atoms can be repartitioned; exchanges between adjacent temperature replicas were attempted every 1 ps; and conformations were extracted every 0.1 ns. Trp-cage and homeodomain were parameterized by ff14SBonlysc⁴⁴ with GBNeck2¹⁵ and mbondi3 radii.¹⁵ The surface tension parameter was extensively tested in pwSASA²⁰ for proteins and 5, 7, and 10 cal/(mol·Å²) provided reasonable results when compared to simulations of a model system in explicit water. Thus, we implemented the GB/SA simulations for LCPO, pwSASA, and dSASA with 5 or 7 cal/(mol·Å²).

For Trp-cage, two production runs starting from unfolded and the first NMR structure were simulated for 1.0 μs in GB, LCPO, pwSASA, and dSASA, with an REMD ladder of 8 temperatures (247.7, 264.0, 281.4, 300.0, 319.8, 340.9, 363.3, and 387.3 K). The backbone RMSD cutoff of 2.0 Å was applied to calculate the fraction of folded. For the homeodomain variant, two production runs starting from extended and the first NMR structure were simulated with a ladder of 10 temperatures (288.7, 300.0, 311.7, 323.9, 336.6, 349.8, 363.5, 377.7, 392.4, and 407.8 K) for 2.0 μs in GB and pwSASA, and 1.0 μs for LCPO and dSASA. The backbone RMSD cutoff of 5.0 Å was used to calculate the fraction of folded. These cutoffs are consistent with the study in refs 2 and 20. MD simulations were not performed for the RNA systems, since weaknesses in current RNA force fields make it challenging to obtain stable simulations for stable hairpins such as 2KOC even with fully explicit water.⁴⁵ We therefore restrict our analysis to the SASA accuracy for RNA.

RESULTS

Molecular SASA Estimation. In the test set to validate SASA estimation, we selected eight proteins from the set of the previously examined proteins for ab initio protein folding,² in which the set of proteins have diverse topologies: Trp-cage (20 residues), Fip35 (33 residues), NTL9 (39 residues), BBL (47 residues), NuG2variant (56 residues), CSPA (69 residues), lambda repressor (80 residues), and Top7 (92 residues). An ensemble of structures for every protein were extracted from the protein folding trajectories in that work with an even interval, obtaining a set of conformations with diverse atomic and molecular SASA values. The corresponding data were estimated for every conformation using ICOSA, LCPO, and pwSASA and dSASA. The numerical method ICOSA can provide more accurate SASA estimation for molecules than LCPO and pwSASA, so we mainly compared the molecular SASA values from dSASA with ICOSA. The results are shown in Figure 6 and every point represents one conformation in the trajectory. In all systems, dSASA can generate well-correlated SASA values with ICOSA as Pearson correlation coefficients of the linear regression are $R^2 = 0.98$ and 0.99 and the slopes are also close to 1. We note that this close agreement in fact demonstrates the relative accuracy of ICOSA as compared with our exact geometric method. However, ICOSA lacks the derivatives needed for MD simulation.

The molecular SASA values from LCPO and pwSASA were calculated using the same conformations and compared with the ICOSA method. As shown in Figure 7, the discrepancy of LCPO molecular SASA values with ICOSA values varies for different conformations in the systems. Overall, the LCPO tends to underestimate the values, with most of the values falling above the perfect agreement line. The correlation coefficients are around 0.9, but the slope and correlation are lower than 0.9 for NTL9 and CSPA. The results became worse with pwSASA shown in Figure 8, where a similar phenomenon

