Assessing and Improving Large Scale Parallel Volume Rendering on the IBM Blue Gene/P

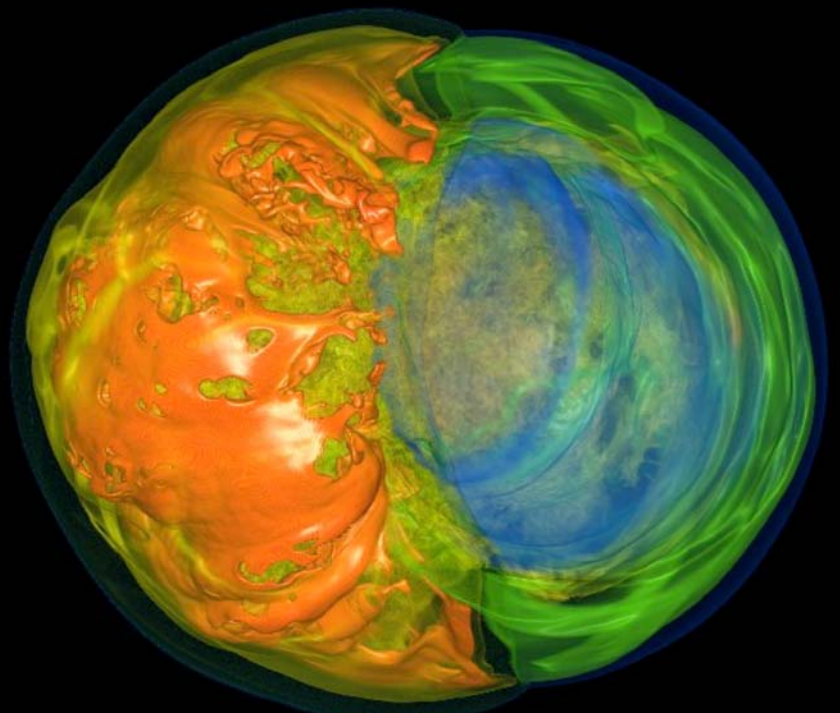

Entropy in core-collapse supernova, time step 1354

Rob Ross - ANL

Hongfeng Yu - SNL California

Kwan-Liu Ma – UCD

Wesley Kendall – UTK

Jian Huang - UTK

Tom Peterka

tpeterka@mcs.anl.gov

Mathematics and Computer Science Division

Argonne NATIONAL LABORATORY

... for a brighter future

UltraVis SciDAC Institute for Ultrascale Visualization

www.ultravis.org

U.S. Department of Energy

UChicago ▶ Argonne LLC

Office of Science U.S. DEPARTMENT OF ENERGY

A U.S. Department of Energy laboratory managed by UChicago Argonne, LLC

Leadership Resources

Computation, communication, and storage

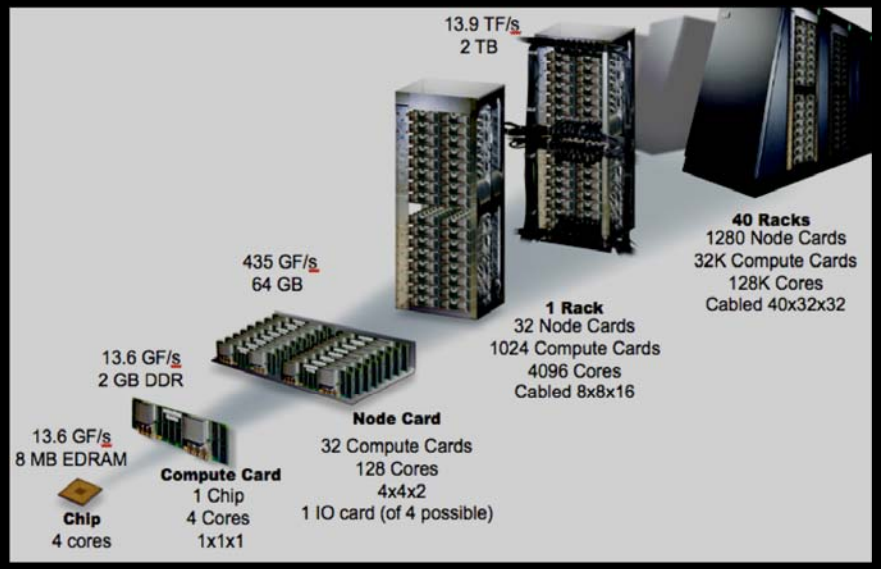

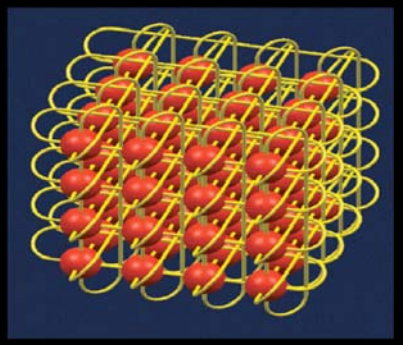

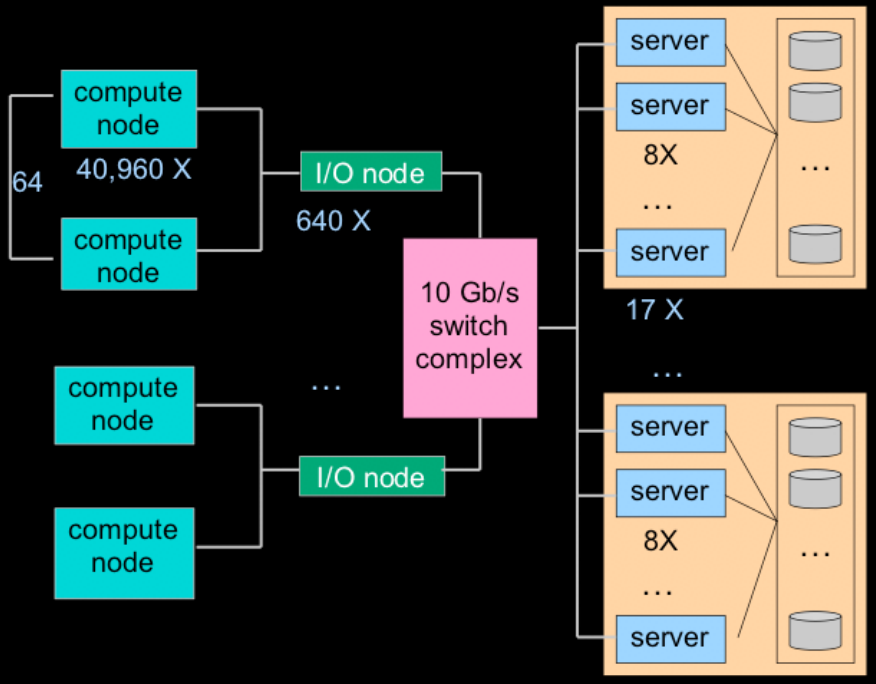

Ever-Increasing Scale of Data and Visualization

Problem sizes are data-dominated. Visualization is no exception.

Computations

