EG UK Computer Graphics & Visual Computing (2017)
Tao Ruan Wan and Franck Vidal (Editors)

Human-in-the-Loop Visualisation Architecture for Monitoring Remote Compute

M. J. Turner[†1,2], S. Nagella[2], R. Fowler[2], R.J. Allan[3], E. Pasca[2] and E. Yang[2]

[1]Research IT Services, The University of Manchester, Manchester M13 9PL
[2] Scientific Computing Department, STFC, Didcot, Oxon, OX11 0QX
[3] Hartree Centre, STFC, Sci-Tech Daresbury Warrington, WA4 4AD

Abstract

This paper describes the timeline of use cases of large and remote display VEs (Virtual Environments), hosted by STFC (Science and Technology Facilities Council), which were linked to HPC (High Performance Computing) systems. Considered is the development and use in the last few years of putting the human back into the HPC loop and clarifying the main types of interaction and collaboration that have been re-explored.
It describes a set of specific common modes of use as well as stages of development, categorising and explaining how best practice may be achieved.

CCS Concepts
•***Computer Graphics*** → *Hardware Architecture; Picture/Image Generation; Graphics Utilities; Application packages; Methodology & Techniques;*

1. Introduction: Survey Time

Over the last two years the STFC's Scientific Computing Department's Visual Analytics and Imaging System Group [Yan17], has reconsidered visualisation needs for computational sciences. To achieve this a series of informal and formal surveys were carried out [TFM16]. Two of these specifically considered the Tomographic Imaging and CFD (Computational Fluid Dynamics) communities [Jon17]. The outcomes stated which software tools were actually being used and proposed some best methods to support them [MFT14].

Open source was not always the most important issue but the easy creation of plug-ins, new readers and writers, as well as analysis tools were requested. There was also indicated a strong growth in the use of the ParaView visualisation system [HAL04] `http://www.paraview.org` that is an open source, multi-platform data analysis and visualisation application where users can build systems including adding qualitative and quantitative techniques. Two other popular tools, repeatedly specified by tomographic imaging users, were Avizo `https://www.fei.com/` and ImageJ `https://fiji.sc/`. To address the communities' needs, the related research software engineering group [Nag17] now run licences and training courses, and have created a plug-in wrapper service [Wor13, NF17] for these three application platforms. This is now part of the Collaborative Computational Project in Tomographic Imaging (CCPi) core activities `http://www.ccpi.ac.uk/`

From the perspective of these applications we will consider in this paper some use cases and experiences, over the past couple of years, of software and hardware built to support the communities with their particular computational and visualisation needs.

2. Visualisation Matters – the need to put the human back in the computational loop; Modes and Development Stages of use

Three common development stages and three modes of interactive use have been a recurring aspect within virtual environment (VE) creation. In a report [LLG*07] while developing and studying the HydroVR software system Lidel et al. described the three stages of development. The first stage is to create an application from a single data set with real-time interaction. Then the second stage adds collaborative features where users even within multiple VEs can interact and view the same data set simultaneously. The final and third stage is the incorporation of multiple remote data sets that can be added as required by the user. A follow on report [MTL08] considered at the same time three modes of operation; the first mode is the single dedicated researcher exploring a particular data set and finding individual inspiration. This is where true personal immersive environments have the ability to allow a user to be inside and

