Software Rasterizer (SWR)

Timothy Rowley, Graphics Software Engineer, Parallel Visual Engineering

Software Rasterization

A Software Rasterizer for OpenGL®

Timothy Rowley - Graphics Software Engineer, Parallel Visual Engineering

Abstract: Introduction to a new open source project – a software implemented OpenGL® pipeline. A combination of reasons in the scientific and large-scale visualization fields make such a component interesting in comparison to GPU or Mesa3D based solutions. This talk will discuss the motivations for this effort, the current state of implementation, and some preliminary benchmark results.

Bio: Tim joined Intel in 2009 to work on the Larrabee graphics architecture, starting on the OpenGL® driver and then moving down the stack to rendering pipeline and shader compiler. Since Larrabee he has developed on the range of Intel graphics solutions. Prior work includes a stint at PowerVR® and project lead for Mozilla Scalable Vector Graphics (SVG).

Agenda

- Motivation: Scientific and Large-scale visualization problems
- Overview of existing implementations
- Solution: SWR (Software Rasterizer)
- Current project status
- Preliminary benchmark results
- Wrap-up
- Q&A

Motivation: Scientific and Large-scale Visualization Problems

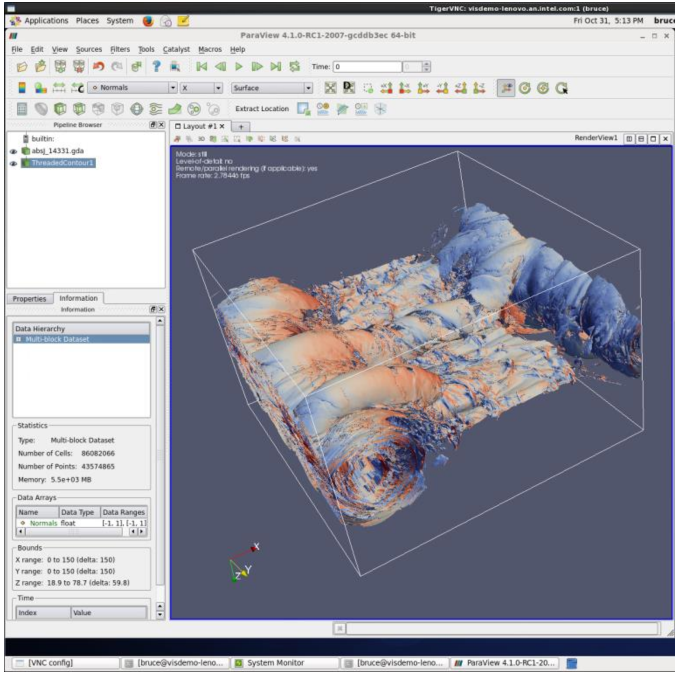

- **Interactive visualization is challenging!** Often requiring a creative workflow and a lot of patience
- **Extremely large datasets are becoming common**, demanding more advanced resources and techniques
- Advanced rendering techniques can be employed, but usually require software rewrite
- Strong **desire to use existing industry-standard software packages** with little or no changes

