



This is a repository copy of *Multiscale modeling and cinematic visualization of photosynthetic energy conversion processes from electronic to cell scales*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/181671/>

Version: Accepted Version

Article:

Sener, M., Levy, S., Stone, J.E. et al. (8 more authors) (2021) Multiscale modeling and cinematic visualization of photosynthetic energy conversion processes from electronic to cell scales. *Parallel Computing*, 102. 102698. ISSN 0167-8191

<https://doi.org/10.1016/j.parco.2020.102698>

© 2020 Elsevier B.V. This is an author produced version of a paper subsequently published in *Parallel Computing*. Uploaded in accordance with the publisher's self-archiving policy. Article available under the terms of the CC-BY-NC-ND licence (<https://creativecommons.org/licenses/by-nc-nd/4.0/>).

Reuse

This article is distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivs (CC BY-NC-ND) licence. This licence only allows you to download this work and share it with others as long as you credit the authors, but you can't change the article in any way or use it commercially. More information and the full terms of the licence here: <https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>



Published in final edited form as:

Parallel Comput. 2020 May ; 102: . doi:10.1016/j.parco.2020.102698.

Multiscale modeling and cinematic visualization of photosynthetic energy conversion processes from electronic to cell scales

Melih Sener^{a,*}, Stuart Levy^{#b}, John E. Stone^{#a}, AJ Christensen^b, Barry Isralewitz^a, Robert Patterson^b, Kalina Borkiewicz^b, Jeffrey Carpenter^b, C. Neil Hunter^c, Zaida Luthey-Schulten^d, Donna Cox^{a,*}

^aBeckman Institute, University of Illinois at Urbana-Champaign

^bAdvanced Visualization Laboratory, NCSA, University of Illinois at Urbana-Champaign

^cDepartment of Molecular Biology and Biotechnology, University of Sheffield, Sheffield, U.K.

^dDepartment of Chemistry, University of Illinois at Urbana-Champaign

These authors contributed equally to this work.

Abstract

Conversion of sunlight into chemical energy, namely photosynthesis, is the primary energy source of life on Earth. A visualization depicting this process, based on multiscale computational models from electronic to cell scales, is presented in the form of an excerpt from the fulldome show *Birth of Planet Earth*. This accessible visual narrative shows a lay audience, including children, how the energy of sunlight is captured, converted, and stored through a chain of proteins to power living cells. The visualization is the result of a multi-year collaboration among biophysicists, visualization scientists, and artists, which, in turn, is based on a decade-long experimental-computational collaboration on structural and functional modeling that produced an atomic detail description of a bacterial bioenergetic organelle, the chromatophore. Software advancements necessitated by this project have led to significant performance and feature advances, including hardware-accelerated cinematic ray tracing and instanced visualizations for efficient cell-scale modeling. The energy conversion steps depicted feature an integration of function from electronic to cell levels, spanning nearly 12 orders of magnitude in time scales. This atomic detail description uniquely enables a modern retelling of one of humanity's earliest stories—the interplay between light and life.

*Correspondence to be addressed to MS: melih@ks.uiuc.edu and DC: donnacox@illinois.edu.

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

Declaration of Competing Interest

Authors declare no conflict of interest.

1. Introduction

Public outreach of science, through efforts such as fulldome planetarium shows, fulfills a crucial educational need by bringing leading edge scientific discoveries to a populace that supported the research yet often remains unable to access it due to its inherent complexity [1]. Such a fulldome show, *Birth of Planet Earth*, introduces the story of life emerging on our planet and, particularly, how it is powered by harvesting the energy of sunlight, i.e., photosynthesis [2]. Here, an excerpt¹ is presented as a stand-alone explanatory visualization (Fig. 1A). The accessible narrative of this visualization constitutes a modern extension of one of the oldest discoveries of humanity, namely, the relationship between sunlight and the proliferation of life. (“The sun is not only the author of visibility in all visible things, but of generation and nourishment and growth.”—Plato, *The Republic*.) Thus, this visualization is a retelling of the ancient story of “nourishment and growth,” but with atomistic accuracy obtained by a decade of experimental and computational research [3, 4, 5, 6, 7, 2], rendered accessible to children as a result of a multi-year collaboration between biophysicists, visualization scientists, and artists.

High performance computing (HPC) exhibits a dual role in the study of the complex biomolecular processes of photosynthesis. On the one hand, HPC is a critical *exploratory* tool for building atomic detail structures and for multi-scale modeling of energy conversion steps from electronic interactions to cell level integration [7, 2]. On the other hand, HPC provides an *explanatory* tool for depicting the cascade of energy conversion steps in a continuous ‘powers of ten’ style visual narrative [8, 9]. Efficient harvesting and conversion of solar energy, thus depicted, may hold the key for a clean energy future for humanity [10] with biological systems outperforming comparable human technologies at system-level [2].

The visualization presented here was awarded the ‘Best Scientific Visualization & Data Analytics Showcase’ at the Supercomputing 2019 (SC19) conference. Several performance and capability improvements were developed for NAMD and VMD as a result of experiences gained during show production, including up to 8× faster hardware-accelerated cinematic ray tracing, and instanced visualizations for efficient cell-scale modeling, described in Sec. 5. The visualization involved the development of novel workflows by a collaborating team of biophysicists (Fig. 1B) and visualization scientists (Fig. 1C), bridging multiple length and time scales as well as the corresponding data modalities and software tools. These workflows present an example for creating new visual narratives that bring complex processes in nature to broad audiences.

2. Summary of bioenergetic processes depicted

Bioenergetic processes in photosynthesis involve integration of function across hundreds of proteins [3, 4, 6] and span length and time scales ranging from electronic excitation transfer (femtoseconds/picoseconds) [5] to organelle-scale diffusion of charge carriers [7] and ATP synthesis [11, 12] (milliseconds) to adaptive response of protein composition to changing environmental conditions at the cell level (hours) [2, 13]. A combination of experimental

¹See a zoomed, flat-screen version of the fulldome show excerpt at: <https://www.youtube.com/watch?v=NTgAok6n7I4>

and theoretical approaches (Fig. 1B) involving multiple data modalities (Fig. 2A,B) are employed to determine structure and function at atomic, supra-molecular, organelle, and cell levels of organization [2, 7, 14].

The visualization here is based on the purple phototrophic bacterium, *Rhodobacter (Rba.) sphaeroides*, resembling evolutionarily primitive light-harvesting organisms [15]. Typically limited to anoxic habitats, *Rba. sphaeroides* performs a non-oxygenic form of photosynthesis that is significantly simpler than found in plants or cyanobacteria [16, 17]. The photosynthetic apparatus in this bacterium is organized in the form of spherical membrane vesicles of approximately 60 nm diameter, called chromatophores (Fig. 1A), comprising several hundred proteins containing an array of up to 3,000 bacteriochlorophylls for light capture and charge separation [3]. The chromatophore structural models permit the computation of electronic interactions between bacteriochlorophyll molecules (Fig. 2C) using effective Hamiltonian formulations [5] and, subsequently, the rate kinetics of electronic excitation transfer, i.e., delivery of the captured solar energy to reaction center proteins that initiate charge transfer and diffusion steps [6]. (For a summary, see Refs. [7, 2].)

The interaction of the protein arrays across an entire chromatophore, integrated through multi-scale rate kinetics formulations [6, 7, 14], determines the rate of formation of the final photoproduct, ATP, the universal energy currency of all life on Earth [12]. The ATP production rate as a function of incident light intensity (Fig. 2D) provides a quantitative measure of efficiency for the chromatophore as an energy conversion device. A related performance metric, energy return-on-investment, permits a comparison with human-made devices of similar purpose, where the chromatophore outperforms photovoltaic or fossil fuel based systems [2]. Combining the collective light-response of all chromatophores in an entire bacterium permits the computation of reproductive efficiency at cell level, related to the cell doubling time as a function of growth conditions [2] (Fig. 2E).

This electronic-to-cell scale organization of energy conversion steps provides a natural event sequence for a storyboard (Fig. 2F) used to construct the present visualization in the form of an accessible narrative (Fig. 3).

3. Constructing an accessible narrative

Communicating a complex biophysical process to a lay audience requires elements beyond those needed in an exploratory visualization aimed at a scientific audience [20]. A balance between accuracy and comprehensibility was sought by the team in conceptualizing, researching, designing, choreographing, writing narration scripts, and testing audiences iteratively for this explanatory visualization. The primary goal was to tell the story of how life is powered by sunlight in a way that was informative yet not burdensome, educational yet not didactic, and visually faithful to scientific complexity yet intuitively simple. All technical work described in Sec. 4 below is rooted in this goal. Fig. 3 illustrates, scene-by-scene, how this narrative balance was maintained throughout the visualization.