† Contact, martin.turner@manchester.ac.uk

DOI: 10.2312/cgvc.20171275

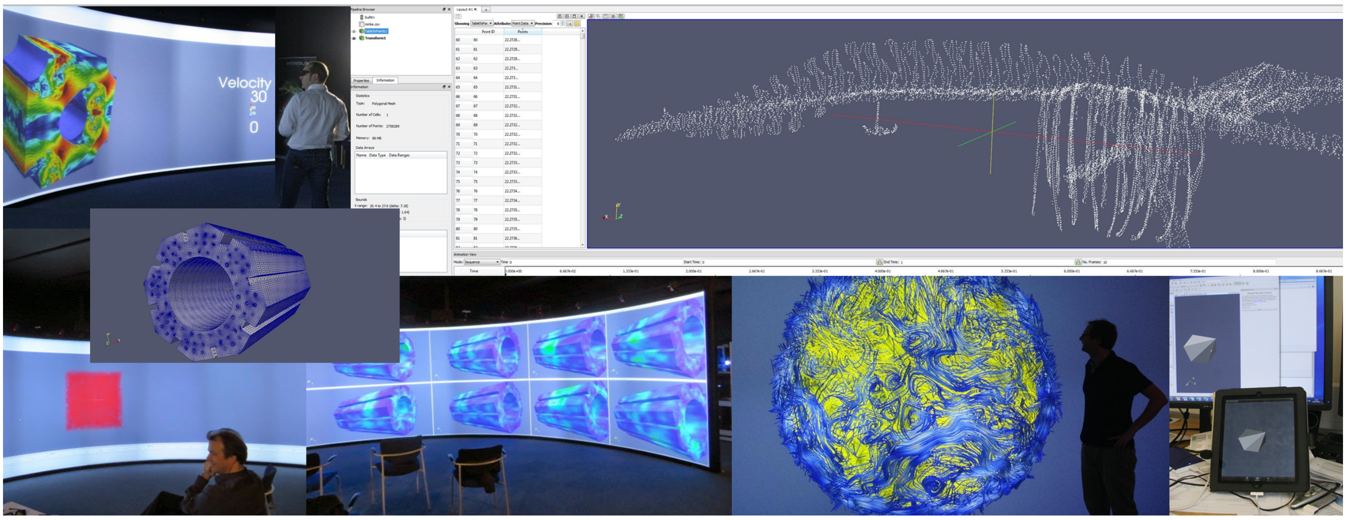

Figure 1: *Interactive use Mode One – ParaView in various forms for individual interactive VE use;- clockwise from top left; CFD monitoring of remote visualisation for HPC dataruns; LiDAR data evaluation and pre-processing visual focussed qualification for smoothed particle hydrodynamics; Remote server from client end for iPad like control panels automating simple interactions; Multi-modality and stereoscopic display for post-Biophysical computational modelling; Tomography and ParaFEM simulation of bricks from a fusion reactor showing tens of simulations simultaneously.*

"at one" with the data set(s) and gain insights that would be difficult otherwise. The second mode is a more common use of a VE where a global, possible elaborate, group presentation occurs and the key part is a one-way conversation. The important feature here is that the presenter(s) can "tell a story" to the audience. The final mode of operation involves small team interaction exploiting all parts of the VE space. This third mode also illustrated a use of VE in the way of a visualization "what if" analysis. This can be considered in the same way as a spreadsheet allows users to experiment with financial "what if" questions, but by dynamically changing parameters a VE can be used to visually verify simulations and manage computational job submissions. The team interaction allows consensus, or otherwise, between different specialists to occur in an efficient manner

We will consider use cases that highlight these three modes and three stages within a specific physical VE that was built to interact with and plan for HPC job control.

3. Use Case 1: Hartree Centre development and building a universal virtual environment platform

The Hartree Centre `https://www.hartree.stfc.ac.uk` created two bespoke and adaptable visualisation suites as multi-use centres to become physical VEs. *Leverhulme* (8025 × 1920, 10.25m by 2.3m) a 15 megapixel stereoscopic curved wall; blended to create one large interactive and collaborative space for data presentation and exploration; and *Crosfield* (3615 × 1880, 5.5m by 2.85m) a 6.6 megapixel stereoscopic rear-projected flat glass wall; blended and integrated into a flexible project boardroom and meeting environment.

One of the best ways to understand and explain complex high resolution data, is to show it on a very high resolution display and personally walk through and around a problem; ParaView's stereo display enabled individuals and groups of users to be immersed within their data; and also walk "inside the problem". Shown in Figure 1 is a range of technical features that were incorporated to aid the single user gaining insight including; remote interaction using an ipad on a client system; kiwi viewer and ParaViewWeb for mobile devices; and the Catalyst method for HPC specific network visualisation streaming. The ParaView code base is designed in such a way that all of its components can be reused to develop so-called vertical applications and it has been successfully deployed as an augment to the IBM Blue Gene/Q HPC system. It has also been used within development projects and user communities; within the ISIS 'Mantid' data flow system and for volume visualisation and within the RCaH (Research Complex at Harwell) community (new tools include `http://www.tomviz.org/`).