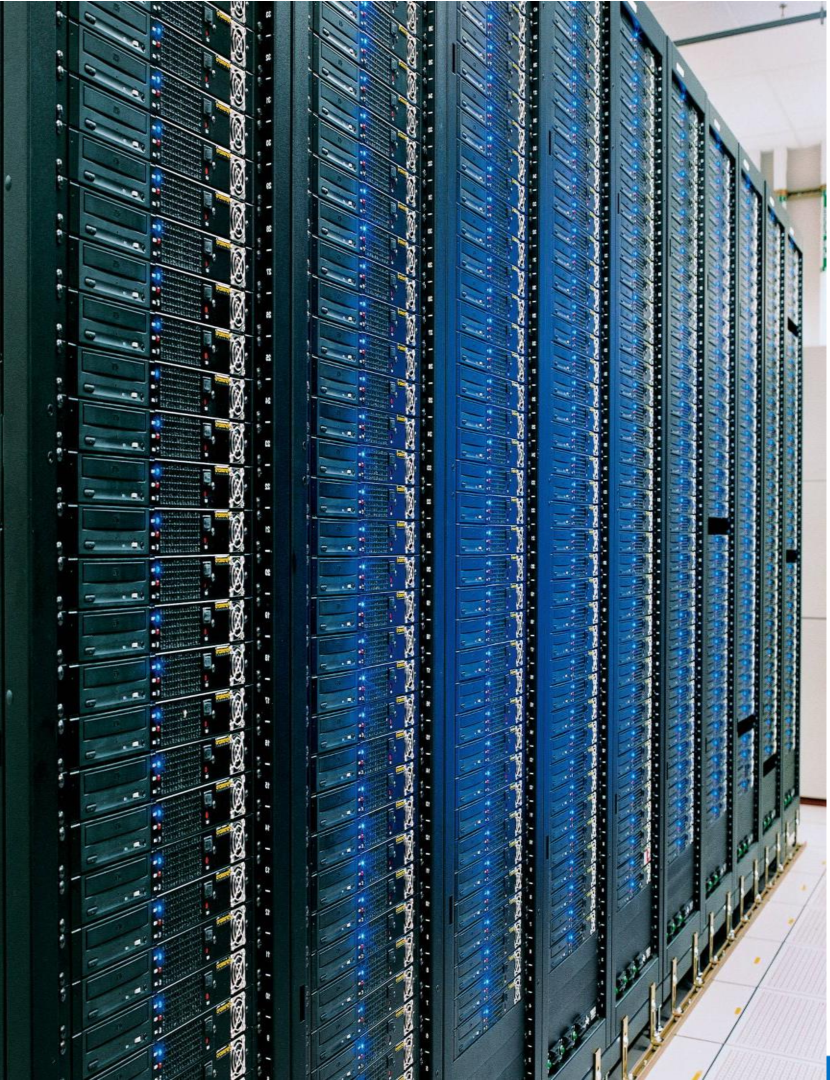

Clusters can be Better Utilized

- HPC clusters have enormous amounts of CPU power
- But GPUs are generally scarce or nonexistent
- Even with GPUs, the dataset is too large to fit in GPU memory
- Software rendering has access to the entire host memory

Overview of existing implementations

Mesa3D

- The standard software OpenGL® implementation
- Vertex throughput is single threaded
- Limited scaling of pixel backend

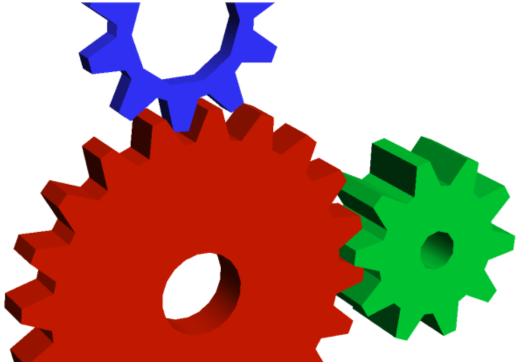

Larrabee (many-core software rendering)

- Targeted different workloads (more emphasis on shading, than geometry)
- Targeted games and consumer-level content
- Targeted unique Larrabee architecture, not Intel big core

Solution: SWR

- Fully scalable OpenGL® taking advantage of SSE, Intel® AVX, Intel®AVX2
- Implemented in C++, x86 intrinsics, and JIT (LLVM) code
- Completely runs on host CPU
- Off-the-shelf software packages just work
- Existing rendering distribution solutions
  can be used to scale across
  compute capacity

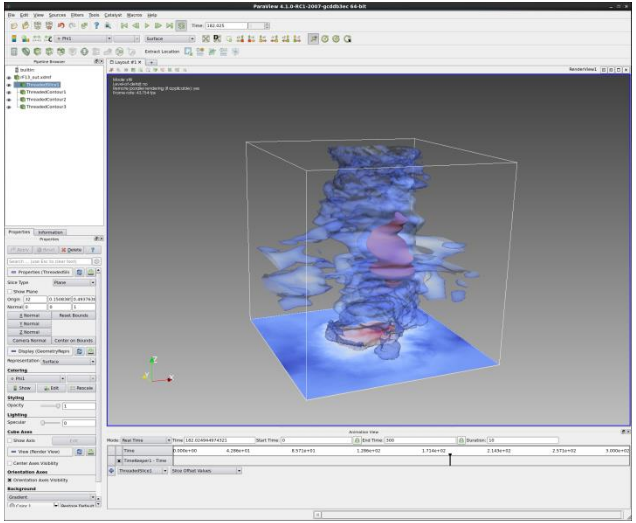

- Intel® Advanced Vector Extensions (Intel® AVX/ Intel®AVX2): Intel® AVX/AVX2 is designed to achieve higher throughput in certain integer and floating point operations. Depending on processor power and thermal characteristics, and system power and thermal conditions, AVX/AVX2 floating point instructions may run at lower frequency to maintain reliable operations at all times. For further details see product data sheet.

