Visualizing Large-Scale Atomistic Simulations in Ultra-Resolution Immersive Environments

Khairi Reda*
University of Illinois at Chicago

Aaron Knoll
Texas Advanced Computing Center

Ken-ichi Nomura
University of Southern California

Michael E. Papka
Argonne National Laboratory

Andrew E. Johnson
University of Illinois at Chicago

Jason Leigh
University of Illinois at Chicago

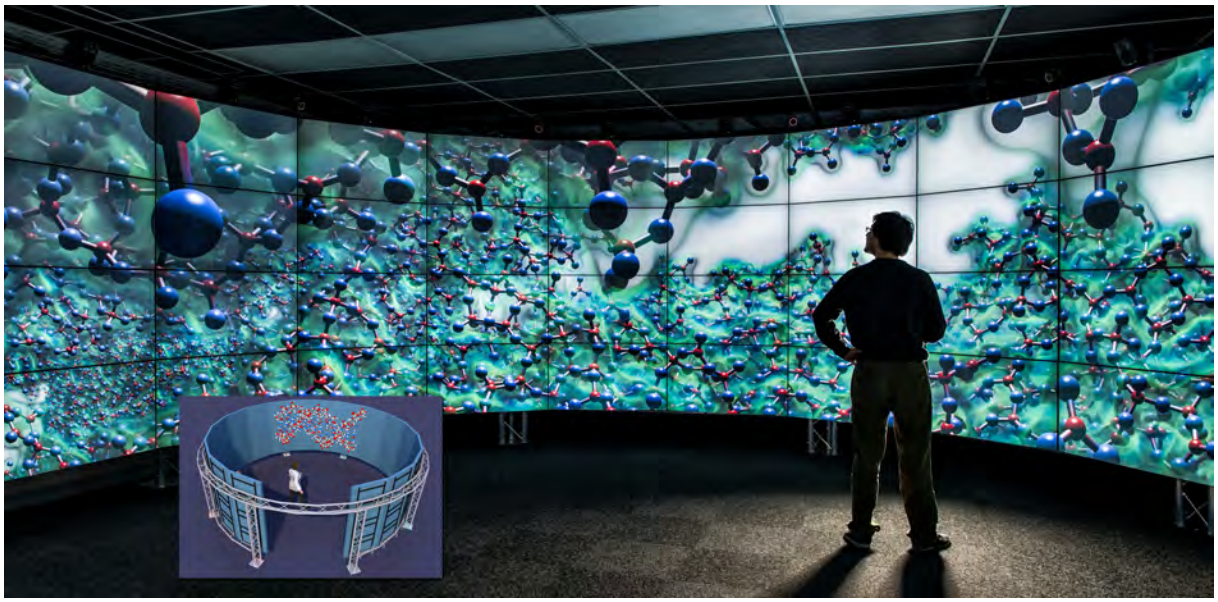

Figure 1: Interactive visualization of an amorphous glass fracture computed in a 5 million atoms Molecular Dynamics simulation in *CAVE2*. *CAVE2* is composed of 72 micro-polarized LCD panels arranged cylindrically to provide a 320-degrees stereoscopic, user-centered view at a total resolution of 74 Megapixels. The user explores large-scale simulations by flying through molecules using a wireless 'wand' or by physically walking in the space, causing the view to be re-rendered from his/her perspective.

ABSTRACT

Molecular Dynamics is becoming a principle methodology in the study of nanoscale systems, paving the way for innovations in battery design and alternative fuel applications. With the increasing availability of computational power and advances in modeling, atomistic simulations are rapidly growing in scale and complexity. Despite the plethora of molecular visualization techniques, visualizing and exploring large-scale atomistic simulations remain difficult. Existing molecular representations are not perceptually scalable and often adopt a rigid definition of surfaces, making them inappropriate for nanostructured materials where boundaries are inherently ill-defined. In this paper, we present an application for the interactive visualization and exploration of large-scale atomistic simulations in ultra-resolution immersive environments. We employ a hybrid representation which combines solid ball-and-stick glyphs with volumetric surfaces to visually convey the uncertainty in molecular boundaries at the nanoscale. We also describe a scalable, distributed GPU ray-casting implementation capable of rendering complex atomistic simulations with millions of atoms in real-time.

Index Terms: I.3.8 [Computer Graphics]: Applications—; H.5.1 [Information Interfaces and Presentation]: Multimedia Information Systems—Artificial, augmented, and virtual realities;

1 INTRODUCTION

Molecular Dynamics (MD) is becoming a principle methodology in the study of nanoscale systems, paving the way for many innovations in energy storage and alternative fuel applications. Modern MD simulations involve complex structures with tens to hundreds of millions of atoms, posing many challenges for visualization and analysis. To handle the sheer number of atoms in such structures, the majority of molecular visualizations employ simple representations, such as space-filling and ball-and-stick models. Such low-level representations, however, are not *perceptually scalable* and are often ineffective at portraying emergent, high-level features in the simulated structure (e.g. channels and cavities).

