SeedMe Preview: Your Results from Disk to Device

Amit Chourasia
amit@sdsc.edu

Mona Wong-Barnum
mona@sdsc.edu

Michael L. Norman
mlnorman@sdsc.edu

San Diego Supercomputer Center, University of California San Diego
9500 Gilman Drive, MC 0505, La Jolla, CA – 92093

ABSTRACT

Computational simulations have become an indispensible tool for a wide variety of science and engineering investigations. With the rise in complexity and size of computational simulations, it has become necessary to continually and rapidly assess simulation output. Visualization could play an even more important and critical role for qualitative assessment of raw data. The result of many visualization processes is a set of image sequences, which can be encoded as a movie and distributed within and beyond the research group. The movie encoding process is a computationally intensive, manual, serial, cumbersome and complicated process as well as one that each research group must undertake. Furthermore, sharing visualizations within and outside the research groups requires additional effort. On the other hand, the ubiquity of portable wireless devices has made it possible and oftentimes desirable to access information anywhere and at anytime, yet the application of this capability for use in computational research and outreach has been negligible. We are building a cyberinfrastructure ***SeedMe*** (**S**tream **E**ncode **E**xplore **D**isseminate **M**y **E**xperiments) to fill these gaps that will enable seamless sharing and streaming of visualization content on a variety of platforms from mobile devices to workstations making it possible to conveniently view and assess the results thus provide an essential yet missing component in computational research and current High Performance Computing infrastructure.

Categories and Subject Descriptors

H.3.5 **[Online Information Services]** Data sharing, Web-based services; H.3.7 **[Digital Libraries]** Collection, Dissemination; H.5.1 **[Multimedia Information Systems]** Video; H.5.3 **[Group and Organization Interfaces]** Asynchronous interaction, Collaborative computing, Web-based interaction

General Terms

Management, Documentation, Design, Human Factors, Standardization

Keywords

Visualization, Video Encoding, Ubiquitous Computing

1. INTRODUCTION

Increased computation capability as well as ubiquity of numerous sensing devices has enabled researchers to compute and collect vast amounts of data. The latest NSF funded High Performance Computing (HPC) resources including Stampede at Texas Advanced Computing Center; Blue Waters at National Center Supercomputing Applications and Gordon at San Diego Supercomputer Center are providing an unprecedented computing capability to a wide range of researchers [1, 2, 3]. Researchers are gearing up to use the unique features of these systems to conduct massive simulations. However, analysis of simulated data has become a daunting challenge due to its exponential growth [4, 5, 6]. We begin to see a shift in the simulation community where **the focus is rapidly moving towards the joint analysis of simulated and observational data**. There is a new push towards analyzing the big data as fast as possible through in-situ or co-located methods.

Visualization is a key and efficient method to gather qualitative insights about underlying data. Interactive visualization capabilities continue to improve at a modest rate and many researchers are taking advantage of this in their work. However, the exponential increase in raw data coupled with only modest improvements in input-output rates has widened the gap for data analysis in general including visualization. Interactive visualization is impractical for very large data sets, considering the amount of time required to load data from disk into memory, and the time to perform the necessary exploration and manipulation. In practice interactive visualization may be used to identify settings to capture features of interest and then batch visualization is performed to generate images or image sequences. Scientists rely on these images for qualitative assessment of data. Furthermore, visualization serves a critical role during development, debugging and test phase of simulations.

In recent years web technologies have made significant strides to improve interaction, performance and presentation of content. These improvements are a result of concerted effort towards standards compliance. With the imminent release of HTML5 specification [7], which is expected to support features including canvas tag, video tag and potentially WebGL the web browser vendors are rapidly adopting these features in addition to improving interaction performance. Moreover, the rise of social media [8] has fueled the growth of web content not just limited to desktop and laptops but also on smart phones and tablets devices [9]. The web browsers have become an integral feature in most computing devices and have made it possible to consume a variety of content anywhere and at any time. However, the adoption of these capabilities to augment and facilitate research within HPC community has been quite limited. The purpose of this project is to change the status quo by providing HPC researchers a ubiquitous, easy to use cyberinfrastructure for encoding and sharing of scientific visualizations.

