

# State of In Situ Visualization in Simulations: We are fast. But are we inspiring?

Axel Huebl  
Arianna Formenti  
Marco Garten  
Jean-Luc Vay  
{axelhuebl,jlvay}@lbl.gov  
Lawrence Berkeley National Laboratory  
Berkeley, California, USA




Figure 1: WarpX in situ visualization with Ascent and VTK-m. The figure shows a laser-wakefield accelerator stage (left), driven by a laser pulse (right) that travels to the bottom right. Iso-contours show the longitudinal electric field used in particle acceleration. The scenario is modeled and visualized in a Lorentz-boosted reference frame [15–17]. If this was not *in situ*, the authors would be able to transform to a different reference frame, add multiple light sources, cast hard and soft shadows, select some iso-contours for semi-transparent representation, and would add more triangles to smooth the generated iso-contours.

## ABSTRACT

Visualization of dynamic processes in scientific high-performance computing is an immensely data intensive endeavor. Application codes have recently demonstrated scaling to full-size Exascale machines, and generating high-quality data for visualization is consequently on the machine-scale, easily spanning 100s of TBytes of input to generate a single video frame. In situ visualization, the technique to consume the many-node decomposed data in-memory, as exposed by applications, is the dominant workflow. Although in situ visualization has achieved tremendous progress in the last

decade, scaling to system-size together with the application codes that produce its data, there is one important question that we cannot skip: is what we produce insightful and inspiring?

## CCS CONCEPTS

• **Human-centered computing** → **Visualization techniques**; *Visualization toolkits*; • **Computing methodologies** → **Rendering**; • **Applied computing** → *Physics*; • **Software and its engineering** → Software libraries and repositories.

## KEYWORDS

in situ visualization, high-performance computing, particle-in-cell, reflections, directions, lightning presentation submissions

## ACM Reference Format:

Axel Huebl, Arianna Formenti, Marco Garten, and Jean-Luc Vay. 2023. State of In Situ Visualization in Simulations: We are fast. But are we inspiring?. In *Proceedings of In Situ Infrastructures for Enabling Extreme-scale Analysis and Visualization (ISAV '23)*. ACM, New York, NY, USA, 3 pages. <https://doi.org/XXXXXXXX.XXXXXXX>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

ISAV '23, November 13, 2023, Denver, CO

© 2023 Association for Computing Machinery.  
ACM ISBN 978-1-4503-XXXX-X/18/06...\$15.00  
<https://doi.org/XXXXXXXX.XXXXXXX>

## 1 INTRODUCTION

In situ visualization is a tremendously powerful workflow to generate insight into the largest simulations run today. Recently, the 2022 Gordon Bell Prize-winning application WarpX [2] was used to run in situ visualization on 552 nodes of the Frontier supercomputer [9].

Immediate visualization of simulation dynamics at scale, from various camera angles, is powerful and helpful, providing answers to domain-science questions such as: Is a simulation evolving as planned? Are numerical options and resolution sufficiently set? Are any hardware or software issues/bugs appearing at scale? Yet, the scientifically most important question is: Does the visualization develop insight?

Gaining scientific insight from simulations is a complex and iterative process, with domain scientists connecting existing theory, empirical evidence and data from experiments and simulations. Visualizations can produce qualitative and quantitative representations of the dynamics at play. These representations can solidify understanding, guide the theoretical model building, help testing approximations and assumptions. An attractive visualization does help to communicate results and might inspire new scientific ideas.

Particularly for the latter part, domain scientists and audiences will compare the quality of their visualization with the state-of-the-art seen in everyday life: movies, games, advertising, etc. That is a high bar, given photo-realistic capabilities in these industries at high frame rates. Based on these expectations, can we produce in situ visualizations of scientific data that can be awe-inspiring and stimulate our minds? And - how much costs and/or scalability penalty are we willing to trade for this in high-performance computing?

## 2 SCALABLE METHODS WANTED

Many algorithms offered in contemporary visualization frameworks [5, 6, 10, 12] are able to exploit some locality, e.g., by domain decomposing ray traces and iso-contour searches, composing results later on [11]. Yet, advanced visualization techniques for casting shadows, tracing reflections, sorting collisions with objects, etc. are notoriously non-local and are thus challenging for multi-GPU implementations. Even volume-rendering more than one spatially overlapping source is non-trivial to do *in situ*, since established methods depend on a sampling technique that is hard to scale [8]. Additionally, many visualization techniques that scientists can use on single-node implementations would be highly desirable as distributed implementations for in situ frameworks: Taking Figure 1 as an example, if this was not *in situ* generated, the authors would add multiple light sources, cast hard and soft shadows, select some iso-contours for semi-transparent representation, and would smooth the generated iso-contours, by adding additional triangles that interpolate beyond the original resolution of the data source.