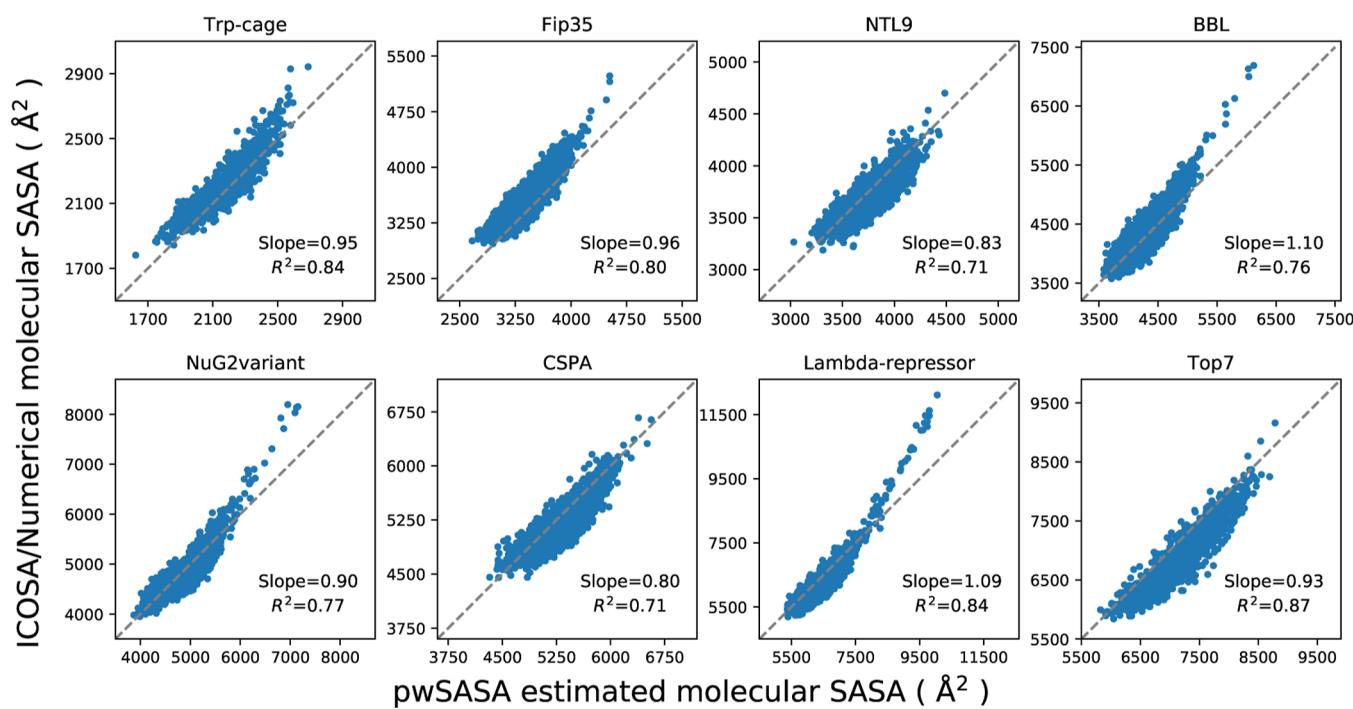


Figure 8. Comparison of molecular SASA between the pwSASA and ICOSA numerical method for the proteins. Each point represents one conformation of the protein. The diagonal dashed lines indicate perfect agreement.

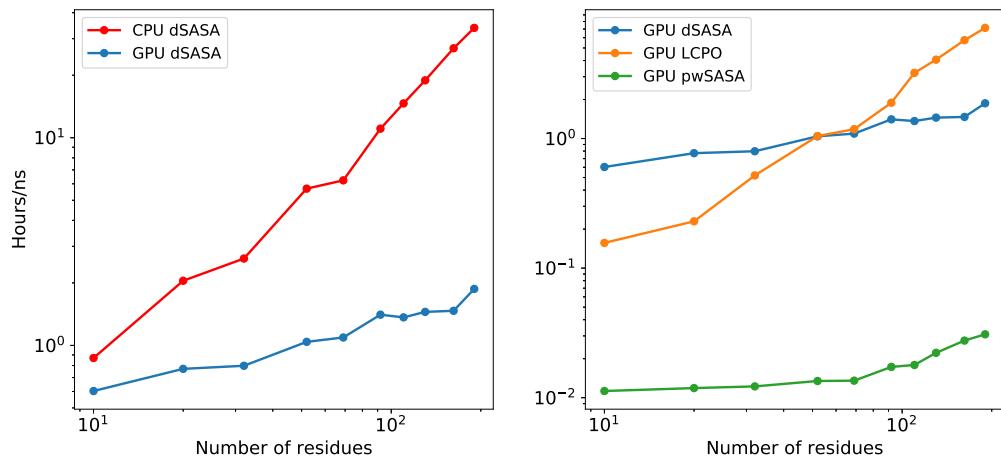


Figure 9. Left: Performance of dSASA for various sizes of proteins on CPU and GPU. The y-axis shows the wall-clock time in hours to complete a 1 ns simulation. Right: Performance of different methods for proteins on GPU.

arose: the discrepancy of pwSASA molecular SASA values with ICOSA values varies for conformations in the systems; furthermore, all correlation coefficients are less than 0.9 and the coefficients can be less than 0.8 for NTL9, BBL, NuG2variant, and CSPA. This reflects on the goal of pwSASA being speed rather than accuracy.

Both LCPO and pwSASA use pairwise overlaps, which can be relatively faster, but the inexact calculation will lead to inaccuracy, especially for some types of atoms. For example, both methods occasionally produce unphysically negative SASA values for individual atoms, while the atomic values from dSASA are all non-negative. LCPO includes the higher-order correction terms to account for the overlap between two neighbors of the central atom. Then, several atom types are defined by the environment, such as the atomic number, the number of bonded neighbors, and the state of hybridization. Every atom type has predefined parameters through training.

Because of less data for certain atom types, the parameter values for these atoms may generate unexpected values, which explains the discrepancy for the proteins above. As noticed in ref 28, the method works better for more exposed atoms and tends to overestimate the surface area of some buried atoms, such as the carboxyl carbons of the amino acids Asp, Asn, Gln, and Glu. pwSASA only considers the overlap between neighboring pairs and is compensated by training highly specific parameters applicable only to the atomic environments present in standard amino acids. Even though pwSASA introduced many additional empirical parameters, it suffers similar problems as LCPO, likely due to weak transferability of the pair parameters, and accounting for 3-body and higher terms only in an average way. In fact, the pwSASA values are not the original SASA values because they introduce an empirical adjustment to the SASA values to compensate for a