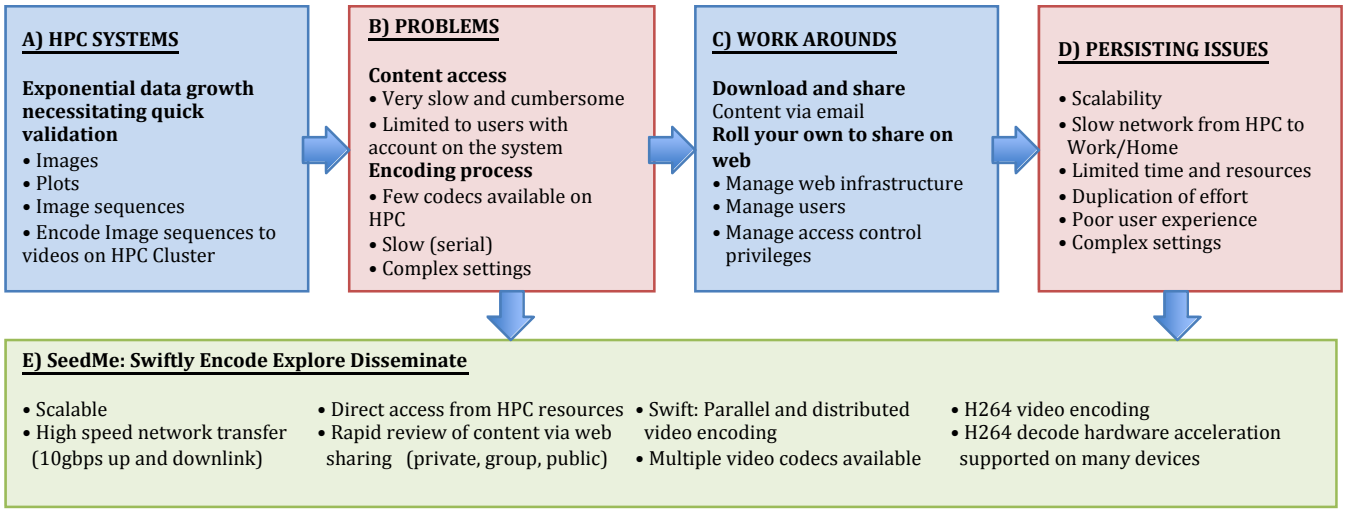

Figure 1. Illustrating problems with current infrastructure box (a) indicates the need of derived content from raw data for rapid assessment, box (b) highlights the problems of sharing and video encoding, box (c) lists the common mitigation methods applied, box (d) highlights persisting problems after mitigation effort. Box (e) lists the SeedMe solution and its merits that include swift, accessible and automatic publishing with granular rights management for large number of users.

2. EXISTING PROBLEMS

Sharing visualization results within and beyond research groups is a slow, cumbersome and error prone task (Figure 1). Currently, scientists resort to viewing individual images or downloading them often on a slow network and share them via email to their colleagues. Images sequences can be turned into a movie through a slow serial video encoding process, which could take many hours to complete.

Moreover, the movie encoding process is riddled with several arcane video codec choices and video compression settings. To our knowledge the Linux-like operating systems have only few tools (e.g. FFmpeg [10]) and do not support parallel video encoding. Furthermore, the security model used on existing HPC infrastructures is not suitable for direct public sharing. Thus each research group needs to roll their own solution to share videos, determine adequate video encoding settings, manage and configure access control, and allocate resources for download and streaming. These efforts result in varying degree of success, which lacks scalability and leads to duplication of effort.