SWR: Overview

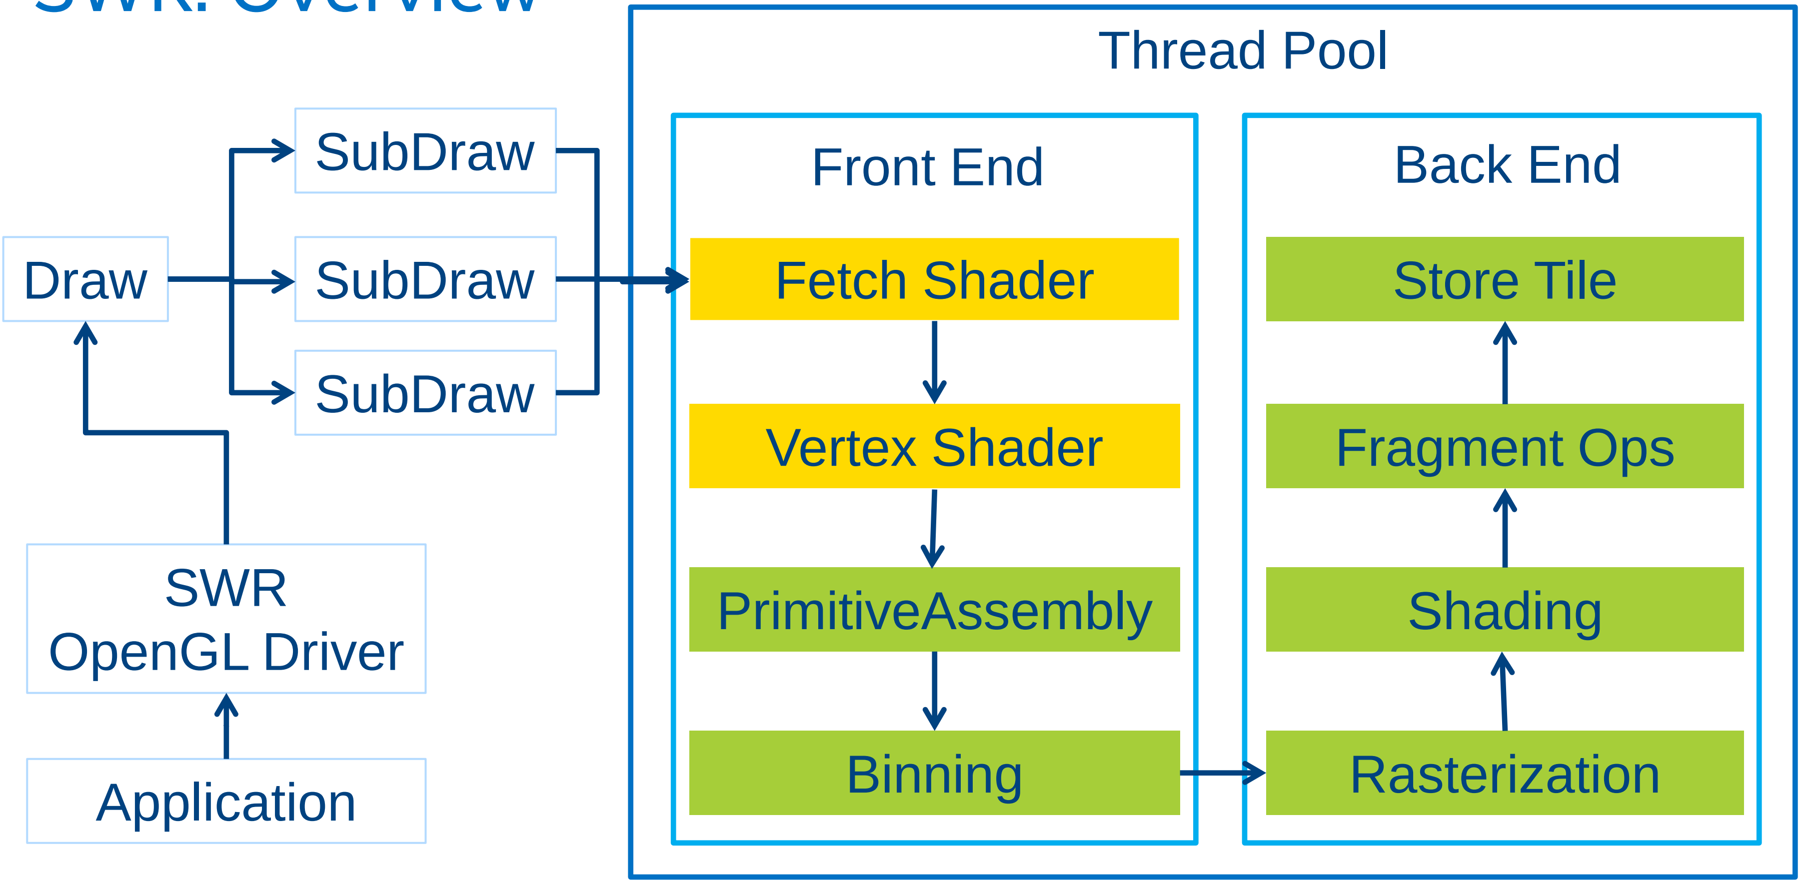

SWR: Key Concepts

- High-performance rendering core with an OpenGL® driver
- Thread pool which dynamically chooses front-end (vertex) or backend (fragment) work
- Large draws are split to allow parallelism, ordering enforced in backend
- JIT code for both array-of-structures to structure-of-arrays conversion and vertex shading

SWR: Usage

- Drop-in OpenGL®
  - Shipped as a shared library
  - Set LD_LIBRARY_PATH and run normally
- Also supports OSMesa API
  - Standard off-screen API from Mesa3D
  - One extra API for flushing output

SWR: Current Status

- OpenGL® 1.4 class implementation
