The Desktop is Dead, Long Live the Workstation!

Aaron Knoll, Carson Brownlee, Khairi Reda, Brian Haymore, Michael E. Papka, Valerio Pascucci

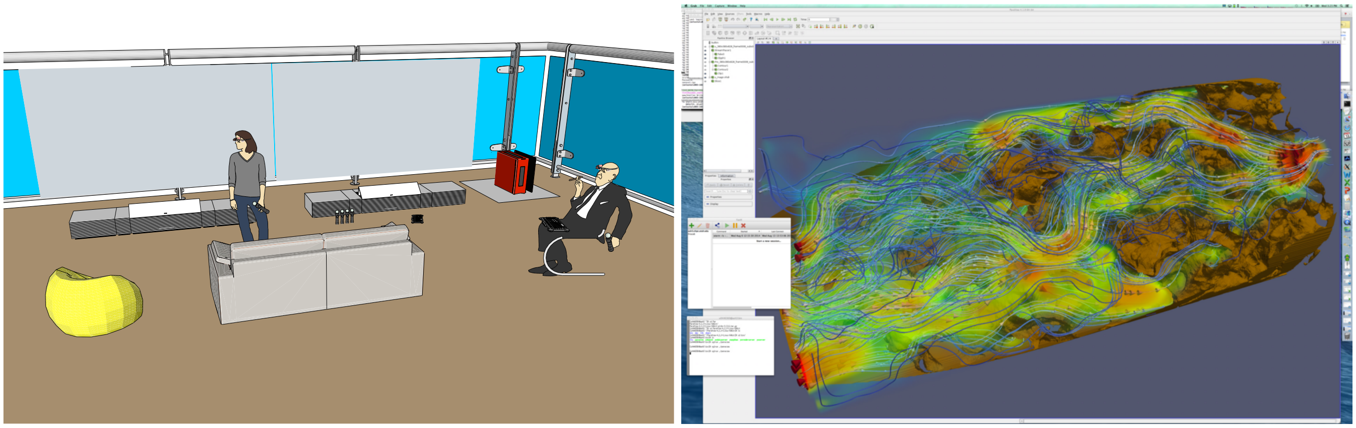

Fig. 1. **Left:** "Usable" vis lab concept. A SMP workstation with 4–8 manycore CPUs and several GPUs could singlehandedly drive a tiled display or multi-projector system with a single OS, with on the order of 100 megapixels, 20-50 teraflops, and 1–2 terabytes of main memory. Tethered to a supercomputer filesystem and globus tools, it is also a multi-user remote visualization resource. Instead of dedicated dark rooms, the workstation-driven vis lab is public office space: well-lit and at least as comfortable to work in as a café, library, airport, or home sofa. **Right:** Remote visualization from a "fat node" on the Ash cluster at the University of Utah. This node has 768 GB RAM, dual Intel Xeon CPUs, an NVIDIA GPU, and cost roughly $12,000. The resource requires no scheduling – user state is persistent until logout. FastX [9] enables seamless remote visualization – in this case at 4K resolution over WAN on a wireless network. While the desktop may die, personal computing will live on in the form of workstations and fat nodes such as these.

Abstract— We explore shared-memory workstations as compelling alternatatives to desktops and small clusters, for purposes of scientific visualization. With new manycore CPU hardware on the horizon and the current popularity of large-memory "fat nodes" in HPC, SMP workstations are poised to make a comeback. These machines will augment, not replace, HPC and cloud resources, providing both remote visualization and more personalized vis labs. They will be accessible anytime, anywhere on any device, running a single operating system, capable of handling all but the absolute largest scientific data. We describe current state of the art, emerging trends, and use cases that could make the SMP workstation the dominant driver of high-end scientific visualization in the next decade.

Index Terms—VR, HCI, HPC, supercomputing, desktop, workstation

1 Introduction

The era of desktop computing is nearly over. Thanks to mobile and cloud computing, monitors and towers sit at desks gathering dust. High-end interactive visualization, traditionally dependent on fast GPUs with direct video output, represents the "last stand" for the desktop outside of the home gaming market. Even for interactive visualization, there has been a broad movement toward remote rendering on supercomputers, and (in practice) replacement of desktops with laptops. With the increasing importance of cloud computing and shared resources, and a new array of promising visual and haptic interfaces, it is not surprising that many are heralding the demise of the desktop.

