San Jose Downtown

August 25-27, 2008

NVIDIA.

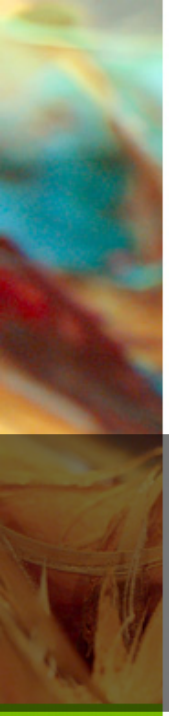

Michele D. Cash, University of Washington

Multi-Fluid Modeling with GPUs

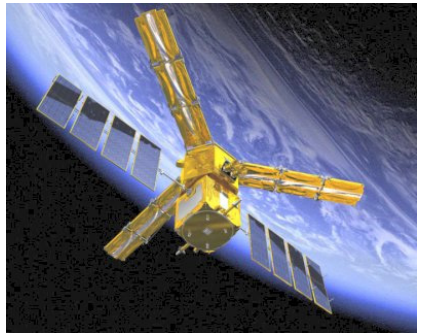

University of Washington

- Michele Cash
- Robert Winglee
- Erika Harnett

Eagle Harbor Technologies

- Tim Ziemba
- John Carscadden
- Dan O'Donnell

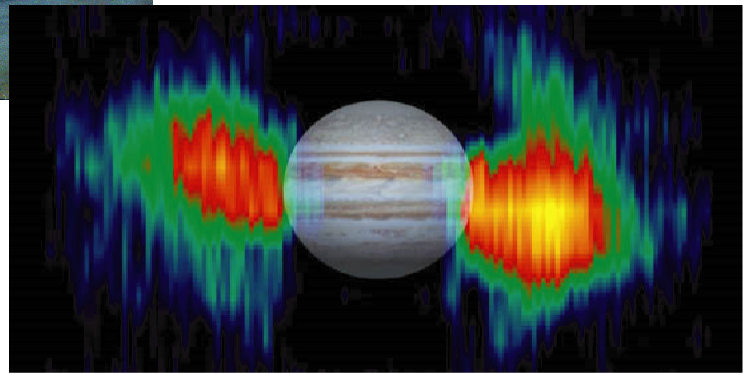

Multi-Fluid Modeling with GPUs