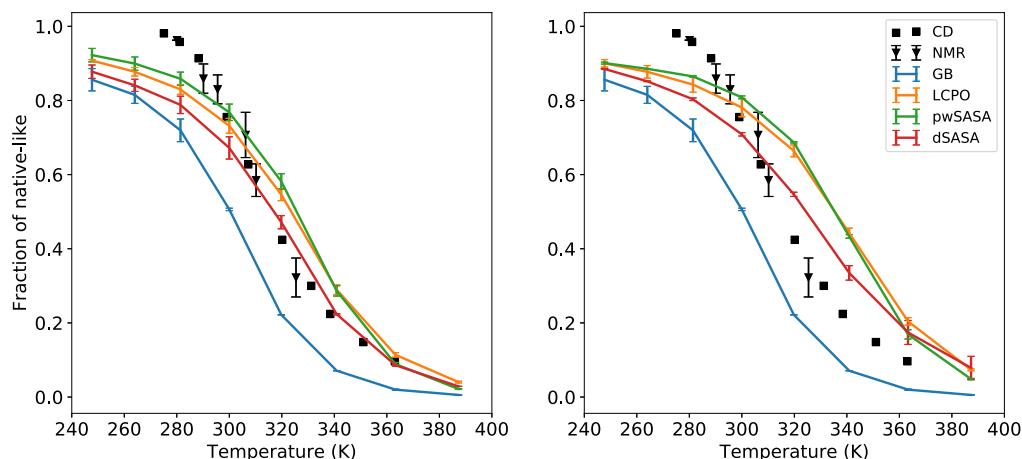


Figure 10. Thermal stability profiles for Trp-cage with surface tension 5 (left) and 7 $\text{cal}/(\text{mol}\cdot\text{\AA}^2)$ (right) in GB and GB/SA simulations, respectively, including experimental data.

systematic divergence from the numerical ICOSA calculations, which may adversely affect the overall accuracy as well.

Speed Comparison in GB/SA MD. The original version of dSASA was implemented on CPUs, and the transformation to a GPU version speeds up the computation by taking advantage of parallel computing. The program written in CUDA was integrated into the AMBER20 version. In sander, pmemd, or pmemd.cuda, setting the *gbsa* flag to 4 in the GB suite will activate GB/SA simulations with dSASA. We first compared the speed of GB/SA MD on GPU with CPU and next with other methods using 10 proteins with various sizes ranging from 10 residues to 190 residues. Four of them were selected from above and the PDB code of the rest is given: 10 (PDB code, SAWL), 20, 33, 52 (2P6J), 69, 92, 110 (1BYW), 130 (1E6K), 162 (2B75), and 190 (1BK7).

We benchmarked the method on the CPU Intel Platinum 8268 2.90 GHz and GPU Nvidia RTX 2080 Ti. The estimated wall-clock time for one-step SASA calculation on CPU is from 12 ms (10 residues) to 90 ms (190 residues), while the time trend on GPU is from 12 to 30 ms. The wall-clock time here and below is averaged over several simulations. The speed on CPU is comparable to the speed on GPU when the size of the protein is small, and the speed on GPU becomes relatively faster as the size of the protein increases. The estimated wall-clock time in hours required to complete 1 ns GB/SA simulation in AMBER is given in the left of Figure 9. The wall-clock time required to finish 1 ns simulation on the CPU increases from 0.87 h (10 residues) to 33.80 h (190 residues), while the corresponding trend of timing on the GPU is from 0.6 to 1.87 h. When the size of the protein is small, the improvement of GPU over CPU version is small with a speedup of 1.44 times for 10 residues, while the speedup grows as the size of the molecule increases to 18.07 times for 190 residues. For the system with 190 residues, we tested the parallel version with multiple CPUs. The GPU version is still much faster than the parallelized CPU version. The speedup of the GPU version over the parallel version with 4, 8, and 16 CPUs is 7.6, 9.1, and 4.8, respectively.

We next compared the speed with other methods in AMBER using the same set of systems described above. The wall-clock time in hours needed to complete a 1 ns simulation is provided in the right of Figure 9. The speed of pwSASA is stable and much faster than the other two methods as it is designed to have a fast and simple estimation with more

approximations, so we can confirm that the accuracy of the method is at the cost of speed as LCPO and dSASA can reproduce more accurate SASA values by the testing above. As can be seen in Figure 9, the speed of dSASA is more stable as the size of the system grows, while LCPO has a steeper increasing trend when the size of the molecule becomes larger. When the size of the system is small (10 residues), dSASA is slower than LCPO. While the crossing point is around the size of 69 residues and after that, dSASA starts to gain an advantage over LCPO for 69 and 92 residues and the speedup becomes greater for larger systems. This phenomenon fits the computation complexity for these two methods: the computation complexity of LCPO is $O(N^2)$ while the complexity of dSASA is $O(N \log N)$ (the complexity of tetrahedralization algorithm), so the system with a larger size will have relatively better performance. Another reason is that the GB/SA simulation with LCPO is CPU/GPU hybrid, so the data transfer further hinders the performance. On the other hand, pwSASA is ~ 80 times faster than dSASA for small proteins, reducing to ~ 50 times for larger proteins. While pwSASA is a pairwise method and its computation complexity is also $O(N^2)$, however, its computation is efficiently embedded in the overall energy computation which is of the same order; thus pwSASA adds little to the overall computational cost.