Molecular surfaces have been proposed as a higher-order alternative to simple space-filling and ball-and-stick models [3]. These surfaces provide visual continuity across larger structures and serve to illustrate important structural and functional properties. Nevertheless, computing and visualizing these surfaces in real-time re-

*e-mail: mreda2@uic.edu

IEEE Symposium on Large Data Analysis and Visualization 2013
October 13 - 14, Atlanta, Georgia, USA
978-1-4799-1658-0/13/$31.00 ©2013 IEEE

main prohibitively expensive, limiting their applicability to small datasets. Further, traditional molecular surfaces, such as Solvent Accessible Surfaces, are defined rigidly and often represented with discrete polygonal meshes. At the nanometer level, however, the definition of surfaces takes on a probabilistic character as it is subject to subatomic and quantum phenomena, leading to uncertainty in what constitutes a 'boundary' or a 'surface'. Such uncertainly cannot be faithfully portrayed with discrete surfaces.

The spatial complexity of today's atomistic simulations yet poses another challenge. To correctly interpret the results of these simulations, scientists need not only see the overall structure but also need to identify emergent large and small-scale features that are scattered throughout the dataset. Immersive visualizations with stereoscopic depth and viewer-centered perspective can greatly enhance perception and improve the spatial understanding of complex 3D structures [2]. Although immersive molecular visualizations have been around since the 1990s [4], the majority are aimed at decade-old CAVEs and low-resolution stereo projectors. Modern immersive environments on the other hand deliver much higher resolutions while greatly improving the picture quality. To leverage the full potential of these emerging platforms, however, we need new visualization algorithms that are both data and resolution scalable, while capable of achieving interactive frame rates.

In this paper we present an application for the interactive visualization and exploration of large Molecular Dynamics simulations in ultra-resolution immersive environments. Our visualization employs a hybrid visual encoding which combines solid ball-and-stick glyphs with volumetric molecular surfaces to visually convey the uncertainty in boundaries at the nanoscale. We demonstrate the expressiveness of this technique in generating richer molecular visualizations and its ability to visually classifying marked regions that emerge during the simulation. We also describe a distributed GPU ray-casting implementation that is capable of rendering millions of atoms in real-time at resolutions of up to 74 Megapixels.

The rest of the paper is divided as follows. Section 2 surveys the literature. We describe our visualization approach in section 3, motivating the discussion by outlining challenges in visualizing nanostructured materials. In section 4, we describe the implementation and algorithm details. We illustrates example use cases and gives performance and preliminary user evaluation results in section 5. We sum up and conclude the paper in section 6.

2 Related Work

Molecular visualization has a rich history that dates back to the 1800s. A variety of visual representations has been devised to represent molecular systems in ways that convey meaningful structural and functional properties. Despite the diversity of metaphors, we can generally recognize three types of molecular representations: solid glyphs (such ball-and-stick and space-filling models), molecular surfaces (such as Solvent Accessible Surfaces), and volumetric representations, which are often used to depict quantum properties around the molecule.

2.1 Solid glyphs

The vast majority of molecular visualizations employ a solid glyph metaphor to represent the molecular structure. Many of the recent techniques deal specifically with large atomistic simulations. For instance, Nakano et al. employ multi-resolution, polygon-based rendering to interactively visualize few millions of atoms [16]. Sharma et al. uses probabilistic occlusion culling to achieve near-interactive frame rates when rendering hundreds of millions of atoms [19]. Several techniques have also been proposed to deal with temporally large simulations. For instance, Stone et al. describe an out-of-core method to quickly render long MD simulations by leveraging solid-state disks [22]. GPU-based techniques are also abundant. For example, Grottel et al. employ a GPU ray-casting algorithm with occlusion culling and deferred shading to attain interactive frame rates [7].

Although ball-and-stick and space-filling models are universally recognized and are quite rendering-efficient, they do not visually aggregate as good as surfaces. Often, they are too low-level of a metaphor to convey structural intricacies in complex nanostructured materials.

2.2 Molecular surfaces

Molecular surfaces have become a popular alternative to space-filling and ball-and-stick glyphs as a somewhat higher-order visual abstraction, particularly in biomolecular applications [3]. Solvent Accessible Surfaces, for instance, are computed by rolling a 'probe' atom on the molecule of interest. As the center of the probe rolls along the atoms, it defines a surface that is equidistant from each atom in the molecule. This surface can be visualized as a polygonal mesh engulfing the molecule, making certain structural properties more obvious (such as the portion of a protein that is accessible to a solvent). They are sometimes rendered semi-transparently so that one can still see the molecule through.

Molecular surfaces have traditionally been computed offline, but the introduction of programmable pixel shaders in modern GPUs made interactive rendering feasible. For example, Krone et al. describe a GPU ray-casting method to render molecular surfaces using their inherent mathematical representations [13]. Techniques have also been proposed to accelerate their construction on parallel architectures [15]. Yet, scaling these algorithms beyond a few hundred thousand atoms remains difficult even with modern GPU architectures. Moreover, since conventional molecular surfaces are defined and visualized deterministically, they could potentially mislead the interpretation of molecular boundaries. This is particularly limiting in nanoscale materials design where the correct interpretation of material interfaces requires consideration of fuzzy subatomic properties, such as the underlying electronic structure.