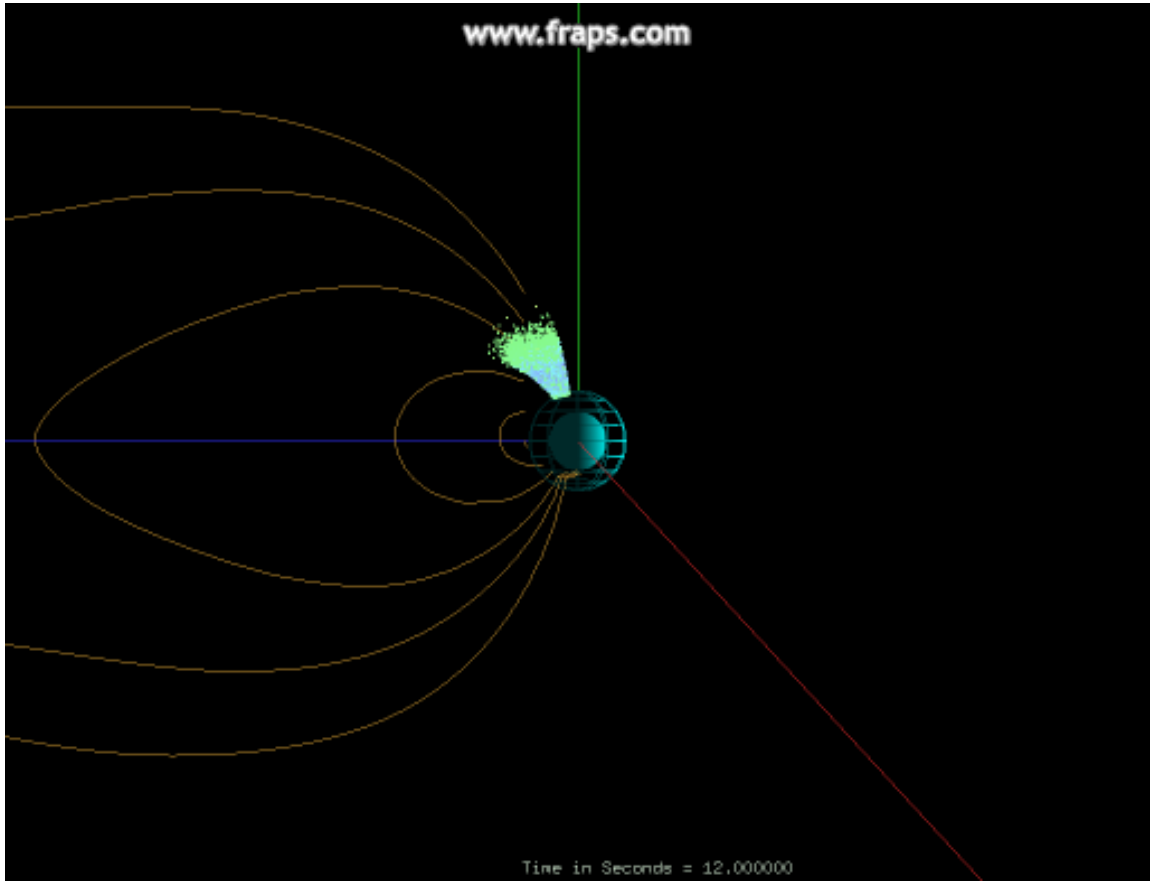

Multi-Fluid Modeling with GPUs

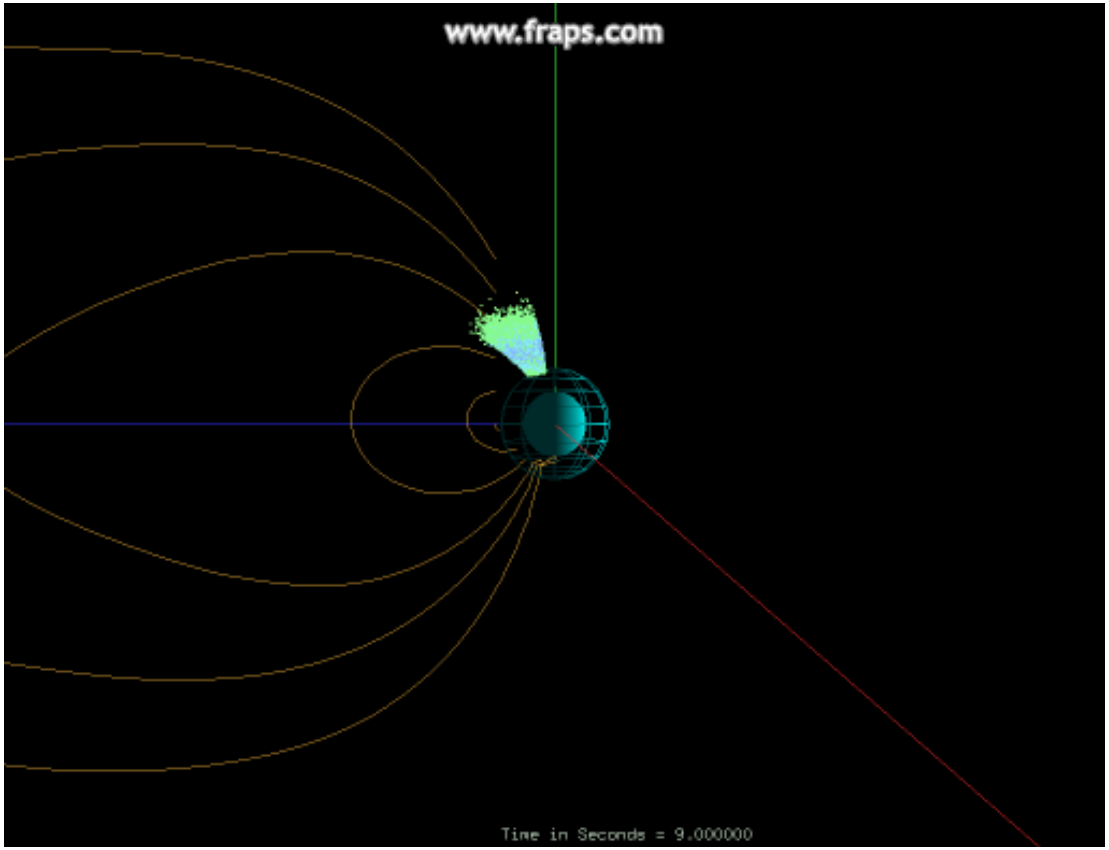

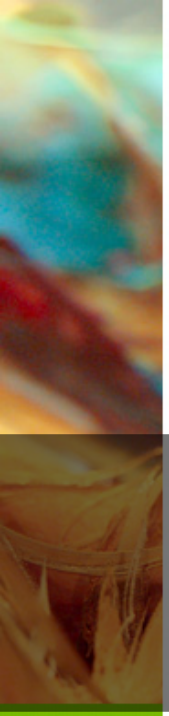

Gregory Diamos, GA Tech

Harmony

An Execution Model and Runtime for Heterogeneous Many Core Systems

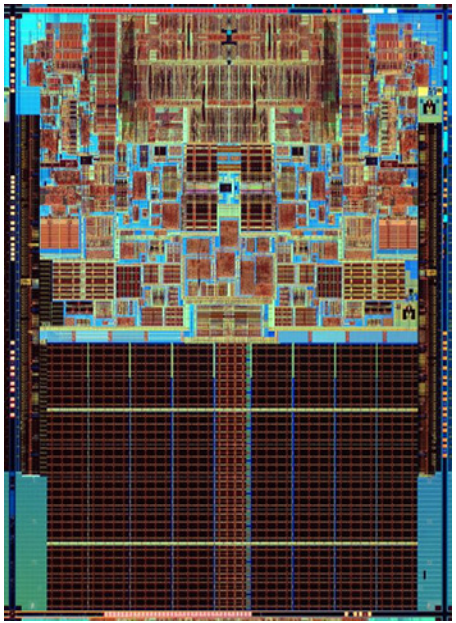

Gregory Diamos and Sudhakar Yalamanchili
Computer Architecture and Systems Lab
Georgia Institute of Technology