Stability Analysis of Proteins. dSASA can estimate the molecular SASA accurately in proteins with diverse shapes based on the testing above. With the speedup on GPUs, it is now possible to examine their performance in GB/SA simulations and evaluate their effect on the stability of protein structures. However, a comparison of the simulation results to experiments should be done with the understanding that deviations can arise from inaccuracy in any of the multiple components of the model (such as the biopolymer force field or GB model, which typically contributes more than the SASA term). Additionally, the implicit solvation models may become less reliable away from 300 K. Given these caveats, we simulated two proteins with dSASA and compared the results with experiments and GB-only and other SASA methods. The trends in the simulations can provide useful information, albeit with these limitations. The thermal stability profiles of the Trp-cage and the fraction of folded on the homeodomain are calculated from REMD at various temperatures. The temperature ladder for Trp-cage contains 8 replicas (247.7–387.3 K,

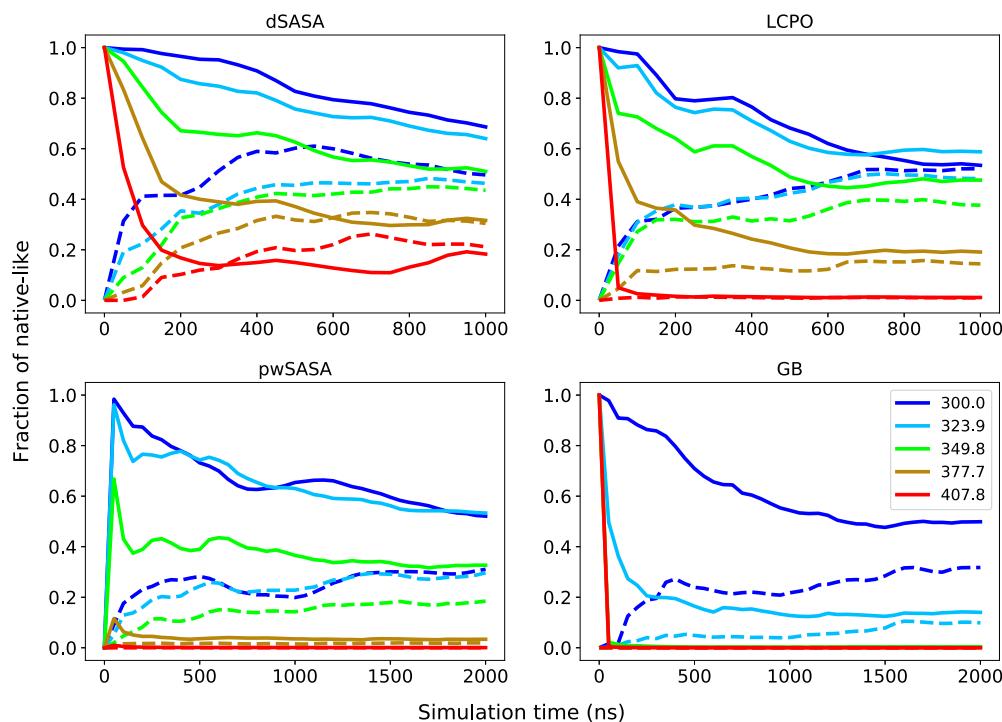


Figure 11. Fraction of native-like structures for the homeodomain starting from native (solid lines) and extended (dashed lines) at 5 selected temperatures. Top left: dSASA; top right: LCPO; bottom left: pwSASA; and bottom right: GB-only method.

see the [Methods](#) section for the full list). The ladder for the homeodomain has 10 replicas (288.7–407.8 K).

The default surface tension value in the AMBER GB/SA module is 5 cal/(mol·Å²)⁴ and the recommended value is 7 cal/(mol·Å²)⁵ for pwSASA.²⁰ For the Trp-cage, we tested surface tension values of 5 and 7 to examine the performance of the method. For the homeodomain, we only used the value 7, which can produce more near-native states in the trajectories.

The thermal stability profiles of Trp-cage with surface tension 5 and 7 cal/(mol·Å²) were first computed and the fractions of the near-native conformations are given in [Figure 10](#) for all methods. The left shows the results with a surface tension of 5 cal/(mol·Å²), while the right shows the results with a surface tension of 7 cal/(mol·Å²). With a surface tension of 5, all three methods incorporating the nonpolar term achieve better agreement with experiments by creating more stable thermal stability profiles at various temperatures. At 300 K, simulations using dSASA generate ~75% fraction of folded states, while other methods with the nonpolar term and experimental results are ~80% and the GB-only method leads to ~50%. When the temperature is greater than 310 K, the fraction of folded from dSASA is still slightly smaller than other methods, but dSASA can produce closer results to the experiments. The predicted melting temperature T_m is 316.8 K and the experimental result is around 317 K, which is better than the GB-only (predicted T_m is around 300 K) and other methods (predicted T_m is greater than 320 K for LCPO and pwSASA). Moreover, it is expected that a larger surface tension can generate more native-like conformations through which we can demonstrate that the nonpolar term does not disrupt the system. As shown on the right of [Figure 10](#), all three methods can generate more native-like structures with a surface tension of 7 cal/(mol·Å²).

To show the extent of dynamics and variation in sampled RMSD, we provide the detailed trajectories at several temperatures for Trp-cage starting from extended and native states with surface tension 5 cal/(mol·Å²) in the [Supporting Information](#).

For the homeodomain variant, we first carried out REMD with surface tension 5 cal/(mol·Å²) using pwSASA for 2 μ s, but we noticed that the fraction of folded was far from the experimental results at higher temperatures. The experimentally measured T_m is greater than 372 K, so we expect that a larger surface tension can achieve better agreement with experiments. We then simulated a surface tension of 7 cal/(mol·Å²) for all GB/SA methods. All three methods with the nonpolar term can improve the fraction of folded at various temperatures compared to the GB-only method. The fractions of native-like starting from native and unfolded states at selected temperatures are shown in [Figure 11](#). Since the experimental T_m is greater than 372 K, we expect that the system will not collapse at high temperatures. The two more accurate SASA methods, LCPO and dSASA, produce more native-like conformations at higher temperatures, as shown in [Figure 11](#). Moreover, dSASA generated more native-like conformations than LCPO at various temperatures, demonstrating that our method can have a better effect on stabilizing the system. The detailed trajectories of the homeodomain at selected temperatures starting from extended and native states with surface tension 7 cal/(mol·Å²) are given in the [Supporting Information](#). Furthermore, at the temperature of 300 K, the difference between the native-like fractions of the simulations starting from native and extended is around 20% for dSASA and LCPO after 500 ns simulations, while the difference is around 40% for GB and pwSASA even after 1000 ns simulations, indicating that the inclusion of an accurate SASA term might lead to improved sampling and faster convergence (or conversely, that the SASA approximations

