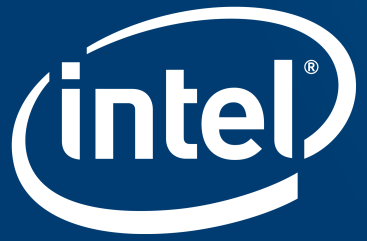

IXPUG SDVis In-Situ Workshop

Welcome and Intel SDVis Update

May 22, 2017

Jim Jeffers

Special Thanks!

Dan “The Man” Stanzione – Director, TACC

Joe ”The Boss” Curley – GM, Code Modernization and SDVis, Intel

Melyssa “The Glue” Franklin – Industry Programs Director, IXPUG and TACC

Lisa “The Bucks” Smith – Director, Ecosystem Programs, Intel

Ben “The Sweat” Han – Ecosystem Programs and IXPUG, Intel

Valerie ”The Food!” Wise – Sr. Events Program Coordinator, TACC

SDVis Quick Recap

http://software.intel.com/sdvis www.sdvis.org

Intel supported community effort enabling high performance, high fidelity and in-situ CPU based rendering

- High, interactive performance for even very large (TB+) datasets
- Fully Open Source with liberal Apache 2 and MIT License
- Single Node and Cluster-wide Scalability in 1) Render Time; 2) Render Quality; 3) 3D Model Data Size
- OpenGL and Ray Tracing Support with shadows, ambient occlusion, up to photorealistic quality

Now broadly Integrated and tested with both general and targeted domain applications

ParaView

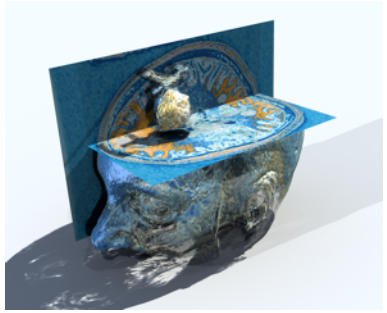

Data courtesy Kitware. Visualization, Carson Brownlee, Intel

VisIt

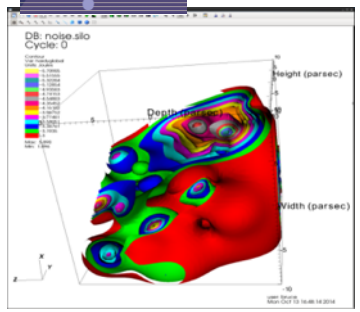

Courtesy Hank Childs, U Oregon, Jian Huang and Alok Hota, UTenn

EnSight
Extreme Visualization Software

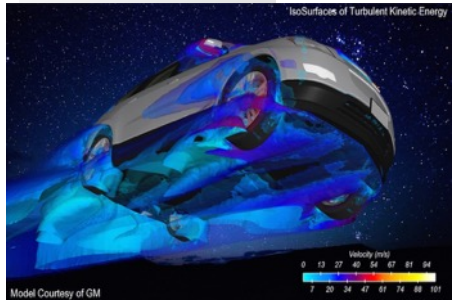

Courtesy Sean Ahern, CEI and General Motors

VMD
Visual Molecular Dynamics

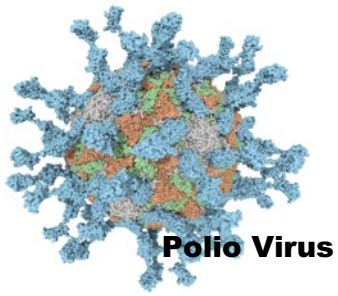

Courtesy John Stone, Beckman Institute, Univ. Illinois at Urbana-Champaign

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

Brayns

Data and Visualization courtesy Cyrille Favreau, EPFL

Argonne NATIONAL LABORATORY

VL3

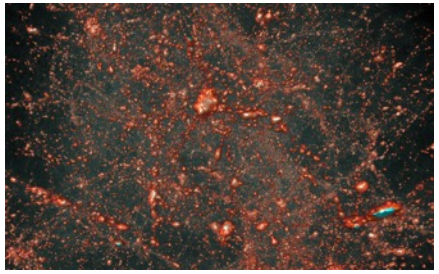

Data: Salman Habib, Katrin Heitmann, and the HACC team.
Visualization: Joe Insley, Slivio Rizzi, ANL

Mesa/OpenSWR overview

OpenGL compatibility

- Compatibility profile 3.0
- Core profile 3.3
- ES profile 3.0

Drop-in replacement, no code changes or recompilations needed

- Set a couple environment variables (LD_LIBRARY_PATH, GALLIUM_DRIVER) and run!