Over the last 5 years, however, annual predictions of the desktop's demise have proven premature. Thanks to the consumer PC game market, GPU-based visualization continues strong. However, large-scale visualization, analysis and database problems increasingly depend on efficient memory access as opposed to FLOPS. The needs of large-scale visualization increasingly resemble those of data analysis and server applications addressing terabytes of data, as opposed to computer games optimized for gigabytes on a GPU. Perhaps the desktop is not dying, but evolving with the rest of computing. We believe the survivors of the "destocalpyse" will be larger SMP workstations: bigger, faster, and more connected than ever, providing supercomputing levels of FLOPS but interactive access to large-scale data, without the limitations of current distributed approaches.

2 Rebirth of the Workstation

We see workstation-driven visualization driving the following trends:

- **Fewer, bigger, shared resources.** Users migrating from a laptop will want a resource that provides clear advantages over their personal machine.
- **Large main memory.** Unlike HPC workloads, visualization and analysis tools benefit from having a large amount of usable memory relative to FLOPS.

• *Aaron Knoll is with the SCI Institute, University of Utah. E-mail: knolla@sci.utah.edu. Carson Brownlee is with the Texas Advanced Computing Center, University of Texas at Austin. E-mail: carson@tacc.utexas.edu. Khairi Reda is with the Electronics Visualization Laboratory at University of Illinois Chicago, and Argonne National Laboratory. Email: mreda2@uic.edu. Brian Haymore is with the University of Utah Center of High Performance Computing. Email: brian.haymore@utah.edu Michael E. Papka is with Argonne National Laboratory. Email: papka@anl.gov Valerio Pascucci is with the SCI Institute, University of Utah. Email: pascucci@sci.utah.edu*

Manuscript received 31 Mar. 2014; accepted 1 Aug. 2014; date of publication xx xxx 2014; date of current version xx xxx 2014.
For information on obtaining reprints of this article, please send e-mail to: tvcg@computer.org.

- **Manycore CPUs.** Manycore CPUs (and likely future GPUs) will access main memory directly without the latency of current hardware. Ultimately, these machines will be programmed using thread-parallel paradigms that take advantage of shared memory.

- **Real-time Ray Tracing.** In contrast to current visualization approaches employing OpenGL rasterization, rendering workloads will use scalable real-time ray tracing algorithms [10, 5] designed for CPU and manycore architectures. Ray tracing enables direct rendering of scientific data, including volume data. These methods thrive on cache-oblivious, large memory hardware, and exhibit linear weak scaling (i.e., roughly the same rendering performance for small and large data) [4, 3].

- **In situ workbench computing.** With the power of these machines, small to midsize HPC jobs can be run directly on the workstation, and visualized directly in software such as SCIRun [2]. This will aid a variety of smaller-scale scientific studies.

- **No explicit scheduling.** Users can log into the resource anytime. User state is persistent, never lost when the session ends. To fully utilize computational resources, a backfill scheduler will operate as a persistent service on the operating system, adjusting backfill tasks to the current workload.

- **Turnkey application support.** As single machines driving a single OS, workstations will run a wide array of software in parallel. Current applications include data analysis tools (Mathematica, Matlab, Tableau), software renderers (Embree, Arnold, Renderman), and computational packages that encourage shared-memory and in situ visualization (NWChem, VASP).

- **Supercomputer filesystems.** The workstation will mount parallel filesystems (e.g., LUSTRE [8]) directly, sharing them with the supercomputer and enabling manipulation of large HPC data on-site. Thanks to parallel IO tools (e.g., Globus [1]), data movement across WAN will be similarly painless.

- **Out-of-core streaming.** For the largest HPC data, the workstation can function as both server and client for out-of-core and streaming methods such as Visus [6].

