

FP7 Support Action - European Exascale Software Initiative

DG Information Society and the unit e-Infrastructures



Addressing the Challenge of Exascale

European Exascale Software Initiative EESI Towards Exascale roadmap implementation

EESI2 – Recommendations

In situ Data Processing

Toward real time Extreme Data Processing and better science through I/O avoidance in High-Performance Computing systems

Philippe Ricoux (TOTAL, EESI2 Coordinator)



EESI2 recommendations



This recommendation is fully in the core of Exascale development, in the Data Centric Pillar, for the best compromise between Extreme Computing and Extreme Data.

Motivation

This approach is really disruptive by Changing the classical sequential paradigm which is: First Simulation ,Then "Off Line" post processing, Visualization and then feedback to new Simulations. This paradigm imposes to move data in memories and in storage.



New large simulations (Turbulence, Combustion, Oil & Gas Reservoir, Neutronics, Electromagnetism, Astrophysics, ...) generate huge amount of simulated data Combustion Turbulent : 1PB each 30 min Giga Model Reservoir: 350 TB/Run

Analysis, Reduction and Visualization on the Fly , In situ Data Processing

New Paradigm: In Parallel Compute and Analyse

Typical Reservoir simulations





Exascale Combustion Turbulent



DNS of Autoigniting Jet Flames

- 3D turbulence coupled with Chemistry (>100 species)
- Fundamental insights into stabilizing a lifted flame in hot ignitive environment
- Predict behavior of new fuels in different combustion scenarios at realistic pressure and turbulence conditions
 - Develop new combustor concepts
 - Design new fuels
- Co-design center is focusing on highfidelity direct numerical simulation methodologies
 - Need to perform simulations with sufficient chemical fidelity to differentiate effects of fuels where there is strong coupling with turbulence
 - Need to address uncertainties in etherma-cohemical properties











generates PB of data

Full DNS with coupling all

physics and chemistry

Exploring the design space of new workflows



Location of analysis compute resources

- Same cores as the simulation (in-situ)
- Dedicated cores on the same node (in-situ)
- Dedicated nodes on the same machine (in-transit)
- Dedicated nodes on external resource (in-transit)

Data access, placement, and persistence

- Shared memory access via hand-off / copy
- Shared memory access via non-volatile near node storage (NVRAM)
- Data transfer to dedicated nodes or external resources
- Synchronization and scheduling
 - Execute synchronously with simulation every nth simulation time step
 - Execute asynchronously





















primary compute resources

- Works well for data-parallel analyses with short run times
- For more complex analyses, impact to the simulation becomes too great



simulation

analysis































Hybrid analysis requires decomposition of algorithms into 2 stages

In-situ	In-transit
Data-parallel	More forgiving of complex communication needs
Short run time with respect to simulation	Can have longer run times while minimizing impact to simulation
Limited amount of memory; minimize cache impacts	Limited to memory and processing constraints of secondary resources
Should minimize the amount of data sent in-transit	Can only require data sent by in-situ stage

EESI2-Final-Conf-Dublin-2015-May-28

In situ techniques and methods

Data Reduction

The transform by itself should be reversible

Subsample, Single precision or double precision,

Quantization Direct scalar quantization,, Adaptive scalar quantization,, Vector quantization (VQ) or block quantization groups (LBG, k means, Codebook training) (VQ is computationally

expensive as an in-situ processing method)

Transform-based compression (FT, wavelet, ...)

Feature Extraction

Large-scale scientific simulations generate massive amounts of data that must be validated and analyzed for understanding and possibly new discovery.

Figure 3. Feature tracking on a turbulent vortex data set [14]. Four selected time steps are shown. Each tracked feature is rendered with a distinct color.

Basic visualization algorithms exist for feature identification, extraction, and tracking, which incorporate principles from image processing, computer vision, and machine learning (especially new deep learning).

Quality Assessment

If data has to be reduced, the corresponding information loss must be conveyed to the user of the reduced data set. Quality assessment thus plays a crucial role in large-scale data analysis and visualization since in many cases the reduced versions of data rather than the original version are used for evaluating the simulation and modeled phenomena.

Issues In Situ Visualization

Compared with a traditional visualization task that is performed in a postprocessing fashion, in-situ visualization brings some unique challenges.