More info…

- www.mesa3d.org; www.openswr.org ; openswr@googlegroups.com

OpenSWR Releases

Mesa 12.0.0

- 8 July 2016
- Initial officially integrated upstream OpenSWR release

Mesa 13.0.0

- 1 November 2016
- Initial tuning of SWR integration for HPC workloads
- SWR core optimizations

Mesa 17.0.0

- 13 February 2017
- OpenGL conformance fixes

Mesa 17.1.0

- *Just Released 9 May 2017*
- Geometry shader support
- True MSAA available as an experimental feature

Benchmarks

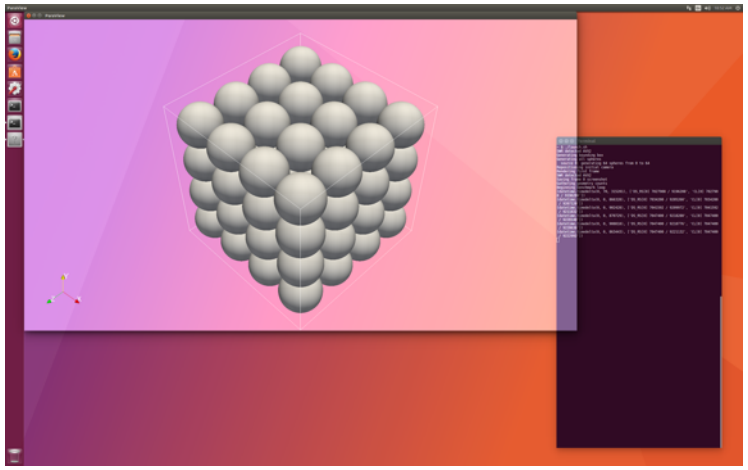

manyspheres.py
67 MiPolys

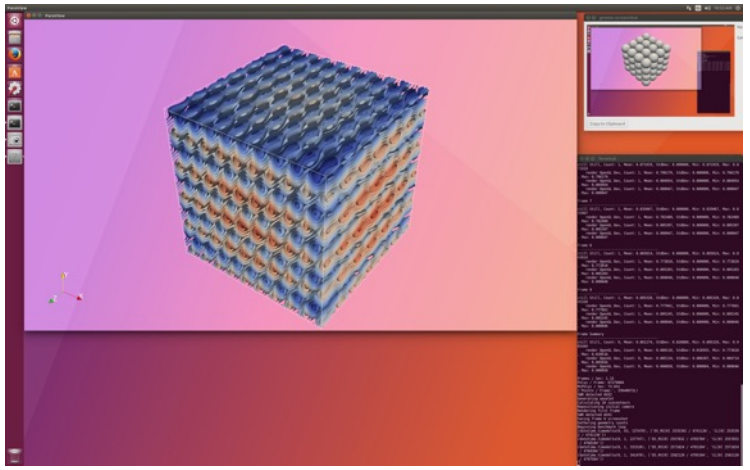

wavelets.py
11 MiPolys

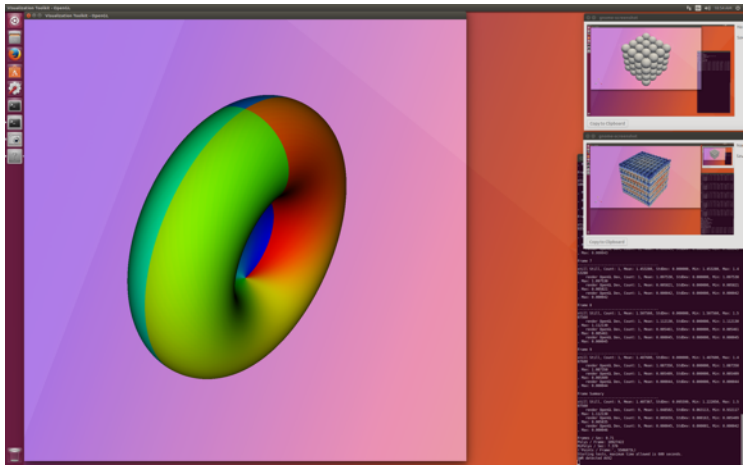

TimingTests
30 MiTris

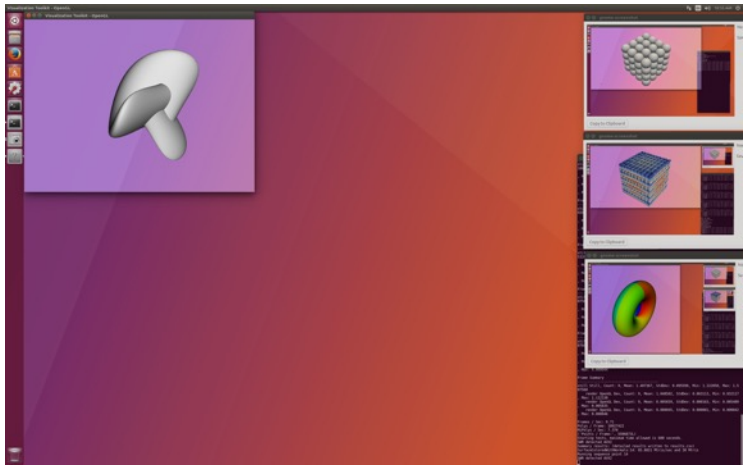

GLBenchmarking
30MiTris

KNL SWR/LLVMPIPE PERFORMANCE RATIO

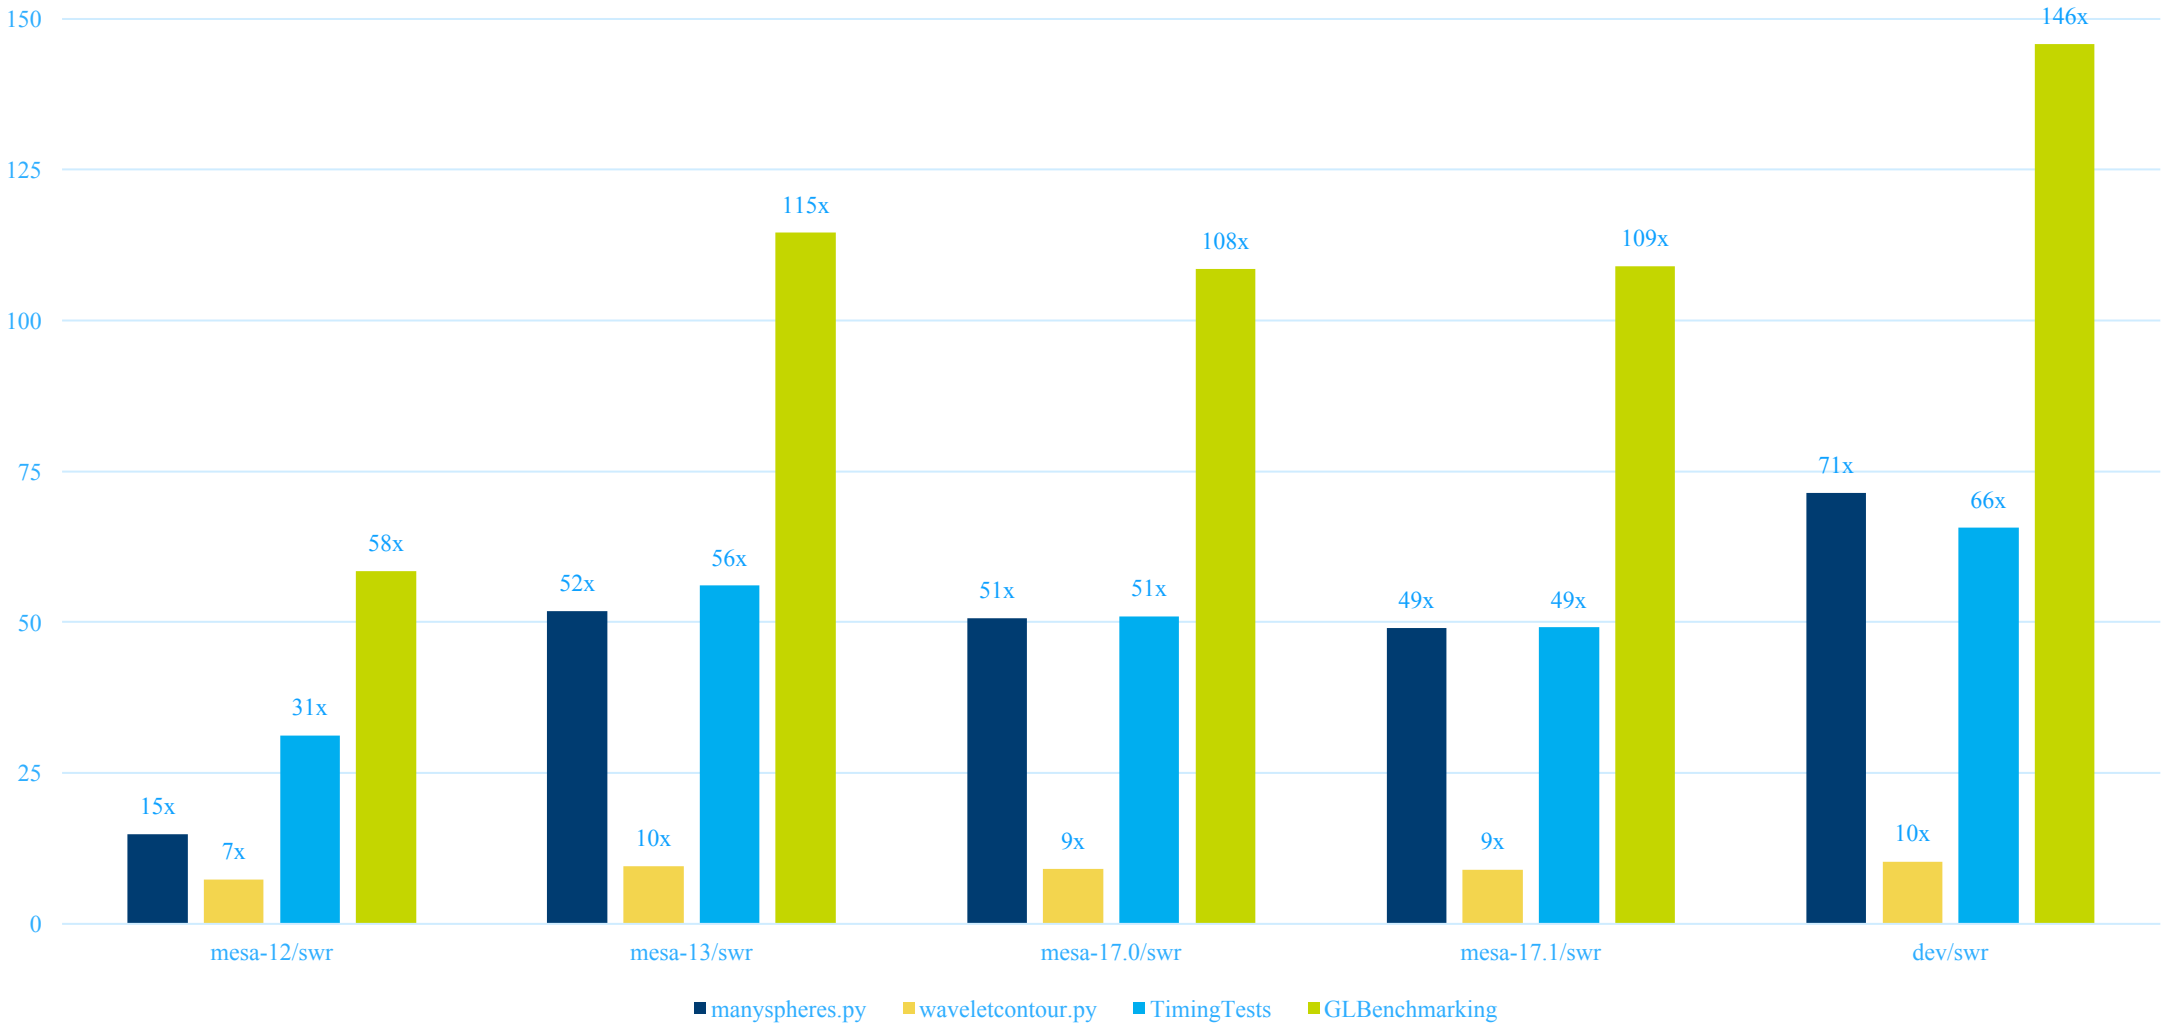

Applications

Kitware/VTK/ParaView

- Close collaboration
- ParaView ships with OpenSWR (run with “—mesa-swr”)

VisIt

- OpenSWR works with current version
- Looking forward to them moving to VTK 7.0 or later

Ensight

- OpenSWR works with current version
- Looking forward to their rendering rewrite

VL3 (from Argonne National Laboratory)

- Required implementing point sprites, otherwise OpenSWR dropped in

With our OpenGL compatibility, we’re fairly confident in supporting new customers without heroics