Generality

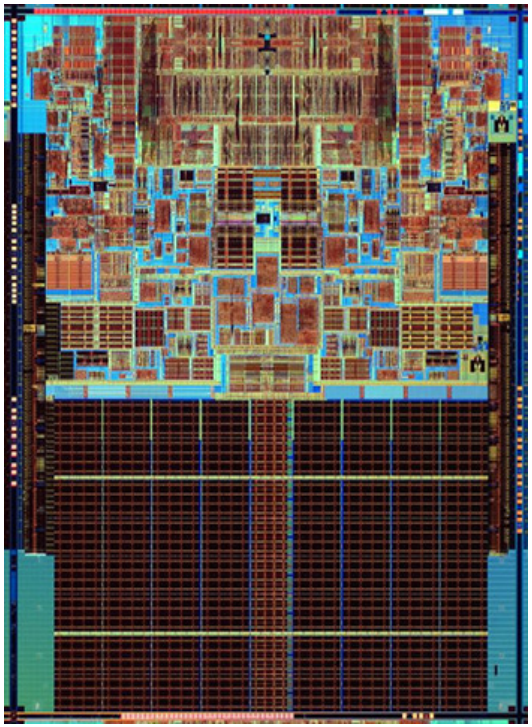

Generality is bought with efficiency

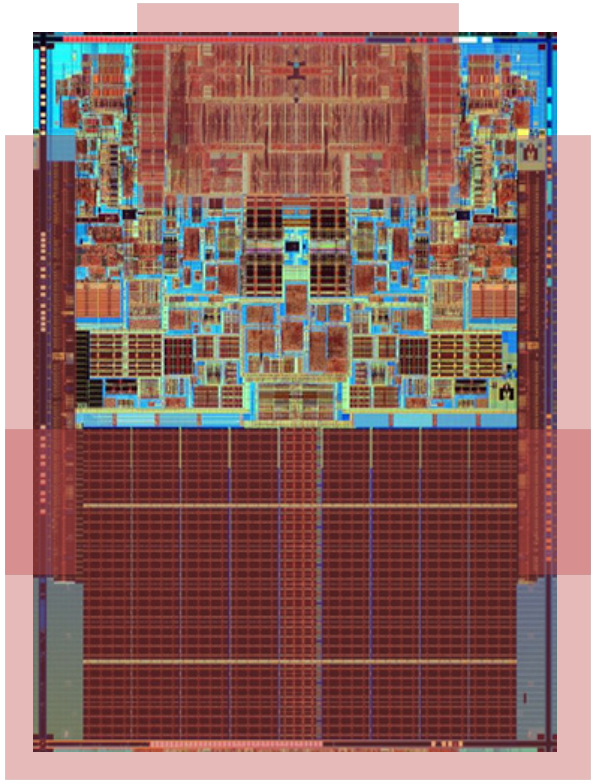

Applications have diverse requirements

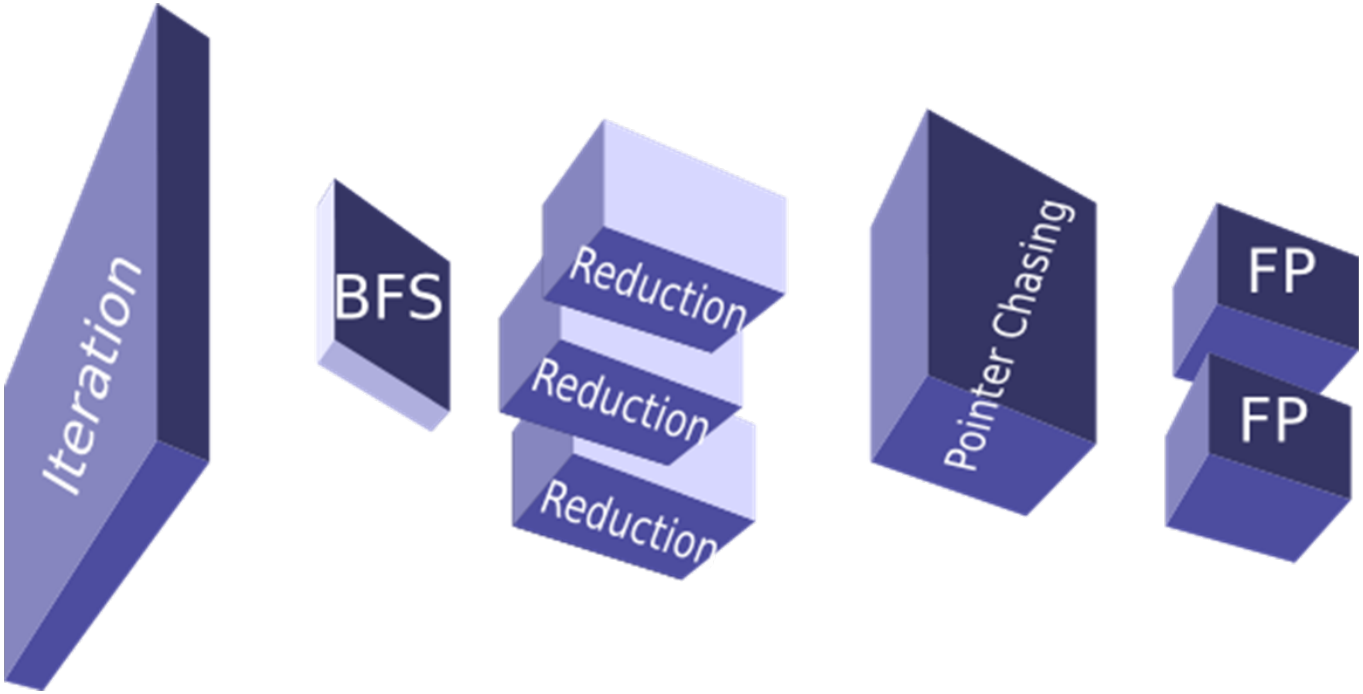