- **Vis labs with fewer megapixels, greater usability.** When driven by a single operating system, the need for a "control workstation" from which to drive a tiled display vanishes. The role of the visualization lab changes from a place for demos and tours, to a place to physically conduct work.

- **Remote visualization.** In addition to driving a physical vis labs, workstations can serve their users anytime, anyplace using video streaming technologies. As WAN network speeds improve, the separation between these workstations and thin clients (laptops, tablets) will vanish. However, the physical vis lab will remain useful for high-resolution, interactive use.

Running a single operating system, they will singlehandedly drive tiled displays and CAVEs, and simultaneously support seamless remote visualization with laptop, tablet and mobile devices. In many cases, these machines will be part of existing HPC landscapes, tethered directly to supercomputing filesystems or Globus tools. Their massive main memory will be used for co-processing of massive data using tools such as Tableau and Matlab, and real-time scientific rendering using ray tracing algorithms that thrive on cache-oblivious, large memory hardware.

3 The Talk

This talk will describe industry trends already underway that herald the rise of the workstation. We consider current "fat nodes" and large memory workstations already used in HPC facilities and visualization laboratories at the University of Utah, Texas Advanced Computing Center, and Argonne National Laboratory. We examine how they are used today for a variety of visualization workloads, and how they may improve in the future. We will describe the current workflows of visualization users: production visualization experts, materials scientists, and visualization researchers. Building on our experiences with tiled displays [4] and CAVE2 [7], we describe how single-OS environments might improve use and interactivity. While we note the potential of VR devices (Oculus Rift, touch), we are cognisant that much production visualization (and even interactive use) requires mouse and keyboard for effective navigation, and the best solutions for scientific visualization are those that enable command-line use, scripting, and programming. The focus of our talk will be on the *backend* hardware replacing desktops and desktop clusters.

4 Timeline

"Rebirth of the Workstation" envisions the next 30 years as follows:

- **now: "fat nodes" and multicore SMP workstations** Large memory workstations are prevalent in HPC infrastructure, enabling remote visualization for virtualized mouse-and-keyboard sessions.

- **1–2 years: manycore SMP workstations** New technology (Intel KNL, wide-SIMD Xeon CPUs) enable cost-effective shared-memory, dramatically increasing the computational power of single-OS solutions. The gulf between CPU and GPU performance narrows. Large-memory SMP nodes become cheaper and more powerful than ever, with a single multi-KNL workstation able to compete with a small GPU cluster in FLOPS, while offering a unified programming model.

- **1–5 years: new VR devices** The next generation of headsets (Oculus Rift) reaches mass-market. Remote visualization at low resolutions (1080p to 4K) looks effectively "as good" as a CAVE. Interaction will suffer due to latency; physical vis labs will still needed for interactive and collaborative work.

- **5+ years: effective distributed shared memory** The distinction between data-parallel and shared-memory will be blurred, with effective in-kernel remote memory access in GPU architectures such as NVIDIA Pascal. The question remains: single or distributed OS?

- **5–10 years: petascale multi-node SMP workstations.** Fueled by the advances in QPI and interconnects, vendors will start to offer compelling single-OS shared memory resources for the first time since the 2000's. Fat nodes will be merged together into a single virtual system, running a single OS, and delivering performance competitive with far-larger clusters. The divide between cluster and SMP computing will begin to vanish. Visualization algorithms will continue to prioritize memory over compute, ensuring demand for fat nodes.

- **7–10 years: CPUs and GPUs merge.** As parallel programming models evolve, the gap between CPU and GPU programming will narrow. Visualization will continue to depend on fast access to large memory – workstations and fat nodes will continue to replace small clusters in the institutional and small to midsize HPC market.

- **7–10 years: fiber everywhere.** Residential fiber will be commonplace in many cities in the US and Europe. Remote visualization, driving either laptop displays or headsets, is faster and more interactive than ever before. Physical installations will be limited to a handful of use cases where interactivity and super-high resolution are paramount.