Future development

Features done pending customer/partner requests

MSAA to be finished off and promoted to non-experimental

Build changes

- Improve integration of swr to reduce disk space
- Add KNL and SKX specific swr library builds

Performance, performance, performance

- Dynamically pick the frontend parallelization split based on machine topology
- Widen pipe to 16-wide for AVX512
- Use KNL AVX512ER instructions for shader math
- Add back early rasterization
- Allow early-Z testing to handle z-write shaders (ParaView targeted optimization)
- Optimize MSAA triangle rejection
- Investigate better lane utilization in backend
  - Small triangles (typical HPC visualization workloads) results in tiny triangles = poor lane utilization
  - Quad-packing or pixel-packing triangles (book keeping overhead, render ordering)

test configuration

KNL: Ninja Developer Platform Pedestal, Intel® Xeon Phi™ CPU 7210 @ 1.30GHz, Ubuntu 17.10, 4.10.0-20-generic, gcc 6.3.0

BDW: Intel® Xeon 2699v4 x 2 (44 cores) @ 2.20GHz, Ubuntu 17.10, 4.10.0-19-generic, gcc 6.3.0

Mesa-12: mesa-12.0.6/llvm-3.9.1

Mesa-13: mesa-13.0.6/llvm-4.0

Mesa-17.0: mesa-17.0.4/llvm-4.0

Mesa-17.1: mesa-17.1.0-rc1/llvm-4.0

Development: mesa-master(April 22)/llvm-4.0, -march=knl, patches scheduled to land, KNOB_MAX_DRAWS_IN_FLIGHT=256

pvpython manyspheres.py -s 64 -r 726 -v 1920,1080

pvpython waveletcontour.py -d 256 -v 1920,1080

TimingTests -width 1536 -height 1536 -regex SurfaceColoredWithNormals -nochart -ss 14 -se 14

GLBenchmarking --start 14 --end 14

Ray Tracing Update

Embree and OSPRay

Ray Tracing Foundation: Embree Ray Tracing Kernel Library

Provides highly optimized and scalable ray tracing kernels

- Acceleration structure build and ray traversal
- Single Ray, Ray Packets(4,8,16), Ray Streams(N)

Targets up to photorealistic professional and scientific rendering applications

Highest ray tracing performance on CPUs

- 1.5–6× speedup reported by users

Support for latest CPUs / ISAs

- Intel® Xeon Phi™ Processor (codenamed *Knights Landing*) – AVX-512
- Intel® Xeon® Processor (codenamed Skylake) - AVX-512 (coming soon!)

API for easy integration into applications

Free and open source under Apache 2.0 license

- http://embree.github.com

Embree Adoption* - 50+ Studios, Professional Rendering, Gaming Tools, and SIM Apps

UBISOFT MOTION PICTURES

AUTODESK.

DREAMWORKS

FluidRay RT
Physically Based Renderer

SURVICE ENGINEERING COMPANY

Simlab Soft
Bringing Art closer

corona

3DEXCITE

Courtesy of Jeff Patton, Rendered with Corona Renderer

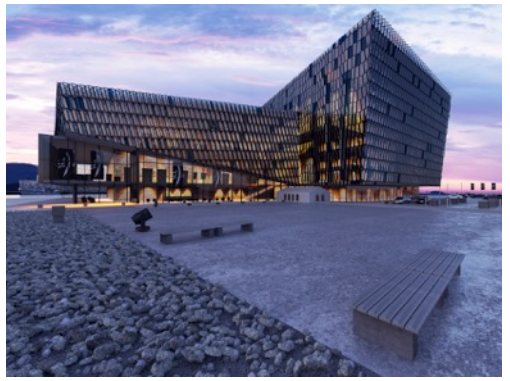

Image rendered with FluidRay RT

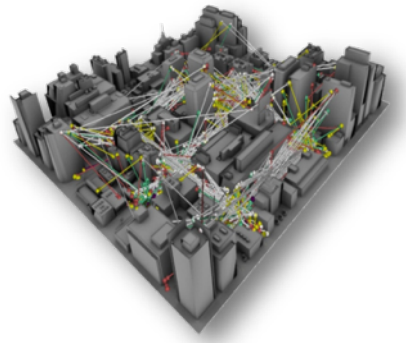

Rendered with StingRay, SURVICE Engineering

pCon.planner rendered courtesty EasternGraphics

*Many other announced users incl.: Pixar, Weta Digital, Activision, Chaos V-Ray, Ready At Dawn, FrostBite, EpicGames UnReal, High Moon, Blue Sky, UBISoft MP, Framestore, Illumination, ….