2.3 Volumetric representations

Volumetric representations have traditionally been used to depict subatomic properties of the 3D space surrounding molecules [10, 6, 8]. They are also employed sometimes in analysis and classification tasks. For instance, Krone et al. employ direct volume rendering (DVR) to visually track structural features over time such as internal cavities in proteins [14]. Knoll et al. classifies material interfaces by volume rendering approximate charge density fields [11]. Although this volumetric approach to rendering nanoscale boundaries appear to be a promising alternative to conventional molecular surfaces, it has only been tested on relatively small datasets with less than a million atoms.

In this paper, we adopt Knoll et al's volumetric formulation and visualize approximate electron charge density fields as proxies for molecular surfaces and apply it to large-scale atomistic structures. Our main goal is obtaining flexible representations of molecular surfaces to support the probabilistic interpretations of boundaries at the nanoscale. Crucially, this requires our visualization to support both solid and fuzzy representations of these boundaries. However, volume and isosurface rendering have traditionally been treated as separate visualization modalities with different techniques, though some methods allow the explicit blending of the two (e.g. [18]). Yet, there is still separation between volumes and surfaces at the algorithmic or representational levels. Nevertheless, this separation is rather artificial as both techniques are aimed at the same data type (i.e. scalar fields). An isosurface in fact can be produced by volume rendering the scalar field with a Dirac impulse in the transfer function at the desired iso value. Knoll et al. have previously shown that a single DVR pass is capable of producing high-quality results using a combination of peak finding and differential sam-

Figure 2: Volume rendering of (approximate) electron densities allows for a flexible interpretation of molecular boundaries and enables a variety of visual representations which can be obtained by simply varying the transfer function. Here, the same molecule is shown under three transfer functions emphasizing different electron density ranges. The left and center views employ the same transfer function (shown in the middle inset), but with different volume opacities. An overview of the dataset is shown in the left-most inset.

pling [12]. This unified volume-surface formulation fits our design goals because it gives us the ability to smoothly transition from the traditional rigid definition of molecular surfaces to the nebulous representation of boundaries we advocate here. Our rendering algorithm therefore does not distinguish between an isosurface and a volume cloud, but leaves it up to the user to modulate the transfer function to obtain the desired surface representation.

In addition to volume rendering, we also visualize ball-and-stick glyphs and embed them in the electron density volume. Many techniques have been devised to accelerate the rendering of quadratic surfaces (e.g. spheres and cylinders). On the GPU, splatting remains one of the most efficient techniques [20]. Methods for ray tracing arbitrarily complex surfaces from their mathematical representation have also been proposed [21]. We take the more straightforward approach of rendering implicit ball-and-stick surfaces by performing world-space intersections along the ray. While more computationally intensive than splatting, the overhead is minimal compared to the time it takes to perform volume rendering. Further, this approach is algorithmically easier to combine with volume rendering and results in a smaller shader code, which minimizes the footprint on the limited SIMD registers in the GPU. To accelerate the ray-casting process, we embed the glyphs in a uniform grid and use the 3DDA algorithm to traverse the structure [1].

3 A HYBRID *volume + glyphs* APPROACH

We first motivate the discussion by making the case for new visual metaphors that can cater better to the requirements of materials scientists. We then describe our *volume + glyphs* approach, highlighting its perceptual scalability and its suitability for visualizing nanostructured materials. We also discuss the rising need for immersive molecular visualizations at high-resolutions.

3.1 The elusive metaphor

When looking at results of atomistic simulations, scientists often want to see familiar representations of atoms and bonds (e.g. ball-and-stick models). But they also want the visualization to be scalable and integrative so that they can recognize higher-order structures that emerge during the simulation. Furthermore, scientists need to be able to visually approximate the boundaries of molecules in order to understand how they might react with their surroundings.

Molecular surfaces provide visually contiguous shapes that serve to illustrate geometric boundaries. Compared to standalone glyphs, these surfaces are more perceptually scalable and are often better at conveying high-level features in the molecular structure. Yet, two issues arise from the conventional, discrete interpretation of molecular surfaces:

- At the nanoscale, molecular boundaries suffer from uncertainty due to underlying subatomic and quantum effects. Instead of portraying this uncertainty, conventional molecular surfaces convey a false sense of confidence, potentially leading to incorrect interpretation in what constitutes a 'boundary' or 'surface'. Although subatomic uncertainty is less relevant in large biomolecules (such as proteins), it impacts the chances of interaction between interfacing materials at the nanometer level, making it germane to many energy catalysis and materials design applications.

- The discrete, geometrical definition of conventional molecular surfaces is also rather inflexible. While several definitions exist, including Solvent-Excluded, Solvent-Accessible, and Molecular Skin surfaces, each of these have to be computed and visualized separately, and there is no easy way of visually 'interpolating' between them in a meaningful way. Yet, there is no single appropriate definition of a surface that can capture boundaries in all possible applications. Instead, scientists often want to visually explore multiple interpretations of surfaces and transition smoothly and interactively between them during analysis. Tools such as VMD [9] allow one to explore different boundaries by specifying different iso values. However, this cannot be always performed fast enough in real-time due to the complexity of constructing isosurfaces.

