# In Situ and Post Processing Workflows for Asteroid Ablation Studies

J. M. Patchett[1,2] and B. Nouanesengsy [1] and G. Gisler [1] and J. Ahrens [1] and H. Hagen [2]

[1]Los Alamos National Laboratory, USA
[2]TU Kaiserslautern, Germany

Figure 1: *Volume renderings of two thresholds of partial asteroid density. One colored yellow to red by temperature and showing the ablation material the other colored by partial density in green shows only the cohesive asteroid flattened from the entry.*

Abstract

Simulation scientists need to make decisions about what and how much output to produce. They must balance their ability to efficiently ingest the analysis with their ability to get more analysis. We study this balance as a tradeoff between flexibility of saved data products and accessibility of saved data products. One end of the spectrum is raw data that comes directly from the simulation, making it highly flexible, but inaccessible due to its size and format. The other end of the spectrum is highly processed and comparatively small data, often in the form of imagery or single scalar values. This data is typically highly accessible, needing no special equipment or software, but lacks flexibility for deeper analysis than what is presented. We lay out a user driven model that considers the scientists' output needs in regards to flexibility and accessibility. This model allows us to analyze a real-world example of a large simulation lasting months of wall clock time on thousands of processing cores. Though the ensemble of simulation's original intent was to study asteroid generated tsunamis, the simulations are now being used beyond that scope to study the asteroid ablation as it moves through the atmosphere. With increasingly large supercomputers, designing workflows that support an intentional and understood balance of flexibility and accessibility is necessary. In this paper, we present a new strategy developed from a user driven perspective to support the collaborative capability between simulation developers, designers, users and analysts to effectively support science by wisely using both computer and human time.

Categories and Subject Descriptors (according to ACM CCS): I.3.8 [Computer Graphics]: Applications—

1. Introduction

Asteroids are potentially deadly objects racing through our solar system with one, at least a meter in diameter, entering the earth's atmosphere every other week [NAS16b]. Small asteroids are common, while larger are more rare. Asteroid TC$_3$ [JSN*09] entered the atmosphere over the Northern Sudan in 2008. It marked the first time scientists were able to detect an asteroid prior to it entering

Simulation	Reduction	VDA
10s – 100s GB	10s-100s MB	KBs - 10s MB
Larger		Smaller
General		Specific

More	Flexibility	Less
Data Intensive		Less Data Intensive

Less	Accessibility	More
Special tools		WYSIWYG

Figure 2: *A summary of the supercomputing workflow operations. Reduction refers to data reducing operations. VDA stands for visualization and data analysis.*


Figure 3: *A visualization of the initial conditions. The asteroid is red, water is blue and air is everything else in the bounding box. The axis values are in centimeters.*

the earth's atmosphere. Scientists did not see the 2013 Chelyabinsk meteor [PJE*13] which disturbed a substantial population area, damaging buildings, mostly broken windows, in the middle of the Russian winter and sending many to the hospital for related injuries. If scientists are fortunate enough to detect an asteroid prior to its entering the earth's atmosphere, scientists need to have studied the problem and provide decision makers with potential solutions.

Variables to consider when studying an asteroid impact include size, composition, speed, angle of entry and whether or not there is an airburst. Our work supports scientists studying asteroid impacts using ensembles of simulation results generated with in situ capabilities to make decisions about how, how much and which data to save. We go beyond the simple in situ versus post processing argument and contribute a study of an actual workflow that leads to a user driven model that can support decision making with domain scientists to meet user needs over time. We observe that scientist analysis needs change over time and the requirement of systems to support the balance of flexibility and accessibility of results needs to be maintained.

2. Background

The data product summary in Figure 2 shows typically expected output sizes and their effects from the three main operations of the supercomputing workflow that we study here: Simulation, Reduction, and Visualization and Data Analysis (VDA). Simulation output is the largest and least refined, VDA output is the smallest and most refined and reduction output is somewhere in between, on the spectrum generated by these two extremes. Reduction operation examples include compression or feature extractions. They typically are lossy operations, sacrificing precision or exclusion of certain subsets.

Flexibility refers to the total information content in the data product. For instance, a simulation dump contains the most general information while a single image contains very specific information. Accessibility refers to the ease of access of the data product. A very general data product coming directly from a simulation usually requires sophisticated tools to access usable information, while

a highly refined data product from a VDA operation produces data that can be viewed in a shell or web browser. It is easily and quickly ingested by the domain scientist.