Performance: Embree vs. NVIDIA* OptiX*

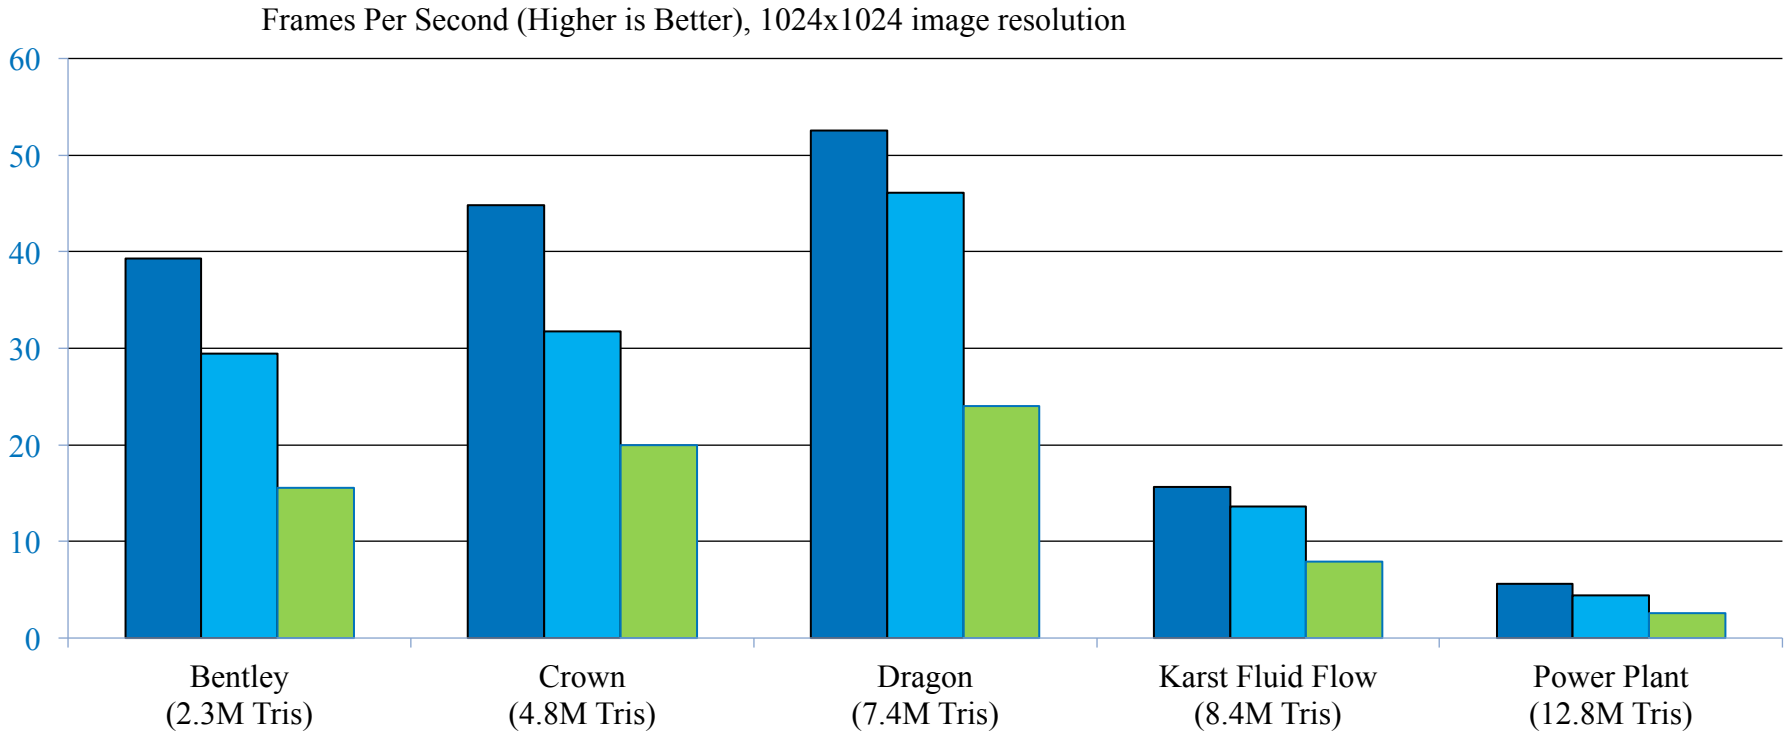

Intel® Xeon® E5-2699 v4 Processor
2 x 22 cores, 2.2 GHz

Intel® Xeon Phi™ 7250 Processor
68 cores, 1.4 GHz

NVIDIA Tesla P100 Coprocessor
PCIe, 16 GB RAM

Embree 2.16.0-alpha, Intel® C++ Compiler 17.0.2, Intel® SPMD Program Compiler (Intel® ISPC) 1.9.1