3. PROPOSED SOLUTION

We are building *SeedMe*, a turnkey cyberinfrastructure to ***S****wiftly* ***E****ncode,* ***E****xplore* and ***D****isseminate* derived data in the form of visualization images, plots and movies on the web. We aim to address the problems described above (see Figure 1) by creating a cyberinfrastructure. We must note that sharing of raw HPC data is impractical just by its sheer size, moreover large datasets are of little utility without corresponding compute power thus beyond the scope of this project. However, this is not the case with derived visualization content in most circumstances. The key characteristics and strengths of *SeedMe* cyberinfrastructure are enumerated as follows:

- **Publish and Disseminate.** *SeedMe* will provide a capability to publish scientific plots and visualization images on the web with granular access control directly from the HPC resource. A web interface will enable users and wider audience to upload, download, browse, search, tag, and rate the movies. A wiki-like interface will allow content owners to edit meta-data and provide scientific explanation. A rating method will allow web users to rate the movies and provide comments while a community based flagging method will be implemented to purge inappropriate content.
- **Near Real Time Monitoring.** *SeedMe* will be connected with two 10 Gbps up and downlink making it possible to transfer content rapidly. **Simulation monitoring** can be performed via periodic submission of text, images or plots for ongoing simulations as they become available.
- **Swift**. *SeedMe* will encode movies in a parallel-distributed manner (See Figure 2). This will result in significant speeding up of the video encoding process thus making results available in a very short time span.
- **Share and Stream**. *SeedMe* will provide easy web access to the images and movies. Automatic encoding of movies with multiple settings will make it possible to stream movies on many devices including smart phones, tablets, laptops and desktops. We postulate that this will enable significant improvement in accessibility and time saving over current practice.
- **Short Term Storage**. *SeedMe* will store and host the images and movies for at least 90-days from date of creation and up to the duration of the project. Longer duration may be possible, but will be contingent on storage availability.
- **Easy Access.** We propose a low barrier to access the *SeedMe* resource. We will provide direct access from XSEDE resources via command line and from other systems and web.
- **H.264 Compliant Encoding**. *SeedMe* will encode movies using a H.264 compliant video codec [11, 12] that has both hardware encoding and decoding support [13, 14] built into many devices including smart phones, tablets, laptops and desktops. Compression and video quality of H.264 compliant codecs is generally higher than other video codecs [15].

