

LA-UR-17-23903

Approved for public release; distribution is unlimited.

Title:	A Pipeline for Large Data Processing Using Regular Sampling for Unstructured Grids
Author(s):	Berres, Anne Sabine Adhinarayanan, Vignesh Turton, Terece Feng, Wu Rogers, David Honegger
Intended for:	Report
Issued:	2017-05-12

Disclaimer:

Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by the Los Alamos National Security, LLC for the National Nuclear Security Administration of the U.S. Department of Energy under contract DE-AC52-06NA25396. By approving this article, the publisher recognizes that the U.S. Government retains nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or to allow others to do so, for U.S. Government purposes. Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy. Los Alamos National Laboratory strongly supports academic freedom and a researcher's right to publish; as an institution, however, the Laboratory does not endorse the viewpoint of a publication or guarantee its technical correctness.

A Pipeline for Large Data Processing Using Regular Sampling for Unstructured Grids

Anne Berres, Vignesh Adhinarayanan, Terece L. Turton, Wu Feng, David H. Rogers

Wednesday 10th May, 2017, 20:12




Figure 1: Kinetic energy plotted for the entire ocean.

1 Motivation

Large simulation data requires a lot of time and computational resources to compute, store, analyze, visualize, and run user studies. Today, the largest cost of a supercomputer is not hardware but maintenance, in particular energy consumption. Our goal is to balance energy consumption and cognitive value of visualizations of resulting data. This requires us to go through the entire processing pipeline, from simulation to user studies.

To reduce the amount of resources, data can be sampled or compressed. While this adds more computation time, the computational overhead is negligible compared to the simulation time.

We built a processing pipeline at the example of regular sampling. The reasons for this choice are two-fold: using a simple example reduces unnecessary complexity as we know what to expect from the results. Furthermore, it provides a good baseline for future, more elaborate sampling methods.

We measured time and energy for each test we did, and we conducted user studies in Amazon Mechanical Turk (AMT) for a range of different results we produced through sampling.

2 Data

As test data, we used climate simulation data from a resolution run of the Model for Prediction Across Scales – Ocean (MPAS Ocean), a multi-resolution approach to climate modeling [RPH⁺13, MM05], which is part of the Accelerated Climate Model for Energy (ACME) project []. For the purpose of this study, we chose Kinetic Energy, a variable that is commonly used to visualize eddies in the data. Eddies are circular currents which occur along streams, and which contribute heavily to global climate. We applied a colormap that is optimized for good contrast [SPG⁺15, SA15, SPA⁺15]. An example can be seen in Figure 2.




Figure 2: Eddies, seen as dark green rings, originate from the Agulhas Retroflection at the southern tip of Africa. They travel westwards across the Atlantic towards South America.


3 Sampling and Testing Setup

We use a latitude/longitude projection of the simulated data, as shown in Figure 1. The original grid is a hex-dominant Voronoi grid on a sphere, so this projection results in very uneven cells. In a preliminary step, we sample this grid to a regular grid which faithfully represents the original resolution. The original simulation has cells that are 15×15 km in size. In the projection, cells at the equator have the same size, whereas cells that are closer to the poles are strongly stretched. We apply a regular sampling that oversamples the original grid by a factor of about 1.5, resulting in 3750×1500 cells.

3.1 Subsampling the High-Resolution Grid

To subsample the original data, we apply the `ResampleToImage` filter which is part of the ParaView [FMT⁺11] suite. We gradually increase the grid’s cell size, thereby gradually lowering the resolution.

For example, the original grid resolution corresponds to 1×1 units. Doubling the cell size to 2×2 units equates cutting the resolution of both dimensions in half, and reducing the overall resolution to a quarter of the original resolution. As the filter interpolates, one can also use decimals, like 1.5×1.5 , etc. Figure 3 shows an example of the results achieved using different sampling resolutions.



(a) Unstructured grid. (b) Supersampled grid. (c) Subsampled 3x3 grid. (d) Subsampled 9x9 grid.

Figure 3: Comparison of different sampling resolutions, from the original hex-dominant Voronoi grid, supersampled regular grid, and two different sampling densities.

3.2 Rendering

In the raw data, each grid cell has one value, resulting in a flat color when rendered directly. When sampling the data to larger grid cells, this has the appearance of increasingly large pixels. An advantage of this rendering approach is that it displays the data exactly as-is. However, it comes with the drawback that color/value changes between neighboring cells can look very abrupt and seem like feature.

4 Results

To determine the differences between various sampling rates and rendering styles, we conducted a crowdsourced perceptual user study, and we conducted energy measurements.