CPUs are one-size-fits-all solutions

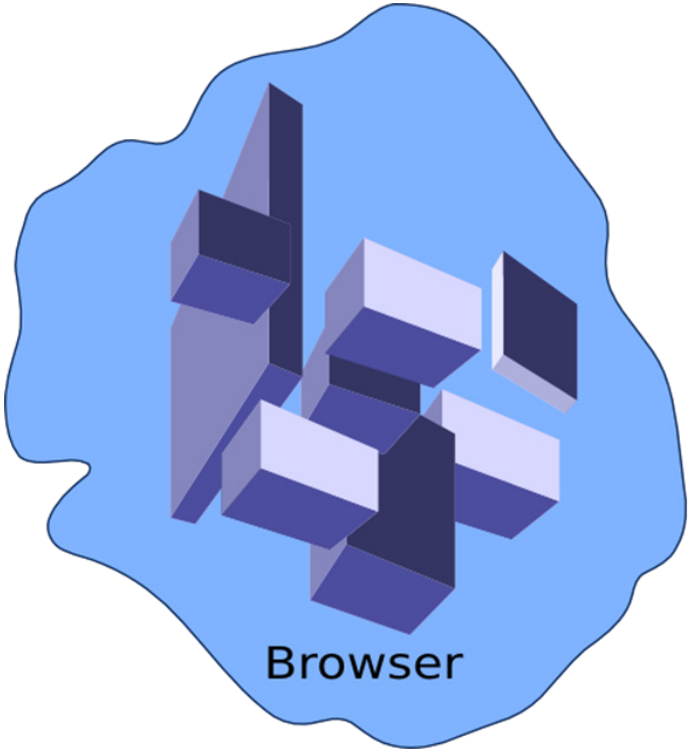

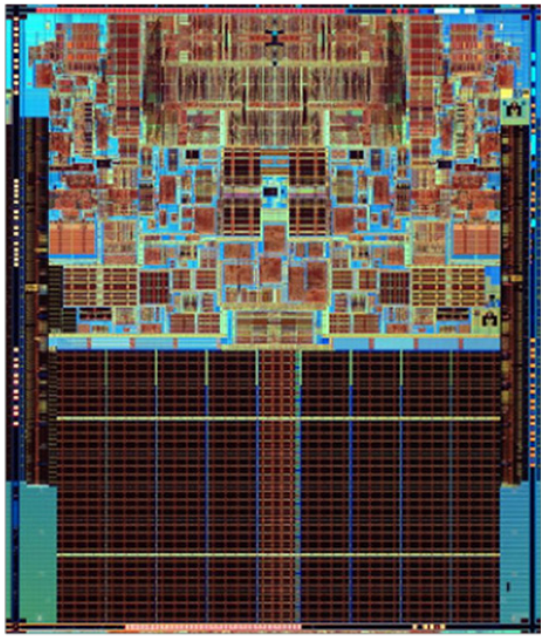

Specialization – Improves Efficiency

Heterogeneous systems can match algorithms to architectures