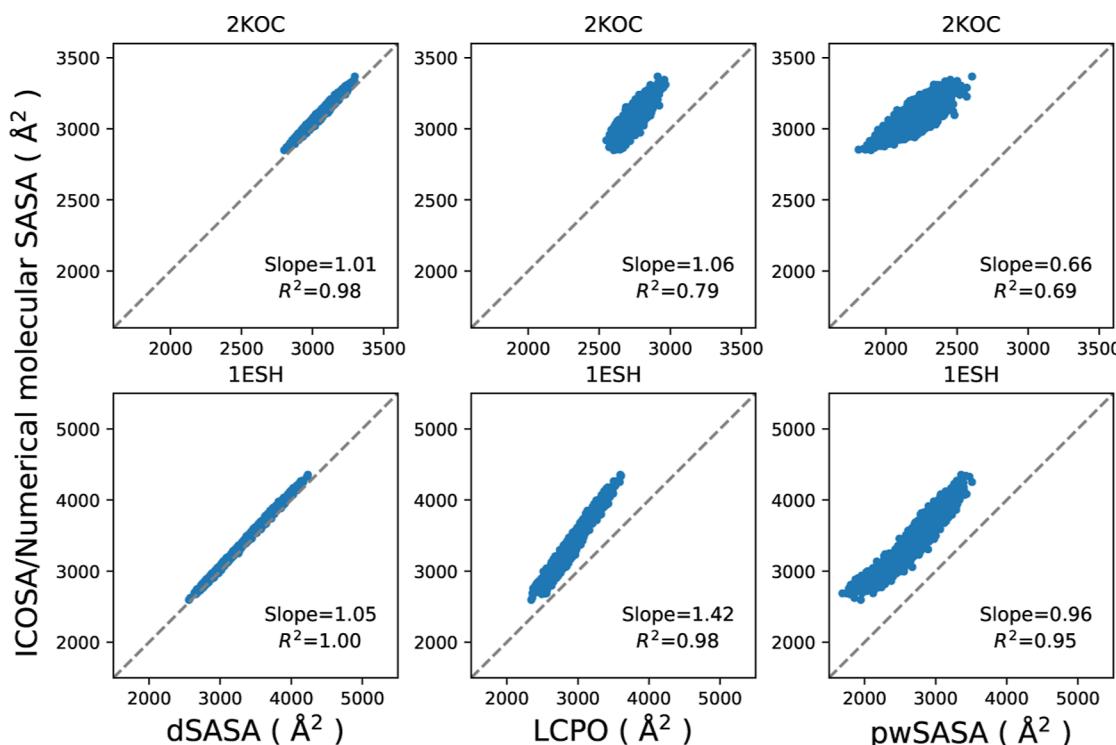


Figure 12. Comparison of molecular SASA values from various methods; each point represents one conformation of the RNA. The first row is for 2KOC and the second is for 1ESH. The diagonal dashed lines indicate perfect agreement.

may lead to increased kinetic traps). The improved performance from dSASA still has variance from the experimental results; for example, the fraction of the native-like conformations at 377.7 K is around 35% while it is 50% in the experiment. Such discrepancy may come from limitations of the SASA-based nonpolar term, or from the force field and the polar GB term; comparison of simulation results to the experiment typically does not reveal which component is responsible, and further model system development is needed.

Performance on RNA Systems. Another advantage of dSASA over pwSASA is that it can be applied to other types of systems such as nucleic acids or small-molecule ligands because only the radii of the atoms are considered without defining diverse atom types. Here, we test this aspect using two RNA systems, 2KOC and 1ESH, and compare the molecular SASA values with ICOSA and other methods to see if the excellent performance on proteins can be transferred to other biopolymers. Simulation results for RNA are not presented in detail, since accurate force fields and implicit solvent models for RNA are still a matter of active research.^{45,46}

To validate the SASA estimation, we generated alternate conformations for these RNA systems with the GB-only method for 1 μ s and extracted 10^4 conformations for each system, respectively. The comparison of the molecular SASA values between all methods is shown in Figure 12. dSASA can reproduce SASA values from ICOSA with more than 98% accuracy for both systems, while LCPO and pwSASA produced divergent results for the two systems, and they had an overall underestimation of the SASA values. The SASA estimates from pwSASA have large discrepancy from the ICOSA method, likely because it was trained exclusively on protein data so that the parameters do not accurately model the specific atom types and local environments present in RNA. Overall, the results

confirm the expected outcome that dSASA accuracy is transferable across diverse molecular systems.