3.1. Accessibility of the narrative

Dome shows can reach millions of viewers between the target ages of 8 and 80, have a lifespan of over a decade, are often translated into a dozen or more foreign languages, and are distributed worldwide. Creating a long-lasting visualization accessible to diverse audiences requires an understanding of human perception, education, psychology, and visual communication practices. An inspiring and engaging visualization additionally requires storytelling, cinematography, and production.

The first step in planning an accessible narrative is the storyboard (Fig. 2F). The storyboard is the basis for scene design and camera treatment to ensure accessibility for the target audience as they are exposed to multiple size and time scales in a complex process. A script for narration was co-developed alongside the storyboard with writer Thomas Lucas along with the biophysics and visualization teams, including the late Klaus Schulten. The final script was recorded by a professional actor, Richard Dormer, to ensure that the narration is performed with empathetic, emotional delivery that connects with the audience. Advance screenings for audience testing were held for the in-progress documentary where psychology researchers evaluated whether the visual narrative was enjoyable, met learning goals, and raised awareness of computational science. The show was revised accordingly.

3.2. Inclusion within a larger story

The present visualization must fit within a larger story as it is the culminating scene in a documentary explaining how Earth came to be the cradle of life—starting with the remnants of a supernova, accretion of a molecular cloud that lends materials to the formation of the Solar System, the creation of our Earth-Moon system, and, at last, how early life was powered by photosynthesis.

The immediately preceding scene shows a pool of bubbling water surrounded by rocky, prehistoric-looking terrain created from reference photography in Iceland. Visual cues are made to flow smoothly between these two scenes as the camera smoothly zooms from the macroscopic to the atomic scale. To cross this scale gap between that of a pool and a solitary chromatophore, a larger contextual environment is constructed, consisting of (i) a colony of hundreds of bacterial cells (Fig. 3A) and (ii) a bacterial cell model containing hundreds of chromatophores (Fig. 3B), providing a background for the individual ‘hero’ chromatophore upon which the action later focuses (Fig. 3C).

3.3. The writer’s dilemma and the visual dilemma

The compromise between accuracy and comprehensibility brought a number of difficult design challenges. First, relevant size scales are smaller than the wavelength of visible light; therefore, intuitive meanings of color and opacity become inapplicable. Second, scene complexity of a cellular environment is overwhelming, especially for a child who gets only brief exposure to it. Third, the processes depicted span a vast variety of time scales (femto-/picoseconds to milliseconds, as discussed in Sec. 2). A simple linear scaling of time would have implied that some processes appear at standstill, while others become blurs. Fourth, some processes are causally connected yet are distant from each other, thus are not obviously associated. Fifth, the primary processes involved in the interaction of light

with biological matter are quantum mechanical (Fig. 3D), which human perception has not evolved to process. Any accurate depiction of probability dynamics, as would be customary for a scientific presentation involving quantum processes, would be incomprehensible to a lay audience.

Creating an accessible narrative for an eight year old in the face of these challenges requires careful simplification of the story, while still retaining the underlying structural model as a basis for the visualization. For example, picosecond scale excitation transfer steps are depicted alongside millisecond scale charge diffusion steps. However, a qualitative ordering of event speeds is employed so that fast events do appear faster to the viewer, though not commensurately with their actual time scales. The chromatophore and its host cell are depicted in terms of protein scaffolds only, where lipids are abstracted as transparent membranes. Charge carrier molecules (quinol/quinone) are visualized as glowing point particles instead of molecular models (Fig. 3E), and their dynamics are abstracted as a semi-directed flow instead of stochastic diffusion [7]. The protein complex responsible for ATP synthesis (ATP synthase) is animated as a combination of ‘turnstile’ and ‘mushroom’ components (F_0 and F_1 units, respectively), undergoing directed, cohesive motions (Fig. 3F–G), instead of a stochastic machine [12]. Components of the ATP synthase and its interaction partners (protons, ADP/ATP) were synchronized to guide the perception of the viewer to the culmination of the story in the form of ATP release (Fig. 3G), powering life.

3.4. Choreography of elements in the fulldome environment

The motion of photons, protons, quinones/quinols, and ATP/ADP molecules were meticulously choreographed so that the viewer could follow the path of energy directly through a chain of conversion steps, creating a narrative visual thread through the chromatophore (Fig. 3A–G). Movements of many elements in this visualization were designed to appear random, but their motions are guided by the viewpoint of the camera.

A ‘hero’ quinol and a ‘hero’ proton were separately choreographed and given a brighter starburst appearance to guide the viewer’s eye, without being too visibly different from the surrounding elements. The ‘hero’ chromatophore was shown with transparent surfaces, allowing the viewer to see the internal structures (Fig. 3C). In a fulldome environment, the viewer is free and encouraged to look around at the action happening all around them, but their vision must be guided to a central location to where the story is evolving, including focusing one-by-one on successive steps of a multi-step process. This is done through the design of the scene as well as the motion of the camera.

Some energy conversion steps are causally related yet spatially separate—such as proton passage at the ATP synthase ‘turnstile’ (F_0) causing the synthesis of ATP from ADP and phosphate at the ‘mushroom’ (F_1) (Fig. 3F, G). To establish this crucial causal relationship in the narrative, we have employed a *rhythm*, coordinated with both visual and auditory cues, in the synchronization of the corresponding animation components. Specifically, as the ATP synthase is being climbed by the camera as a tower from the bottom, arrival of protons and the lurching forward of the ‘turnstile’ are locked onto one another in alternating steps of motion and stillness. Sound effects (‘thrum–thump, thrum–thump’) further emphasize this same rhythm, which is then continued further onto the animation of the ‘mushroom’. This

rhythm does not correspond to a physical process, which instead is stochastic [11], and is employed as a narrative tool only.

4. Visualization and rendering using HPC

The basis of the visualization presented here is the structural model of the chromatophore (Fig. 1A) [3, 4, 5, 6, 7] as well as corresponding atom selections, molecular representations, and animation scripts developed for exploratory visualizations [8, 9]. The molecular visualization software VMD [18] is unparalleled at generating visual abstractions of large biomolecular complexes and was the tool used to build the chromatophore structural models [7]. No single visualization tool provided the required range of rendering capabilities and visual effects, so several were integrated to create an effective production pipeline (Fig. 1C) [24], starting with computational models and ending with the fulldome visualization (Fig. 1A).

4.1. Visualization Tools (Virtual Director, Houdini, VMD, Nuke)

Visualization of organelle-scale biomolecular complexes presents many challenges in terms of I/O, host and GPU memory capacity, and efficient rendering. VMD uses sophisticated I/O mechanisms and molecular data structures that permit large molecular structures and long simulation trajectories to be processed efficiently [25, 26]. VMD incorporates a built-in GPU-accelerated ray tracing engine that excels at high-fidelity rendering of large biomolecular complexes with standard molecular representations, but particularly semi-transparent molecular surfaces, ambient occlusion lighting, and depth of field focal blur (Fig. 3C–G) [27, 28, 29]. The visualization algorithms in VMD generate compact geometry critical for rendering large biomolecular scenes within available GPU memory capacity [29]. To meet the needs of the project, VMD was extended with special material shaders for visual effects (e.g., Fig. 3D electronic excitation), and direct ray tracing of fulldome images (Fig. 1A) [9, 19].

VMD scripting enabled the camera choreography software, Virtual Director, to control all VMD camera parameters, and allowed effective interchange of scene geometry and motion effects with the visual effects software, Houdini. VMD was the tool of choice for portraying the ‘hero’ chromatophore, seen in close up, as it efficiently renders highly-detailed biomolecular data. But not all rendering could be handled by VMD, which lacked some visual effects, e.g., to create an “under water” appearance, needed for the larger bacterial cluster environment (Fig. 3A–B,D). VMD was used to generate the geometric data used in the other software tools.

The visual-effects package Houdini was a capable authoring system for constructing the larger environment—a pond, incoming photons of sunlight, and a cluster of bacteria (Fig. 3A–B)—and for rendering the resulting computationally demanding scenes. Houdini also served as an authoring platform for choreography of dynamic elements such as electronic excitations (Fig. 3D), protons and quinols/quinones (Fig. 3E), ATP molecules (Fig. 3G), and so on.

4.2. Previsualization and camera choreography

Virtual Director [30] software provided fluid cinematic camera path design (Fig. 2G) in a pre-visualized version of the scene. It is our preferred tool for camera design, supporting recording and editing of tours performed with an intuitive six-degree-of-freedom input device.

The pre-visualization of the scene included a simplified static chromatophore model that was exported from VMD into meshes which Virtual Director could display at interactive rates. Choreographed phosphate/ADP/ATP elements (Fig. 3G) from Houdini were incorporated as animated glyphs, providing an interactive representation of the scene for camera design. Further, using VMD scripting, a live connection was made from a running Virtual Director instance to a running VMD session. Interactive camera motion in Virtual Director promptly updated the fully-detailed VMD view (Fig. 2H).