NVIDIA* OptiX* 4.0.2, CUDA* 8.0.44

Source: Intel

OSPRay: A Ray-Tracing based Rendering Engine for *High-Fidelity* Visualization

www.ospray.org

- Build on top of Embree; Launched June 2016

- Scalable Visualization targeted features
  - Surfaces (both polygonal and non-polygonal)
  - Volumes, and volume rendering
  - *High-Fidelity* rendering/shading methods
  - *Scalable* Cluster Wide Rendering

- Packed it up in an 'easy-to-use' rendering library for visualization
  - Same "spirit" as OpenGL, but different API

- 10+ Application adoption in 9 months, more under development!

Brayns

PowerCT

NASA

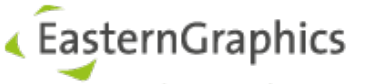

VL3

OSPRAY FEATURES v1.2

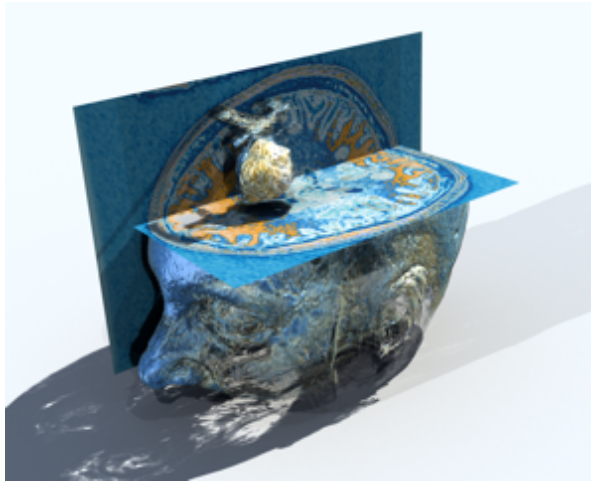

What OSPRay can do today:

- Ray tracing based scientific visualization
  - Extendable via in-source ‘modules’ and/or installed SDK
  - Great performance on both Intel® Xeon®, Intel® Xeon Phi™ Processors
    - Embree for very fast ray intersections
    - TBB for great thread scalability
    - Intel® SPMD Program Compiler (ISPC) for rendering/shading good vectorization

- Many ‘built-in’ rendering features
  - Geometry types (+instancing): Cylinders, Spheres, Streamlines, Triangles, Implicit isosurfaces
  - Structured volumes (regular grids) with configurable transfer functions
  - Various cameras, lights, materials, etc
  - Ambient Occlusion, Shadows, Progressive Refinement,….
  - Simple renderers (raycast, AO), to more advanced renderers (scivis, pathtracer)

- Multiple ways to scale rendering compute
  - Local rendering (single application’s CPUs + memory)
  - MPI cluster offload (single application offloading to N worker nodes via MPI)

Future development

Gravational Waves : GR-Chombo AMR Data, Stephen Hawking CTC, UCambridge; Queens College, London; Visualization, Carson Brownlee, Intel, ParaView)

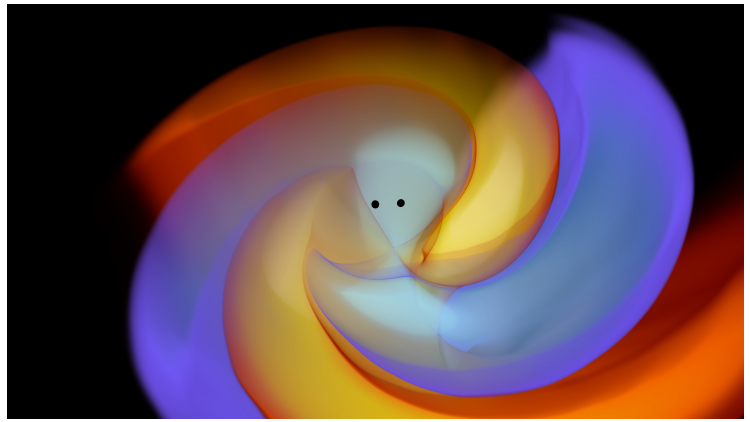

OSPRay's 2017 Feature Roadmap:

- Distributed API for MPI applications (Q2)
  - Built for in-situ rendering of running simulations
  - Both data-distributed geometry and volumes
- AMR "Chombo" volumes (Q2-early Q3)
  - Native rendering support with ParaView integration

Early Access This Week!