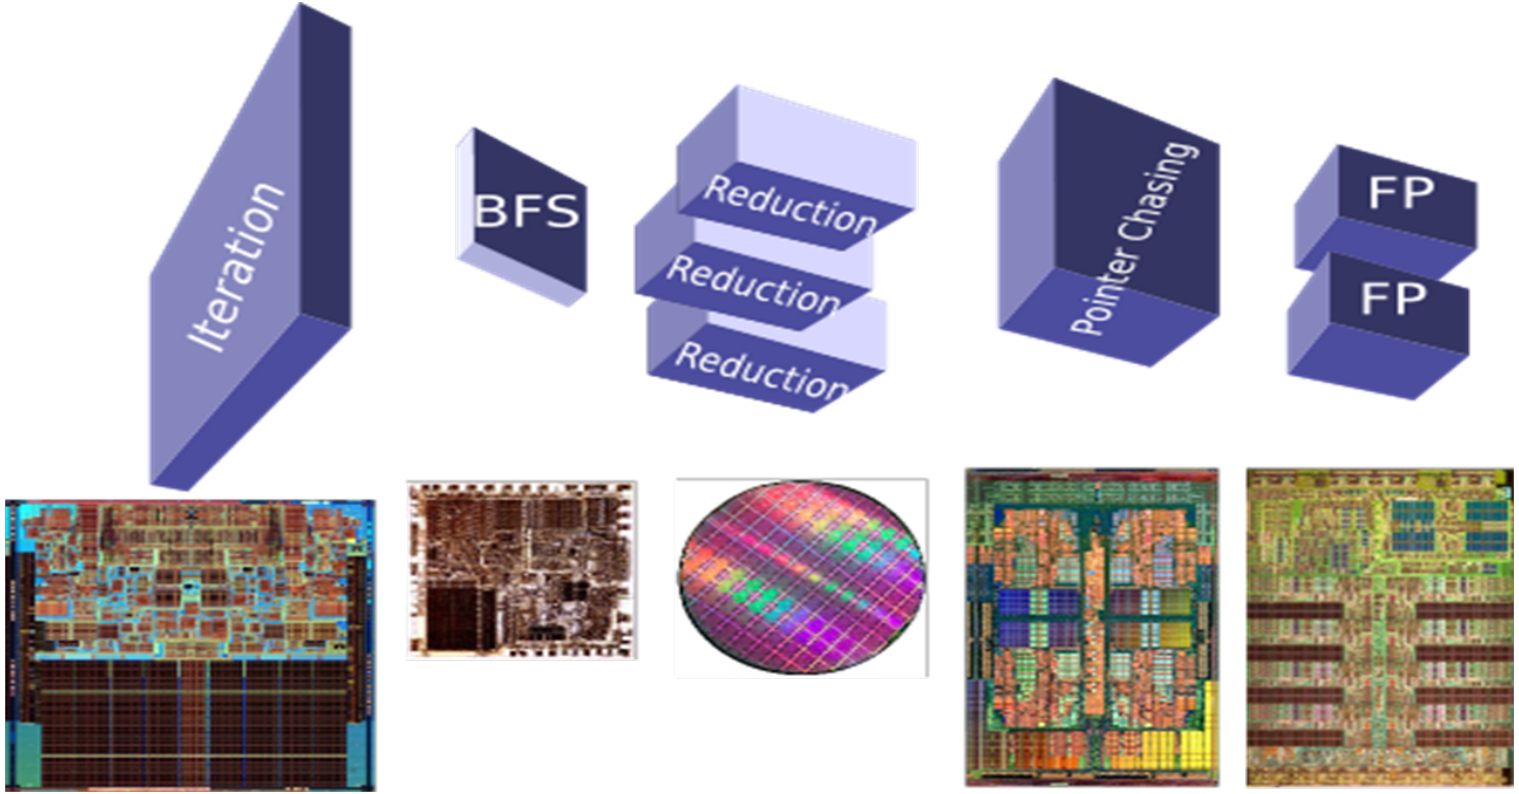

Specialization – Compounds Complexity

The user must handle partitioning, parallelism, and scheduling using diverse programming models.