3. Related Work

Work presented here is empirically derived directly from work on in situ for supercomputing applications [FMT*11, CBB*05], particularly while preparing a domain scientist for the the Second International Workshop on Asteroid Threat Assessment: Asteroid-generated Tsunami (AGT) and Associated Risk Assessment [NAS16a] which supported decision making by the attendees and was also presented at Supercomputing 2016 [PGN*]. We seek to improve on the analysis capabilities afforded by saving defensive checkpoint restarts at some regular interval [Dal06]. We also search for ways to move toward a smaller footprint on file systems and are inspired by research like the Cinema project [AJO*14], in situ feature detection and preservation [WPB*11, WPS*16, WHA*11], and compression [DC16, BHM*16] all of which enables users to make tradeoffs between accessibility and flexibility and still manage practical constraints like storage, space and time. Our models and strategies also attempt to consider, or at least not exclude, emerging hardware and software technologies like burst buffers [LCC*12, BFA*12] and VTKm [MSU*16]. Many observations have been influenced by the data, information, knowledge and wisdom hierarchy described in [Ack89].

4. Approach

We attempt to limit our specifics to a single simulation of a 500-meter diameter asteroid, initially at a 20-kilometer elevation ripping through the atmosphere at 17 kilometers per second on a simulation grid containing three materials: air, water and asteroid as seen in Figure 3. The simulation begins with a modest 150 million cells, but quickly ramps up to 500 million cells when the asteroid impacts the water and then quickly goes to 1.3 billion cells. Our initial goal is to analyze the effect on the deep ocean water after the asteroid impact. A later goal, generated after the initial simulation run, is to then study the asteroid itself, prior to impact.

The computational simulation is performed by xRage [GWC*08], a parallel multi-physics Eulerian hydrodynamics code that is developed and maintained by the ASC program at the Los Alamos National Laboratory. xRage uses a continuous adaptive mesh refinement (AMR) technique that allows

Figure 4: *A simple workflow that begins with a simulation, extracts a feature and performs visualization and data analysis on that feature with the decision to store or not store data between each processing element.*

Figure 5: *Volume rendering that shows ocean surface extracted.*

smaller computational cells in areas of interest and larger, thus fewer, cells in other areas, which enables more efficient use of the supercomputer. The simulation is outfitted with an integrated ParaView Catalyst capability that translates the computational grid into a VTK unstructured grid representation, then hands control to ParaView which is capable of executing a Reduction or VDA pipeline before returning control back to the simulation code.

We use a supercomputer that contains 2 GB of memory per processing core with a high performance interconnect between nodes (16-36 cores per node). We also have access to a VDA cluster which has many fewer nodes but contains 196 GB of memory and 12 cores per node. Both of these machines have access to large multi petabyte shared file systems. We also have many desktop computers with modern graphics hardware, 64 GB of RAM and a single terabyte SSD storage.

Figure 4 shows the basic workflow that is being studied here. The simulation produces data that can either be persistently stored or can be handed directly to a feature extraction algorithm. The algorithm will produce a reduced data set which can be either persistently stored or passed directly into the visualization and data analysis algorithm. This algorithm, in turn, will produce easily accessible data products like imagery or numbers. There are four paths through this workflow. The full in situ stores no intermediate data and produces only final visualization and data analysis products. The full post processing stores intermediate data products at each juncture. There are two hybrid approaches. The first passes the data directly to the reduction algorithm in situ and stores the much smaller output from that which can be quickly and efficiently read by the visualization and data analysis operation. The second starts as a traditional post processing algorithm with a large data set dropped from the simulation and the feature extraction and visualization and data analysis are combined with no intermediate writes.

$$T_{total} = S + S_{out} + R_{in} + R + R_{out} + V_{in} + V \qquad (1)$$

Formula 1 shows the cost model derived from the possible workflows. Let T_{total} be the sum of all costs, S be the cost of the simulation operation, R is the data reduction operation and V be the cost of the visualization and data analysis. S_{out} is the cost of the simulation's data output, which could be a disk write or an in situ adapter cost. R_{in} is the cost for the reduction operation to access the data, this could be near zero in the case of in situ or the time to read from disk for post processing. R_{out} is the cost of the reduction operation producing data and either storing it or passing it to the VDA operation. Finally V_{in} is the cost of ingesting data for the visualization and data analysis operation.