CONCLUSIONS

In this work, we present dSASA, an analytical SASA evaluation method for molecules, and its implementation on GPUs. In this approach, the weighted Delaunay tetrahedralization is first computed with the atomic coordinates and radii. Next, the atomic and molecular SASA values are estimated using the inclusion–exclusion method based on exterior simplices from tetrahedralization, resulting in a more accurate estimation than LCPO and pwSASA in reproducing numerical ICOSA SASA values. The pretrained parameters for atoms in LCPO and pwSASA can speed up the calculation, but these parameters can have low transferability. This is confirmed by the comparison of computed molecular SASA values, which for these approximate methods had smaller correlation coefficients with the values from the ICOSA method, such as in the protein systems CSPA and NuG2variant and RNAs. In contrast, estimation from dSASA can reproduce ICOSA values with correlation coefficients greater than 98%. Moreover, the current GPU version speeds up the calculation substantially compared with the CPU version, making it applicable for longer GB/SA MD simulations, especially for larger systems. The gain of speed arises from the complexity of the algorithm itself which will scale well for larger systems and from the full implementation on GPU devices removing the data transfer with CPUs. dSASA has a stable trend of speed as the size of molecules increases, but it is still relatively slower than the highly approximate pwSASA. However, dSASA started to outperform LCPO when the system contains around 70 residues, and the performance for larger molecules would be more stable than LCPO because the computation in LCPO is

pairwise and its implementation is CPU/GPU hybrid in AMBER.

In the GB/SA simulations, two proteins (Trp-cage and homeodomain) were simulated and compared with other methods and experimental results. dSASA achieved a comparable performance with LCPO and pwSASA on the small protein, and the performance became better for the larger system. The simulated melting curve with the nonpolar term for the Trp-cage was more consistent with the experimental measures compared with that without the nonpolar term. In the homeodomain variant, the melting temperature is greater than 372 K so we expect that the system can maintain some amount of native-like conformations at higher temperatures. The simulations generated more extended conformations with other methods at high temperatures, while dSASA can produce more native-like conformations in the trajectories, which is closer to the experimental data. However, these comparisons to experiments must be considered in the context of the other force-field components as well.

As the program has been rigorously examined, we anticipate to extend its applications to other functions in the AMBER software, such as postprocessing of the trajectories in MM-PBSA or MM-GBSA. The calculation in dSASA is geometry-based, obtaining accurate results with the provided atomic coordinates and radii. Therefore, it can be used for diverse types of systems without restrictions, such as the simulations of larger proteins, protein–ligand complexes, and protein–nucleic acids complexes. One ns GB/SA simulation for 200-residue proteins with dSASA takes around 2 h. Even though it is slower than pwSASA, the accuracy of the method will allow us to explore the impact of the nonpolar term on simulations in the future. Given the accuracy of dSASA, the calculation of atomic and molecular SASA values can be benchmark data set for the training of the parameters in the pwSASA approach for RNAs in the future, which will combine the advantage of its speed on GPUs and the accurate calculation from dSASA. The SASA is certainly an approximation to nonpolar solvation, but many studies have shown that including it improves agreement with the experiment for things like protein stability or binding affinities.¹² In addition, dSASA, a more accurate SASA with derivatives implemented on GPUs, will also enable fast polar solvation methods for MD. Furthermore, the calculation of SASA is based on the diagram of Laguerre intersection cells, so it can be easily extended to the computation of molecular volumes along with the corresponding atomic derivatives. Inclusion of the volume term provides a more complete description for the nonpolar solvation term. With the volume derivatives, we can further examine the effect of the term on the stability of molecules in the MD simulations. Moreover, the current algorithm being used for weighted Delaunay tetrahedralization, gReg3D, is designed to work on a set of random points, and the size of the workspace depends on the distribution of points in the workspace. As it consumes nearly 70% of the wall-clock time of our surface area calculation, improvement on this algorithm will further speed up the simulations. The program dSASA written in CUDA will be freely available from the authors.

ASSOCIATED CONTENT

Supporting Information

The Supporting Information is available free of charge at <https://pubs.acs.org/doi/10.1021/acs.jctc.3c01366>.

Examination of energy conservation after including the dSASA term in MD; detailed trajectories with dSASA, GB, LCPO, and pwSASA for Trp-cage and homeodomain at selected temperatures; and fractions of folded for the second test set of homeodomain (PDF)

AUTHOR INFORMATION

Corresponding Author

Evangelos A. Coutsias – Department of Applied Mathematics and Statistics, Stony Brook University, Stony Brook, New York 11794, United States; Laufer Center for Physical and Quantitative Biology, Stony Brook University, Stony Brook, New York 11794, United States;  orcid.org/0000-0003-2910-9125; Email: evangelos.coutsias@stonybrook.edu

Authors

Xin Cao – Department of Applied Mathematics and Statistics, Stony Brook University, Stony Brook, New York 11794, United States; Laufer Center for Physical and Quantitative Biology, Stony Brook University, Stony Brook, New York 11794, United States

Michelle H. Hummel – Sandia National Laboratories, Albuquerque, New Mexico 87123, United States

Yuzhang Wang – Department of Chemistry, Stony Brook University, Stony Brook, New York 11794, United States

Carlos Simmerling – Department of Applied Mathematics and Statistics, Stony Brook University, Stony Brook, New York 11794, United States; Laufer Center for Physical and Quantitative Biology, Stony Brook University, Stony Brook, New York 11794, United States;  orcid.org/0000-0002-7252-4730

Complete contact information is available at:

<https://pubs.acs.org/10.1021/acs.jctc.3c01366>

Notes

The authors declare no competing financial interest.

ACKNOWLEDGMENTS

This work was partly supported by NIH grants RM1 GM135136 and R01 GM107104 and by NSF grant DMS 2054251. We gratefully acknowledge support from the Laufer Center for Physical and Quantitative Biology.