3.2 Volumetric molecular surfaces

Molecular boundaries are largely influenced by subatomic particles (e.g. electrons), which ultimately impact the probability of

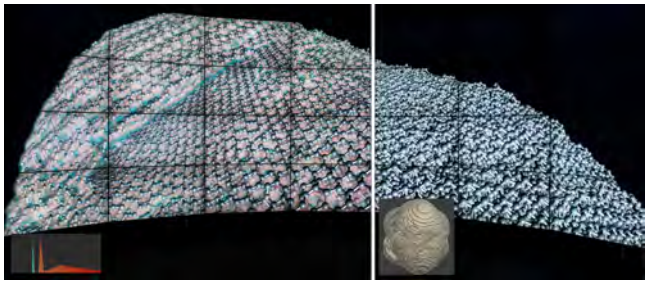

Figure 3: An amorphous carbon nanosphere used as battery material. The left side illustrates a semi-transparent volume + glyphs combination, where as the right side shows a standalone ball-and-stick model. Structural features, such as open cavities where Lithium ions tend to diffuse (shown with blue outline), become much more evident with volume rendering. The left inset illustrates the transfer function used to the surfaces. The right inset shows an overview of the entire nanosphere structure.

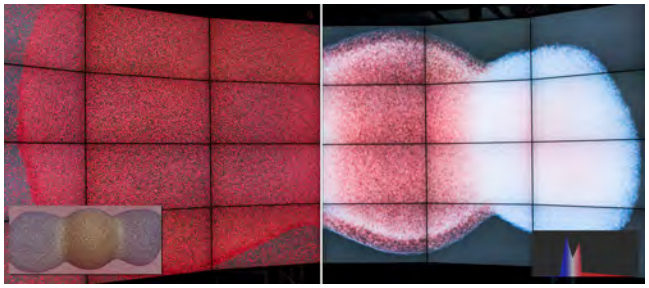

Figure 4: Simulated combustion of three aluminum nanoparticles coated in amorphous alumina computed in a 15 million atoms Molecular Dynamics simulation. The right side illustrates volume rendering of approximate electron densities, where as the left side shows sphere glyphs only (red for oxygen and black for aluminum). With volume rendering we can visually classify the amorphous alumina shell from the core aluminum atoms and convey differences in core atom densities between the different nanoparticles.

interaction between adjacent molecules. Accordingly, electron density presents a reasonable source of information when determining molecular surfaces.

Our method employs direct volume rendering to visualize approximate electron densities as proxies for surfaces. We use a conventional transfer function to assign distinct colors and opacities to different electron density ranges, resulting in different surface representations. This also allows us to smoothly transitioning between different surface interpretations by simply varying the transfer function. Traditional discrete molecular surfaces can also be obtained using a transfer function with a *Dirac* impulse at the desired iso value. Figure 2 illustrates the expressiveness of this volumetric formulation.

The ability to interactively transition between different representations of boundaries is also useful when one is trying to visually classify distinct regions in the molecular structure. For example, in Figure 4, we were able to visually classify the amorphous alumina shell from the core aluminum atoms in the three nanoparticles.

3.3 The need for high-resolution immersion

To explore and make sense of complex simulations, materials scientists are increasingly turning to immersive environments which provide stereoscopic depth, user-centered perspective, and embodied navigation in six degrees-of-freedom (e.g. [22, 19]). Immersion can greatly enhance spatial understanding of 3D structures [2], and facilitate the exploration of complex molecular simulations [4]. However, achieving good perceptual scalability requires high-resolution displays capable of delivering high visual acuity [23]. Yet, the majority of immersive molecular visualizations are intended for use on low-resolution immersive platforms. While these platforms offer an added degree of immersion compared to conventional desktop and laptop monitors, they supply limited resolution and poor image quality. These technical limitations restricts the amount of data that can be shown at a time and limits the complexity of visual encodings that can be employed. Volumetric representations, for instance, are used scarcely because they are inherently difficult to perceive in traditional, projectors-based immersive environments.

Modern LCD-based immersive systems on the other hand provide nearly 20/20 visual acuity. The *CAVE2* system, for instance, provides a 320-degrees panoramic view at 74 Megapixels aggregate stereoscopic resolution (37 Megapixels per eye) [17]. The use of stereoscopic LCD panels in place of projectors makes for a very bright display environment with high contrast ratio. This enables scientists to explore complex 3D representations of massive molecular structures that are orders of magnitude larger, compared to what can be visualized with older immersive environments. The challenge, however, is to render such structures in high-resolutions at interactive frame rates.

We employ an image-parallel ray-casting algorithm and use implicit mathematical representations of the ball-and-stick glyphs. To generate high-quality molecular surface representations, we employ higher-order filtering functions (B-spline and Catmull-Rom splines) in addition to supporting grid-independent computation of approximate electron densities.