Efficient In situ Data Processing: Coupling 3 Algorithms with different characteristics





In-situ Visualization



Parallel volume rendering

- Design grid adaptor mechanism
- Visualization directly takes data regions from grid adaptor
- Highly scalable parallel volume rendering, particle rendering and image compositing
- Down sampling

Reduced topology computation

- Complete characterization of level-set behavior of simulation variables
- Used to define features of interest
- Compute local merge trees
- Integrate to resolve features spanning multiple cores
- Adjust local merge trees

21

🛛 minimum 🤁 saddle 🖸 maximur





full resolution







Fig. 13. Features extracted from the HCCI simulation for a diff-OH threshold of 0.12 (left) and 0.09 (right).





Digital Rocks Properties

On the fly Visualization





In-Situ Visualization : Collaboration between simulation and visualization On the fly Visualization : Accessing to In-situ Visualization at "simulation time"



On the fly Visualization





at "simulation time"

In-Situ Visualization : Collaboration between simulation and visualization On the fly Visualization : Accessing to In-situ Visualization at "simulation time"

On the fly $\bigcirc \bigcirc \bigcirc$

On the fly Visualization





In-Situ Visualization : Collaboration between simulation and visualization On the fly Visualization : Accessing to In-situ Visualization at "simulation time"







Courtesy of Jackie Chen, EXaCT, Sandia



Data Centric Approaches

Proposal : Fund R&D programs in order to explore

- The design space of future workflows: Real Time (In situ) data-related energy/performance trade-offs for end-to-end simulation workflows running at scale on current high-end computing systems
- Take into account future multi tiered storage architectures and asynchronous data transfers
- Design and implement new analysis techniques typically performed on large-scale scientific simulations: topological analysis, descriptive statistics, surrogate data model, filtering, compression, active learning, pattern/feature discovery, error analysis ... Identify metrics required to characterize classes of analyses Identify which classes of algorithms perform best under which workflow designs Algorithmic shifts – subsampling with quantification of error
- Explore disruptive approaches like sub-linear algorithms addressing the fundamental mathematical problem of understanding global features of a data set using limited resources.
- The use of in situ data analysis for tracking and checking fault or error propagation into the simulation, associating a resilience aspect with the execution of workflows on parallel systems.

Some Biblio on In situ Data Processing & Extreme Computing



- □ In situ data processing for extreme-scale computing Scott Klasky1, All ORNL, LBNL, Sandia, ..., 2013
- **Combining In-situ and In-transit Processing to Enable Extreme-Scale Scientific Analysis** Janine C. Bennett, Hasan Abbasiy, Peer-Timo Bremer, ... Sandia, LBNL, ..., Nov 2012
- Parallel In Situ Indexing for Data-intensive Computing Jinoh Kim, Hasan Abbasi, Luis Chac ´on Ciprian Docan, LBNL, ORNL, 2013
- Parallel In Situ Coupling of Simulation with a Fully Featured Visualization System Brad Whitlock, Jean M. Favre, Jeremy S. Meredith, LLNL, Swiss National Supercomputing Center (CSCS), ORNL, 2011
- Enabling Real-time In-Situ Processing of Ubiquitous Mobile-Application Workflows Hariharasudhan Viswanathan, Eun Kyung Lee, and Dario Pompili, NSF, Rutgers University, New Brunswick, NJ, 2013
- Fast Multiresolution Reads of Massive Simulation Datasets Sidharth Kumar, Cameron Christensen1, John A. Schmidt, Peer-Timo Bremer, Eric Brugger, Venkatram Vishwanath, Philip Carns, Hemanth Kolla,, Ray Grout, Jacqueline Chen, Martin Berzins, University of Utah, LBNL, Sandia, ..., 2014
- Multivariate Volume Visualization through Dynamic Projections Shusen Liu, Jayaraman J. Thiagarajan, Peer-Timo Bremer, Valerio Pascucci, Utah, LBNL, 2014
- Gaussian Mixture Model Based Volume Visualization Shusen Liu, Joshua A. Levine, Peer-Timo Bremer, Valerio Pascucci, Utah, LBNL, Nov. 2012
- Sublinear Algorithms for In-situ & In-transit Data Analysis at Exascale Janine Bennet, Seshadhri Comandur, Ali Pinar, David Thompson, Sandia, 2013 -2104
- In-Situ Processing and Visualization for Ultrascale Simulations Kwan-Liu Ma, Chaoli Wang, Hongfeng Yu, Anna Tikhonova, Department of Computer Science, University of California, 2014
- In-situ Visualization: State-of-the-art and Some Use Cases
 Marzia Rivia, Luigi Caloria, Giuseppa Muscianisia, Vladimir Slavnicb, CINECA, University of Belgrade, 2013



THANK YOU FOR ATTENTION