REFERENCES

- (1) Fennell, C. J.; Dill, K. A. Physical modeling of aqueous solvation. *J. Stat. Phys.* **2011**, *145*, 209–226.
- (2) Nguyen, H.; Maier, J.; Huang, H.; Perrone, V.; Simmerling, C. Folding simulations for proteins with diverse topologies are accessible in days with a physics-based force field and implicit solvent. *J. Am. Chem. Soc.* **2014**, *136*, 13959–13962.
- (3) Perez, A.; Morrone, J. A.; Brini, E.; MacCallum, J. L.; Dill, K. A. Blind protein structure prediction using accelerated free-energy simulations. *Sci. Adv.* **2016**, *2*, 1601274.
- (4) Sitkoff, D.; Sharp, K. A.; Honig, B. Accurate calculation of hydration free energies using macroscopic solvent models. *J. Phys. Chem.* **1994**, *98*, 1978–1988.
- (5) Still, W. C.; Tempczyk, A.; Hawley, R. C.; Hendrickson, T. Semianalytical Treatment of Solvation for Molecular Mechanics and Dynamics. *J. Am. Chem. Soc.* **1990**, *112*, 6127–6129.
- (6) Kollman, P. A.; Massova, I.; Reyes, C.; Kuhn, B.; Huo, S.; Chong, L.; Lee, M.; Lee, T.; Duan, Y.; Wang, W.; Donini, O.; Cieplak, P.; Srinivasan, J.; Case, D. A.; Cheatham, T. E. Calculating structures and free energies of complex molecules: combining molecular

mechanics and continuum models. *Acc. Chem. Res.* **2000**, *33*, 889–897.

(7) Levy, R. M.; Zhang, L.; Gallicchio, E.; Felts, A. K. On the nonpolar hydration free energy of proteins: surface area and continuum solvent models for the solute-solvent interaction energy. *J. Am. Chem. Soc.* **2003**, *125*, 9523–9530.

(8) Honig, B.; Nicholls, A. Classical Electrostatics in Biology and Chemistry. *Science* **1995**, *268*, 1144–1149.

(9) Eisenberg, D.; McLachlan, A. D. Solvation energy in protein folding and binding. *Nature* **1986**, *319*, 199–203.

(10) Wagoner, J. A.; Baker, N. A. Assessing implicit models for nonpolar mean solvation forces: the importance of dispersion and volume terms. *Proc. Natl. Acad. Sci. U.S.A.* **2006**, *103*, 8331–8336.

(11) Allison, J. R.; Boguslawski, K.; Fraternali, F.; van Gunsteren, W. F. A refined, efficient mean solvation force model that includes the interior volume contribution. *J. Phys. Chem. B* **2011**, *115*, 4547–4557.

(12) Genheden, S.; Ryde, U. The MM/PBSA and MM/GBSA methods to estimate ligand-binding affinities. *Expert Opin. Drug Discovery* **2015**, *10*, 449–461.

(13) Onufriev, A.; Bashford, D.; Case, D. A. Exploring protein native states and large-scale conformational changes with a modified generalized Born model. *Proteins: Struct., Funct., Bioinf.* **2004**, *55*, 383–394.

(14) Mongan, J.; Simmerling, C.; McCammon, J. A.; Case, D. A.; Onufriev, A. Generalized Born Model with a Simple, Robust Molecular Volume Correction. *J. Chem. Theory Comput.* **2007**, *3*, 156–169.

(15) Nguyen, H.; Roe, D. R.; Simmerling, C. Improved Generalized Born Solvent Model Parameters for Protein Simulations. *J. Chem. Theory Comput.* **2013**, *9*, 2020–2034.

(16) Nguyen, H.; Perez, A.; Bermeo, S.; Simmerling, C. Refinement of generalized Born implicit solvation parameters for nucleic acids and their complexes with proteins. *J. Chem. Theory Comput.* **2015**, *11*, 3714–3728.

(17) Case, D. A.; Cerutti, D. S.; Cheatham III, T. E.; Darden, T. A.; Duke, R. E.; Giese, T. J.; Gohlke, H.; Goetz, A. W.; Greene, D.; Homeyer, N. e. a. AMBER 2016; University of California: San Francisco, 2017.

(18) Hawkins, G. D.; Cramer, C. J.; Truhlar, D. G. Pairwise solute descreening of solute charges from a dielectric medium. *Chem. Phys. Lett.* **1995**, *246*, 122–129.

(19) Lang, E. J.; Baker, E. G.; Woolfson, D. N.; Mulholland, A. J. Generalized Born implicit solvent models do not reproduce secondary structures of de novo designed Glu/Lys peptides. *J. Chem. Theory Comput.* **2022**, *18*, 4070–4076.

(20) Huang, H.; Simmerling, C. Fast Pairwise Approximation of Solvent Accessible Surface Area for Implicit Solvent Simulations of Proteins on CPUs and GPUs. *J. Chem. Theory Comput.* **2018**, *14*, 5797–5814.

(21) Lee, B.; Richards, F. M. The Interpretation of Protein Structures: Estimation of Static Accessibility. *J. Mol. Biol.* **1971**, *55*, 379.

(22) Connolly, M. L. Analytical molecular surface calculation. *J. Appl. Crystallogr.* **1983**, *16*, 548–558.