3.4 Interactive features

CAVE2 provides tether-less 3D navigation using a wireless 6 degrees-of-freedom joystick coupled with head-tracking. We implemented a kite interaction metaphor; the user flies through the simulation by pressing the trigger button on the joystick and physically moving and/or rotating the joystick in the desired direction, which translates and rotates the user's vantage point accordingly. Wireless head tracking also provides a user-centered perspective, allowing the tracked user to explore the simulation by physically moving in the immersive environments and by turning his/her head.

The user can also edit the transfer function interactively using the wireless wand, or using a separate user interface on a laptop or tablet display. By changing the shape of the transfer function or the colors assigned to electron density ranges, the user indirectly affects the resulting surface representation (see Figure 2). Modifying the transfer function updates the visualization in real-time, allowing the user to smoothly transition through various surface interpretations. The user also has control over all aspects of the visualization, including the overall opacity of the volume and the quality of the rendering, all of which can be changed interactively.

4 Implementation

We give an overview of the implementation, illustrate the data structures, and then describe the ray-casting algorithm.

4.1 Overview

Our implementation uses a distributed, image-parallel algorithm to perform volume rendering of electron density fields and to ray-cast ball-and-stick glyphs in one pass. To speed up rendering, we employ an in-core solution making sure at least one time step of the simulation is replicated and loaded into the GPU memory of all rendering nodes. For longitudinal MD simulations with a large number of time steps, we page-in time steps from the main memory to the GPU just before rendering.

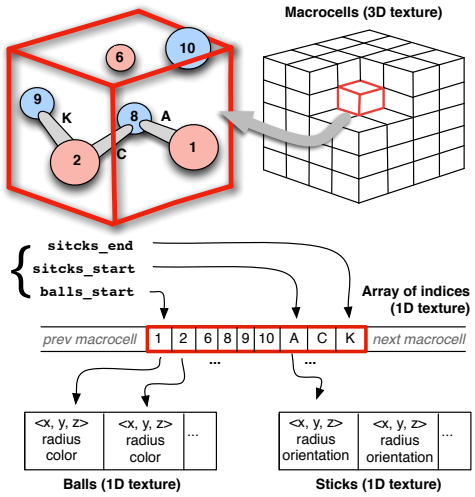

Figure 5: Illustration of the Macrocells data structure used to store analytical representations of ball-and-stick glyphs in GPU memory.

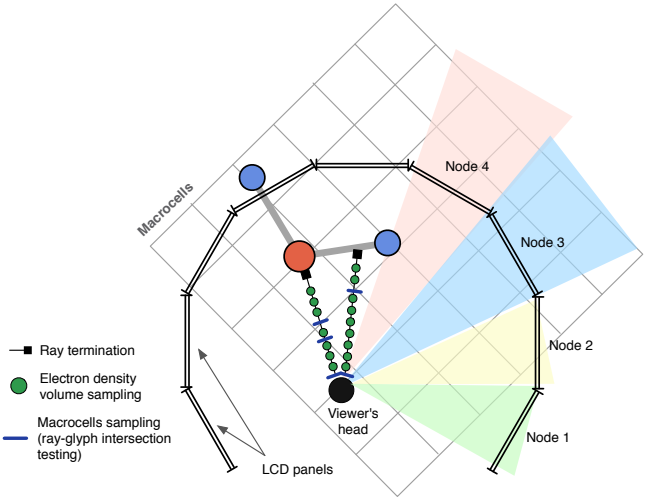

Figure 6: Illustration of the immersive, distributed ray-casting process.

To scale up the visualization to high-resolution multi-panels immersive systems, we employ multiple GPU nodes for rendering with each node assigned a rectangular portion of the scene for load balancing. In the *CAVE2* system, we use a 36-node cluster to render the immersive visualization on 72 LCD panels in parallel. To provide a viewer-centered perspective, the algorithm cast rays from the position of the user's head to each of the panels. Figure 6 illustrates this process.

Stereoscopic rendering is achieved by rendering and interleaving two images separated by the average inter-pupillary distance (63 mm). To support correct stereoscopic viewing for multiple co-located individuals (e.g. Figure 7), we normalize the head orientation when rendering the view of each panel so that the eyes are always parallel to the screen's surface. While this results in some distortion between the tiles, this distortion is minimal in practice, and a good compromise to enable multiple collaborators to simultaneously view the visualization (albeit, from the perspective of a single tracked user).

4.2 Data structures

We employ a uniform 'Macrocells' 3D grid to represent ball-and-stick glyphs analytically in the GPU memory. Each voxel (typically 64 $\AA^3$) in the Macrocells grid stores information about the atoms and bonds that are within the extent of that cell, or those that partly overlap with the borders of the cell. A single ball or stick spanning multiple cells will therefore be represented in all of them. However, instead of replicating atom and bond data across several cells we store pointers to this information; each cell contains three pointers into a contiguous 'indices array'. The indices array contains the global ID of every ball and stick in the simulation. The actual physical properties of the balls (center, atom type, radius, color) and sticks (center, cylinder length, cylinder radius, orientation) are stored in a separate array indexed by a global primitive ID. Figure 5 illustrates this scheme.