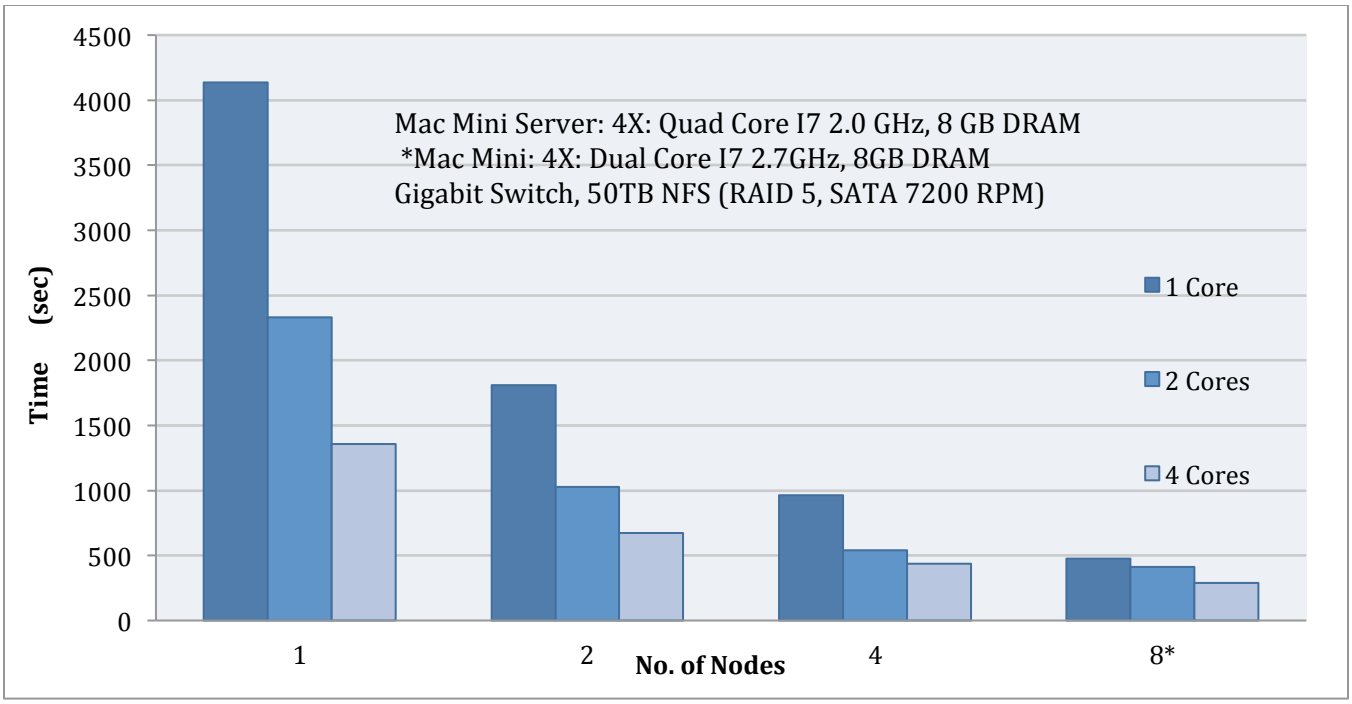

Figure 2. Video encoding scalability benchmark (Lower is better). A sequence of 1660 images at 4096x2048 pixels was encoded using H.264 codec in a distributed manner using Compressor 4 software. Multiple encoding nodes provide performance gain.

4. APPLICATION SCENARIOS

The following exemplary scenarios illustrate application of the *SeedMe* infrastructure in variety of scientific contexts. In each case the user submits a job comprising of image or image sequences along with a setting file that includes email addresses and desired privacy settings. In all cases the primary user is granted administrative rights to their content.

- A researcher conducting analysis of his simulation data creates plots and image sequences from temporal data on a HPC resource. A job is submitted to *SeedMe* with specific content using a simple command line interface. In a short time the researcher and specified collaborators will be notified by an email indicating the URL where the plots and visualization movies can be accessed on the web from desktop or other portable devices.
- A researcher would like to monitor and share progress of his simulation in form of plots and visualization images with his colleagues who do not have access to the HPC resource. The researcher sets up a periodic submission to *SeedMe* using a script. A one-time email is sent out to specified users with a URL location where the content can be accessed on web in near real-time as it becomes available.
- An environmental science researcher would like to share a set of image streams from cameras recording a wild fire or a tornado from several locations with public safety officials and even a wider audience. The user can set up a periodic submission as above and share the URL with others without risking downtime on his limited yet critical infrastructure due to heavy load.
- A researcher or science museum staff would like to transcode a set of videos from workshops or planetarium movies into a different format for better compression, archival or dissemination. They can submit a transcoding job to *SEED*. Once the transcoding completes the user is notified by email with a URL location to download results.
- A research team conducting a simulation of potential natural disaster for example earthquake would like to share simulation results in form of visualization to a large number of people for education and outreach. They can use this resource easily.
- We will also develop plug-in for visualization tools like ParaView [16] and VisIt [17] so users can utilize *SeedMe* from within and outside these tools.

5. CONCLUSIONS

The production system will be widely accessible to all users of XSEDE systems and beyond. We anticipate *SeedMe* to have a very broad impact by serving a diverse range of communities and providing them many benefits as follows

- **Collaboration and Sharing**

  Large computational efforts are usually collaborative and require coordination with many individuals. Currently, there is a serious lack of tools and processes for collaboration. *SeedMe* will enable swift sharing of analysis results and visualizations within and beyond a group of researchers. We believe this will lead to rigorous and robust data analysis.

- **Ubiquity**

  SeedMe will provide direct access from many resources and will offer a portable solution to view/review visual content on multitude of devices.

- **Near Real Time Monitoring**

  Large simulations often run the risk of continued run when errors may have occurred early on due to bad parameter setting, bad IO operation, network issues or processor failure. These risks increase significantly on large systems. *SeedMe* will allow instrumenting the simulations with user set periodic update of visualization or analysis results during the simulation run. The continuous monitoring mechanism enabled by *SeedMe* will allow researchers to potentially catch simulation errors early on saving precious compute time and effort.

- **Video Encoding for Visualization**

  Visualization images are significantly different than other media content like films, natural scenery and other commonly recorded video, where lossy compression does not seriously impact the viewer's assessment. However, this is not the case with scientific content, which needs to be encoded at higher bitrates and ideally also offered as a lossless stream. *SeedMe* cyberinfrastructure is being developed to offer these derived products.