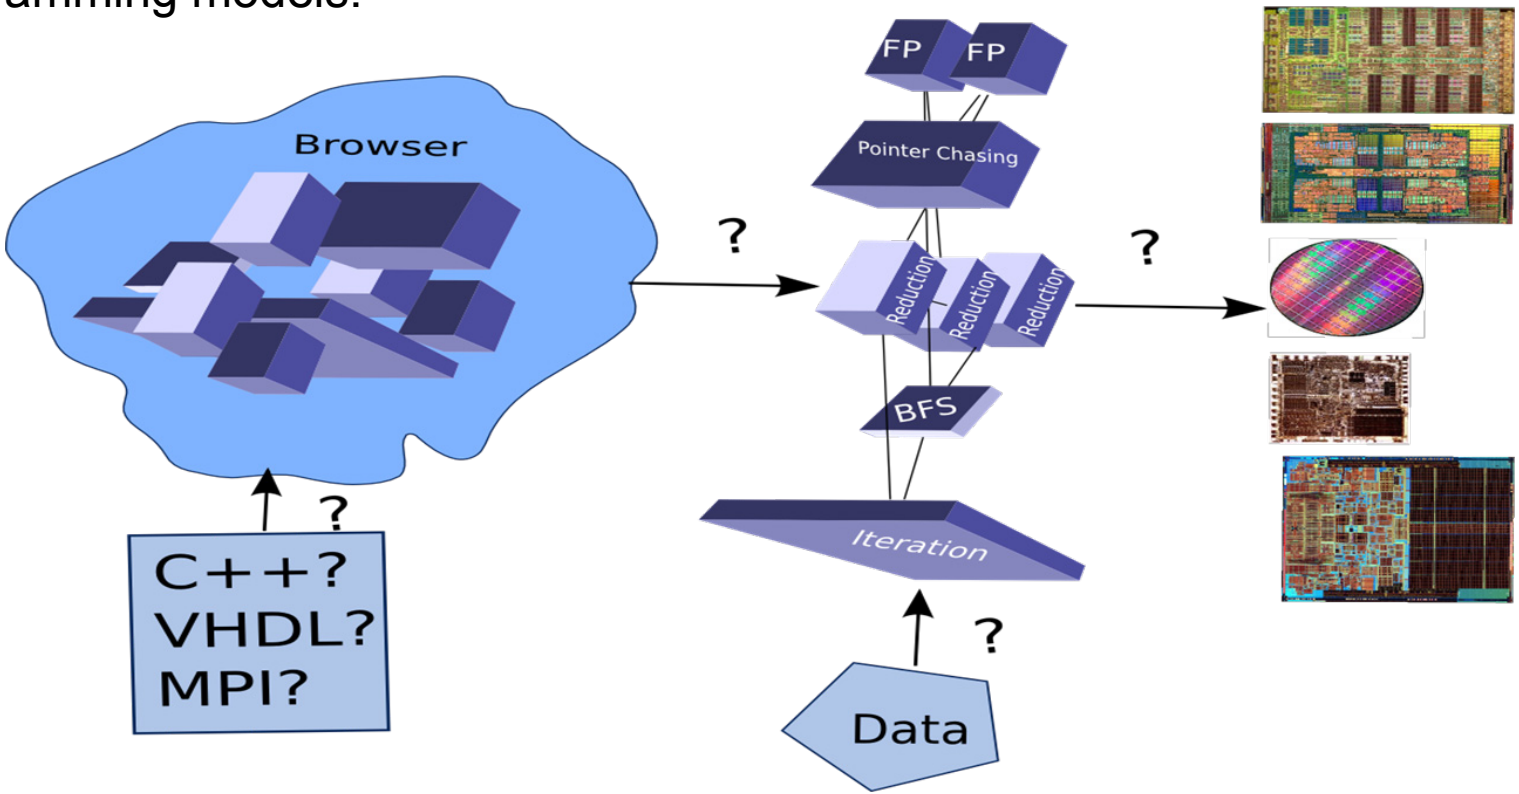

Runtime support can allow heterogeneous cores to cooperate harmoniously using dynamically available information

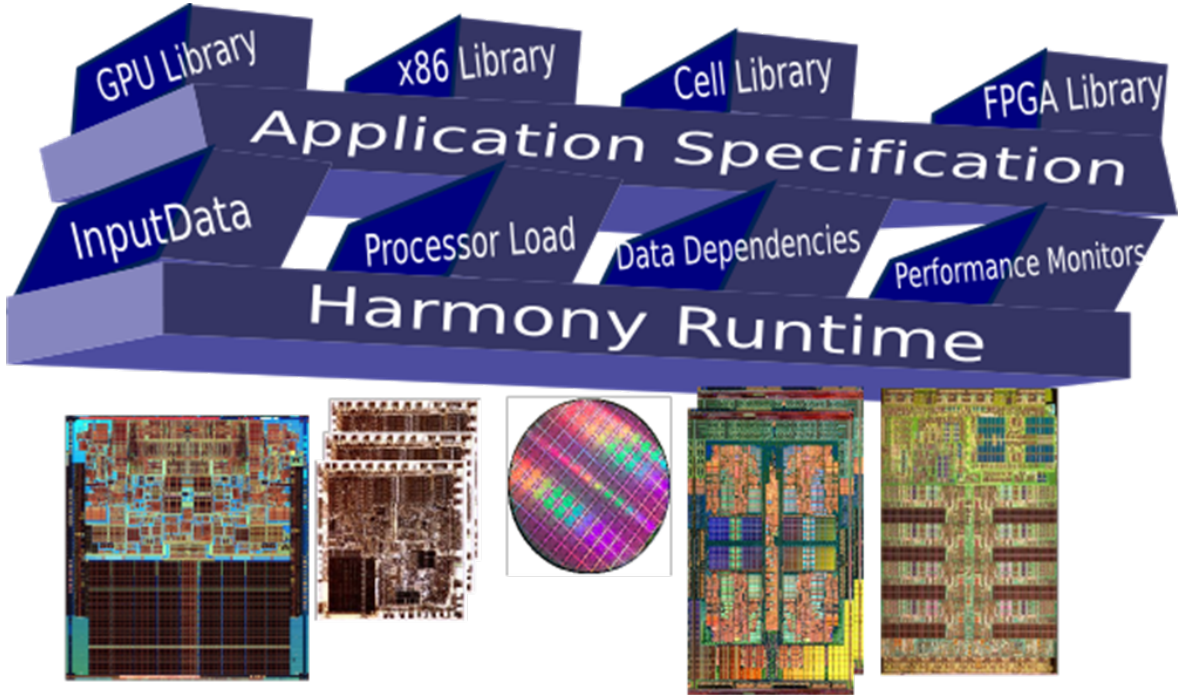

Contact Us

http://www.ece.gatech.edu/research/labs/casl/harmony.html

CASL – Georgia Institute of Technology

Milos Hasan, Cornell University

Large Matrix Problems in Rendering

- Inherently quadratic problems:
  - Indirect illumination
  - Environment map lighting
  - Many direct lights
- Treat them as large matrices