5. Results

We identify four fundamental use cases based on our model and our experience. We present them in a natural order that they appear to occur that aligns with the maturing of a simulation coupled with a topic and a group of domain scientists. A goal would be to achieve the fourth use case as quickly as possible.

S, R, and V are static costs, regardless of workflow. The computational expense does not change and we can simply remove them from our analysis, leaving only the details of the data movement decisions. $S_{outInsitu}$ and $S_{outFile}$ are the two main classes of expense coming out of the simulation. $R_{inInsitu}$ and R_{inFile} represent the cost of getting data into the reduction operation, and $R_{outInsitu}$ and $R_{outFile}$ are the costs of getting data out of the reduction operation. Finally $V_{inInsitu}$ and V_{inFile} are the costs for getting data into the VDA operation.

5.1. Full Post Processing: $S_{outFile} + R_{inFile} + R_{outFile} + V_{inFile}$

The first data product that is typically and persistently stored from simulations is the checkpoint restart (CR). This is the entire state of the simulation required to restart from that state. Analysis can be performed with the CR as a source. Then some other type of output that is more accessible usually comes next: the viz dump. We classify the viz dump in this bin as it usually contains the full grid and a selection of scalar values. This workflow maximizes flexibility at the cost of disk space and user/computer time. We sought to perform interactive volume rendering. This is a difficult task using production software on unstructured grids with hundreds of millions of cells. The solution was to read the visualization dumps, sample them onto structured grids and save those, which effectively turned 30-100 GB per time step data sets into .6-1.2 GB per time step data sets. This data reduction not only made interactive exploration feasible, it made movement of the data to local disks that didn't require supercomputers feasible. This is a vast improvement in accessibility. An output from such data can be seen in Figure 5, which shows the ocean surface after the impact, which is important when analyzing tsunami generation.

5.2. Hybrid: $S_{outFile} + R_{inFile} + R_{outInsitu} + V_{inInsitu}$

Improving the reduction workflow to increase the accessibility and specificity of the output, starting from the vis dump all the way to distinct data products, is a natural progression in scientific inquiry. Output files which were produced to answer more general questions generated more specific questions. These specific questions required more specialized data in order to obtain an answer.
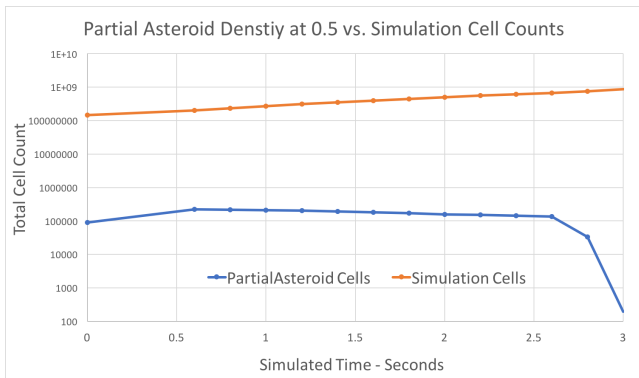
Figure 6: *A plot showing total cells in the full simulation grid versus the total cells in an asteroid partial density threshold greater than 0.5.*

(a) *1.0 seconds*   **(b)** *1.2 seconds*   **(c)** *1.4 seconds*   **(d)** *1.6 seconds*

Figure 7: *A sequence of asteroid at the temporal resolution saved from the simulation partial density threshold greater than 0.5.*

In our case, volume rendering the entire mesh generated questions of how much of the asteroid was actually being lost while it went through the atmosphere. The sampled data clearly showed that a large quantity of asteroid material was left behind and disintegrated in the atmosphere before the impact. The sampled data, although great for making macro level, qualitative imagery accessible, was not fully appropriate for measuring the size of the solid asteroid. The native grid needed to be accessed again. This required loading the time step data, which were in the tens of gigabytes. Once the data was loaded on the supercomputer using hundreds of processors, the asteroid could be extracted, requiring a fraction of the original grid, on the order of tens of megabytes, which is easily accessible for interactive volume rendering. Since we were not fully confident we knew the exact values required to extract the solid asteroid, we made sure to include density values that were well into asteroid dust range. Since basalt density is on the order of 2.7 g/cm^3 we chose to threshold on computational grid cells of .5 g/cm^3 this produced data sets that are sufficiently small to transfer with high accessibility. Results of this reduction in size can be seen in Figure 6. Full resolution for the volume of interest is preserved in the reduction. The image in Figure 1 used such a refined workflow and the relatively small asteroid in Figure 3 helps to explain such a reduction.