4.3 Approximating electron density

Electron charge densities can be plotted from density functional theory (DFT) computations. For most MD simulations and particularly large classical MD, DFT computation is infeasible. Therefore, we provide a method to approximate electron density using Gaussian distributions, either based on the Van der Waals radii of atoms or averaged from bulk DFT computation, if that data is provided.

The approximate electron density at point X in the 3D space is the summation of the electronic contribution from all atoms that are within a limited distance of X. The contribution of a single atom centered on C to X is a radial basis function:

$$\phi(X,C) = \phi(\|X-C\|) \tag{1}$$

For a point X in macrocell M, the total electronic density is:

$$\psi(X) = \sum_{C \subset M} \phi(\|X-C\|) \tag{2}$$

We use a Gaussian formulation of $\phi(x) = \sqrt{w}e^{\frac{-x^2}{w^2}}$ where w is the covalent radius of the atom (assuming a Van der Waals radius of roughly twice the covalent bond length). This approximation can be precomputed offline over a uniform 3D grid. Alternatively, it can be done in real-time by evaluating equation 2 along they ray. While the latter solution is more attractive as it gives us a resolution-independent solution, it could potentially slow down the visualization. We therefore report rendering times for each separately in Table 1.

4.4 Ray-casting

The ray-casting algorithm traverses the Macrocells grid at discrete intervals using the 3DDA algorithm [1], and reads the list of geometry indicated by the `balls_start`, `sticks_start`, and `sticks_end` pointers (Figure 5). Because the list of glyphs cannot be depth-ordered within a macrocell for all possible perspectives, the algorithm has to test every primitive (sphere or cylinder) in the current macrocell against the ray. If a hit is detected, the surface of the closest glyphs is Phong shaded. After testing all the glyphs and determining the hit point, the charge density volume is sampled within the current macrocell and the color is accumulated according to the transfer function.

The algorithm terminates the ray when it hits one of the glyphs (we only support fully opaque glyphs at the moment). Otherwise, the algorithm proceeds along the ray to the next macrocell and repeats the above process. We also skip empty Macrocells (electron density is zero without atoms) and use an opacity termination scheme to stop the ray when the accumulated volume opacity saturates the pixel. Figure 6 illustrates this scheme.

Table 1: Data characteristics and rendering speeds in frames per second (FPS) achieved in the *CAVE2* environment

Dataset	Atoms	Precomputed density grid			Real-time density computation	
		Density grid size	**Mono** FPS *min - max (avg)*	**Stereo** FPS *min - max (avg)*	**Mono** FPS *min - max (avg)*	**Stereo** FPS *min - max (avg)*
Octyne reactivity study	200	301 x 301 x 301	6 - 70 (15)	5 - 64 (12)	4 - 14 (7)	2 - 7 (2)
Carbon Nanosphere	740K	219 x 219 x 219	7 - 31 (17)	8 - 23 (12)	2 - 10 (5)	1 - 4 (2)
Amorphous SiO_2 Glass fracture	5M	1601 x 567 x 107	5 - 28 (15)	4 - 30 (13)	1 - 3 (3)	1 - 3 (2)
Al_2O_3 Nanoparticles	15M	1221 x 611 x 611	3 - 12 (7)	2 - 32 (5)	< 1	< 1

5 USE CASES

We illustrate our techniques with four example use cases in nanoscale materials design and computational chemistry. We also report on preliminary validation results with users from these domains. For each of the example datasets, we generated approximate electron density volumes, except for the Octyne Reactivity study which came with a DFT-computed charge density volume.

1. The first dataset is a small-scale DFT simulation. The system is composed of 5 amine molecules (Octylamines) and 1 hydrocarbon (Octyne) molecule adsorbed together on the same face of a 35 atoms Platinum pyramidal cluster. The goal of the study is to investigate the reactivity of the Octyne molecule and how it is affected by the presence of amines (Figure 7).

2. The second dataset represents a quasi-spherical amorphous carbon structure used as the anode material for intercalation in Lithium-ion energy storage. The structure has been obtained from an MD simulation involving approximately 740,000 atoms (Figure 3).

3. The third dataset is a nanoscale simulation of an amorphous glass (SiO_2) fracture computed in a 5 million atoms MD simulation (Figures 1 and 2).

4. The forth datasets is a MD study of the combustion of three aluminum nanoparticles coated in amorphous alumina (Al_2O_3). The modeled system comprises 15 million atoms densely pack into three nanoparticles. The study aims to understand the stages in the burning process (Figure 4).

Figure 7: Our technique can also be applied to visualize electronic structures obtained from Density Functional Theory (DFT) simulations. High-quality rendering of electron and molecular orbitals can be achieved with an appropriate transfer function. The inset illustrates an overview of the structure.

5.1 Rendering performance

We evaluated our implementation in the *CAVE2* immersive environment [17]. *CAVE2* is a cylindrical immersive environment measuring 24 feet in diameter and 8 feet in height. It consists of 72 micro-polarized, passive stereo LCD panels that are arrange on the circumference of the cylinder, providing a 320-degrees panoramic view at a total stereoscopic resolution of 74 Megapixels (37 Megapixels per eye). The system is driven by a 36-nodes cluster with each node driving two of the 72 LCD panels at a resolution of 2 Megapixels. Each of the 36 nodes boasts an *Nvidia GTX 680* graphics card with 2GB of video memory, 64GB of main memory, and a 16-core 2.9GHz *Intel Xeon E5-2690* processor.