The motion of the camera in an immersive fulldome theater can create a visceral impact on viewers, which can be used to create engagement and excitement, but risks causing motion sickness if not employed well. The camera's path is key to constructing the visualization narrative, guiding the viewer's attention throughout the scene from macroscopic to atomic scale. The camera treatment here was performed in two sections with different spatial scales: the flight through the bacterial cell colony (Fig. 3A–C) and the flight through the 'hero' chromatophore (Fig. 3C–G). The two performances were merged and blended to create a seamless, immersive journey that allows the viewers to maintain a sense of scale. The measured pacing and smooth motion of the camera enable the viewer to focus on the unfolding story.

For the camera performance, we approach and fly over the chromatophore, looking down on it as if it were a planet landscape. A photon is captured, represented as a column of light (Fig. 3D). We follow the generated excitation to the reaction center and then descend (Fig. 3E) into the chromatophore interior with the audience. A 'hero' proton released inside is followed toward the towering ATP synthase (Fig. 3F). The audience rises up to see the complex transforming machinery of the ATP production process (Fig. 3G), and then finally, the camera pulls away to have a grand view of the chromatophore environment.

4.3. Model construction in Houdini

Electronic excitations (Fig. 3D) were choreographed in Houdini using a random walk representing transfers between neighboring light-harvesting proteins in resemblance of excitation transfer kinetics [3, 5, 7]. This random walk ends with arrival at a 'hero' light-harvesting protein (LH1, Fig. 3C), followed by a jump inward to the so-called 'special pair' pigments at the reaction center, where quinols are released.

The quinols (Fig. 3E), storing captured light energy in their protonated state, traverse between membrane proteins in apparent stochasticity, with docking positions determined from the chromatophore structure [6]. Protons (Fig. 3E) were choreographed as five separate populations, with synchronized arrival/departure times at constituent proteins, and discarded when not visible on screen.

Houdini was further used to build the wider environment: a pond, incoming photons of sunlight, and a cluster of bacterial cells (Fig. 3A). Each bacterium comprised a volumetrically-rendered membrane and many copies of the chromatophore model (Fig. 3B). Chromatophores were packed into the cell interior in a perturbed hexagonal grid, to approximate the density of organelles in extreme low light adapted *Rba. sphaeroides* cells [2, 8]. The ‘hero’ chromatophore was placed near the outer edge of a bacterium, such that the camera move approach could see the featured chromatophore clearly. This placement facilitated a dissolve from a Houdini-rendered chromatophore to the VMD ‘hero’ render, and also enabled unobstructed lighting of the ‘hero’, since internal chromatophores were dimmed by ambient occlusion shading.

4.4. Exporting geometry for Houdini and reading into VMD

Dynamic elements designed using Houdini were exported as tables of particles for each animation frame, describing their configuration—position, orientation, brightness, unique id, ‘hero’ status. A VMD driver script created and destroyed corresponding VMD scene geometry elements as needed for each frame.

Electronic excitations within the chromatophore were embedded as transparent surfaces rendered by VMD. VMD was adapted so that the visualization team could add a *glow* shader to VMD’s shader library. The blue-glow excitation animations, as well as embedded proton and quinol elements, use this new feature.

4.5. Simulation/animation cross-over

Depicting ATP synthase (Fig. 3C, F–G) function combined molecular dynamics simulation data [11] for the ‘mushroom’ (F_1) part (Fig. 3G), a pulsed, stepwise rotation of the ‘turnstile ring’ (F_0) part (Fig. 3F), choreography of diffusing molecules, and matching timing to give the appearance of a cohesive biological machine, as mentioned in Sec. 3. No molecular dynamics simulation was available of the entire ATP synthase system, but one did exist of a related F_1 -like V-ATPase unit as it performed a complete catalytic cycle [11], so it was adapted and installed in dynamic relationship to the animated F_0 unit. A rhythmic synchronization was enforced between each pulsed motion of the ‘turnstile’ animation and the F_1 (‘mushroom’) simulation results, which was effective at visually tying together the mechanical connectivity of the F_0 and F_1 components.

4.6. Rendering Layers in Blue Waters

The visualization was rendered in 14 separate layers from both VMD and Houdini as hemispherical fulldome images in 4096×4096 resolution. A VMD batch rendering pipeline imported the Houdini-generated particle effects (Fig. 3D–E) and followed the Virtual Director camera path, rendering the structural model as a high-quality foreground element. Five versions of the VMD scene were rendered, with various combinations of moving elements, allowing those elements’ brightnesses to be adjusted in the Nuke compositing step. Nine more layers were rendered in Houdini’s Mantra, including instanced chromatophores (Fig. A-B), particle effects (Fig. D-E), photons, and cell membranes, along with depth passes to control relative luminance levels and depth effects. Houdini scene rendering was done on the Blue Waters supercomputer with a custom dome lens shader for

fulldome imagery, permitting an extended geometric complexity that could be represented in a scene—i.e., 300 1.3GB volumetric bacterial domains (Fig. 3A) containing more than 500,000 900MB chromatophore surface geometries (Fig. 3B). Rendering these layers was computationally intensive—the final run taking 27,763 node-hours, with many more used in preceding iterations.

A new process was developed for rendering on Blue Waters. The Houdini standard method of doing such large-scale rendering created over 100TB of intermediate data, which was not feasible to work with. This amount was reduced to 0.5TB by selectively creating minimal packed disk primitives from the geometry and linking these to the scene description files rather than including the data directly.

4.7. Compositing

Rendering of layers provided flexibility in the image composite to control luminance levels of the various elements and the magnitude of depth effects across the scene. It also provided efficiency in rendering by reducing the number of iterations for some layers as well as segmenting the computational demands of rendering. The ray tracing engine in VMD was augmented to generate alpha buffers for compositing. The 14 render layers from Houdini and VMD were composited with Nuke, where the scene was made to appear integrated and under water (Fig. 2I and 3D).

5. Post-Project Simulation and Visualization Advancements

After the completion of the *Birth of Planet Earth* visualizations and fulldome show as a whole, the team has continued development of the NAMD molecular dynamics simulation software and VMD molecular visualization software to increase their respective performance, and to implement missing features and software limitations that arose during the project.

Two of the most broadly valuable new developments in the VMD software motivated by this project are the addition of high level support for so-called *instancing* of molecular geometry, and a new built-in ray tracing engine implementation with support for hardware ray tracing acceleration on state-of-the-art GPUs.

5.1. Fast Single-Node GPU-Accelerated Molecular Dynamics

The NAMD parallel molecular dynamics simulation software [31, 32, 33] was used for all of the molecular dynamics studies [11] that informed the visualizations. The depicted ATP synthase rotation was based on 28 μ s of so-called string method with swarms of trajectories [34] supercomputer simulations of a 490,000-atom V-ATPase model. The simulations were used to calculate a minimal free energy pathway of hydrolysis stalk rotation, along with 37 μ s of free-energy and related supercomputer calculations to determine initial, less-refined rotation pathways and to characterize rotation simulation results [11].

A new fully GPU-resident version of NAMD 3 [35] allows so-called replica-exchange or multiple-copy simulations of million-atom biomolecular complexes that formerly required large-scale parallel computers to be performed effectively using more modest GPU-

accelerated workstations, cloud instances, or compute nodes. The new GPU-resident NAMD 3 achieves over 9× the simulation throughput over NAMD 2.13 on the NVIDIA DGX-A100 (8× A100 GPUs) for a standard millionatom STMV virus benchmark. The string method V-ATPase simulations described above use replica-exchange and multiple-copy methods and required a total 160 million SU on the OLCF Titan supercomputer; these simulations could be performed with 200 ns/day (estimated) throughput on a DGX-A100 system. If run today, the 28 μ s of string method simulations could be completed with two DGX-A100 machines — or with a cluster containing 16× A100 GPUs — in 70 days.

Ongoing work on exploiting next-generation GPU hardware platforms promises to further improve performance and decrease both runtime and cost of simulations, making it feasible for molecular dynamics simulations to play a greater role in facilitating, informing, and inspiring biomolecular visualizations for broad audiences.

5.2. Hardware Accelerated GPU Ray Tracing

At the very end of full-dome show development, the state-of-the-art for high performance ray tracing underwent radical change, with the introduction of the first commodity GPU hardware incorporating dedicated architectural extensions and machine instructions for acceleration of ray tracing algorithms. New GPU functional units provide hardware acceleration for bounding volume hierarchy (BVH) construction, BVH intersection and traversal, and ray-triangle intersection, all of which are low-level performance-critical operations for ray tracing of large biomolecular scenes. By routing hardware-accelerated ray tracing work items to dedicated GPU functional units, general purpose arithmetic units are left available to process remaining software-based ray tracing and shading algorithms. First generation GPU hardware ray tracing acceleration features enable ray casting rates as high as 10 billion rays per second when performance is not bound by shading calculations and the hardware is fully saturated with work. Second generation hardware acceleration for GPU ray tracing is just beginning to emerge at the time of writing, and has incorporated acceleration for additional curved tube-like geometric primitive types, and improved concurrent utilization of ray tracing-specific functional units and general purpose arithmetic hardware used for shading.