Dataset	Problem size (billion elements)	Year	PI
Lifted H2 air	0.9	2008	Grout
Lifted C2 H4 air	1.3	2008	Grout
Supernova	1.3	2008	Blondin
Turbulence	8.0	2005	Yeung

Visualizations

Dataset	Problem size (billion elements)	System size (CPUs)	Year	Reference (et al.)
Taylor-Raleigh	1.0	128	2001	Kniss
Molecular Dynamics	0.1	256	2006	Childs
Earthquake	1.2	2048	2007	Ma
Supernova	0.6	4096	2008	Peterka

2008 INCITE projects

Domain	Data size (TB)	PI
Fusion	54.0	Klasky
Materials	100.0	Wolverton
Astrophysics	300.0	Lamb
Climate	345.0	Washington

Parallel Volume Rendering

Divide, conquer, and reunite

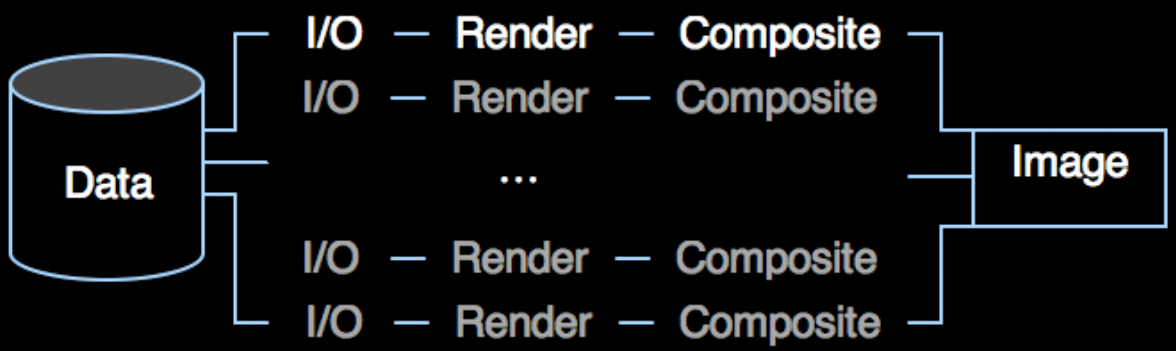

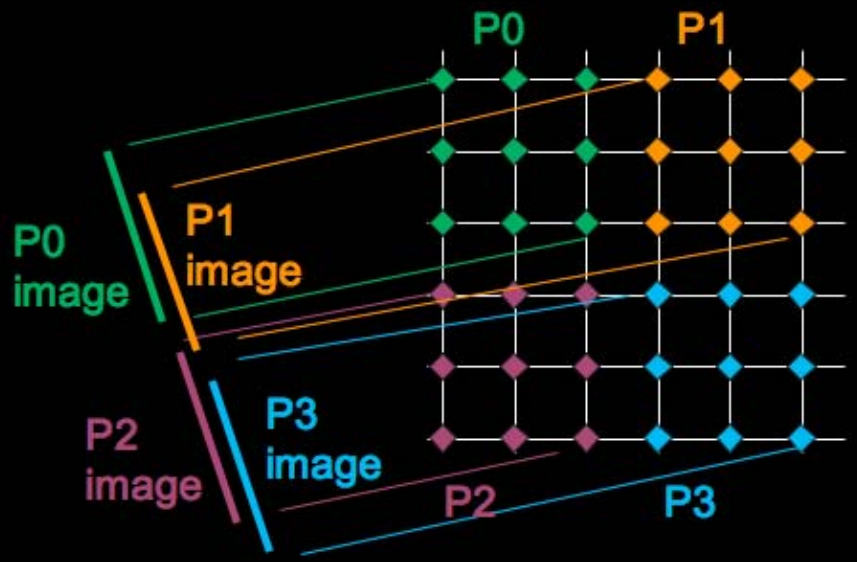

Some Parallel Rendering Parameters

Knobs to turn, switches to flip, buttons to press

Argument	Sample Values
DataSize	1120x1120x1120
ImageSize	1600x1600
ImageType	ppm, rgb, rgba
IP, port	137.72.15.10, 5000
Stereo	y, n
NumProcs	16384
NumPipes	16
BlockingFactor	8
NumWriters	64
NumThreads	1