To get a sense of the performance under real usage conditions, we generated plausible fly-through trajectories, moving through the datasets to generate different view configurations ranging from closeups to overviews. Table 1 lists the worst, best, and average frame times achieved in the *CAVE2* system, breaking them down by the volume rendering method (precomputed electron density grid vs. real-time density computation in the GPU). We also report separate results for stereoscopic and monoscopic rendering.

In the four datasets we evaluated, we were able to achieve real-time interactive rendering using precomputed electron density grids (columns 4-5 of Table 1). The rendering algorithm scales solidly up to 5 million atoms with an average frame rate of 15 FPS. The frame rate halves by the time we get to 15 million atoms, but we still achieve interactive rendering at 7 FPS. Non-surprisingly, rendering is about 2-3x faster with a precomputed electron density grids, compared to real-time density computation (columns 6-7). Yet, we were able to achieve reasonably interactive speeds with the latter, except for our largest dataset. These results are generally competitive with published SPH rendering techniques [5].

5.2 Preliminary user evaluation

The development of the proposed technique was a collaborative effort with computational chemists, physicists, and materials scientists at Argonne National Laboratory and the University of Southern California. As discussed in section 3, our goal was perceptual scalability and the desire to highlight emergent high-level features in the molecular structure. Secondly, we wanted a flexible definition of molecular surfaces to enable scientists to explore different boundary interpretations. To provide some user validation of these goals, we report our informal observations on how scientists utilized the visualization to view and analyze their molecular datasets. These observations shed a light on the advantages of the hybrid visual metaphor and the analytical tasks it supports. However, further evaluation is needed to formally validate these observations and further refine the technique.

For large-scale MD datasets (such as the simulation of amorphous glass fracture and the combustion of aluminum nanoparticles), volumetric surfaces were particularity successful at abstracting low-level details in the structure. For instance, in the glass fracture dataset, one researcher was able to identify the fracture mode in the amorphous silica system by flying into the fracture. The task

was simplified by rendering an opaque surface with a transfer function that emphasizes low electron densities, which consolidated the structure. This in turn allowed the researcher to visually compare the spatial configuration of the two sides of the fracture with relative ease. The researcher was also interested in understanding the morphology of the fracture surface. Is it sharp and clean, or is it rough and bumpy? And how does it vary across the fracture? A semi-transparent volumetric surface allowed the researcher to investigate those questions. Additionally, the researchers was able to see through the electron charge cloud and investigate atom arrangement patterns near the surface of the fracture. The electron cloud also served to naturally reduce visual clutter by producing a 'fog' effect, which reduced interference with distant atoms.

Volume rendering was also useful as a visual classification tool when the number of atoms in the system was too large to render them individually. For example, in the simulation of aluminum nanoparticles combustion, we able to visually segment the amorphous alumina shell from the aluminum core in the center of the nanoparticle using an appropriate transfer function (see Figure 4). We used a cross-sectional clipping plane to slice through the data, which allowed the researcher to see how oxygen and aluminum migrate at the interface of the particles by looking at changes in the electron density cloud.

The hybrid volume + glyphs technique was also useful for computational chemists who need to see the molecular as well as the electronic structure. Volume rendering proved more flexible compared to traditional isosurface-based techniques, allowing researchers to plot a range of charge density values as uncertainty clouds, in addition to the ability to emphasize known electron orbitals with sharp peaks in the transfer function (see Figure 7).

6 Conclusion

Large-scale Molecular Dynamics simulations pose a challenge for current visualization techniques. For many nanoscale materials design and alternative energy applications, conventional molecular surfaces are semantically inappropriate as they cannot portray the uncertainty in boundaries at the nanoscale. Real-time rendering of molecular surfaces is also prohibitively expensive. In this paper, we argued for a hybrid visual metaphor for the visualization of complex nanostructured materials via volume rendering of approximate electron densities. We demonstrated the expressiveness of this approach and its scalability with respect to data size and the resolution of the output framebuffer. We also described a distributed GPU ray-casting scheme for visualizing large-scale atomistic simulations in ultra-resolution immersive environments. While orthogonal, the coupling of these two components result in clear analytical advantages over existing molecular visualization techniques.

Acknowledgements

The authors wish to thank Lance Long, Mark Hereld, Adarsh Shekhar, Priya Vashishta, Aslihan Sumer, Julius Jellinek, Kah Chun Lau, and Maria Chan. This work was supported by National Science Foundation award CNS-0959053 and U.S. Department of Energy contract DE-SC005067. Funding was also provided by the Office of Advanced Scientific Computing Research, Office of Science, U.S. DOE contract DE-AC02-06CH11357. CAVE and CAVE2 are trademarks of the University of Illinois Board of Trustees.

References

[1] J. Amanatides, A. Woo, et al. A fast voxel traversal algorithm for ray tracing. In *Proceedings of EUROGRAPHICS*, volume 87, pages 3–10, 1987.