Molecular scenes are composed of diverse geometric primitives, including triangle meshes, arrays of spheres, cylinders, cones, discs, and others. These geometric primitives are used by VMD to construct higher level graphical representations of molecular structure, such as molecular surfaces, secondary structure “ribbons” and “tubes”, vector fields, and electrostatic field lines. Since only triangle meshes benefit from ray-triangle intersection acceleration hardware, the VMD/OptiX ray tracing engine must use a hybrid hardware- and software-based ray tracing approach to render complete scenes. Other geometric primitives such as spheres, cylinders, cones, and discs, are processed in software, by custom VMD CUDA kernels for ray-object intersection, surface normal computation, and so on [28]. When hardware acceleration is available, VMD funnels triangle mesh geometry through hardware-acceleration-specific software APIs that enforce hardware-specific vertex buffer data types, memory layouts, and other constraints that do not apply to the software-based ray tracing path. Software-based OptiX ray tracing allows completely application-

defined geometric primitives, with custom choice of vertex, normal, and color data types, and memory layout. VMD uses compressed surface normals (octahedron normal vector encoding) and quantized color representations to facilitate rendering of very large biomolecular structures in limited GPU on-board memory [29].

Table 1 and Table 2 show VMD performance results for pure software-based ray tracing and hardware-accelerated ray tracing for scenes with a high fraction of triangle mesh content (Table 1), and a low fraction of triangle mesh content (Table 2), respectively. Both test cases involve spherical and cylindrical geometry, so the results show the outcome of hybrid software and hardware ray tracing performance for varying degrees of hardware acceleration. Speedup results for Table 1 are normalized vs. the Quadro M6000 GPU hardware that was used for all of the VMD production renderings for *Birth of Planet Earth*. We note that the factor of 15× GPU hardware accelerated ray tracing performance gain relative to the Quadro M6000 result was achieved in just a three year time period. The Quadro GV100 performance results shown in both Tables are nearest in general purpose arithmetic and memory performance to the Quadro RTX 6000 GPUs with hardware ray tracing acceleration. These results show that hardware ray tracing acceleration provides roughly 5–8× the performance using VMD’s hybrid software and hardware accelerated ray tracing approach implemented in OptiX on first generation NVIDIA RTX GPUs.

Detailed tests of VMD’s OptiX hybrid ray tracing engine on the Quadro RTX 6000 GPU show that it can achieve ray casting rates as high as 9 billion rays/sec for scenes using the simplest lighting and shading model, and rates ranging from 1–3 billion rays/sec for full complexity scenes with depth cueing (fog), ambient occlusion lighting, and large numbers of transparent surfaces, and depth-of-field focal blur, all of which involve significantly more shading arithmetic, and generation of secondary rays for shadow tests and transmission rays. The VMD hybrid ray casting performance results are comparable to publicly reported results for commercial cinematic rendering tools tested on the same hardware. The molecular scenes rendered in VMD make much greater use of software-based ray-object intersections (for spheres, cylinders, etc) than would be typical in cinematic rendering. Conversely, VMD surface shading workloads are usually simpler and more streamlined relative to typical cinematic rendering cases that involve many types of texture mapping and complex shaders that emulate complex surfaces such as car paint containing metallic flakes.

Ongoing VMD development efforts aim to further optimize the software-based shading and primary ray generation components of the OptiX ray tracing engine to raise the average ray casting rates for the most challenging cases through targeted use of C++ template specialization in the CUDA-based OptiX shaders. We expect that the use of second generation GPU ray tracing acceleration hardware will increase the peak ray casting rates achieved for complex VMD scenes, as a result of improvements to concurrent scheduling of ray tracing and shading operations, and reductions in software overheads through new ray tracing APIs that operate at a lower level of abstraction and provide greater programmer control (OptiX 7, Vulkan Ray Tracing, and DirectX Ray Tracing). Although the fulldome renderings in this project used the custom-written OptiX-based ray tracing engine designed specifically for VMD on NVIDIA GPUs, multiple vendors are developing GPUs with varying degrees of ray tracing hardware acceleration and we expect those to become

interesting platforms for future work of this type. Similarly, high-level APIs for scientific visualization increasingly target ray tracing and path tracing techniques, easing the task of incorporating hardware optimized cinematic rendering features directly into domain specific scientific visualization applications such as VMD [36, 37].

5.3. Interactive Preview with VR HMDs

A key benefit of the performance gains achieved through hardware accelerated ray tracing is the ability to perform high resolution interactive ray tracing of large molecular scenes at frame rates sufficient to support not only traditional windowed viewing on desktop displays, but also panoramic, omnidirectional, and stereoscopic viewing, using immersive displays or commodity VR HMDs. Table 3 summarizes interactive ray tracing performance results for live stereoscopic viewing of omnidirectional projections with an Oculus Rift VR head mounted display (HMD). The results shown for the NVIDIA VCA cluster represent an array of 8× Quadro M6000 GPUs in a dense single-node configuration (the configuration providing highest frame rate and lowest latency), as previously reported [38]. The hybrid software and hardware-accelerated ray tracing results shown for the GeForce RTX 2080 demonstrate a significant performance gain relative to the best previously published VCA results, yielding much higher interactive frame rates. The performance levels achieved are sufficient to allow immersive flythroughs of large biomolecular scenes such as those in *Birth of Planet Earth*, enabling both high quality exploratory scientific visualizations and cinematic pre-visualization to be performed in VR HMDs. The combination of high performance interactive ray tracing with VR HMDs allows VMD to emulate the field of view and immersive viewing experience in a panoramic or fulldome theater such as a planetarium, providing authors with an accurate impression of the audience experience in the target venue.

5.4. Instancing of Molecular Geometry

As described in Sec. 4, one of the features that VMD previously lacked that motivated the use of Houdini for rendering of the bacterial cluster environment, was support for efficient rendering of up to thousands of *instances* of chromatophores, as shown in (Fig. 3A–B,D and Fig. 4).

VMD visualizations are composed from a collection of graphical representations, each built from selected atomic structure components, using a molecular structure drawing method and a coloring method. Each graphical representation drives underlying visualization algorithms to generate molecular scene geometry (geometric primitives, and associated vertices, normals, color, textures, and materials), that form the input to built-in VMD rendering engines, e.g., interactive OpenGL rasterization, interactive progressive refinement GPU ray tracing with OptiX, or export to files for rendering by external tools such as Houdini. The addition of instancing support in VMD required adaptation using two major schemes, one for streaming-oriented renderers such as OpenGL rasterization, and another for retained-mode renderers such as ray tracing engines and for export to files read by external renderers.

By design, streaming renderers have a minimal memory requirement, so instanced geometry must either be submitted to the renderer multiple times, or the instanced geometric elements

can be cached in moderate sized vertex buffers in on-board GPU memory, with a series of subsequent draw calls emitted for each instance. To facilitate GPU-resident caching of molecular geometry in moderate-sized vertex buffers, VMD associates instances with individual graphical representations. This choice naturally limits the size of individual geometry buffers, while balancing user convenience and the ability to create instances of very large biomolecular complexes or organelles such as the chromatophore.

Retained mode renderers that store all geometry for render-time random access do not benefit from any particular choice of geometric instancing granularity since they must make all of the scene geometry available at once, whether in host memory, or GPU-resident memory, e.g., for the built-in VMD OptiX ray tracing engine. Ray tracing engines make use of acceleration structures, e.g. a BVH, to eliminate a large fraction of their ray-geometry intersection test workload, so it is important to provide a molecular scene graph hierarchy that leads to efficient acceleration structure construction and render-time traversal. Typically, this is achieved by grouping spatially co-resident geometry together in a single acceleration structure, rather than building separate acceleration structures for individual geometry buffers. This strategy favors an approach that groups together the graphical representations associated with the same molecular complex and/or those that are involved in the same set of destination instances in the molecular scene, such that they share the same acceleration structure.

The addition of support for instancing of molecular geometry has enabled the use of VMD for more effective visualization and modeling of cell-scale structures. Further work in this direction will make it much easier to incorporate VMD into cinematic rendering pipelines for cell-scale visualization tasks. However, we note that artistic tools such as Houdini will remain key to accomplish visualization and storytelling tasks that require the application of human choreography and illustrative creativity over purely simulation and data-driven approaches provided by scientific tools such as VMD. Having improved VMD's performance and capabilities will however ease the incorporation of atomic structures and simulation trajectories into cinematic visualizations for broad audiences.

As an application of the aforementioned *instancing* capabilities in VMD, a structural model of a low-light grown *Rba. sphaeroides* cell containing 985 chromatophores was developed (Fig. 4). This model corresponds to a packing density of chromatophores observed in cryoEM tomography images [39], including a nonuniform radial distribution. The cellular environment thus depicted contains 2,431,965 bacteriochlorophylls for light-harvesting, distributed across approximately a hundred thousand proteins.