Consequently, there is a continued need for new, innovative, scalable in situ visualization methods. Both fast, low-overhead and higher-overhead (yet scalable), high-quality methods are needed. With respect to scalability, maybe there are tricks one can lend from other communities to generate artificial locality: occlude far-from-focus parts with mist as in gaming, simplify shadow masks and reflections, or aggressively exploit the adaptive resolution of

mesh-refined data sources. Additionally, successful in situ implementations and workflows can likely be enhanced and benefit from evolution through standardization of APIs, vendor abstractions, render scene control and data descriptions, e.g., [3, 13, 14].

## 3 SELECTED IN SITU VISUALIZATION NEEDS

Adding to the challenges of addressing expectations set from offline rendering for in situ visualization, we surveyed the Beam, Plasma & Accelerator Simulation Toolkit (BLAST) [1, 2, 4, 7] codes and identified three selected needs specific to in situ visualization.

First, we noticed that domain scientists have to relearn how to express rendering scene descriptions for each in situ tool. Standardization is needed [13]. Another approach might be domain-specific options in the simulation input language, automating the creation of visualization-configuration templates with mostly defaulted options - ready to be configured further for details by the inclined scientists when needed.

Second, video generation of iso-contours, glyphs (e.g., vectors placed in space), etc. often create “flicker” effects for surfaces and pointing of objects, simply based on the roughness of simulation data and steps selected for visualization. Research into transitions (or animations) between key/data frames with low memory overhead for HPC could be beneficial to reduce such effects.

Third, we also identified a commonly used algorithmic and simulation pattern for which in situ visualization would be ideally suited, but are not aware of any implemented solution yet: rendering of spatially-sliced data pipelines. In a large class of modeling codes, efficient solutions can be calculated by splitting the 3D domain over one axis. Instead of advancing the whole domain by an update, algorithms update a slice of the domain, e.g., from the back to the front of the 3D domain, and parallelize for the third spatial axis in *time*. Without spatially sliced rendering tools, a large number of algorithms and codes currently need to fall back to costly data output to “reconstruct” the spatial data domain that is required at once in offline visualization. Examples in laser-plasma and accelerator physics are the boosted frame technique [15–17] as shown in figure 1 (a more meaningful representation would transform slice-wise to the laboratory frame), the quasi-static method [1], or representations in reference trajectory space instead of time and space [4].

We believe addressing these challenges is timely and resulting in situ visualization will provide insight and inspiration for scientists.

## ACKNOWLEDGMENTS

This research was supported by the Exascale Computing Project (17-SC-20-SC), a collaborative effort of the U.S. Department of Energy Office of Science and the National Nuclear Security Administration. This research was supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research and Office of High Energy Physics, Scientific Discovery through Advanced Computing (SciDAC) program. This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.