- **SeedMe Invitation**

  We anticipate this infrastructure will serve a large and wide user base and would like to extend and invitation to sign up for early access or get notification when the system goes in production on our project website at **http://www.seedme.org**

6. ACKNOWLEDGMENTS

We would like to thank Apple Inc. for providing the test hardware and software for two weeks during conceptualization phase, Andrew Ferbert and Doug Weimer for software and hardware setup, Thomas Hutton for identifying network switch for the infrastructure. This material is based upon work supported by the National Science Foundation under Grant No. OCI-1235505.

7. REFERENCES

[1] Texas Advanced Computing Center. 2013. *Stampede*. Retrieved Apr 25, 2013 from http://www.tacc.utexas.edu/stampede

[2] National Center for Supercomputing Applications. 2013. *About the Blue Waters project*. Retrieved Apr 25, 2013 from http://www.ncsa.illinois.edu/BlueWaters

[3] San Diego Supercomputer Center. *Gordon User Guide: Technical Summary*. Retrieved Apr 25, 2013 from http://www.sdsc.edu/us/resources/gordon

[4] T. Hey and A. Trefethen. 2003. *The data deluge: an e-Science perspective*. In F. Berman, G. C. Fix, and A. J. G. Hey, editors, Grid Computing: Making the Global Infrastructure a Reality, pages 809–824. Wiley, 2003.

[5] A. Szalay and J. Gray. 2006. *Science in an exponential world*. Nature, 440, March 23 2006.

[6] C. Anderson. 2008. *The End of Theory: The Data Deluge Makes the Scientific Method Obsolete*. Wired Magazine 16.07 (June 23, 2008)

[7] World Wide Web Consortium. 2013. *HTML5*. Retrieved Apr 25, 2013 from http://www.w3.org/TR/html5

[8] Nielson. 2011. *Social Media Report: Q3 2011*. Retrieved Apr 25, 2013 from http://blog.nielsen.com/nielsenwire/social

[9] Nielson Wire. 2011. *Report: The Rise of Smartphones, Apps and the Mobile Web (Dec 15, 2011)*. Retrieved Apr 25, 2013 from http://blog.nielsen.com/nielsenwire/online_mobile/report-the-rise-of-smartphones-apps-and-the-mobile-web/

[10] FFmpeg. 2013. *FFmpeg FAQ*. Retrieved Apr 25, 2013 from http://ffmpeg.org/faq.html#Which-are-good-parameters-for-encoding-high-quality-MPEG_002d4_003f

[11] ITU-T and ISO/IEC JTC 1. 2007. *Advanced video coding for generic audio- visual services*, ITU-T Recommendation H.264 and ISO/IEC 14496-10 (MPEG-4 AVC), Version 1: May 2003, Version 2: May 2004, Version 3: Mar. 2005, Version 4: Sep. 2005, Version 5 and Version 6: June 2006, Version 7: Apr. 2007, Version 8 (including SVC extension): Consented in July 2007.

[12] T. Wiegand, G. J. Sullivan, G. Bjøntegaard, and A. Luthra. 2003. *Overview of the H.264/AVC video coding standard*. IEEE Transactions on Circuits and Systems for Video Technology, vol. 13, no. 7, pp. 560-576, July 2003.

[13] Intel. 2013. *Intel Quick Sync Video*. Retrieved Apr 25, 2013 from http://www.intel.com/content/www/us/en/architecture-and-technology/quick-sync-video/quick-sync-video-developer.html

[14] TI. 2013. *DaVinci Digital Media Processors | TI.com*. Retrieved Apr 25, 2013 from http://www.ti.com/lsds/ti/dsp/video_processors/overview.page

[15] D. Kulikov and A. Parshin. *MPEG-4 AVC/H.264 Video Codecs Comparison*. April 2010. Retrieved Apr 25, 2013 from http://compression.ru/video/codec_comparison/h264_2010/

[16] Kitware. 2013. ParaView. *Open Source Scientific Visualization*. Retrieved Apr 25, 2013 from http://www.paraview.org

[17] VisIt. 2013. *VisIt Visualization Tool*. Retrieved Apr 25, 2013 from https://wci.llnl.gov/codes/visit