(23) Durham, E.; Dorr, B.; Woetzel, N.; Staritzbichler, R.; Meiler, J. Solvent accessible surface area approximations for rapid and accurate protein structure prediction. *J. Mol. Model.* **2009**, *15*, 1093–1108.

(24) Juba, D.; Varshney, A. Parallel, stochastic measurement of molecular surface area. *J. Mol. Graphics Modell.* **2008**, *27*, 82–87.

(25) Dynerman, D.; Butzlaff, E.; Mitchell, J. C. CUSA and CUDE: GPU-accelerated methods for estimating solvent accessible surface area and desolvation. *J. Comput. Biol.* **2009**, *16*, 523–537.

(26) Lee, K. H.; Chen, J. Optimization of the GBMV2 implicit solvent force field for accurate simulation of protein conformational equilibria. *J. Comput. Chem.* **2017**, *38*, 1332–1341.

(27) Gong, X.; Chiricotto, M.; Liu, X.; Nordquist, E.; Feig, M.; Brooks, C. L.; Chen, J. Accelerating the generalized Born with molecular volume and solvent accessible surface area implicit solvent model using graphics processing units. *J. Comput. Chem.* **2020**, *41*, 830–838.

(28) Weiser, J.; Shenkin, P. S.; Still, W. C. Approximate atomic surfaces from linear combinations of pairwise overlaps (LCPO). *J. Comput. Chem.* **1999**, *20*, 217–230.

(29) Gotz, A. W.; Williamson, M. J.; Xu, D.; Poole, D.; LeGrand, S.; Walker, R. C. Routine Microsecond Molecular Dynamics Simulations with AMBER on GPUs. 1. Generalized Born. *J. Chem. Theory Comput.* **2012**, *8*, 1542–1555.

(30) Edelsbrunner, H. The union of balls and its dual shape. *Discrete Comput. Geom.* **1995**, *13*, 415–440.

(31) Cazals, F.; Kanhere, H.; Loriot, S. Computing the volume of a union of balls: a certified algorithm. *ACM Trans. Math Software* **2011**, *38*, 1–20.

(32) Bryant, R.; Edelsbrunner, H.; Koehl, P.; Levitt, M. The area derivative of a space-filling diagram. *Discrete Comput. Geom.* **2004**, *32*, 293–308.

(33) Hummel, M. H. Delaunay-Laguerre geometry for macromolecular modeling and implicit solvation. Ph.D. thesis, The University of New Mexico, 2014.

(34) Hummel, M. H.; Yu, B.; Simmerling, C.; Coutsias, E. A. Laguerre-Intersection Method for Implicit Solvation. *Int. J. Comput. Geom. Appl.* **2018**, *28*, 1–38.

(35) Koehl, P.; Akopyan, A.; Edelsbrunner, H. Computing the volume, surface area, mean, and Gaussian curvatures of molecules and their derivatives. *J. Chem. Inf. Model.* **2023**, *63*, 973–985.

(36) Shrake, A.; Rupley, J. A. Environment and Exposure to Solvent of Protein Atoms - Lysozyme and Insulin. *J. Mol. Biol.* **1973**, *79*, 351–371.

(37) Vasilyev, V.; Purisima, E. O. A fast pairwise evaluation of molecular surface area. *J. Comput. Chem.* **2002**, *23*, 737–745.

(38) Nanjappa, A. Delaunay Triangulation in R3 on the GPU. Ph.D. thesis, Singapore National University, 2012.

(39) Neidigh, J. W.; Fesinmeyer, R. M.; Andersen, N. H. Designing a 20-residue protein. *Nat. Struct. Biol.* **2002**, *9*, 425–430.

(40) Shah, P. S.; Hom, G. K.; Ross, S.; Lassila, J.; Crowhurst, K.; Mayo, S. Full-sequence computational design and solution structure of a thermostable protein variant. *J. Mol. Biol.* **2007**, *371*, 1–6.

(41) Nozlinovic, S.; Fürtig, B.; Jonker, H. R. A.; Richter, C.; Schwalbe, H. High-resolution NMR structure of an RNA model system: the 14-mer cUUCGg tetraloop hairpin RNA. *Nucleic Acids Res.* **2010**, *38*, 683–694.

(42) Tinoco, I.; Kim, C. H.; Kao, C. C. RNA motifs that determine specificity between a viral replicase and its promoter. *Nat. Struct. Biol.* **2000**, *7*, 415–423.

(43) Hopkins, C. W.; LeGrand, S.; Walker, R. C.; Roitberg, A. E. Long-Time-Step Molecular Dynamics through Hydrogen Mass Repartitioning. *J. Chem. Theory Comput.* **2015**, *11*, 1864–1874.

(44) Maier, J. A.; Martinez, C.; Kasavajhala, K.; Wickstrom, L.; Hauser, K. E.; Simmerling, C. ff14SB: improving the accuracy of protein side chain and backbone parameters from ff99SB. *J. Chem. Theory Comput.* **2015**, *11*, 3696–3713.

(45) Mráziková, K.; Mlýnský, V.; Kührová, P.; Pokorná, P.; Kruse, H.; Krepl, M.; Otyepka, M.; Banáš, P.; Sponer, J. UUCG RNA tetraloop as a formidable force-field challenge for MD simulations. *J. Chem. Theory Comput.* **2020**, *16*, 7601–7617.

(46) Liebl, K.; Zacharias, M. The development of nucleic acids force fields: From an unchallenged past to a competitive future. *Biophys. J.* **2023**, *122*, 2841–2851.