  - Not complete API coverage – focused on optimizing existing visualization applications
  - Old style OpenGL®, fixed function
- ParaView 4.2 and VisIt 2.8.1 work both single-node and distributed

SWR: Current Status

ParaView

- FIU dataset
- One contour with clipping plane and 4,000 streamlines.
- 47 million triangles.
- Performance on Stampede is 2 fps on 1-node and 7 fps on 32-nodes.

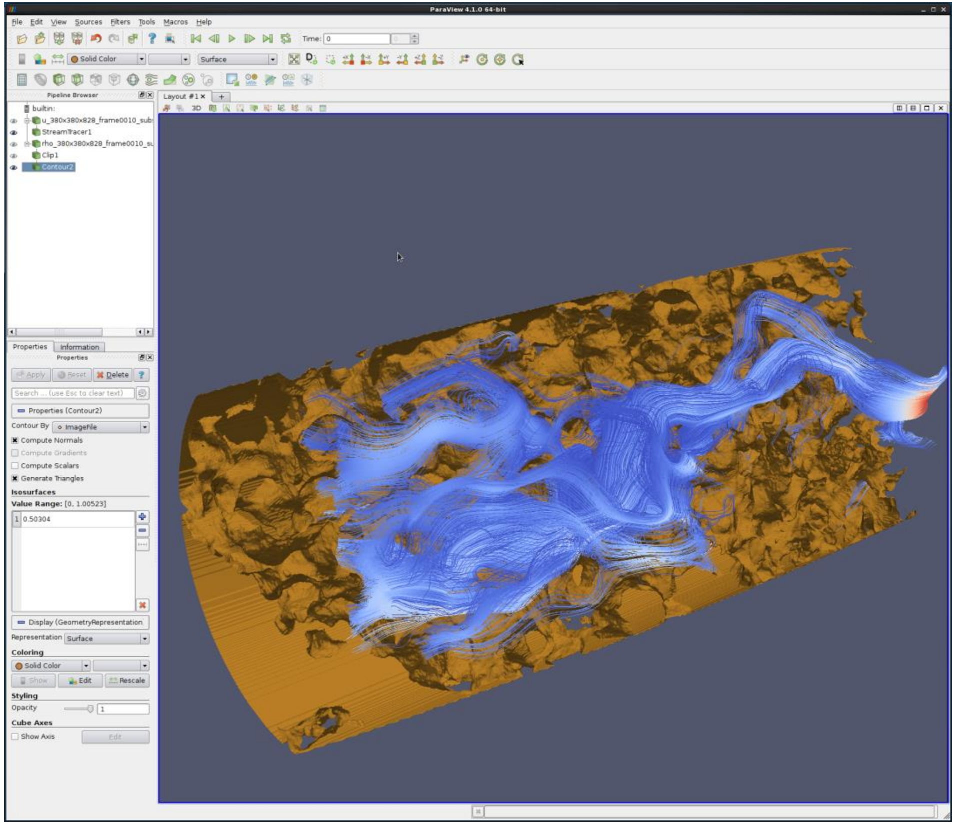

Results measured on TACC Stampede cluster (dual-socket Intel Xeon® E5-2680, 32GB per node) and based on internal Intel analysis and are provided for informational purposes only. Any difference in system hardware or software design or configuration may affect actual performance.