Larger Datasets and Images

Another measure of scalability

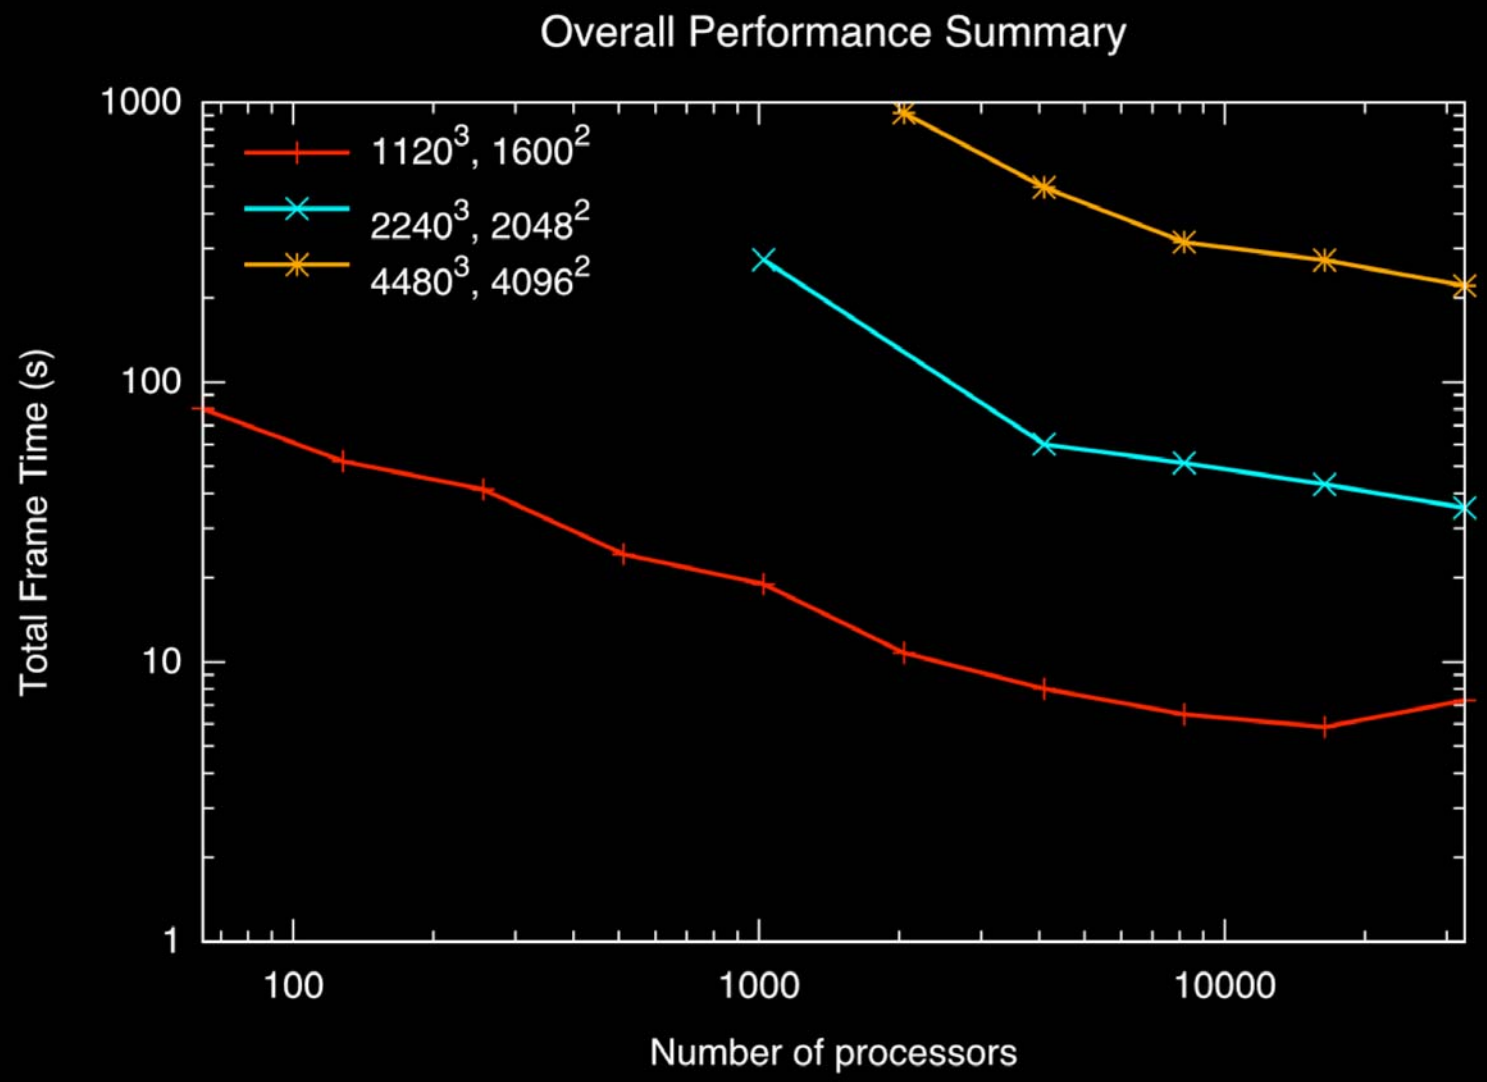

Time Distribution

Reading the data from storage dominates the total cost of a time step.