## REFERENCES

- [1] Severin Diederichs, Carlo Benedetti, Axel Huebl, Rremi Lehe, Andrew Myers, Alexander Sinn, Jean-Luc Vay, Weiqun Zhang, and Maxence Thévenet. 2022. HiPACE++: A portable, 3D quasi-static particle-in-cell code. *Computer Physics Communications* 278 (2022), 108421. <https://doi.org/10.1016/j.cpc.2022.108421>
- [2] Luca Fedeli, Axel Huebl, France Boillod-Cerneux, Thomas Clark, Kevin Gott, Conrad Hillairet, Stephan Jaure, Adrien Leblanc, Rémi Lehe, Andrew Myers, Christelle Piechurski, Mitsuhisa Sato, Neil Zaim, Weiqun Zhang, Jean-Luc Vay, and Henri Vincenti. 2022. Pushing the Frontier in the Design of Laser-Based Electron Accelerators with Groundbreaking Mesh-Refined Particle-In-Cell Simulations on Exascale-Class Supercomputers. In *SC22: International Conference for High Performance Computing, Networking, Storage and Analysis*. IEEE, Dallas (TX), USA, 1–12. <https://doi.org/10.1109/SC41404.2022.00008>
- [3] Cyrus Harrison, Matthew Larsen, Brian S. Ryujin, Adam Kunen, Arlie Capps, and Justin Privitera. 2022. Conduit: A Successful Strategy for Describing and Sharing Data In Situ. In *2022 IEEE/ACM International Workshop on In Situ Infrastructures for Enabling Extreme-Scale Analysis and Visualization (ISAV)*. IEEE, Dallas (TX), USA, 1–6. <https://doi.org/10.1109/ISAV56555.2022.00006>
- [4] Axel Huebl, Remi Lehe, Chad Mitchell, Ji Qiang, Robert Ryne, Ryan Sandberg, and Jean-Luc Vay. 2022. Next Generation Computational Tools for the Modeling and Design of Particle Accelerators at Exascale. *Proceedings of the 5th North American Particle Accelerator Conference NAPAC2022 (2022)*, USA. <https://doi.org/10.18429/JACOW-NAPAC2022-TUYE2>
- [5] Matthew Larsen, James Ahrens, Utkarsh Ayachit, Eric Brugger, Hank Childs, Berk Geveci, and Cyrus Harrison. 2017. The ALPINE In Situ Infrastructure: Ascending from the Ashes of Strawman. In *Proceedings of the In Situ Infrastructures on Enabling Extreme-Scale Analysis and Visualization (Denver, CO, USA) (ISAV'17)*. Association for Computing Machinery, New York, NY, USA, 42–46. <https://doi.org/10.1145/3144769.3144778>
- [6] Matthew Larsen, Eric Brugger, Hank Childs, and Cyrus Harrison. 2022. Ascent: A Flyweight In Situ Library for Exascale Simulations. In *In Situ Visualization For Computational Science*. Mathematics and Visualization book series from Springer Publishing, Cham, Switzerland, 255 – 279.
- [7] LBNL, LLNL, SLAC, CEA-LIDYL, DESY, U Hamburg, CERN, TAE, et al. 2023. Beam, Plasma & Accelerator Simulation Toolkit (BLAST). <https://blast.lbl.gov>
- [8] Nicole Marsaglia. 2023. personal communication.
- [9] Nicole Marsaglia, Cyrus Harrison, Matthew Larsen, Axel Huebl, and Jean-Luc Vay. 2023. *KPP-3 Artifact for WarpX ↔ ALPINE: Ascent + VTK-m*. DOE Exascale Computing Project. <https://doi.org/10.5281/zenodo.8226853>
- [10] Felix Meyer, Benjamin Hernandez, Richard Pausch, René Widera, David Groß, Sergei Bastrakov, Axel Huebl, Guido Juckeland, Jeffrey Kelling, Matt Leinhauser, David Rogers, Ulrich Schramm, Klaus Steiniger, Stefan Gumhold, Jeff Young, Michael Bussmann, Sunita Chandrasekaran, and Alexander Debus. 2023. Hardware-Agnostic Interactive Exascale In Situ Visualization of Particle-In-Cell Simulations. In *Proceedings of the Platform for Advanced Scientific Computing Conference (Davos, Switzerland) (PASC '23)*. Association for Computing Machinery, New York, NY, USA, Article 9, 14 pages. <https://doi.org/10.1145/3592979.3593408>
- [11] Kenneth Moreland. 2009. IceT Users' Guide and Reference. Sandia Report SAND2009-3170. <https://www.osti.gov/servlets/purl/970256>
- [12] Kenneth Moreland, Christopher Sewell, William Usher, Li-ta Lo, Jeremy Meredith, David Pugmire, James Kress, Hendrik Schroots, Kwan-Liu Ma, Hank Childs, Matthew Larsen, Chun-Ming Chen, Robert Maynard, and Berk Geveci. 2016. VTK-m: Accelerating the Visualization Toolkit for Massively Threaded Architectures. *IEEE Computer Graphics and Applications* 36, 3 (2016), 48–58. <https://doi.org/10.1109/MCG.2016.48>
- [13] Pixar. 2016. Universal Scene Description. <https://openusd.org>
- [14] John E. Stone, Kevin S. Griffin, Jefferson Amstutz, David E. DeMarle, William R. Sherman, and Johannes Günther. 2022. ANARI: A 3-D Rendering API Standard. *Computing in Science & Engineering* 24, 2 (2022), 7–18. <https://doi.org/10.1109/MCSE.2022.3163151>
- [15] Jean-Luc Vay. 2007. Noninvariance of Space- and Time-Scale Ranges under a Lorentz Transformation and the Implications for the Study of Relativistic Interactions. *Phys. Rev. Lett.* 98 (Mar 2007), 130405. Issue 13. <https://doi.org/10.1103/PhysRevLett.98.130405>
- [16] Jean-Luc Vay, Cameron G.R. Geddes, Estelle Cormier-Michel, and David P. Grote. 2011. Numerical methods for instability mitigation in the modeling of laser wakefield accelerators in a Lorentz-boosted frame. *J. Comput. Phys.* 230, 15 (2011), 5908–5929. <https://doi.org/10.1016/j.jcp.2011.04.003>
- [17] Jean-Luc Vay, Cameron G. R. Geddes, Estelle Cormier-Michel, and David P. Grote. 2011. Effects of hyperbolic rotation in Minkowski space on the modeling of plasma accelerators in a Lorentz boosted frame. *Physics of Plasmas* 18, 3 (03 2011), 030701. <https://doi.org/10.1063/1.3559483>

Received 11 August 2023; revised 30 September 2023