SWR: Current Status

- Sample noise.silo dataset
- 15 contour surfaces with a clipping plane
- Performance on Core™ i7-4790 workstation is 16 fps @ full-screen 2560x1600.

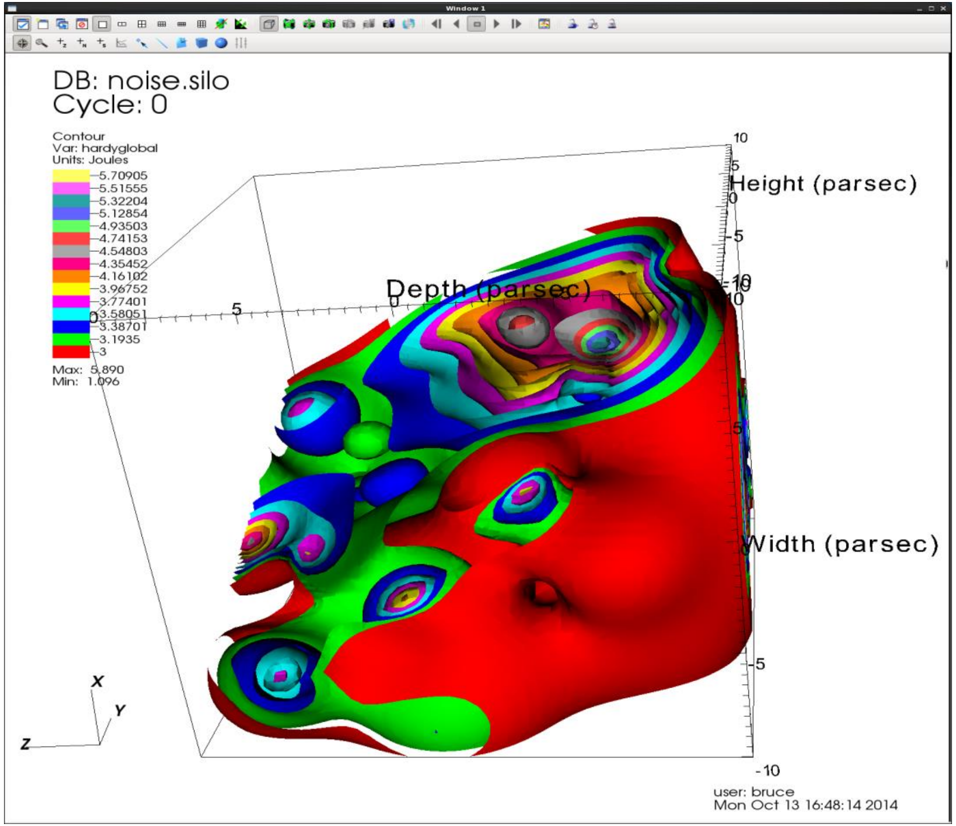

Results measured on an Intel Core™ i7-4790 PC with 32GB) and based on internal Intel analysis and are provided for informational purposes only. Any difference in system hardware or software design or configuration may affect actual performance.

SWR: Performance Scaling

- SWR performance scales by adding threads on a single host

VizOS (internal benchmark rendering 19 million triangles)