6. Discussion and Broader Impact

Full-dome shows typically reach millions of viewers over lifespans exceeding a decade. The visualization presented here as part of a full-dome show provides a workflow that can be employed for creating other accessible narratives that bring complex natural processes to broad audiences. A visualization that successfully engages its target audience requires: (i) choice of a problem of immediate and intuitive relevance to many—e.g., sunlight powering life on Earth—with abundant visualization-ready data on key processes; (ii) close

collaboration between computational/experimental and visualization scientists, neither of whom could produce such content alone; (iii) consultation with psychology researchers, writers, and actors, as well as audience tests, since the challenges of human perception are of equal importance to technical challenges. Accessible presentations of science need to obey the rigors of a performance art in addition to the rigors of research.

But why should a scientist undertake such an ordeal, when one could instead confine presentations to other experts? Arguably, scientists have a responsibility exceeding that of others for performing *science outreach*. Today, data visualization has become akin to empowering people to read [1]; numerical literacy has become the new literacy upon which democratic institutions depend; and the ability to visualize information has become a basis for empathy in a society where misinformation is rampant.

Specifically, the subject of solar energy harvesting is associated directly with the prosperity of human civilization: The total power demand of humanity (~ 15 TW) is a tiny fraction of the solar power incident on Earth ($\sim 120,000$ TW), representing an abundant clean energy solution in the long term [10]. Energy efficiency is shown to correlate with economic, environmental, and social benefits to a society [40], whereas biological systems provide examples of high energy efficiency surpassing comparable human technologies when considered at system-level [2]. Yet, public support for the development of such clean energy solutions remains elusive. It is the responsibility of scientists to make a case for our relevance to the public interest by making research accessible.

Acknowledgments

The visualization on bacterial photosynthesis presented here as well the entire fulldome show *Birth of Planet Earth* are dedicated to Klaus Schulten (1947–2016), whose lifelong passion for quantum biology was a source of inspiration. The authors acknowledge the collaborators and colleagues that contributed to this work: Thomas Lucas, Andrew Hitchcock, Michal Cartron, Pu Qian, Philip Jackson, Amanda Brindley, John Olsen, Peter Adams, Jen Hsin, Johan Strümpfer, Angela Barragan, Ivan Teo, James Phillips, Abhishek Singharoy, Mahmoud Moradi, Christophe Chipot, Lizanne DeStefano, Christine Shenouda, Andre Schleife, Emad Tajkhorshid. The research and visualization presented here was supported through 2007–2020 by NSF grants MCB1157615, MCB1616590, PHY0822613, OCI-0725070, ACI-1238993, ACI-1445176 (CADENS), NIH grants P41-GM104601 and 5R01GM098243-02, the CUDA Center of Excellence at the University of Illinois, and the Blue Waters sustained-petascale computing project supported by NSF awards OCI-0725070, ACI-1238993, and the state of Illinois, NSF PRAC awards OCI-0832673 and ACI-1440026, DOE grant DE-SC 0001035. C.N.H. acknowledges research grant BB/M000265/1 from the Biotechnology and Biological Sciences Research Council (UK) and an Advanced Award 338895 from the European Research Council.

References

- [1]. Borkiewicz K, Christensen AJ, Stone J, Communicating science through visualization in an age of alternative facts, SIGGRAPH ACM-10.1145/3134472.3134488.
- [2]. Hitchcock A, Hunter CN, Sener M, Determination of cell doubling times from the return-on-investment time of photosynthetic vesicles based on atomic detail structural models, *J. Phys. Chem. B* 121 (2017) 3787–3797. [PubMed: 28301162]
- [3]. Sener MK, Olsen JD, Hunter CN, Schulten K, Atomic level structural and functional model of a bacterial photosynthetic membrane vesicle, *Proc. Natl. Acad. Sci. USA* 104 (2007) 15723–15728. [PubMed: 17895378]
- [4]. Sener M, Strumpfer J, Timney JA, Freiberg A, Hunter CN, Schulten K, Photosynthetic vesicle architecture and constraints on efficient energy harvesting, *Biophys. J* 99 (2010) 67–75. [PubMed: 20655834]

- [5]. Sener M, Strümpfer J, Hsin J, Chandler D, Scheuring S, Hunter CN, Schulten K, Förster energy transfer theory as reflected in the structures of photosynthetic light harvesting systems, *Chem. Phys. Chem* 12 (2011) 518–531. [PubMed: 21344591]
- [6]. Cartron ML, Olsen JD, Sener M, Jackson PJ, Brindley AA, Qian P, Dickman MJ, Leggett GJ, Schulten K, Hunter CN, Integration of energy and electron transfer processes in the photosynthetic membrane of *Rhodobacter sphaeroides*, *Biochim. Biophys. Acta, Bioener* 1837 (2014) 1769–1780.
- [7]. Sener M, Strümpfer J, Singharoy A, Hunter CN, Schulten K, Overall energy conversion efficiency of a photosynthetic vesicle, *eLife* 10.7554/eLife.09541 , (30 pages). doi:10.7554/eLife.0954110.7554/eLife.09541, (30 pages). doi: 10.7554/eLife.09541 .
- [8]. Sener M, Stone JE, Barragan A, Singharoy A, Teo I, Vandivort KL, Isralewitz B, Liu B, Goh BC, Phillips JC, Kourkoutis LF, Hunter CN, Schulten K, Visualization of energy conversion processes in a light harvesting organelle at atomic detail, in: *Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis, SC '14, IEEE Press, 2014*, (4 pages).
- [9]. Stone JE, Sener M, Vandivort KL, Barragan A, Singharoy A, Teo I, Ribeiro JV, Isralewitz B, Liu B, Goh BC, Phillips JC, MacGregor-Chatwin C, Johnson MP, Kourkoutis LF, Hunter CN, Schulten K, Atomic detail visualization of photosynthetic membranes with GPU-accelerated ray tracing, *Parallel Computing* 55 (2016) 17–27. 10.1016/j.parco.2015.10.015. [PubMed: 27274603]
- [10]. Blankenship RE, Tiede DM, Barber J, Brudvig GW, Fleming G, Ghirardi M, Gunner MR, Junge W, Kramer DM, Melis A, Moore TA, Moser CC, Nocera DG, Nozik AJ, Ort DR, Parson WW, Prince RC, Sayre RT, Comparing photosynthetic and photovoltaic efficiencies and recognizing the potential for improvement, *Science* 332 (6031) (2011) 805–809. [PubMed: 21566184]
- [11]. Singharoy A, Chipot C, Moradi M, Schulten K, Chemomechanical coupling in hexameric protein-protein interfaces harnesses energy within V-type ATPases, *J. Am. Chem. Soc* 139 (2017) 293–310. [PubMed: 27936329]
- [12]. Junge W, Nelson N, ATP synthase, *Ann. Rev. Biochem* 84 (2015) 631–657. 10.1146/annurev-biochem-060614-034124. [PubMed: 25839341]
- [13]. Niederman RA, Membrane development in purple photosynthetic bacteria in response to alterations in light intensity and oxygen tension, *Photosyn. Res* 116 (2013) 333–48.
- [14]. Singharoy A, Maffeo C, Delgado-Magnero KH, Swainsbury DJK, Sener M, Kleinekathofer U, Vant JW, Nguyen J, Hitchcock A, Isralewitz B, Teo I, Chandler DE, Stone JE, Phillips JC, Pogorelov TV, Mallus MI, Chipot C, Luthey-Schulten Z, Tieleman DP, Hunter CN, Tajkhorshid E, Aksimentiev A, Schulten K, First principles to phenotypes: Design principles of cellular energy metabolism, *Cell* 179 (2019) 1098–1111. [PubMed: 31730852]
- [15]. Xiong J, Fischer WM, Inoue K, Nakahara M, Bauer CE, Molecular evidence for the early evolution of photosynthesis, *Science* 289 (2000) 1724–1730. [PubMed: 10976061]
- [16]. MacGregor-Chatwin C, Sener M, Barnett SF, Hitchcock A, Barnhart-Dailey MC, Maghlaoui K, Barber J, Timlin JA, Schulten K, Hunter CN, Lateral segregation of Photosystem I in cyanobacterial thylakoid, *Plant Cell* 29 (2017) 1119–1136. 10.1105/tpc.17.00071. [PubMed: 28364021]
- [17]. MacGregor-Chatwin C, Jackson PJ, Sener M, Chidgey JW, Hitchcock A, Qian P, Mayneord GE, Johnson MP, Luthey-Schulten Z, Dickman MJ, Scanlan DJ, Hunter CN, Membrane organization of photosystem I complexes in the most abundant phototroph on Earth, *Nature Plants* 5 (2019) 879–889. [PubMed: 31332310]
- [18]. Humphrey W, Dalke A, Schulten K, VMD – Visual Molecular Dynamics, *J. Mol. Graphics* 14 (1) (1996) 33–38. 10.1016/0263-7855(96)00018-5.
- [19]. Stone JE, Ch. 4 - A planetarium dome master camera, in: Haines E, Akenine-Möller T (Eds.), *Ray Tracing Gems*, Apress, 2019, pp. 49–60.
- [20]. Borkiewicz K, Naiman JP, Lai H, Cinematic visualization of multiresolution data: Ytini for adaptive mesh refinement in houdini, *Astronomical J* 158 (2019) 10. 10.3847/1538-3881/ab1f6f.
- [21]. Adams PG, Hunter CN, Adaptation of intracytoplasmic membranes to altered light intensity in *Rhodobacter sphaeroides*, *Biochim. Biophys. Acta, Bioener* 1817 (2012) 1616–1627.