- Unstructured volume support (Q3-Q4)
  - Primarily focused on natively rendering Hex and Tet meshes
- Various ongoing feature requests (and of course bug fixes!):
  - Support for instancing volumes
  - Intel® Xeon® Processor (codenamed Skylake) AVX-512 optimization

Whose Here?

Sim + Vis Teams Working Together!

Addressing Data Visualization with Intel® Solutions

Open Source Libraries Optimized by Intel®

- Optimized for parallel processing and latest instruction sets
- OpenSWR, Embree, & OSPRay available today!
- Used by ParaView, VisIt, VMD, CEI EnSight and more…

Intel® Xeon Phi™ Provides Better Performance for Visualization

- Provides up to 72 cores / CPU
- Addresses up to 384GB memory vs 16GB GPUs
- AVX-512 Instruction Optimiized

Realize Compelling Value

- Cost of Host vs cost of Host + Card for GPUs
- Plus -> Increased performance for visualization
- Additional use as general purpose compute platform

For in-situ, post-processing, and professional rendering visualization needs

ANNOUNCING…. The Turnkey SDVis Appliance!!

Appliance Configuration

Installed Software Intel® HPC Orchestrator, SDVis Libraries (Embree, OSPRay, Mesa/OpenSWR), Open Source Apps (ParaView, VTK, VisIt, VMD), Intel Parallel Studio Cluster Edition, SW Dev. Tools

Nodes 8x Intel® Xeon Phi™ 7250 compute nodes, Intel® Xeon ® E5-v4 head node + 6 monitor display card, Intel® Xeon ® E5-v4 storage node

Storage 32TB Raid

Network 24 port Omni-Path & Ethernet switches

Targeting Launch Mid-Year

Ordering & More Info

Ordering: sales@colfax-intl.com

Information: http://sdvis.xeonphi.com/

SDVis Appliance	
Master Head Node Node + Display Management	**Software:** CentOS 7.3, HPC Orchestrator
	Hardware: 1 Intel® Xeon ® Node Intel® Xeon ® E5-v4 36 cores, 256GB DDR4, OPA HFI, Up to 6 4K Monitors Display Card, SSD 480GB
Compute Nodes Compute + Render	**Software:** CentOS 7.3, SDVis Software (ParaView, VTK, VisIt, VMD), Intel Parallel Studio Cluster Edition, SW Dev. Tools
	Hardware: 8 Intel® Xeon Phi™ Nodes Intel® Xeon Phi™ 7250 68 cores, 192GB DDR4, OPA HFI, SSD 150GB
Storage System	**Software:** CentOS 7.3
	Hardware: 1 Intel® Xeon ® Node Intel® Xeon ® E5-v4 16 cores, 64GB DDR4, OPA HFI, 32TB RAID
Networking OPA 100G Ethernet	**Hardware:** 24 Port OPA Edge Switch, 24 Port Ethernet Switch

Pre-configured solution for visualization apps and development needs!

ENJOY THE WORKSHOP!

THE NEW CENTER OF POSSIBILITY

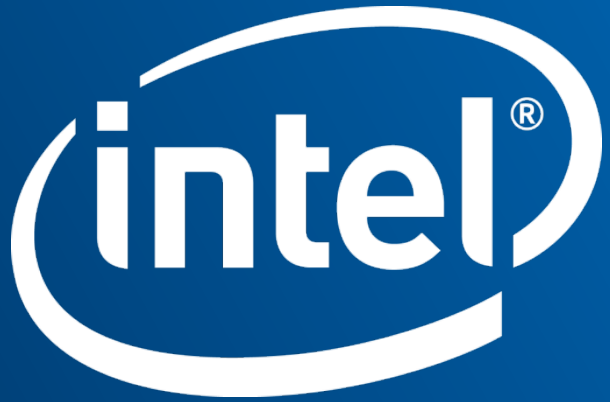
intel