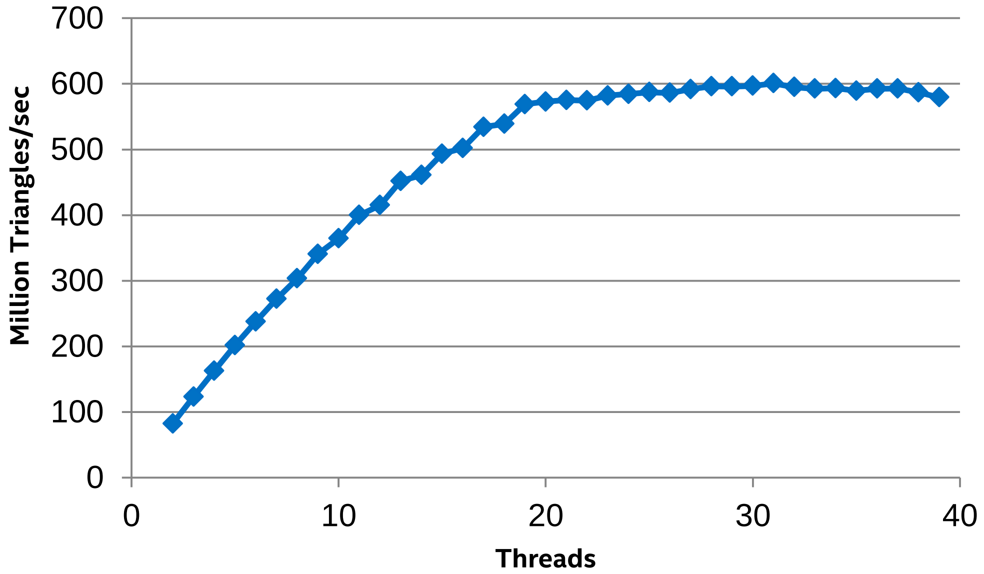

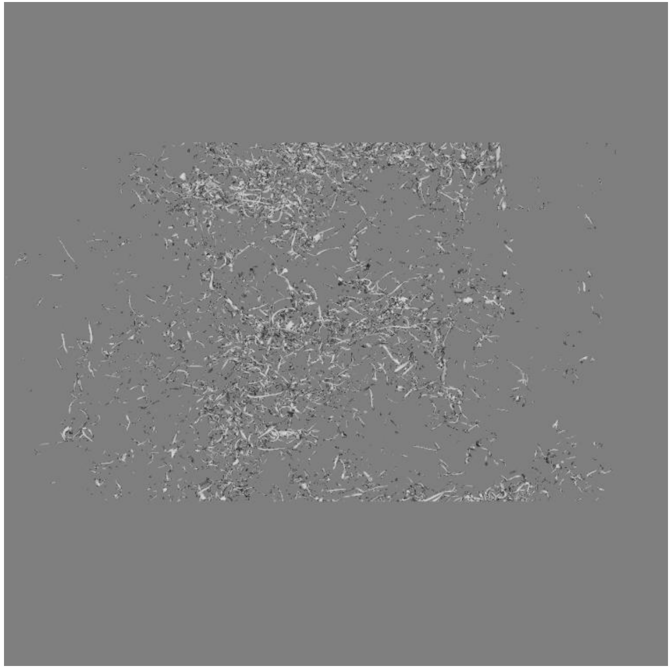

Results measured on a dual-socket Intel Xeon® E5-2687W v3 @ 3.10GHz with 2x10 cores, 128GB) and based on internal Intel analysis and are provided for informational purposes only. Any difference in system hardware or software design or configuration may affect actual performance.

