4th International Workshop on In Situ Visualization (WOIV'19)

Steffen Frey[1], Peter Messmer[2] and Thomas Theußl[3]

[1] University of Stuttgart
[2] NVIDIA
[3] KAUST

1 Introduction

Large-scale HPC simulations with their inherent I/O bottleneck have made *in situ* an essential approach for data analysis. In situ coupling of analysis and visualization to a live simulation circumvents writing raw data to disk. Instead, data abstracts are generated that capture much more information than otherwise possible.

The "Workshop on In Situ Visualization" series provides a venue for speakers to share practical expertise and experience with in situ visualization approaches. This 4th edition of the workshop, WOIV'19, took place as a full-day workshop on 20 June 2019 in Frankfurt, Germany, after two half-day workshops in 2016 and 2017 and a full-day workshop in 2018. The goal of the workshop in general is to appeal to a wide-ranging audience of visualization scientists, computational scientists, and simulation developers, who have to collaborate in order to develop, deploy, and maintain in situ visualization approaches on HPC infrastructures.

For WOIV'19 we again also encouraged submissions on approaches that did not live up to their expectations. With this, we expected to get first-hand reports on lessons learned. Speakers should detail if and how the application drove abstractions or other kinds of data reductions and how these interacted with the expressiveness and flexibility of the visualization for exploratory analysis or why the approach failed.

2 Organization of the Workshop

The workshop content was built on two tracks:

Invited talks experts in the field were invited to share their thoughts and insights
Research paper presentations authors were required to submit a full paper before the workshop, which was then reviewed for inclusion in the conference proceedings

After the submission deadline in mid-May 2019, five manuscripts were submitted. Having the time of a full-day workshop at their disposal, the organizing committee was able to select four of these five submissions for presentation at the workshop. After a full review cycle by an international program committee after the workshop presentations, all four papers were selected for inclusion in these conference proceedings. Additionally, five internationally recognized researchers agreed to each give an invited talk at the workshop. Slides for all presentations can be downloaded from the workshop web page at http://woiv.org.

2.1 Organizing Committee

Steffen Frey	University of Stuttgart
Peter Messmer	NVIDIA
Thomas Theußl	KAUST

2.2 Programm Committee

Hadrien Calmet	BSC-CNS
Jose Camata	Federal University of Juiz de Fora
Hank Childs	University of Oregon
Jens Henrik Goebbert	Jülich Supercomputing Centre
Samuel Li	National Center for Atmospheric Research
Kenneth Moreland	Sandia National Laboratories
Benson Muite	University of Tartu
Guido Reina	University of Stuttgart
Joachim Pouderoux	Kitware
Tom Vierjahn	Westphalian University of Applied Sciences
Max Zeyen	Kitware

3 Workshop Summary

3.1 Invited Talks

Burlen Loring and Silvio Rizzi presented the in situ visualization framework SENSEI, which addresses the challenge of toolchain dependent application instrumentation: In order to use in situ visualization the HPC applications needs to be instrumented to expose the simulation data to the visualization application. So far, no single standard interface has emerged, requiring the developers to instrument their applications with different interfaces depending on the targeted in situ visualization toolchain. SENSEI tries to consolidate these different APIs, offering a single mechanism for application instrumentation independent of the selected visualization toolchain. This should help to lower the adoption threshold for in situ visualization and enable a broader range of applications to benefit from it.

In a next presentation, Peter Messmer gave an update on visualization efforts at NVIDIA. Specifically, the arrival of Turing GPUs with hardware accelerated ray tracing put a new emphasis on ray tracing for scientific visualization. Not only does this technology help to produce powerful images for outreach and education, but it also provides better visual cues for e.g. depth and therefore helps a scientist in the day to day work to better understand the spatial relationship of objects in a scene. With the integration of these technologies into popular scientific visualization tools like ParaView, these technologies become available to a broad range of scientific visualization users. In addition to these latest developments on hardware accelerated ray tracing, Peter raised the question whether there is a need for a standardization efforts for ray tracing in scientific visualization in order to free ISV's from developing their own implementations while offering optimal performance on each platform.