- **10+ years: BOINC everywhere.** Thanks to fiber and improved internet backbones, distributed computing over WAN will enable HPC workloads over multiple university networks. New schedulers will enable SMP workstations to operate as fat nodes in these new environments, while simultaneously retaining their function as personal computers and remote vis / virtualization resources.

- **12+ years: "public cloud" of workstations.** The cloud will no longer be restricted to large providers such as Amazon, and HPC will no longer be comprised of a handful of large centers. Advances in operating systems will allow distributed and cloud computing to operate everywhere on any machine. All visualization, from the smallest to the largest scale, will be remote and on-demand.
- **15+ years: ubiquitous cloud computing.** All computing will be automatically distributed and parallel, we won't even know whose system we're operating on anymore. Augmented reality will be a part of everyday life.
- **20+ years: end of FLOPS as a metric.** Advancements in system-on-a-chip technology, 3D circuitry, and/or quantum computing will result in dramatically different hardware with different algorithmic needs. Computational throughput will be a non-issue compared to cost of data movement.
- **30+ years: "The Matrix".** The Connectome project will have successfully reverse-engineered the human brain. Neural interfaces will connect both visual output and interaction input. Privacy and personal space will be paramount: standalone workstations will be rare, highly specialized, and highly valuable, not unlike aging mainframes today.

5 Discussion

In all, the workstation promises to make remote visualization of large data commonplace, and no less painful than using a laptop. The vis lab will be transformed from an environment for tech demos to a usable interactive workspace. Our perspective is slanted towards the needs of HPC visualization practictioners, and chiefly considers the changes in *computing hardware* over the next several years. We look forward to a discussion with the broader HCI and VR communities on how visual and haptic interfaces will change, and what that portends for a potential "rebirth of the workstation."

References

[1] I. Foster. Globus online: ccelerating and democratizin science through cloud-based services. *IEEE Internet Computing*, 15(3):70–73, 2011.

[2] S. Institute, 2014. SCIRun: A Scientific Computing Problem Solving Environment, Scientific Computing and Imaging Institute (SCI), Download from: http://www.scirun.org.

[3] A. Knoll, I. Wald, P. Navratil, A. Bowen, K. Reda, M. E. Papka, and K. Gaither. Rbf volume ray casting on multicore and manycore cpus. In *Computer Graphics Forum*, volume 33, pages 71–80. Wiley Online Library, 2014.

[4] A. Knoll, I. Wald, P. A. Navrátil, M. E. Papka, and K. P. Gaither. Ray tracing and volume rendering large molecular data on multi-core and many-core architectures. In *Proceedings of the 8th International Workshop on Ultrascale Visualization*, page 5. ACM, 2013.

[5] S. G. Parker, J. Bigler, A. Dietrich, H. Friedrich, J. Hoberock, D. Luebke, D. McAllister, M. McGuire, K. Morley, A. Robison, et al. Optix: a general purpose ray tracing engine. In *ACM Transactions on Graphics (TOG)*, volume 29, page 66. ACM, 2010.

[6] V. Pascucci, G. Scorzelli, B. Summa, P.-T. Bremer, A. Gyulassy, C. Christensen, S. Philip, and S. Kumar. The ViSUS visualization framework. In E. W. Bethel, H. C. (LBNL), and C. H. (UofU), editors, *High Performance Visualization: Enabling Extreme-Scale Scientific Insight*, Chapman and Hall/CRC Computational Science, chapter 19. Chapman and Hall/CRC, 2012.

[7] K. Reda, A. Knoll, K. Nomura, M. Papka, A. Johnson, and J. Leigh. Visualizing large-scale atomistic simulations in ultra-high resolution immersive environments. In *IEEE LDAV*, 2013.

[8] P. Schwan. Lustre: Building a file system for 1000-node clusters. In *Proceedings of the 2003 Linux Symposium*, volume 2003, 2003.

[9] StarNet, 2014. FastX: A Simple X Windows terminal emulator: http://www.starnet.com/fastx.

[10] I. Wald, S. Woop, C. Benthin, G. S. Johnson, and M. Ernst. Embree: A kernel framework for efficient cpu ray tracing. 2014.