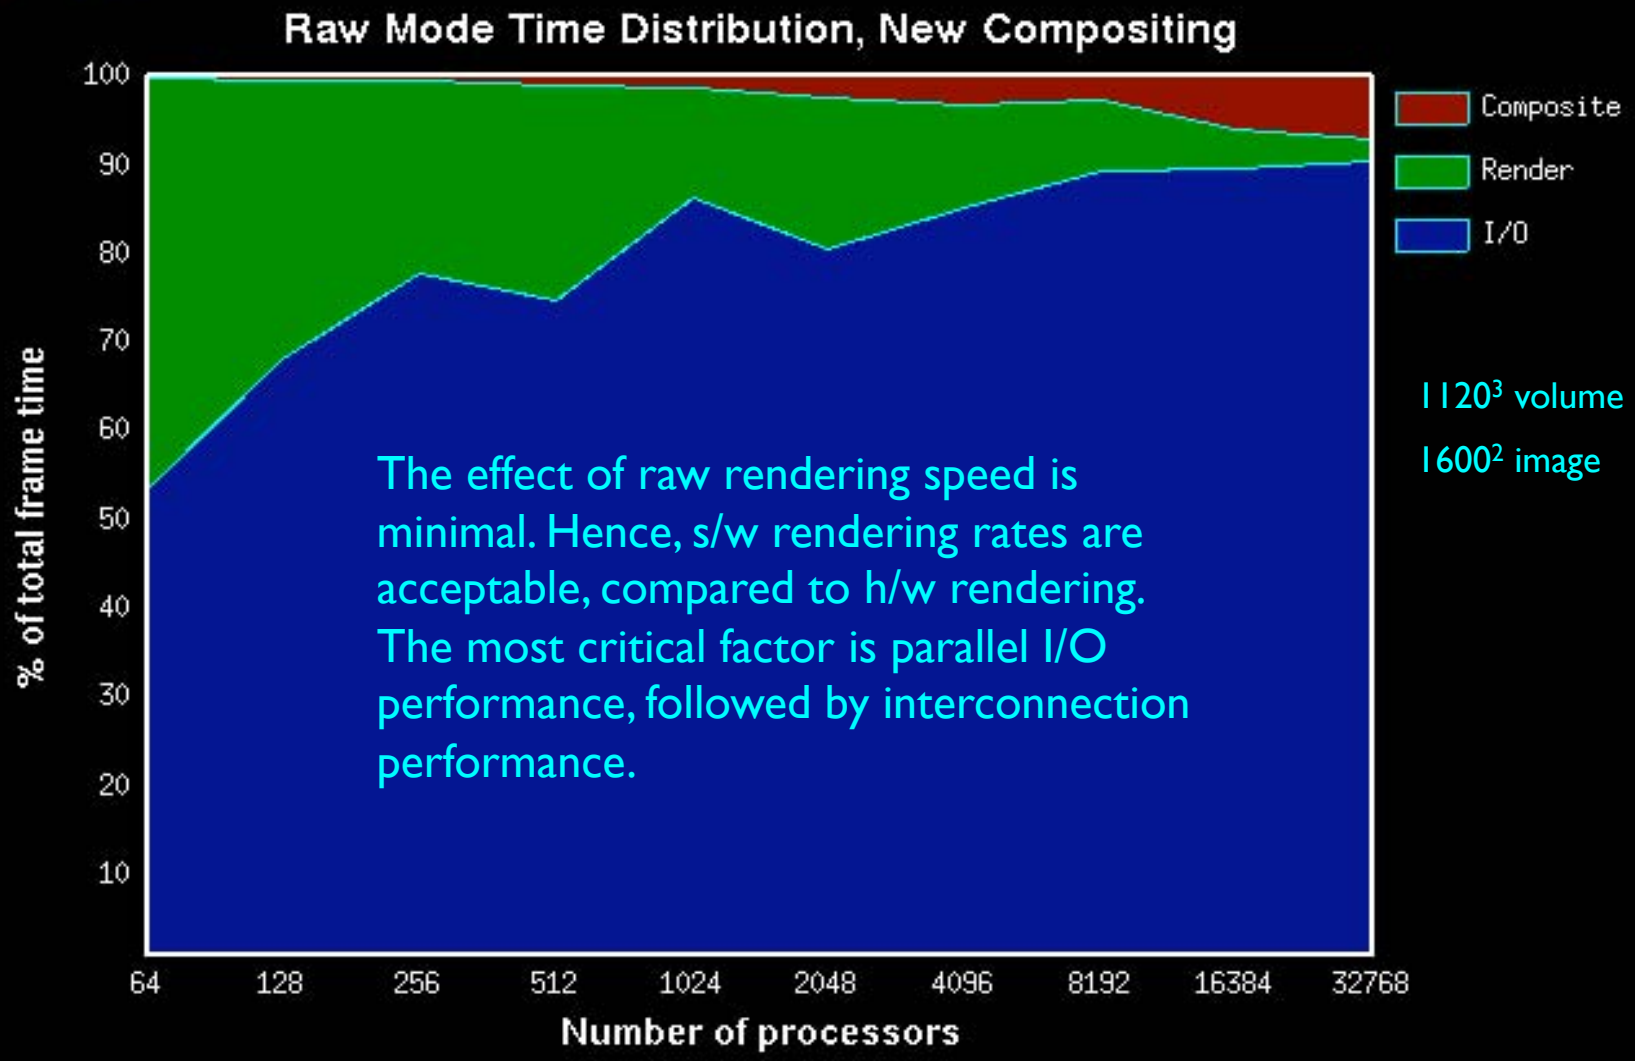

1120^3 volume

1600^2 image

Efficiency

Round robin static block distribution is an inexpensive load balancing scheme that is quite effective.