Figure 2 shown two examples of remote group presentations, where ParaView's linked views and different applications can be controlled for a larger audience to explore singular and multiple datasets. This cinema type mode of use where a story can be presented is in most VE physical systems the most common and publicised mode of use. Interaction between the presenter and the audience can be difficult. but the use of a VE can help make this occur – this has been used within academic classses to public engagement; and from project presentation to proposal planning.

3.1. Use Case 2: The human-in-the-computational-loop: still important for supercomputing

An important role for ParaView is to allow crude computational steering within large batch runs. Although systems like the IBM Blue Gene/Q are highly ranked on the green computing list (the

Figure 2: *Interactive use Mode Two – examples of complex group presentations; top, showing multi-scale linked views in ParaView from Image based Modelling computational analysis of experimental facilities based upon tomographic real-data captured by Culham Centre for Fusion Energy; bottom left, multiple views tomography viewing to explain computational processing techniques within a Nuclear Physics master class; and bottom right showing complete 3D audience CFD immersion.*

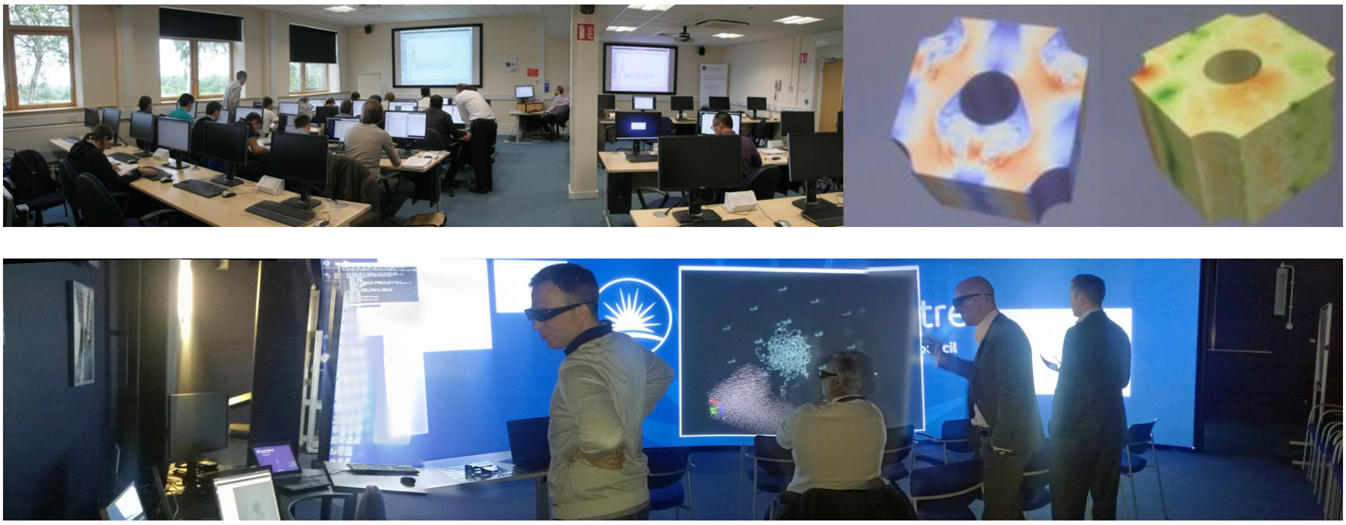

Figure 3: *Interactive use Mode Three – two examples of small team interactions; top, showing a ParaView training and steering example; presenting and running CFD (Code_Saturne) code in the co-located Hartree Centre training rooms and visualisation suites. Bottom, showing discussion space linking with a IBM Meeting system and Skype, combined with curved wall visualisation of chemical structures, deciding on follow-on computational job submissions.*

STFC system was number 26 in the world on the June 2014 Top500 Green List at 2,178 MFlop/s per Watt) the computational cycles are still not free and limited to a certain size of cores. Therefore users have to be careful in selecting the right parameters when running a set of jobs. The Hartree system's GPFS large file store is connected to the visualisation suites using fibre, which has enabled the development of a dynamic monitoring method for long computational jobs. As we had local GPUs and good network this made the process simpler than other systems [LTP*11, DGAW13].