4.1 Perceptual User Study

In order to assess the cognitive impact of the sampling on the visualization – and therefore the impact on the scientist’s ability to see features of interest – we conducted a user study to determine the discrimination threshold. The discrimination threshold is the level at which participants can distinguish between the unsampled data image and a sampled version. A two-alternative force-choice (2AFC) approach was used [Fec89]. Applying the 2AFC method to this work, the baseline (unsampled) image is compared to sampled images spanning some range where one expects to find the threshold. When comparing the baseline to itself, a participant will not be able to detect a difference and random chance should yield a value of 50% for the percentage of participants choosing the baseline as the clearer image. When comparing the baseline to an image well past the threshold, one would expect 100% of the participants to rate the baseline as clearer than the sampled version. The discrimination threshold is set at the point where 75% of participants can detect a difference between the unsampled and sampled images.

In this 2AFC study, participants were shown a series of comparison pairs; one was always the unsampled baseline image and the other had a sampling factor applied. Participants were asked to

focus on the features of interest to the scientist: ocean eddies. We showed the participants images from three regions of interest to climate scientists, allowing us to see the impact of the different scales represented by each region: Agulhas Retroflexion (mid-range view, many eddies), east coast of Australia (very close-up view of a small number of eddies); and the North Atlantic (very far view covering large percentage of globe).

Our study was implemented using Qualtrics survey software [Qua] and participants were recruited from Amazon’s Mechanical Turk [AMT]. For the study, participants were given a brief set of instructions that introduced the concept of eddies and were shown an unsampled and a sampled image as examples of the issue under study. They were told that for each pair of comparison images, they should choose the image that showed the eddies most clearly. Each comparison was presented randomly as A vs. B or B vs. A (left/right presentation). As a validation check to weed out participants who may not have understood the task or who may not have been faithfully attempting the task, we used a comparison between the baseline image and an image that had been heavily compressed using JPEG compression. We required participants to correctly identify the uncompressed image as clearer. Additionally we removed any participant who only clicked on one answer choice (always the “A” or always the “B” choice).

We started with a pilot study using subsamplings of 1x1 against stimuli levels of 3x3 and 9x9 subsampling. With 35 participants, this study indicated that 3x3 was already beyond the discrimination threshold.

In a more detailed study, we used finer steps in the subsampling stimuli levels: from 1.1x1.1 to 3.4x4.4 in steps of 0.1. We collected approximately 200 participants for each sampling level across the three regions. The 2AFC ramps can be seen in Figure 4. Note that the results depend on how zoomed in the view of the eddies, as one would expect. The very close up view around Australia hits the 50% threshold at 1.2x1.2 while the wider geographic views hit the 50% mark farther out. The overall threshold is around 2.0x2.0.

More information on the evaluation techniques can be found in [TBR17, BTP⁺17].

4.2 Energy Measurements

We are interested in the energy consumption for the sampling itself in addition to saving as VTK image data (i.e., `.vti` format) or color-mapped images (i.e., `.png` format). For measuring energy, we use a WattsUp Pro power meter which measures power at one-second intervals. We use the ParaView’s `pvpython` library to read, sample, and write all data. ParaView’s *surface* rendering interpolates the data bi-linearly, which is not only very costly but also introduces diagonal smear as an artifact. The *slice* rendering does not apply this interpolation; it shows the raw data, so one can see the blocks clearly. A comparison of the two rendering methods can be seen in Figure 5.

Figure 6a compares the energy consumption for *surface* rendering against *slice* rendering. One can clearly see that, particularly for high resolutions (i.e., small grid cell sizes), it is far more energy-efficient to use *slices* than *surfaces*. Hence, we use *slice* rendering for our *png* output. However, writing out raw data as VTK image is much more energy-efficient than rendering a color-mapped *png*, as demonstrated in Figure 6b.