1120^3 volume

1600^2 image

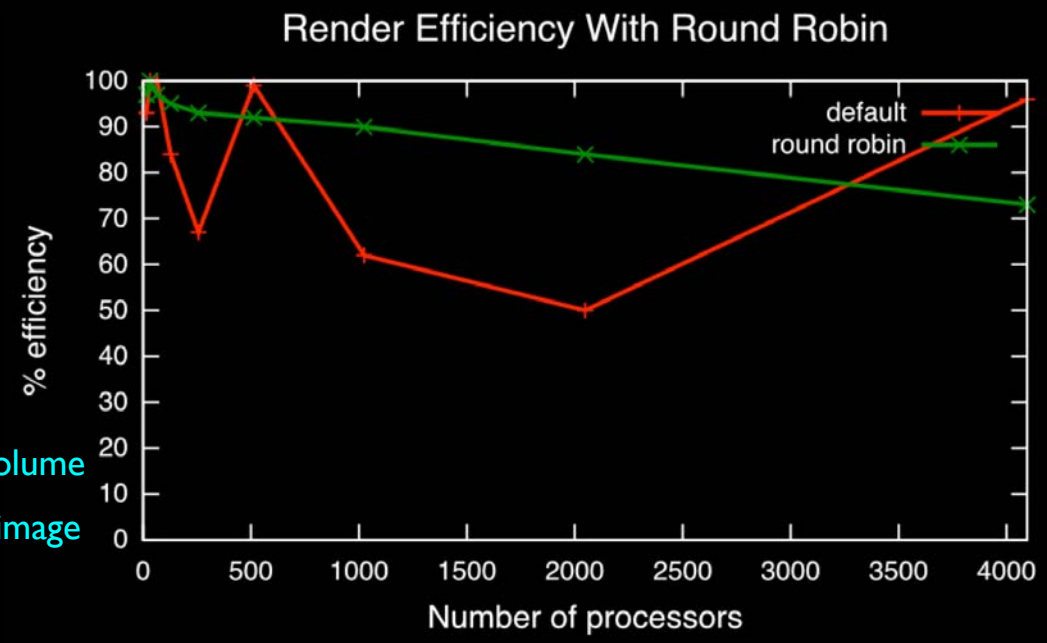

864^3 volume

1024^2 image

Multiple Writers Performance

Improve overall output time by selecting the optimal number of writers.

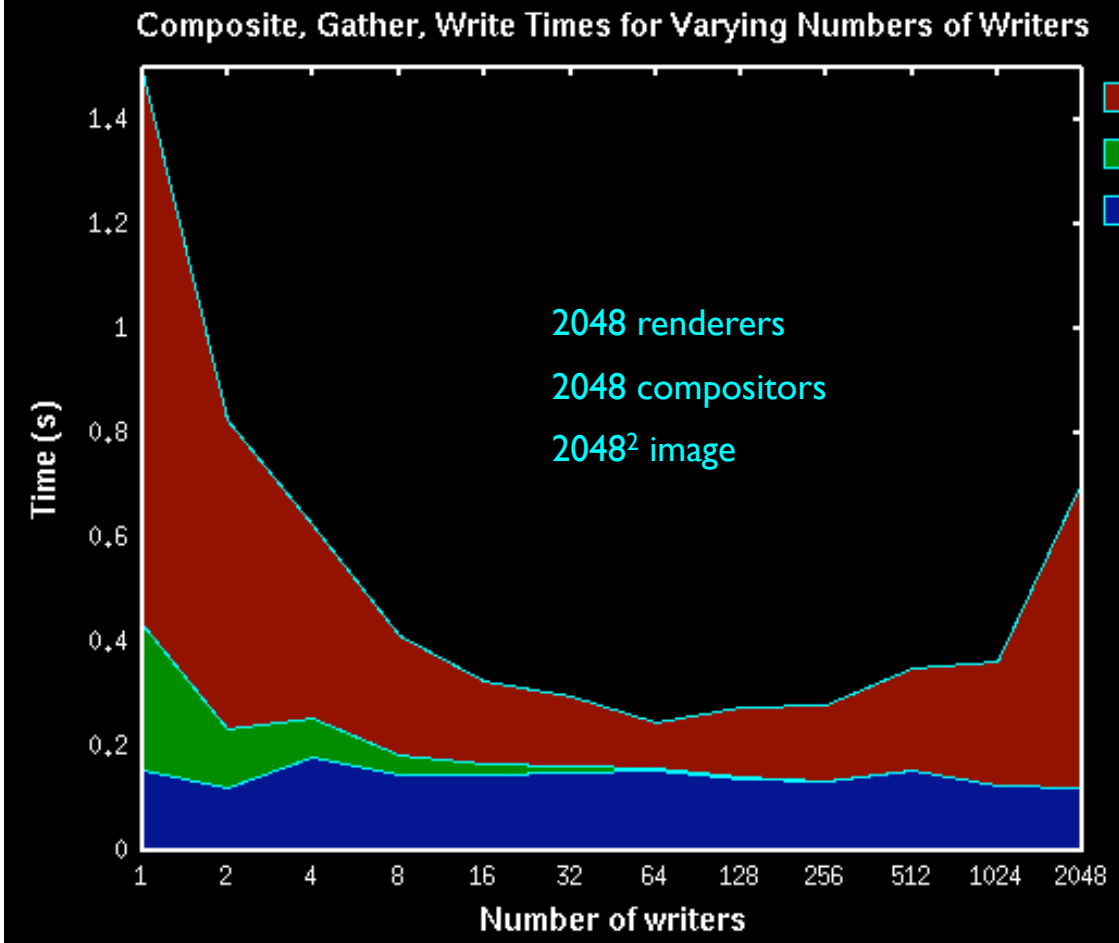

2048 renderers

2048 compositors

2048^2 image

Memory footprint per core =
70MB +
2.5KB * image size / writing_cores +
4 * volume size / rendering_cores