Global illumination …

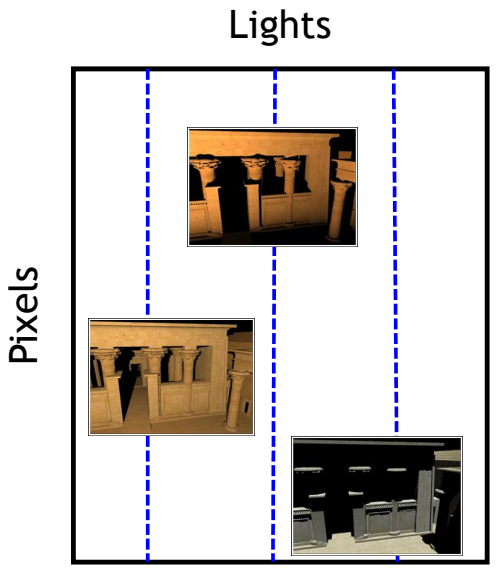

… as a large matrix

Direct-to-Indirect Transfer

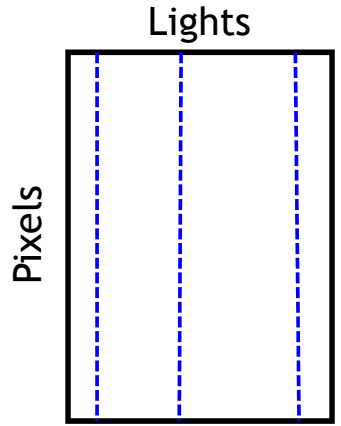

Transport matrix
in wavelet space

Images lit by wavelet lights

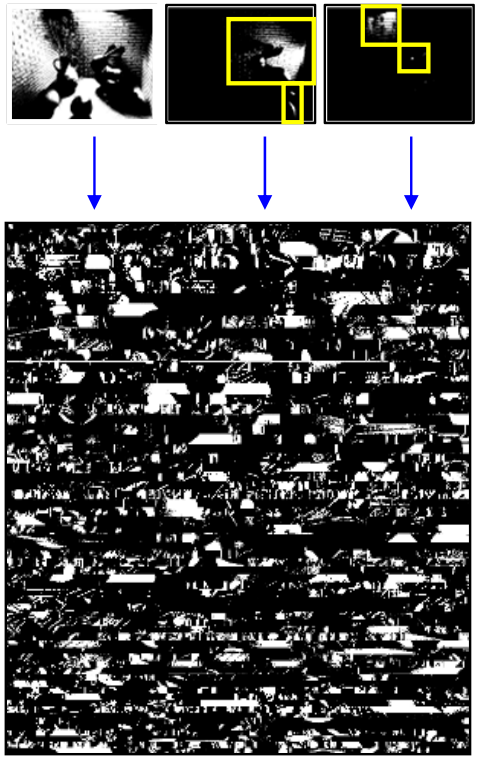

Convert matrix multiplication
to blending of tiny images

SIGGRAPH 2006 - joint work with Fabio Pellacini and Kavita Bala

Direct-to-Indirect Transfer

- Results: 640 x 480, 2x2 AA, up to 2m triangles, running on 7800 GTX

~15 fps
(1.6 hours precomputation)

~15 fps
(2.5 hours precomputation)

SIGGRAPH 2006 - joint work with Fabio Pellacini and Kavita Bala

Matrix Row-Column Sampling

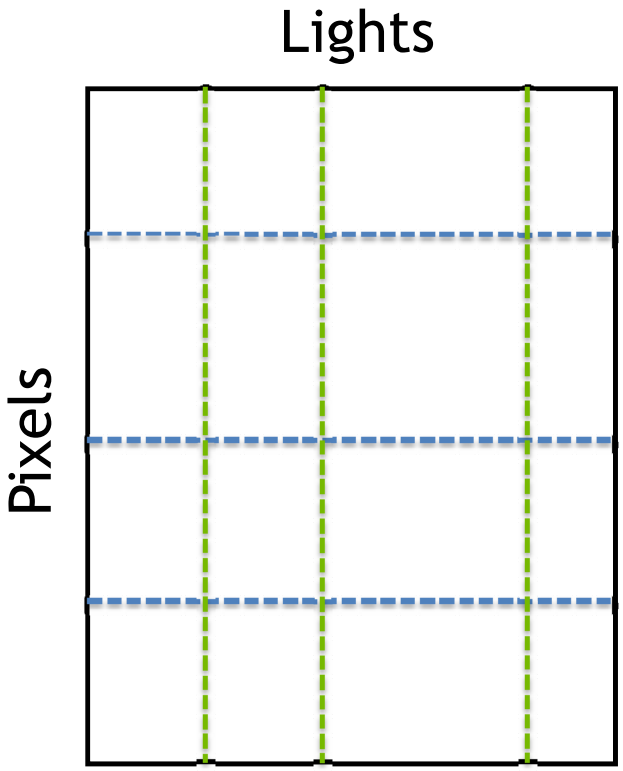

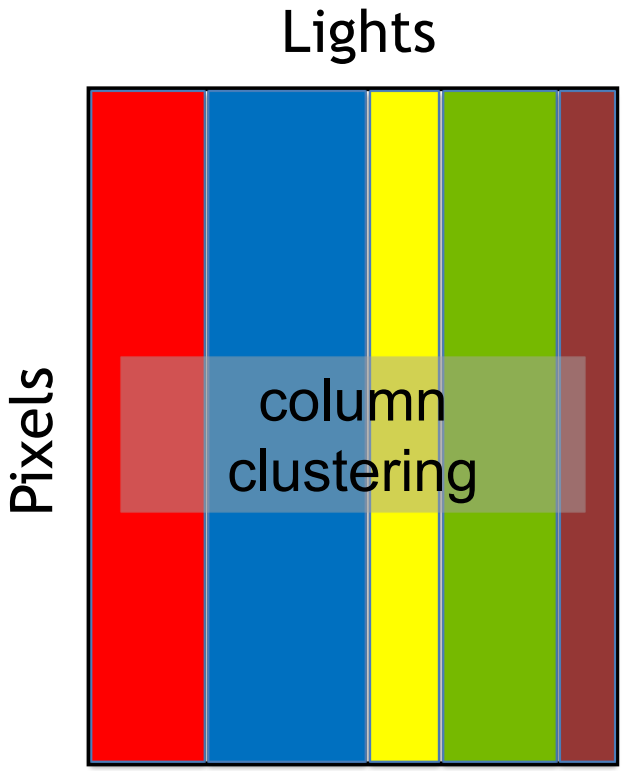

→ A GPU-friendly global illumination renderer

SIGGRAPH 2007 - joint work with Fabio Pellacini and Kavita Bala

Matrix Row-Column Sampling

- Results: 800 x 600, 2x2 AA, 100,000 lights, up to 2m triangles, running on 8800 GTX