So for example in the CFD simulation (Figure 3 top) we see results of small team interaction. There were multiple very long jobs submitted, with each job taking many hours to run but every five to ten minutes an intermediate file was produced that could be interrogated by ParaView. This enabled a scenario where users could return to the visualisation suite and monitor the outputs, interact with and investigate the data sets. If everything was going well then the jobs could be left to complete, otherwise selective jobs could then be terminated and restarted with different parameters.

In a similar format (Figure 3 bottom) the visualisation suites have been used, to discuss past job submissions and plan for further computational tasks. This multi-use space also allows for simple video conferencing to occur enabling external advisers to then be included and share the ParaView visuals. Of the three interaction modes this team mode was perceived as the most important and the one to provoke the most discussion during sessions.

The ability to put the human, locally and remotely, back in the computational-loop can be incorporated very simply, but can easily save a huge amount of computational cycles and energy.

3.2. Use Case 3: Software for the Future: remote ParaView

Tomographic (3D data volumes) data files are becoming too big to be moved efficiently (>20GB). Visualisation results from the dataflow model, for example images or surfaces are often substantially smaller. It is these that can be interactively manipulated remotely and ParaView provides some appropriate tools. It has the ability to run in client-server mode so that the data of a simulation can remain on the parallel file system it was generated on and only the required visualisation geometry is sent to the client. This avoids the need to transfer many GB of data to the client (based in the visualisation suite), which can be slow and inconvenient.

A key component then is to place a new ParaView server node, with good local GPU and multi-core support, in the right location. We have built bespoke high-end and large memory server nodes where the HPC login server nodes are usually mounted; (64GB, 32 core Intel Xeon CPU E5-2670 0 at 2.60GHz with two NVIDIA Quadro 5000). This allows for good response rates, easy access to users as they do not need to retrain for login and job submission, and enables fast access to the datastore without going through extra firewalls. System operation though virtualisation makes software installation independent of the HPC main nodes and data storage service, thus easing the burden considerably.

A disadvantage of this system setup is that the new servers may not be as powerful as the visualisation suite equipment, but the advantage of enabling users to quickly see data has a far higher benefit, and data can always be analysed later in detail off-line if necessary.

The same scenario is being extended to other compute clusters at the Hartree Centre. A recent port of OpenGL to the NVIDIA driver for OpenPower has enabled one of the authors to install VMD `http://www.ks.uiuc.edu/Research/vmd` and ParaView for the computational biology and CFD communities. We believe this is the first time in the UK that it has been done on IBM Power-8 systems with respectively multiple Tesla K80 and Pascal P100 GPUs.

3.2.1. IMAT Neutron Tomography beamline – linking the human-visualisation loop to the imaging-capture process

We now consider a key future use case where we add a national imaging facility to the loop proposed. The IMAT (Imaging and Materials Science and Engineering) beamline, currently in test mode for 2017 [KBK*15], is the first neutron imaging instrument from ISIS that offers unique time-of-flight tomography-driven diffraction techniques. Remote viewable visualisation of the results using ParaView has now been tested, and incorporates fast parallel image reconstruction algorithms so on-the-fly image processing and visualisation can inform and guide experiments, saving experimental time as well as HPC cycles and energy.

Image files (>20GB) are transmitted on a dedicated link to a local STFC HPC cluster `http://www.scarf.rl.ac.uk/`; reconstructed using the latest techniques to create 3D volumes and then processed and viewed remotely via target VMs with high-end GPUs (NVIDIA Grid K2's). This is capitalising on the latest image reconstruction procedures and data collection schemes (CCPi-Core Imaging Library software toolkit within the CCPi [NPF17]).