[2] D. A. Bowman and R. P. McMahan. Virtual reality: how much immersion is enough? *Computer*, 40(7):36–43, 2007.

[3] M. L. Connolly. Analytical molecular surface calculation. *Journal of Applied Crystallography*, 16(5):548–558, 1983.

[4] C. Cruz-Neira, J. Leigh, M. Papka, C. Barnes, S. M. Cohen, S. Das, R. Engelmann, R. Hudson, T. Roy, L. Siegel, et al. Scientists in wonderland: A report on visualization applications in the cave virtual reality environment. In *Virtual Reality, 1993. Proceedings., IEEE 1993 Symposium on Research Frontiers in*, pages 59–66. IEEE, 1993.

[5] R. Fraedrich, S. Auer, and R. Westermann. Efficient high-quality volume rendering of sph data. *Visualization and Computer Graphics, IEEE Transactions on*, 16(6):1533–1540, 2010.

[6] D. S. Goodsell, I. S. Mian, and A. J. Olson. Rendering volumetric data in molecular systems. *Journal of Molecular Graphics*, 7(1):41–47, 1989.

[7] S. Grottel, G. Reina, C. Dachsbacher, and T. Ertl. Coherent culling and shading for large molecular dynamics visualization. In *Computer Graphics Forum*, volume 29, pages 953–962. Wiley Online Library, 2010.

[8] M. Hu, W. Chen, T. Zhang, and Q. Peng. Direct volume rendering of volumetric protein data. In *Advances in Computer Graphics*, pages 397–403. Springer, 2006.

[9] W. Humphrey, A. Dalke, and K. Schulten. Vmd: visual molecular dynamics. *Journal of molecular graphics*, 14(1):33–38, 1996.

[10] Y. Jang and U. Varetto. Interactive volume rendering of functional representations in quantum chemistry. *Visualization and Computer Graphics, IEEE Transactions on*, 15(6):1579–5186, 2009.

[11] A. Knoll, M. K. Chan, K. C. Lau, B. Liu, J. Greeley, L. Curtiss, M. Hereld, and M. E. Papka. Uncertainty classification and visualization of molecular interfaces. *International Journal for Uncertainty Quantification*, 3(2), 2013.

[12] A. Knoll, Y. Hijazi, R. Westerteiger, M. Schott, C. Hansen, and H. Hagen. Volume ray casting with peak finding and differential sampling. *Visualization and Computer Graphics, IEEE Transactions on*, 15(6):1571–1578, 2009.

[13] M. Krone, K. Bidmon, and T. Ertl. Interactive visualization of molecular surface dynamics. *Visualization and Computer Graphics, IEEE Transactions on*, 15(6):1391–1398, 2009.

[14] M. Krone, M. Falk, S. Rehm, J. Pleiss, and T. Ertl. Interactive exploration of protein cavities. In *Computer Graphics Forum*, volume 30, pages 673–682. Wiley Online Library, 2011.

[15] N. Lindow, D. Baum, S. Prohaska, and H.-C. Hege. Accelerated visualization of dynamic molecular surfaces. In *Computer Graphics Forum*, volume 29, pages 943–952. Wiley Online Library, 2010.

[16] A. Nakano, R. K. Kalia, and P. Vashishta. Scalable molecular-dynamics, visualization, and data management algorithms for materials simulations. *Computing in science and engineering*, 1(5):39–47, 1999.

[17] K. Reda, A. Febretti, A. Knoll, J. Aurisano, J. Leigh, A. Johnson, M. Papka, and M. Hereld. Visualizing large, heterogeneous data in hybrid-reality environments. *Computer Graphics and Applications*, 33(4):38–48, 2013.

[18] H. Scharsach, M. Hadwiger, A. Neubauer, S. Wolfsberger, and K. Bühler. Perspective isosurface and direct volume rendering for virtual endoscopy applications. In *Proceedings of the Eighth Joint Eurographics/IEEE VGTC conference on Visualization*, pages 315–322. Eurographics Association, 2006.

[19] A. Sharma, R. K. Kalia, A. Nakano, and P. Vashishta. Large multidimensional data visualization for materials science. *Computing in Science and Engineering*, 5(2):26–33, 2003.

[20] C. Sigg, T. Weyrich, M. Botsch, and M. Gross. Gpu-based ray-casting of quadratic surfaces. In *Proceedings of the 3rd Eurographics/IEEE VGTC conference on Point-Based Graphics*, pages 59–65. Eurographics Association, 2006.

[21] J. M. Singh and P. Narayanan. Real-time ray tracing of implicit surfaces on the gpu. *Visualization and Computer Graphics, IEEE Transactions on*, 16(2):261–272, 2010.

[22] J. E. Stone, K. L. Vandivort, and K. Schulten. Immersive out-of-core visualization of large-size and long-timescale molecular dynamics trajectories. In *Advances in Visual Computing*, pages 1–12. Springer, 2011.

[23] C. Ware and P. Mitchell. Visualizing graphs in three dimensions. *ACM Transactions on Applied Perception (TAP)*, 5(1):2, 2008.