5.3. Hybrid: $S_{outInsitu} + R_{inInsitu} + R_{outFile} + V_{inFile}$

Data representing the asteroid between the originally saved time steps is needed to better visually explain the mushrooming between 1.0 and 1.6 seconds as seen in Figure 7. We don't foresee value in higher temporal, full resolution dumps from the simulation. We are now searching for very specific things, but still want some generality for exploration within the subset for our new study. Improving the in situ dump, to provide only known needed data, is the first push back to the domain scientist who is normally responsible for the simulation.

Having developed a pipeline of asteroid extraction, given the integration of the VDA tool with the simulation, it was trivial to simply rerun the simulation and have the simulation output high resolution asteroid. It is still expensive in terms of time to run the

simulation again, but we don't need to spend time doing defensive checkpoints or vis dumps. The effects on the file system are minimal and we can get hundreds of easily accessible time steps in the disk space consumed by a single vis dump.

This is the first step in really removing flexibility in support of accessibility. If the generality of the dump is needed again it will have to come from the running simulation which would have a greater cost in terms of time, typically, more than simply reading from disk. The improvement in accessibility, though, is large, as the data is much more refined, more information dense in an area of interest, potentially much smaller, and therefore potentially accessed with fewer resources and fewer specialized tools and is definitely faster. Accessibility to an end-user should not be underestimated. Of course they want both accessibility and flexibility.

The decision by a domain scientist to sacrifice the generality found in saving full spatial resolution dumps comes only with an increased maturity of understanding what the implications are of sacrificing that loss for the accessibility. They know what they're looking for and already have a good sense of what they'll be missing.

5.4. Full in situ: $S_{outInsitu} + R_{inInsitu} + R_{outInsitu} + V_{inInsitu}$

The full in situ comes at the end of the pipeline. The simulation is well understood. It is unlikely that any other information will be needed in the near future regarding the simulation. Some number of checkpoint restarts and all of the data necessary to rerun the simulation must still be preserved. We look forward to achieving this point. The risk associated with the full in situ is potentially mitigated by projects like Cinema which can potentially save a large quantity of refined information, which is stored in a searchable database.

6. Conclusion

We presented a model with an example, including the context of a real-world use. The ideas can be leveraged by individuals and groups in designing simulation runs to ensure accessibility and flexibility needs are met. This is important to manage, not only for compute and storage resources, but importantly the scientist time.

7. Acknowledgements

References

[Ack89] ACKOFF R. L.: From data to wisdom. *Journal of applied systems analysis 16*, 1 (1989), 3–9. 2

[AJO*14] AHRENS J., JOURDAIN S., O'LEARY P., PATCHETT J., ROGERS D. H., PETERSEN M.: An image-based approach to extreme scale in situ visualization and analysis. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis* (2014), IEEE Press, pp. 424–434. 2

[BFA*12] BENT J., FAIBISH S., AHRENS J., GRIDER G., PATCHETT J., TZELNIC P., WOODRING J.: Jitter-free co-processing on a prototype exascale storage stack. In *Mass Storage Systems and Technologies (MSST), 2012 IEEE 28th Symposium on* (April 2012), pp. 1–5. doi:10.1109/MSST.2012.6232382. 2

[BHM*16] BAKER A. H., HAMMERLING D. M., MICKELSON S. A., XU H., STOLPE M. B., NAVEAU P., SANDERSON B., EBERT-UPHOFF I., SAMARASINGHE S., DE SIMONE F., ET AL.: Evaluating lossy data compression on climate simulation data within a large ensemble. *Geoscientific Model Development 9*, 12 (2016), 4381. 2

[CBB*05] CHILDS H., BRUGGER E., BONNELL K., MEREDITH J., MILLER M., WHITLOCK B., MAX N.: A contract based system for large data visualization. In *Visualization, 2005. VIS 05. IEEE* (2005), IEEE, pp. 191–198. 2