This project not only brings the human into the computational loop but also brings the human and visualisation system within the imaging facility. This is an exciting development to make efficient use of computation and experimental beamline time.

3.3. Human-in-the-computational-loop

Most physical visualisation systems when connected to HPC results only consider the lone individual gaining insight and the presentation mode to a wider audience. Both these modes can be done off-line and really do not interact with the HPC job submission and procesing pipeline. Over the last few years we have reemphasised that to be efficient a VE visualisation needs to occur at all stages. This may require some investment; in networking, localised and remote GPUs, and install video streaming software so that small team interaction can occur. Once in place, the job subission process can change to include interactions and modifications at the pre-, during and post-computational stages. This has been shown to improve selection of HPC jobs and makes more efficient use of HPC resources. There is still resistance in moving away from the traditional system of submitting many long batch run jobs as this is easier to plan, cost and deliver.

An example of a process that most HPC system would not allow is a complex CFD problem where the user does not know how long convergence will take, perhaps one hour or after three hours. Current solutions would be to run a job for a very long time that wastes CPU ycles but guarantees to converge, or add the addition of checkpointing where all intermediary results need to be stored just in case the job has to be resubmitted to run for longer. A solution with an integrated human in the computational-loop would be to submit a job asking for an infinte run time and then use the monitoring and discussion process during team interactive mode to have a human specify a halting point.

3.4. Conclusions and summary

There are plenty of distributed alternatives to ParaView that are available [CBW*12, NIC17] and some are being used within the STFC groups. Looking at one product and its uses over two years gives an interesting snapshot. We have shown within the Hartree Centre and SCD facilities the same stages of development and use modes defined for VE exploitation, but we also have stated the need to add a small amount of extra computational support and networking when required. This has included careful thought of data transfer, as well as location of new servers which are connected to the data stores. ParaView, supported by Kitware and a large user community, has recent updates including enhanced parallelism (VTK-m), better OpenGL rendering and support for immersive headsets (e.g. HTC Vive) which all helps to keep the software highly relevant in support for HPC visualisation.

The ability to put the human back in the computational-loop has shown itself to save resources and improve the computational dataflow. Future HPC services should consider this at all stages especially during the early procurement. Three objectives are becoming universal; colocate data with the compute; have visualisation applications near the data/compute; and then provide a productive place for collaboration and human to human networking – to quote “create a place where there is a good cup of tea and a heated discussion”.

The next stage we have shown is to consider the integration of large imaging-capture facilities, that create data streams in a similar way to HPC, and then build server and storage infrastructure appropriately.

We acknowledge use of Hartree Centre resources in this work. The STFC Hartree Centre is a research collaboratory in association with IBM providing High Performance Computing platforms funded by the UK’s investment in e-Infrastructure. The Centre aims to develop and demonstrate next generation software, optimised to take advantage of the move towards exa-scale computing. The authors also wish to thank other staff for their time and discussions: Hartree Centre; Terry Hewitt, David Hughes, Thomas Christy and Project IMAT / CCPi; Edoardo Pasca, Winfried Kockelmann, Genoveva Burca; and the many users over the years.

References

[CBW*12] CHILDS H., BRUGGER E., WHITLOCK B., MEREDITH J., AHERN S., PUGMIRE D., BIAGAS K., MILLER M., HARRISON C., WEBER G. H., KRISHNAN H., FOGAL T., SANDERSON A., GARTH C., BETHEL E. W., CAMP D., RÜBEL O., DURANT M., FAVRE J. M., NAVRÁTIL P.: VisIt: An End-User Tool For Visualizing and Analyzing Very Large Data. In *High Performance Visualization–Enabling Extreme-Scale Scientific Insight*. CRC Press, Oct 2012, pp. 357–372. 5

[DGAW13] DEMARLE D., GEVECI B., AHRENS J., WOODRING J.: Streaming and out-of-core methods. In *High Performance Visualization: Enabling Extreme Scale Scientific Insight*. CRC Press, September 2013. 3