- [22]. Şener M, Schulten K, A general random matrix approach to account for the effect of static disorder on the spectral properties of light harvesting systems, *Phys. Rev. E* 65 (2002) 031916, (12 pages).
- [23]. Crofts AR, Proton-coupled electron transfer at the Q_D -site of the bc_1 complex controls the rate of ubiquinone oxidation, *Biochim. Biophys. Acta* 1655 (2004) 77–92. [PubMed: 15100020]
- [24]. Borkiewicz K, Christensen AJ, Levy S, Patterson R, Cox D, Carpenter J, Scientific and visual effects software integration for the visualization of a chromatophore, SIGGRAPH ACM-10.1145/3283289.3283324.
- [25]. Stone JE, Vandivort KL, Schulten K, Immersive out-of-core visualization of large-size and long-timescale molecular dynamics trajectories, *Lect. Notes in Comp. Sci* 6939 (2011) 1–12.
- [26]. Stone JE, Isralewitz B, Schulten K, Early experiences scaling VMD molecular visualization and analysis jobs on Blue Waters, in: *Extreme Scaling Workshop (XSW), 2013, 2013*, pp. 43–50. 10.1109/XSW.2013.10.
- [27]. Krone M, Stone JE, Ertl T, Schulten K, Fast visualization of Gaussian density surfaces for molecular dynamics and particle system trajectories, in: *EuroVis - Short Papers 2012, 2012*, pp. 67–71.
- [28]. Stone JE, Vandivort KL, Schulten K, GPU-accelerated molecular visualization on petascale supercomputing platforms, in: *Proceedings of the 8th International Workshop on Ultrascale Visualization, UltraVis '13, ACM, New York, NY, USA, 2013*, pp. 6:1–6:8.
- [29]. Stone JE, Ch. 27 - Interactive ray tracing techniques for high-fidelity scientific visualization, in: Haines E, Akenine-Möller T (Eds.), *Ray Tracing Gems*, Apress, 2019, pp. 493–515.
- [30]. Lascara CM, Wheless GH, Cox D, Patterson R, Levy S, Johnson A, Leigh J, Kapoor A, Tele-immersive virtual environments for collaborative knowledge discovery, in: *Proceedings of the Advanced Simulation Technologies Conference, 1999*, pp. 416–423.
- [31]. Phillips JC, Braun R, Wang W, Gumbart J, Tajkhorshid E, Villa E, Chipot C, Skeel RD, Kale L, Schulten K, Scalable molecular dynamics with NAMD, *J. Comp. Chem* 26 (2005) 1781–1802. [PubMed: 16222654]
- [32]. Stone JE, Phillips JC, Freddolino PL, Hardy DJ, Trabuco LG, Schulten K, Accelerating molecular modeling applications with graphics processors, *J. Comp. Chem* 28 (2007) 2618–2640. 10.1002/jcc.20829. [PubMed: 17894371]
- [33]. Phillips JC, Stone JE, Schulten K, Adapting a message-driven parallel application to GPU-accelerated clusters, in: *SC '08: Proceedings of the 2008 ACM/IEEE Conference on Supercomputing, IEEE Press, Piscataway, NJ, USA, 2008*, pp. 1–9, (9 pages).
- [34]. Pan AC, Sezer D, Roux B, Finding transition pathways using the string method with swarm of trajectories, *J. Phys. Chem. B* 112 (11) (2008) 3432–3440. [PubMed: 18290641]
- [35]. Phillips JC, Hardy DJ, Maia JDC, Stone JE, Ribeiro JV, Bernardi RC, Buch R, Fiorin G, Hémin J, Jiang W, McGreevy R, Melo MCR, Radak B, Skeel RD, Singharoy A, Wang Y, Roux B, Aksimentiev A, Luthey-Schulten Z, Kalé LV, Schulten K, Chipot C, Tajkhorshid E, Scalable molecular dynamics on CPU and GPU architectures with NAMD, *J. Chem. Phys* 153 (2020) 044130. 10.1063/5.0014475. [PubMed: 32752662]
- [36]. Wald I, Woop S, Benthin C, Johnson GS, Ernst M, Embree: A kernel framework for efficient CPU ray tracing, *ACM Trans. Graph* 33 (4) (2014) 143:1–143:8.
- [37]. Wald I, Johnson G, Amstutz J, Brownlee C, Knoll A, Jeffers J, Gunther J, Navratil P, OSPRay – a CPU ray tracing framework for scientific visualization, *IEEE Transactions on Visualization and Computer Graphics* 23 (1) (2017) 1–1. 10.1109/TVCG.2016.2599041. [PubMed: 27514046]
- [38]. Stone JE, Sherman WR, Schulten K, Immersive molecular visualization with omnidirectional stereoscopic ray tracing and remote rendering, 2016 IEEE International Parallel and Distributed Processing Symposium Workshop (IPDPSW) (2016) 1048–1057doi:10.1109/IPDPSW.2016.121.
- [39]. Noble JM, Lubieniecki J, Savitzky BH, Plitzko J, Engelhardt H, Baumeister W, Kourkoutis LF, Connectivity of centermost chromatophores in *Rhodobacter sphaeroides* bacteria, *Mol. Microbiol* 109 (2018) 812–825. [PubMed: 29995992]
- [40]. Alcorta L, Bazilian M, Simone GD, Pedersen A, Return on investment from industrial energy efficiency: evidence from developing countries, *Energ. Effic* 7 (2014) 43–53.

Highlights

- Development of an accessible visual narrative describing the process of photosynthesis for broad audiences, a cinematic full-dome planetarium show “Birth of Planet Earth”.
- Description of the complete full-dome cinematic rendering pipeline used for “Birth of Planet Earth”, incorporating data from multiscale simulation and modeling on HPC platforms, domain-specific science tools, tools for pre-visualization, camera choreography, artistic modeling tools, rendering on HPC platforms, and tools and techniques used in post production.
- A new VMD ray tracing engine implementing hybrid hardware-accelerated and software based cinematic quality rendering for cell-scale scenes at performance levels up to 8x higher than previously possible, while maintaining high-fidelity rendering of curved biomolecular geometries.
- A new atomic detail structural model of an entire photosynthetic bacterium, featuring 985 light harvesting organelles (chromatophores), containing 2,431,965 chlorophylls for light harvesting, distributed across approximately a hundred thousand proteins. The immense scene complexity of this model required significant feature advances in VMD.
- Supercomputer simulations using NAMD employed for describing the crucial biomolecular mechanisms of energy conversion from light capture to ATP synthesis.