In a subsequent presentation, Mike Ringenburg from Cray talked about the concurrently ongoing workshop series on interactive HPC. This conference series has been established to bring together domain scientist, HPC center managers and middleware developers to address the challenges of using HPC systems in an interactive manner. With visualization benefitting significantly from a high degree of interactivity, it was interesting to see the parallels and differences in the two areas. While the focus of the in situ vis workshop series is on the visualization aspect, both interactive or in a batch fashion, the interactive HPC workshops are more interested in the middleware and center policy aspects of interactivity, specifically in light of the convergence of HPC and data sciences. However, despite the differences in the two target audiences, the talk and the subsequent discussion showed that there clearly are common aspects of interest.

Next, Joao Barbosa from TACC presented a summary of the IXPUG In Situ Hackathon 2019. The Hackathon intends to provide a forum to bring simulation developers together with visualization experts for hands-on implementation of in situ analysis. It took first place in 2017 in Austin, Texas hosted at TACC, then in 2018 in Chicago, Illinois hosted at Argonne National Lab, and finally in 2019 Santa Fe, New Mexico hosted by TACC and LAN. Joao presented some statistics of these events with 30–40 participants each, noting that they were successful because they provided a distraction-free environment and a good mix of new and returning participants. He concluded by presenting selected results of the last Hackathon.

Finally, Silvio Rizzi, on behalf of Jim Ahrens, talked about the development and use of in situ visualization and analysis approaches for the U.S. Exascale Computing Project (ECP), an accelerated research and development project funded by the US Department of Energy (DOE). It is a seven-year, $1.7 B R&D effort that launched in 2016 with the participation of six core DOE National Laboratories (Argonne, Lawrence Berkeley, Lawrence Livermore, Los Alamos, Oak Ridge, Sandia), with staff from most of the 17 DOE national laboratories taking part in the project. While ECP applications target national problems, Silvio focussed on the in-situ aspects. He presented an overiew of ECP Software Technology Data and Visualization projects and an outline of several algorithms and subsystems planned within these projects.

3.2 Research Papers

In the first research presentation, Zhang and Entezari discussed their paper "In-Situ Data Reduction via Incoherent Sensing". In this context, they presented study for data reduction in an in-situ situation based on compressed sensing. The volumetric data is first sparsified using wavelet, curvelet, and surfacelet transforms, then compressed sensing techniques are used for compression and reconstruction. The authors demonstrate the effectiveness of their method with examples from chemistry and astrophysics and examine how different methods for compressing the data influence the quality of the reconstruction.

Next, Sarton et al., in their paper "Distributed out-of-core approach for in-situ volume rendering of massive dataset", present a GPU-based volume rendering approach suitable for out-of-core/in-situ visualization of large volumetric data sets. The volume is divided into bricks and managed with a page table layout. An exemplary 2-CPU and 4-GPU server is used for storing, handling, and rendering the volume, the result is then streamed to a thin client for visualization. By providing each GPU with a local brick cache, and by using multi-resolution volume data the system is capable of providing interactive frame rates, which is demonstrated on two volumetric datasets.

In their paper "In-Situ Processing in Climate Science", Röber and Engels, describe the current efforts of leveraging in situ visualization at the German Climate Computing Centre (DKRZ). The paper describes the Catalyst in situ library being integrated into the ICON model and used to generate images and extract geometric features. It also describes some work to write hierarchical data and to compress data with wavelet decomposition. The authors describe a workflow where a special adaptor is developed to fit ICON model to Catalyst, and then the model output is directly visualized through ParaView. The authors also describe a progressive visualization approach where lower resolution data is used for exploration, whereas the original resolution data is used for detailed examination.

Finally, Hummels and van Kooten, in their paper "Leveraging NVIDIA Omniverse for In Situ Visualization", propose a method that allows interactive, high-quality visualization of distilled simulation geometry. Omniverse is NVIDIA's collaboration platform for 3D production pipelines. It is integrated with a number of commercially available 3D software packages and game engines and enables content creators to work on different aspects of models or entire scenes simultaneously. By integrating ParaView and Catalyst with the Omniverse, the visualization geometry becomes immediately accessible to a number of 3D content authoring and rendering tools without the requirement of invasive software changes or tedious postprocessing and conversion workflows.

The workshop was concluded by a panel dicsussion on the "Future of In-situ at ISC", covering a variety of different topics. Among others, we talked about the different directions the workshop could take in future iterations, considering the landscape of related workshops and events in the field of in situ visualization. Here, especially the strong inclusion of simulation and domain scientists, as well as the focus on sharing experiences gained in the development process were identified as important aspects. In addition, the role of interactivity in in-situ visualization was discussed, to identify collaboration potential with the interactivity workshop in future iterations.