[HAL04] HENDERSON A., AHRENS J., LAW C.: The ParaView guide, 2004. Kitware Inc., Clifton Park, NY. 1

[Jon17] JONES D.: Collaborative Computational Projects portal. STFC / Scientific Computing Department, Jun 2017. http://www.ccp.ac.uk/. 1

[KBK*15] KOCKELMANN W., BURCA G., KELLEHER J. F., KABRA S., ZHANG S.-Y., RHODES N. J., SCHOONEVELD E. M., SYKORA J., POOLEY D. E., NIGHTINGALE J. B., ALIOTTA F., PONTERIO R. C., SALVATO G., TRESOLDI D., VASI C., MCPHATE J. B., TREMSIN A. S.: Status of the neutron imaging and diffraction instrument imat. *Physics Procedia 69* (2015), 71 – 78. Proceedings of the 10th World Conference on Neutron Radiography (WCNR-10) Grindelwald, Switzerland October 5-10, 2014 http://www.sciencedirect.com/science/article/pii/S1875389215006203. doi:http://dx.doi.org/10.1016/j.phpro.2015.07.010. 4

[LLG*07] LIDAL E., LANGELAND T., GIERTSEN C., GRIMSGAARD J., HELLAND R.: A decade of increased oil recovery in virtual reality. *IEEE Computer Graphics and Applications* (2007), 94–97. 1

[LTP*11] LEAVER G. W., TURNER M. J., PERRIN J. S., MUMMERY P. M., WITHERS P. J.: Porting the avs/express scientific visualization software to cray xt4. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences 369*, 1949 (2011), 3398–3412. http://rsta.royalsocietypublishing.org/content/369/1949/3398. doi:10.1098/rsta.2011.0133. 3

[MFT14] MORRIS T., FOWLER R., TURNER M.: Visualisation Matters Survey Results. Wiki on Computer Science, University of Manchester, Dec 2014. http://www.vizmatters.cs.manchester.ac.uk. 1

[MTL08] MCDERBY M. J., TURNER M. J., LEAVER G. W.: Modes of Virtual Environments Integrated within Collaborative Environments. In *Theory and Practice of Computer Graphics* (2008), Lim I. S., Tang W., (Eds.), The Eurographics Association. doi:10.2312/LocalChapterEvents/TPCG/TPCG08/129-135. 1

[Nag17] NAGELLA S.: Collaborative Computational Project in Tomographic Imaging. STFC / Scientific Computing Department, Jun 2017. http://www.ccpi.ac.uk/. 1

[NF17] NAGELLA S., FOWLER R.: Image Quantification Algorithms SourceForge. EPSRC collaborative software development environment, Jun 2017. http://ccpforge.cse.rl.ac.uk/gf/project/iqa/. 1

[NIC17] NICE SOFTWARE: Nice desktop cloud visualization and enginframe, Jun 2017. http://www.nice-software.com/products/dcv. ../enginframe. 5

[NPF17] NAGELLA S., PASCA E., FOWLER R.: Core Imaging Library (CIL) within CCPi. STFC / Scientific Computing Department, Jun 2017. http://www.ccpi.ac.uk/node/CIL. 4

[TFM16] TURNER M. J., FOWLER R., MORRIS T.: Collaborative Computational Projects - Visualisation Applications Survey. In *Computer Graphics and Visual Computing (CGVC)* (2016), Turkay C., Wan T. R., (Eds.), The Eurographics Association. doi:10.2312/cgvc.20161292. 1

[Wor13] WORTH D.: *Developing an image analysis plugin for VolView*. Tech. Rep. 7, STFC Rutherford Appleton Lab.), 6 2013. RAL-TR-2013-007. Persistent URL http://purl.org/net/epubs/work/65673. 1

[Yan17] YANG E.: Visual Analytics and Imaging Systems Group. STFC / Scientific Computing Department, Jun 2017. http://www.scd.stfc.ac.uk/Pages/Visual-Analytics-and-Imaging-Systems-Services.aspx. 1