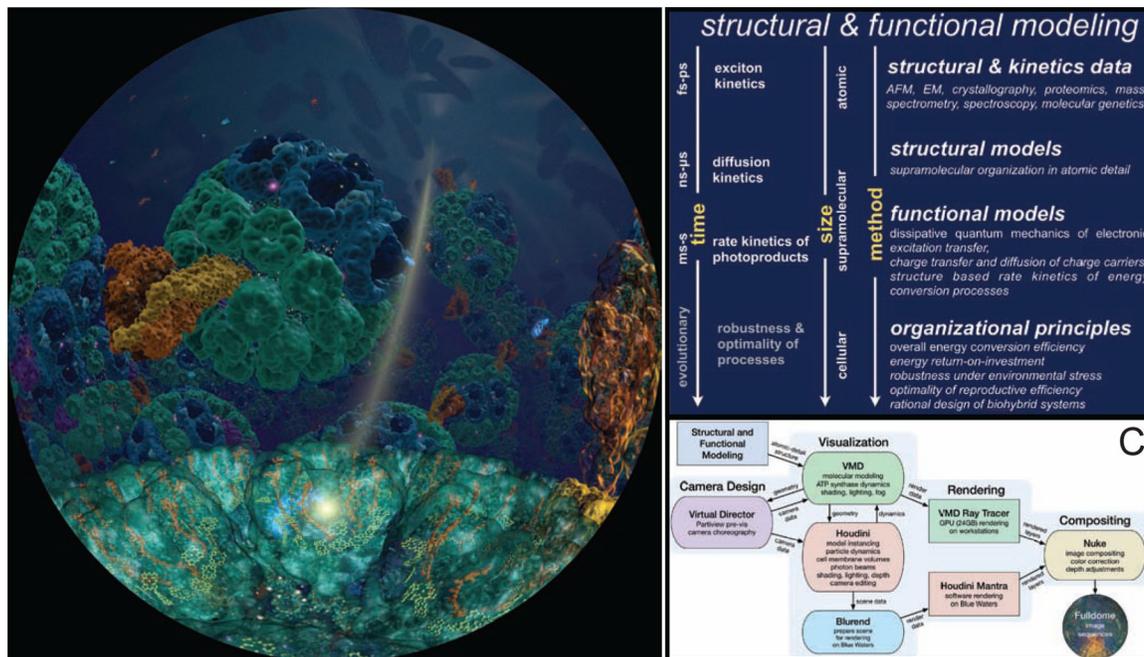


Figure 1: High performance computing as an exploratory and explanatory tool for fundamental energy conversion processes powering life on Earth, from multi-scale rate kinetics models to accessible narratives for lay audiences. (A) Frame from the full-dome show *Birth of Planet Earth*, depicting the capture and transformation of solar energy by an array of proteins inside a biological cell. (B) Structural and functional modeling of photosynthetic energy conversion [3, 4, 6, 7, 2], combining experimental and computational approaches that correspond to time and size scales spanning electronic to cell levels. (C) From computational models to visual narratives [8, 9] using a chain of novel software solutions. Labelled arrows identify the flow and type of data exchanged between applications from scientific computation to full-dome visualization.

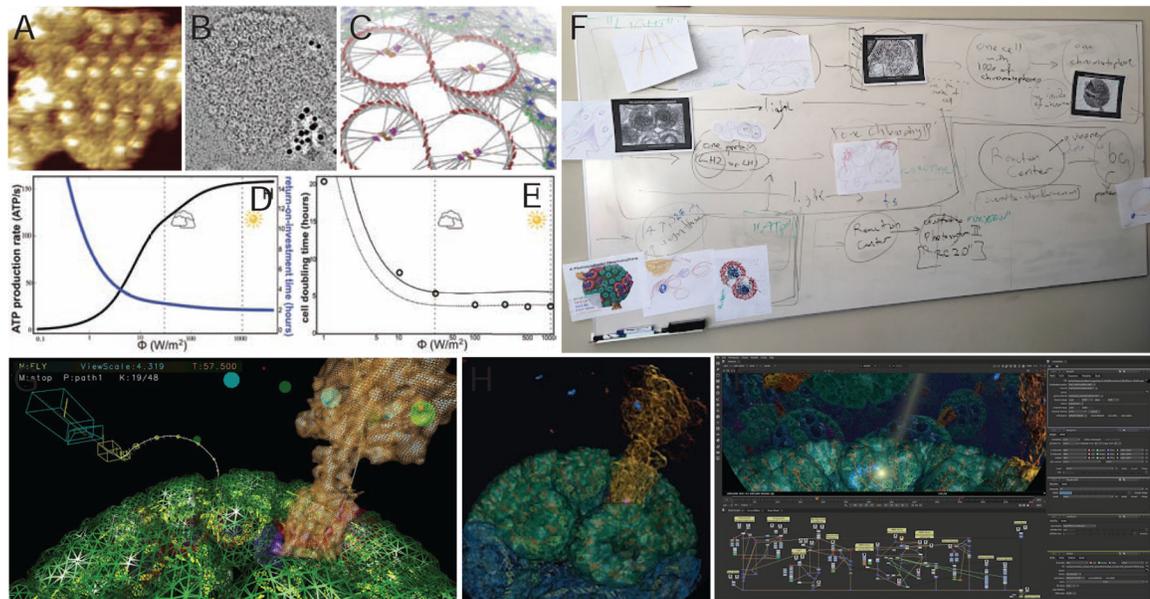


Figure 2:

Workflow from raw data to fulldome visualization represents over a decade of collaborative work. Structural information such as from atomic force microscopy (A) and negative stain electron microscopy (B) is used for the construction of atomic detail structural models that form the basis of quantum mechanical computations for electronic couplings between bacteriochlorophylls (C) [3, 6]. Multi-scale modeling of energy conversion processes [7] culminates in the determination of system-level device performance metrics for the chromatophore (D), such as ATP production rate and return-on-investment time for energy as a function of incident light [2]. Integrated at cell scale, energy conversion efficiency determines the reproductive efficiency of the organism given in terms of cell doubling time (E) [2]. (F) This electronic-to-cell scale description of energy conversion steps translates to a storyboard, which is the basis of an accessible narrative for lay audiences (see Fig. 3). The storyboard is converted to an interactive pre-visualization in Virtual Director, shown with a camera path in (G), and a corresponding ray traced fulldome rendering in VMD [18, 9, 19] (H). Final images are composited using Foundry Nuke software (I), with nodes and connections representing the various Houdini and VMD image sequences being read in, processed, and merged to construct the final image sequence for the visualization, as shown for the frame in (Fig. 1A).

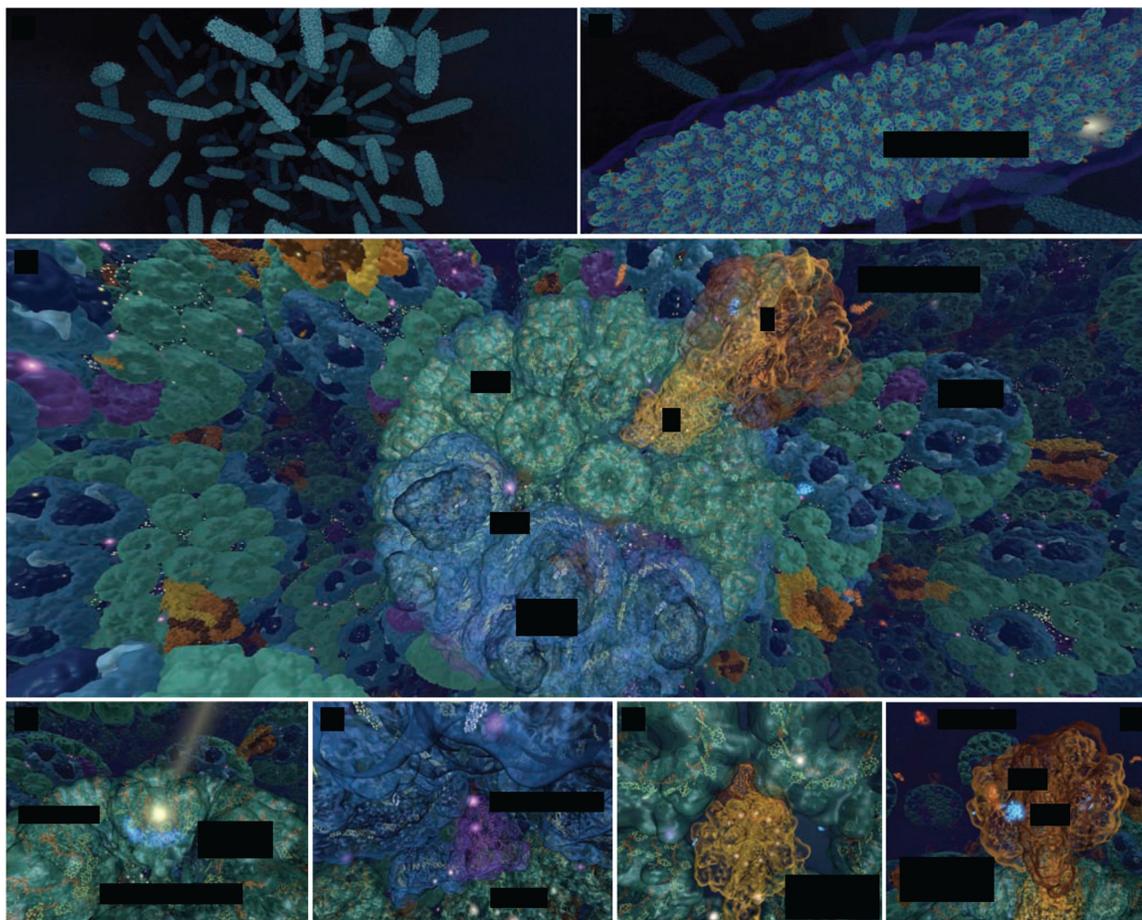


Figure 3:

An accessible visual narrative for the primary energy source of life. The storyboard (Fig. 2F) for energy conversion steps [7, 2] forms the basis of the explanatory visualization spanning cell to electronic length scales. In order to ensure a balance between accuracy and accessibility, some of the scenes shown are based on atomic detail structural modeling and simulation, while others are based on art, animation, and artifice, as detailed below. (A) Approach to a cluster of bacterial cells in a pond. The placement of cells within the cluster and the placement of chromatophores within each cell (B) are not based on data. The overall cell shape reflects *Rba. sphaeroides* electron microscopy images [21]; the dense population of chromatophores corresponds to the low light limit of growth conditions [2]. (C) Approach to the ‘hero’ chromatophore. The atomic structural model of the chromatophore is determined based on atomic force microscopy, cryo-electron microscopy, crystallography, optical spectroscopy, mass spectroscopy, and proteomics data [3, 4, 6, 7]. The sequence of protein complexes responsible for the primary energy conversion steps are visible during approach: light harvesting complex 2 (LH2, green) → light harvesting complex 1 (LH1, light blue) → reaction center (dark blue) → cytochrome bc_1 complex (cyt bc_1 , purple) → ATP synthase (gold/brown) (comprising F_o and F_1 units; see (F) and (G)). This sequence also corresponds to the camera path in the scenes that follow as light energy is converted into chemical energy. (D) Capture and conversion of photon energy into pigment electronic excitations. The network of chromatophore pigments,

namely, bacteriochlorophylls (represented as porphyrin rings only) and carotenoids, is visible inside protein complexes. Quantum mechanical elements, such as a photon or an electronic excitation depicted in this scene, cannot be represented in a literal ‘visual’ sense for human perception. Furthermore, accurate representations of simulations, such as the time evolution of excitonic states [5] do not provide accessible narratives for non-experts. Therefore, intuitive illustrations are instead employed for quantum mechanical elements depicted: (i) incoming photon illustrated as a column of light; (ii) resulting electronic excitation shown as a delocalized halo (delocalization length of five pigments approximately corresponds to room temperature thermal disorder of LH2 excitations [22]); (iii) transfer of excitations between pigment clusters shown as a semi-directed flow to the reaction center, suppressing the degree of stochasticity of the excitation transfer process [3, 4, 5]. (E) The electronic excitation arriving at the reaction center causes the release of a quinol (purple), which diffuses across the membrane to cyt bc_1 , releases protons (yellow), and shuttles back to the reaction center as a quinone (violet). In order to reduce the scene-complexity, quinols/quinones are depicted as glowing point objects only, their trajectory constrained to membrane areas between proteins, but not based on simulation. Quinol/quinone and proton release animations are coordinated with each other, but do not reflect the specific stoichiometry of the Q-cycle at cyt bc_1 [23]. (F) Protons follow a semi-directed flow toward the ATP synthase c-ring (F_o unit) animated as a ‘turnstile’ operated by proton passage. As the camera climbs through the ATP synthase as a tower, the ‘mushroom’ (F_1 unit) of the ATP synthase (G) is animated based on a molecular dynamics trajectory [11]. (See (C) for the relative position of F_o and F_1 units.) The animations of both portions of the ATP synthase as well as the docking and release of ADP (orange), phosphate (red), and ATP (glowing blue) are synchronized with one another. Animation speeds throughout are deliberately not commensurate with physical time scales of the processes shown, which range from picoseconds (electronic excitation kinetics) to milliseconds (ATP synthesis). ATP release represents the finale of this visualization sequence with the proliferation of life thus powered by sunlight as a vital step in the *Birth of Planet Earth*.

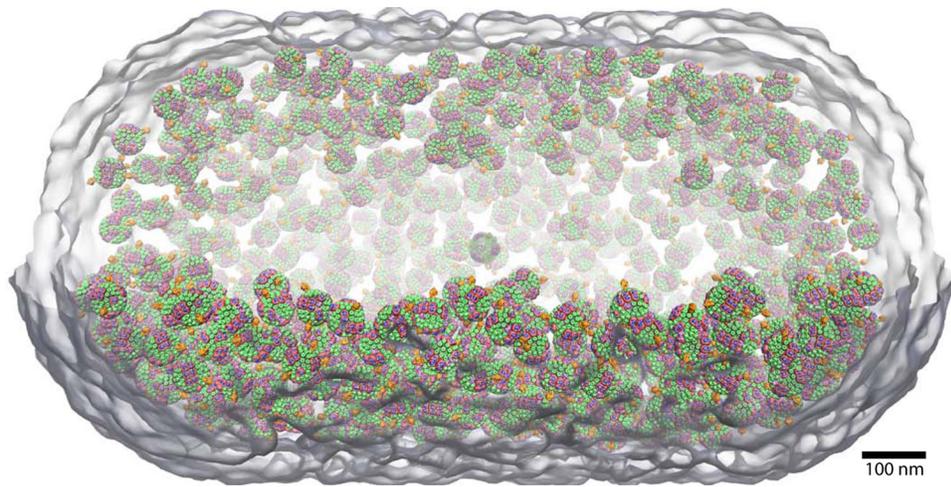


Figure 4: Structural model of a low light adapted *Rba. sphaeroides* cell featuring 985 chromatophores. The cell size ($1.6 \mu\text{m}$) and the radial distribution of chromatophores correspond to cryoEM tomography data [39] and are consistent with the light harvesting capacity of a bacterium grown at 10% of full sunlight intensity [2]. (Non-chromatophore inclusion bodies in the cell are not shown—specifically, DNA and ribosomes are absent.) This packing pattern of chromatophores represents a structural model that was not available at the time of production for the fulldome show *Birth of Planet Earth*, thus improving upon the rudimentary spatial arrangements depicted therein (Figs. 1A, 3B, 3C).

Table 1:

VMD ray tracing performance comparison for a challenging 4096×4096 resolution, cinematic-quality offline rendering of the chromatophore structure with transparent surfaces, depth cueing, shadows, ambient occlusion, and direct lighting. Production VMD renderings for *Birth of Planet Earth* used software ray tracing on a large number of Quadro M6000 GPUs. VMD ray tracing performance achieved with a hardware-accelerated Quadro RTX 6000 GPU is up to 15× that of the GPUs previously used for production rendering.

GPU Hardware and Ray Tracing Method	Chromatophore Ambient Occlusion (Runtime, Speedup)
Quadro M6000 Software Ray Tracing	276s, 1.00×
Quadro GV100 Software Ray Tracing	158s, 1.75×
2× Quadro GV100 Software Ray Tracing	80s, 3.45×
Quadro RTX 6000 Hybrid SW+HW Ray Tracing	18.1s, 15.2×

Table 2:

VMD ray tracing performance comparison for a 7680×4320 resolution, cinematic-quality offline rendering of the chromatophore structure dominated by spherical and cylindrical geometric primitives used to represent the lipid membrane. This case demonstrates a less optimal scene for current generation hardware-accelerated ray tracing, since triangle meshes represent a minority of the scene geometry. Hardware-accelerated BVH traversal still provides a significant performance benefit, yielding an overall speedup over 5× faster than software ray tracing on previous-generation Quadro GV100 GPU hardware with fairly comparable general purpose arithmetic throughput.

GPU Hardware and Ray Tracing Method	Chrom. Memb. Lipids AO, DoF (512 samples) (Runtime, Speedup)
Quadro GV100 Software Ray Tracing	215.8s, 1.00×
2× Quadro GV100 Software Ray Tracing	112.9s, 1.91×
Quadro RTX 6000 Hybrid SW+HW Ray Tracing	40.8s, 5.29×

Table 3:

VMD omnidirectional stereoscopic ray tracing performance for 3072×1536 omnidirectional equirectangular projection. Performance comparison between software ray tracing on a remote VCA GPU cluster [38], and hardware-accelerated ray tracing on a local GeForce RTX 2080 GPU. Progressive refinement ray tracing results update an accumulation buffer as each so-called subframe is rendered, allowing stochastically sampled ray tracing results to be displayed during convergence for greater interactivity. The HMD view is created by texturing the omnidirectional projection onto a spherical projection surface and reprojecting the view for the current HMD head pose using OpenGL rasterization. Oculus Rift HMD head pose and display updates were maintained at the native HMD display refresh rate using a free-running OpenGL rasterization thread with buffer swap vertical retrace synchronization linked to the HMD. Ray tracing workload increases as $(pixel\ count \times light\ count \times sample\ count \times subframe\ count)$.

HIV-1 Scene Lighting Mode and Hardware	Per-subframe Samples AA : AO (AO Per Hit)	RT Update Rate (FPS)
Shadows	1 : 0	15.6
Remote VCA Cluster	2 : 0	12.1
Software Ray Tracing [38]	4 : 0	7.0
Shadows, Ambient Occlusion	1 : 1	13.1
Remote VCA Cluster	1 : 2	9.9
Software Ray Tracing [38]	1 : 4	6.3
Shadows	1 : 0	75.0
GeForce RTX 2080	2 : 0	74.2
Hybrid SW+HW Ray Tracing	4 : 0	37.5
Shadows, Ambient Occlusion	1 : 1	72.0
GeForce RTX 2080	1 : 2	61.0
Hybrid SW+HW Ray Tracing	1 : 4	37.5