SWR: Performance Scaling

- Or adding processors in a cluster

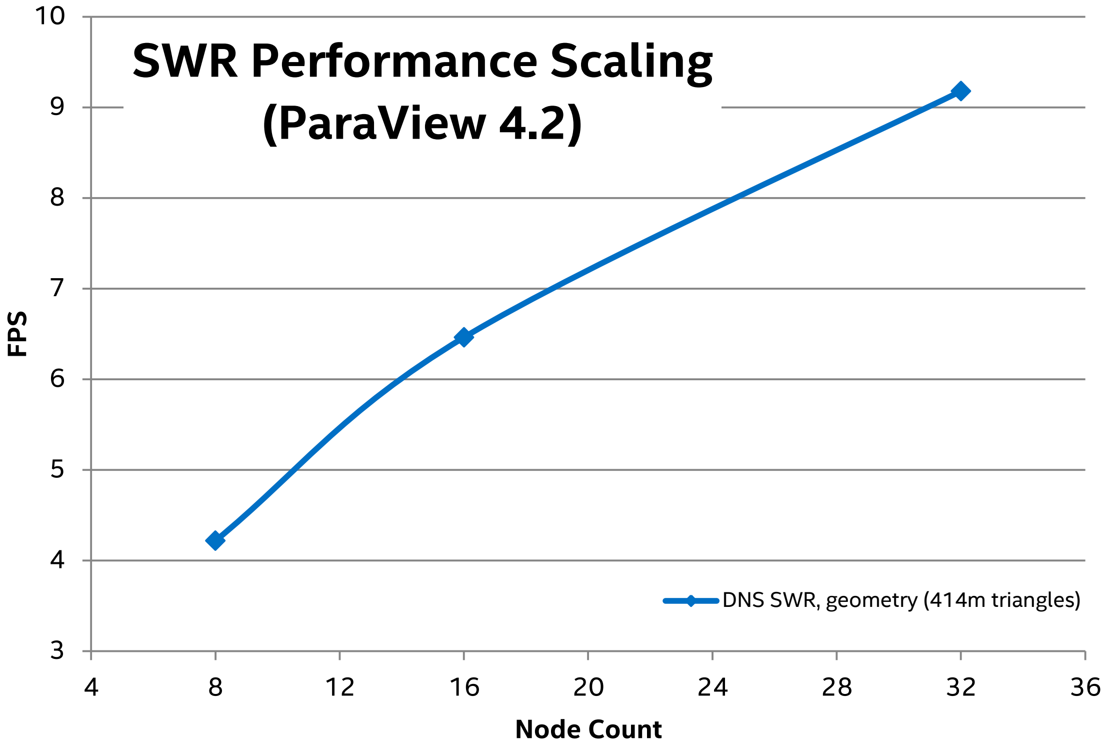

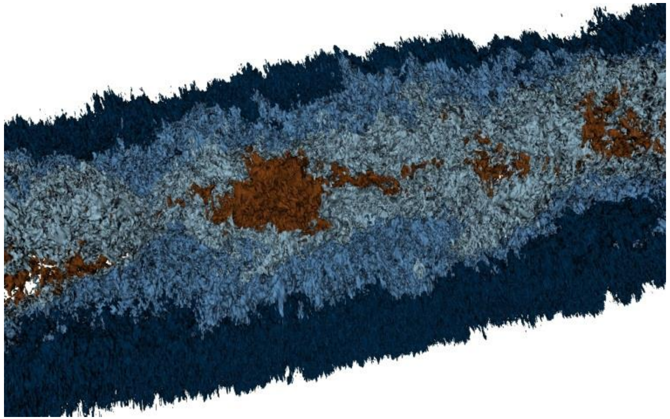

Results measured on TACC Stampede cluster (dual-socket Intel Xeon® E5-2680, 32GB per node) and based on internal Intel analysis and are provided for informational purposes only. Any difference in system hardware or software design or configuration may affect actual performance.