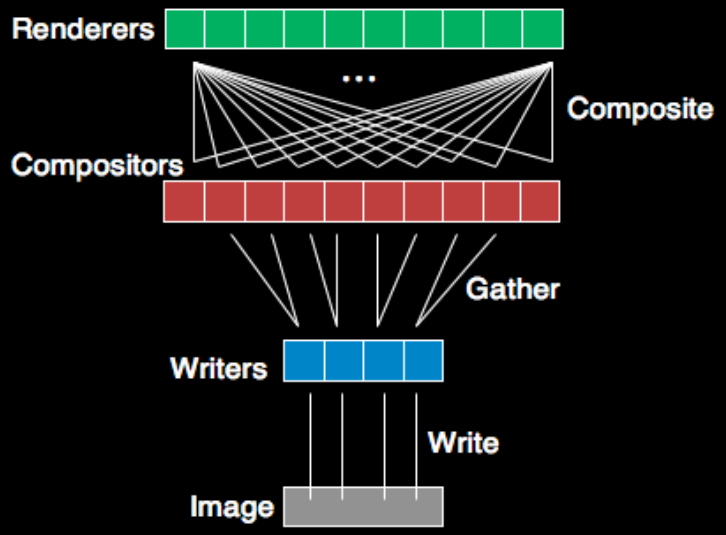

64 writers best for most cases; writers need to be distributed among I/O nodes.