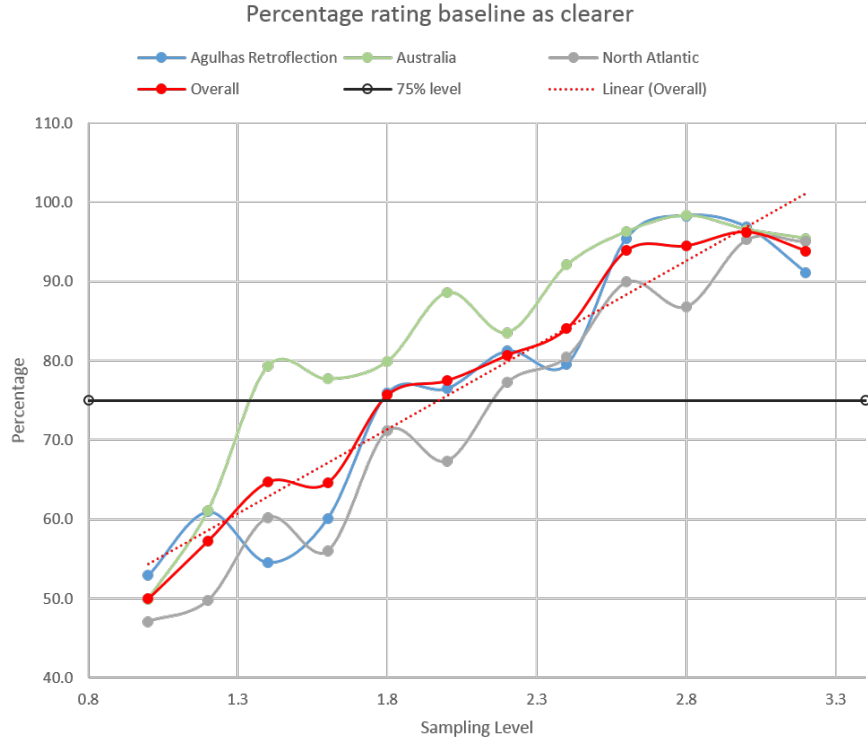


Figure 4: Results of a perceptual evaluation study comparing images with different sampling grid sizes for different ocean regions. In this study, participants were forced to choose which of two images showed the eddies *more clearly*.





Figure 5: Comparison of rendering styles. Surface rendering appears darker due to surface shading. One can see diagonal zigzag artifacts along the diagonal from top left to bottom right but not along the diagonal from top right to bottom left.



(a) Comparison between power for non-interpolated (b) Comparison between power for png output and slices and interpolated surfaces. power for vti output.

Figure 6: Comparison of different sampling grids, from the original hex-dominant Voronoi grid, supersampled regular grid, and two different sampling grid cell sizes.

References

- [AMT] Amazon Mechanical Turk website. www.mturk.com/mturk/welcome.
- [BTP⁺17] Anne S. Berres, Terece L. Turton, Mark Petersen, David H. Rogers, and James P. Ahrens. Video compression for ocean simulation image databases, June 2017. Accepted into to EuroGraphics Workshop on Visualization in Environmental Sciences.
- [Fec89] Gustav Theodor Fechner. *Revision der Hauptpunkte der Psychophysik*, volume 2. Breitkopf & Härtel, Leipzig, 2nd ed. edition, 1889.
- [FMT⁺11] Nathan Fabian, Kenneth Moreland, David Thompson, Andrew C Bauer, Pat Marion, Berk Geveci, Michel Rasquin, and Kenneth E Jan. The ParaView coprocessing library: A scalable, general purpose in situ visualization library. In *Large Data Analysis and Visualization (LDAV), 2011 IEEE Symposium on*, pages 89–96. IEEE, 2011.
- [MM05] M.E. Maltrud and J. L. McClean. An eddy resolving global 1/10° ocean simulation. *Ocean Modelling*, 8(1-2):31–54, 2005.
- [Qua] Qualtrics website. www.qualtrics.com.
- [RPH⁺13] Todd Ringler, Mark Petersen, Robert L. Higdon, Doug Jacobsen, Philip W. Jones, and Mathew Maltrud. A multi-resolution approach to global ocean modeling. *Ocean Modelling*, 69:211 – 232, 2013.
- [SA15] F. Samsel and U. Ayachit. <http://www.kitware.com/blog/home/post/886>, May 2015. Kitware Blog.
- [SPA⁺15] Francesca Samsel, Mark Petersen, Greg Abram, Terece Turton, David Rogers, and James Ahrens. Visualization of ocean currents and eddies in a high-resolution global ocean-climate model. In *Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis*, number 11, 2015.
- [SPG⁺15] Francesca Samsel, Mark Petersen, Terece Geld, Greg Abram, Joanne Wendelberger, and James Ahrens. Colormaps that improve perception of high-resolution ocean data. In *Proceedings of the 33rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems, CHI EA '15*, pages 703–710, 2015.
- [TBR17] Terece L. Turton, Anne S. Berres, and David H. Rogers. ETK: An evaluation toolkit for visualization user studies, June 2017. Accepted into EuroVis 2017: 19th EG/VGTC Conference on Visualization.