SWR: Performance Competition

- Studies on Stampede are very favorable

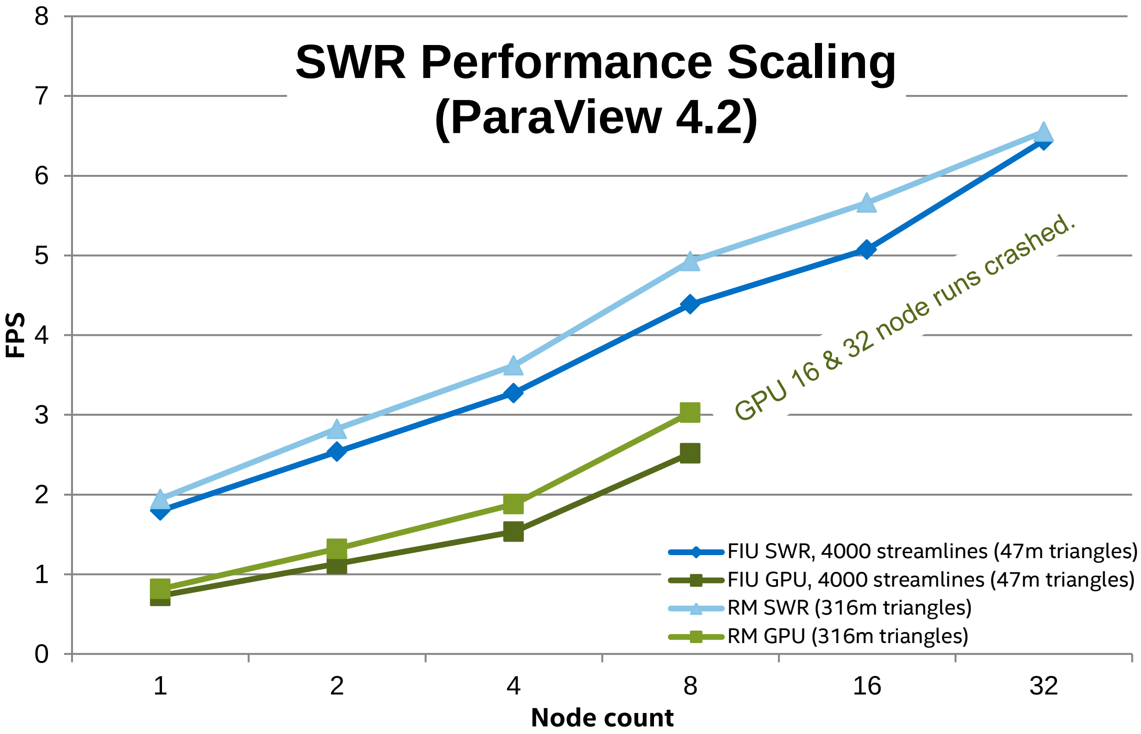

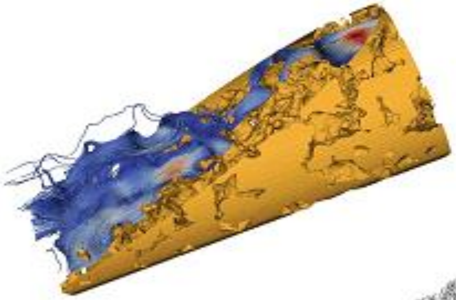

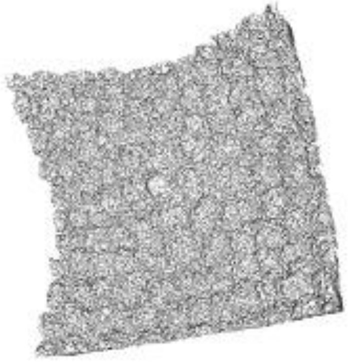

Results measured on TACC Stampede cluster (dual-socket Intel Xeon® E5-2680, 32GB per node, NVIDIA K20 discrete GPUs, 5GB) and based on internal Intel analysis and are provided for informational purposes only. Any difference in system hardware or software design or configuration may affect actual performance.

SWR: Future

- OpenGL® 2.0+
  - Shaders
  - Coordinated schedule to meet release of VTK-next
- JIT more of the pipeline
- Support for AVX-512
- More performance tuning
- More features as workloads demand
- Non-traditional rasterization acceleration

SWR: Obtaining

- Open source
- Apache 2 license
- Targeted release December 2014
- http://openswr.github.io/

Summary

- Interactive visualization of extremely large datasets is driving demand for more advanced resources and rendering techniques
- HPCs have large amounts of CPU power, but are sometimes limited on GPU resources
- SWR is a new open-source highly tuned OpenGL® software rasterizer that supports existing visualization tools
- SWR OpenGL® 2.0+ will be released in coordination with the next major VTK toolkit update

Q&A

Legal Disclaimer

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

A "Mission Critical Application" is any application in which failure of the Intel Product could result, directly or indirectly, in personal injury or death. SHOULD YOU PURCHASE OR USE INTEL'S PRODUCTS FOR ANY SUCH MISSION CRITICAL APPLICATION, YOU SHALL INDEMNIFY AND HOLD INTEL AND ITS SUBSIDIARIES, SUBCONTRACTORS AND AFFILIATES, AND THE DIRECTORS, OFFICERS, AND EMPLOYEES OF EACH, HARMLESS AGAINST ALL CLAIMS COSTS, DAMAGES, AND EXPENSES AND REASONABLE ATTORNEYS' FEES ARISING OUT OF, DIRECTLY OR INDIRECTLY, ANY CLAIM OF PRODUCT LIABILITY, PERSONAL INJURY, OR DEATH ARISING IN ANY WAY OUT OF SUCH MISSION CRITICAL APPLICATION, WHETHER OR NOT INTEL OR ITS SUBCONTRACTOR WAS NEGLIGENT IN THE DESIGN, MANUFACTURE, OR WARNING OF THE INTEL PRODUCT OR ANY OF ITS PARTS.

Intel may make changes to specifications and product descriptions at any time, without notice. Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined". Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them. The information here is subject to change without notice. Do not finalize a design with this information.

The products described in this document may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or go to: http://www.intel.com/design/literature.htm

Larrabee and other code names featured are used internally within Intel to identify products that are in development and not yet publicly announced for release. Customers, licensees and other third parties are not authorized by Intel to use code names in advertising, promotion or marketing of any product or services and any such use of Intel's internal code names is at the sole risk of the user.

Intel, Intel AVX, Intel AVX2 and the Intel logo are trademarks of Intel Corporation in the United States and other countries.

*Other names and brands may be claimed as the property of others.

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel.

Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Notice revision #20110804