Multithread – MPI Hybrid Programming Model

MPI-pthread rendering
MPI-only I/O and compositing

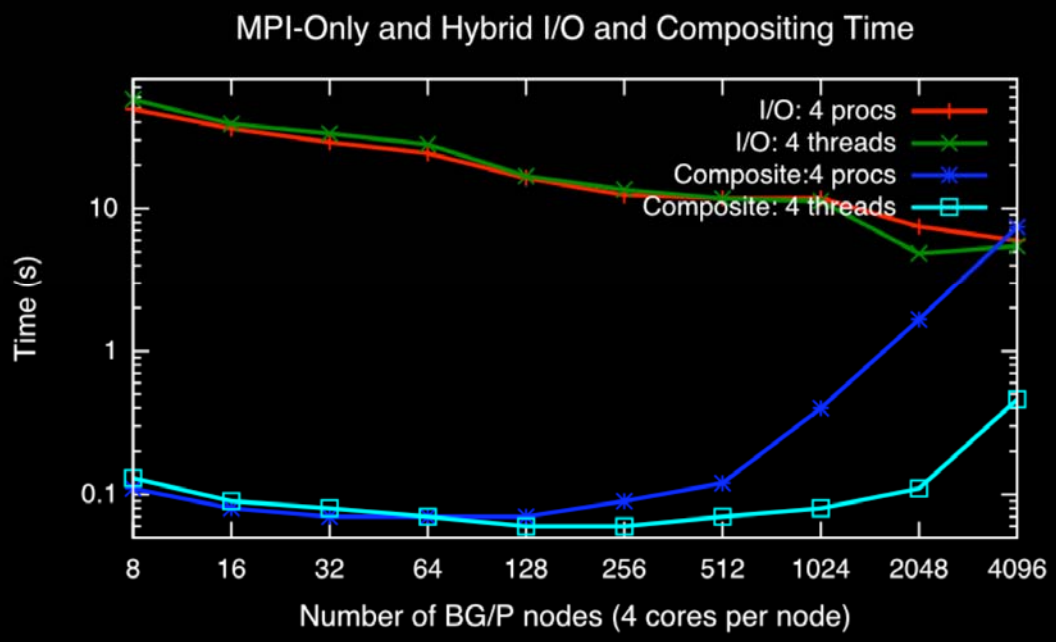

1120^3 volume 1024^2 image

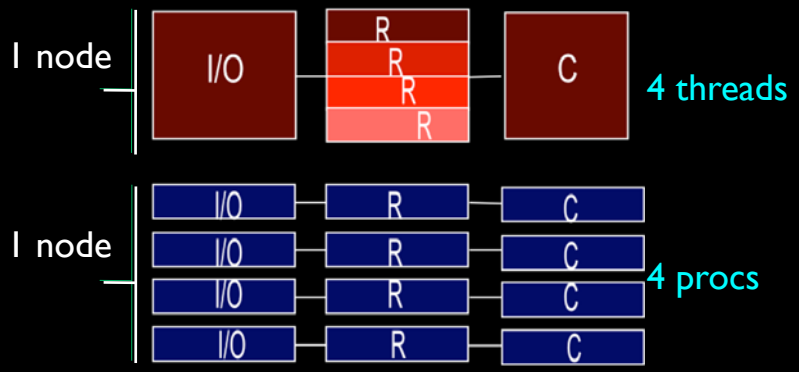

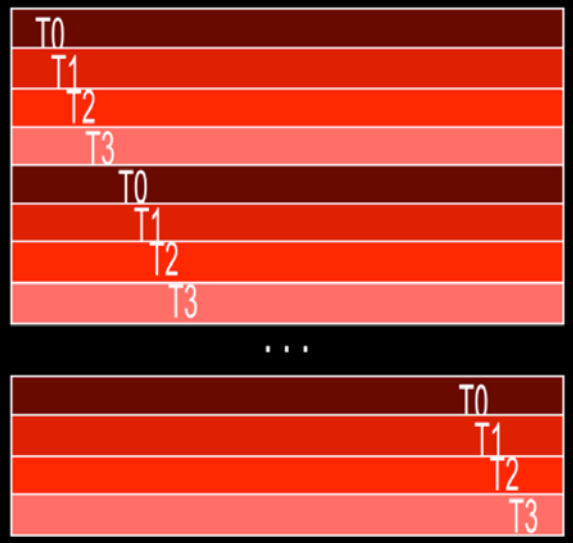

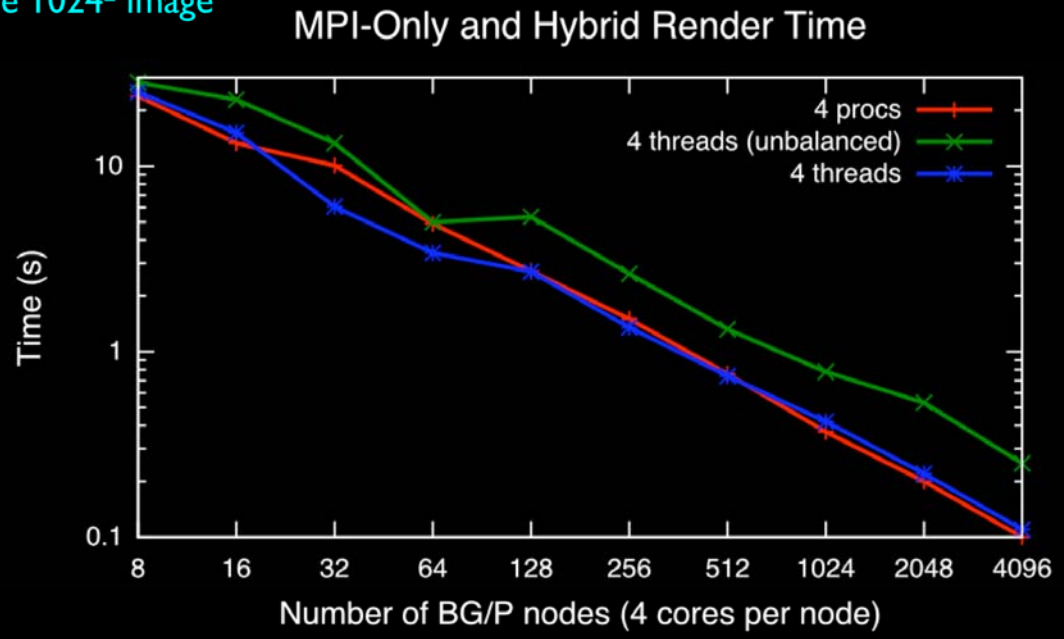

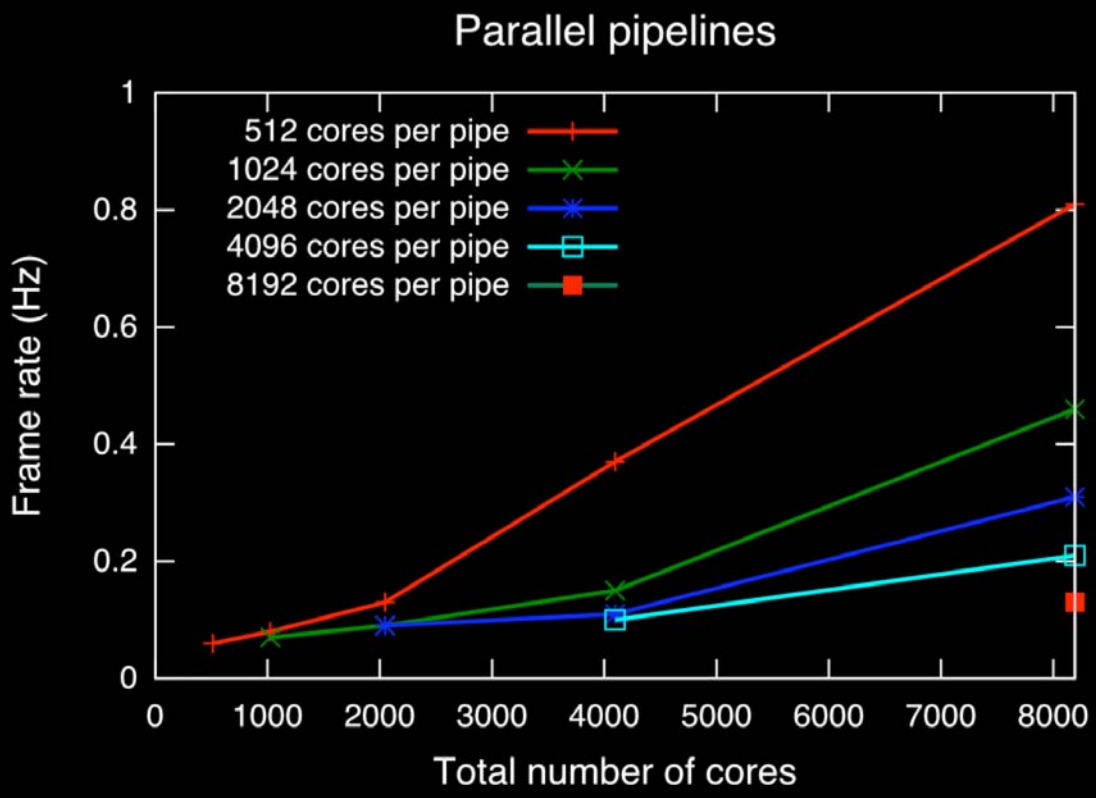

Multiple Parallel Pipelines

Hide I/O latency by extending concurrency between time steps.