[Dal06] DALY J. T.: A higher order estimate of the optimum checkpoint interval for restart dumps. *Future generation computer systems 22*, 3 (2006), 303–312. 2

[DC16] DI S., CAPPELLO F.: Fast error-bounded lossy hpc data compression with sz. In *Parallel and Distributed Processing Symposium, 2016 IEEE International* (2016), IEEE, pp. 730–739. 2

[FMT*11] FABIAN N., MORELAND K., THOMPSON D., BAUER A. C., MARION P., GEVECIK B., RASQUIN M., JANSEN K. E.: The paraview coprocessing library: A scalable, general purpose in situ visualization library. In *Large Data Analysis and Visualization (LDAV), 2011 IEEE Symposium on* (2011), IEEE, pp. 89–96. 2

[GWC*08] GITTINGS M., WEAVER R., CLOVER M., BETLACH T., BYRNE N., COKER R., DENDY E., HUECKSTAEDT R., NEW K., OAKES W. R., RANTA D., STEFAN R.: The RAGE radiation-hydrodynamic code. *Computational Science & Discovery 1*, 1 (2008), 015005. URL: http://stacks.iop.org/1749-4699/1/i=1/a=015005. 2

[JSN*09] JENNISKENS P., SHADDAD M., NUMAN D., ELSIR S., KUDODA A., ZOLENSKY M., LE L., ROBINSON G., FRIEDRICH J., RUMBLE D., ET AL.: The impact and recovery of asteroid 2008 tc3. *Nature 458*, 7237 (2009), 485–488. 1

[LCC*12] LIU N., COPE J., CARNS P., CAROTHERS C., ROSS R., GRIDER G., CRUME A., MALTZAHN C.: On the role of burst buffers in leadership-class storage systems. In *Mass Storage Systems and Technologies (MSST), 2012 IEEE 28th Symposium on* (2012), IEEE, pp. 1–11. 2

[MSU*16] MORELAND K., SEWELL C., USHER W., LO L.-T., MEREDITH J., PUGMIRE D., KRESS J., SCHROOTS H., MA K.-L., CHILDS H., ET AL.: Vtk-m: Accelerating the visualization toolkit for massively threaded architectures. *IEEE Computer Graphics and Applications 36*, 3 (2016), 48–58. 2

[NAS16a] NASA: Asteroid-generated tsunami (AGT) and associated risk assessment. online, August 2016. https://tsunami-workshop.arc.nasa.gov/workshop2016/. 2

[NAS16b] NASA: Newly released map data shows frequency of small asteroid impacts, provides clues on larger asteroid population, 2016. URL: http://neo.jpl.nasa.gov/news/news186.html. 1

[PGN*] PATCHETT J., GISLER G., NOUANESENGSY B., ROGERS D. H., ABRAM G., SAMSEL F., TSAI K., TURTON T.: Visualization and analysis of threats from asteroid ocean impacts. Available at https://youtu.be/yeXcgnj8AG0 LA-UR-16-28934. 2

[PJE*13] POPOVA O. P., JENNISKENS P., EMELÂĂŹYANENKO V., KARTASHOVA A., BIRYUKOV E., KHAIBRAKHMANOV S., SHUVALOV V., RYBNOV Y., DUDOROV A., GROKHOVSKY V. I., ET AL.: Chelyabinsk airburst, damage assessment, meteorite recovery, and characterization. *Science 342*, 6162 (2013), 1069–1073. 2

[WHA*11] WOODRING J., HEITMANN K., AHRENS J., FASEL P., HSU C.-H., HABIB S., POPE A.: Analyzing and visualizing cosmological simulations with paraview. *The Astrophysical Journal Supplement Series 195*, 1 (2011), 11. 2

[WPB*11] WILLIAMS S., PETERSEN M., BREMER P.-T., HECHT M., PASCUCCI V., AHRENS J., HLAWITSCHKA M., HAMANN B.: Adaptive extraction and quantification of geophysical vortices. *IEEE transactions on visualization and computer graphics 17*, 12 (2011), 2088–2095. 2

[WPS*16] WOODRING J., PETERSEN M., SCHMEIβER A., PATCHETT J., AHRENS J., HAGEN H.: In situ eddy analysis in a high-resolution ocean climate model. *IEEE transactions on visualization and computer graphics 22*, 1 (2016), 857–866. 2