864^3 volume
1024^2 image

6X faster for same total system size when 16 pipelines are used instead of one

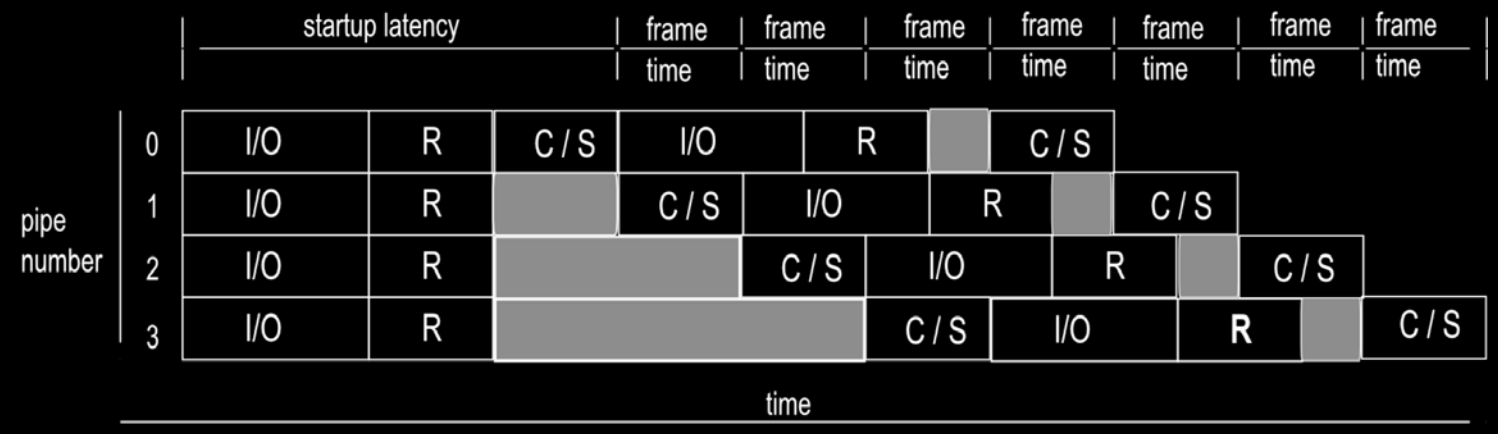

Lessons Learned

and the road ahead

Successes

- Demonstrated scaling
- Large data and image sizes
- Improved compositing
- Improved and benchmarked I/O
- Load balancing
- Memory scalability
- Hybrid programming model
- Parallel pipelines

Challenges, to do

- Other grid topology
- In situ visualization
- Adoption into tools
- Other architectures
- Other vis algorithms

Assessing and Improving Large Scale Parallel Volume Rendering on the IBM Blue Gene/P

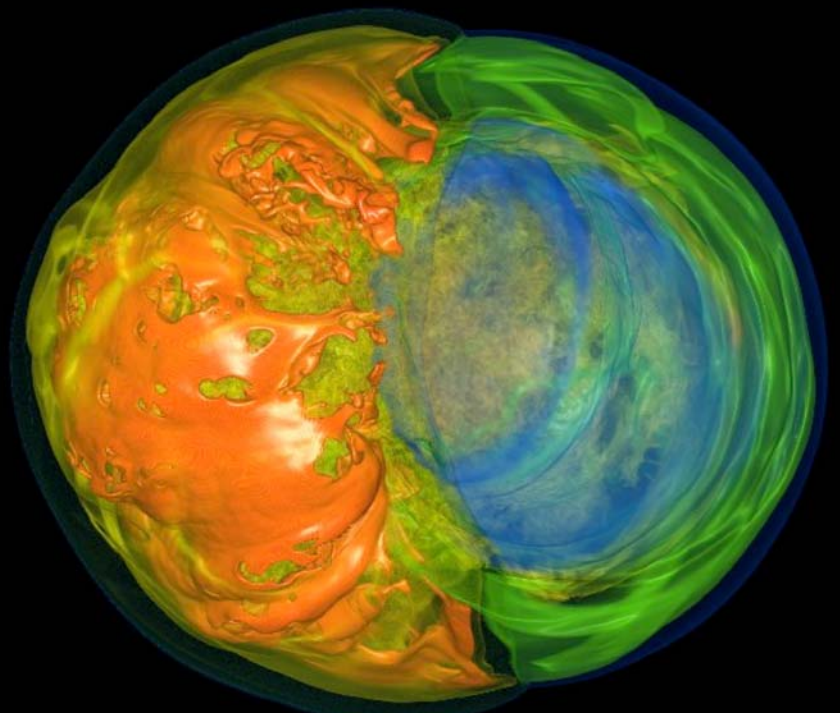

Entropy in core-collapse supernova, time step 1354

Rob Ross - ANL

Hongfeng Yu - SNL California

Kwan-Liu Ma – UCD

Wesley Kendall – UTK

Jian Huang - UTK

Tom Peterka

tpeterka@mcs.anl.gov

Mathematics and Computer Science Division

Argonne NATIONAL LABORATORY

... for a brighter future

UltraVis SciDAC Institute for Ultrascale Visualization

www.ultravis.org

U.S. Department of Energy

UChicago ▶ Argonne LLC

Office of Science U.S. DEPARTMENT OF ENERGY

A U.S. Department of Energy laboratory managed by UChicago Argonne, LLC