Our result: 16.9 sec
(300 rows + 900 columns)

Reference: 20 min
(using all 100k lights)

Our result: 3.8 sec
(100 rows + 200 columns)

Reference: 10 min
(using all 100k lights)

SIGGRAPH 2007 - joint work with Fabio Pellacini and Kavita Bala

Tensor Clustering for Animation

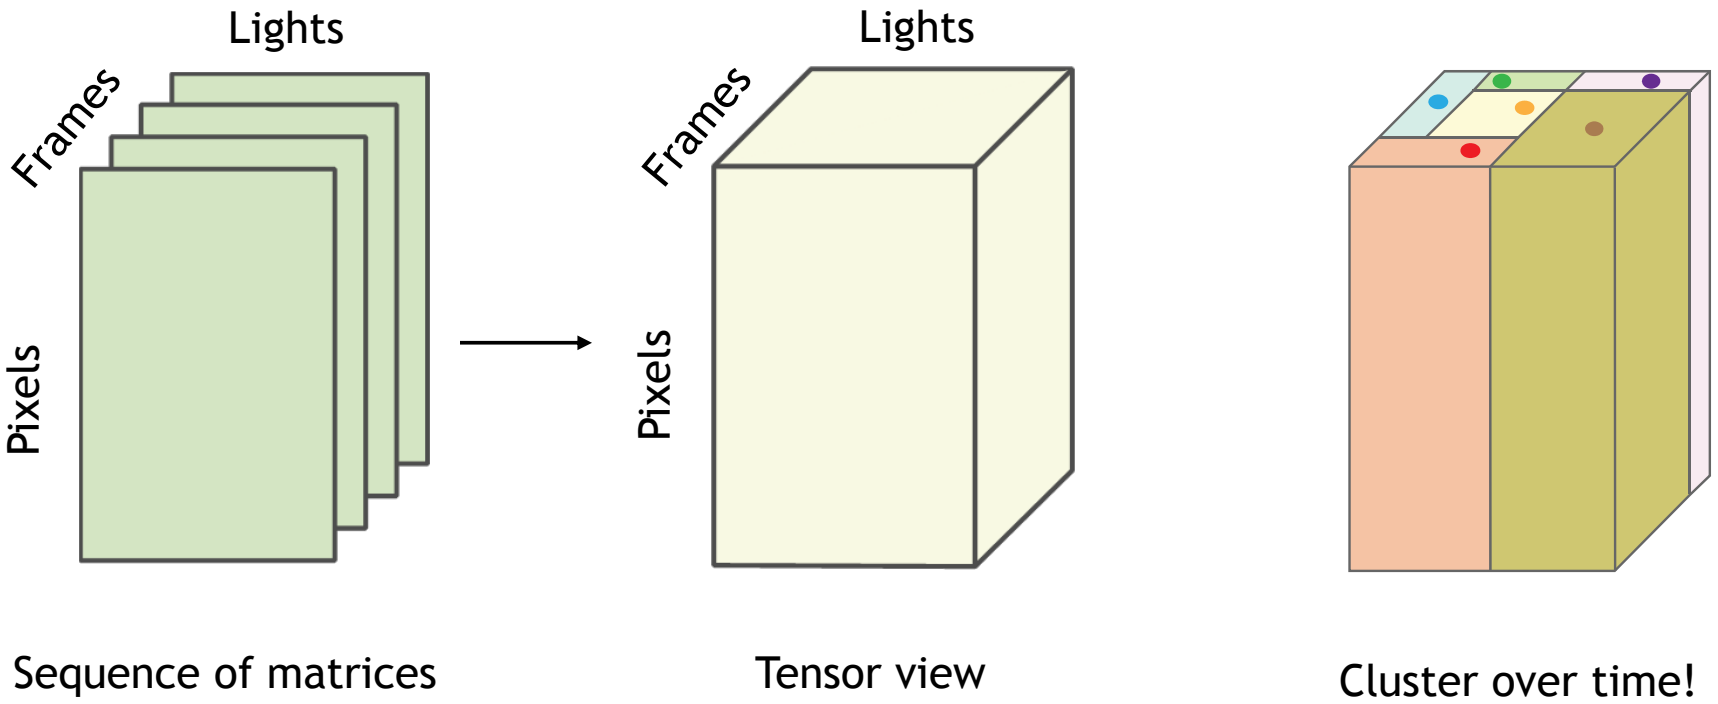

Sequence of matrices

Tensor view

Cluster over time!

EGSR 2008 - joint work with Edgar Velazquéz-Armendáriz, Fabio Pellacini and Kavita Bala

Tensor Clustering for Animation

- Results: 640 x 480, 3x3 AA, 65,536 lights, up to 8.2m triangles, running on 8800 GTX

Our result: 26 sec/frame Reference: 33min/frame

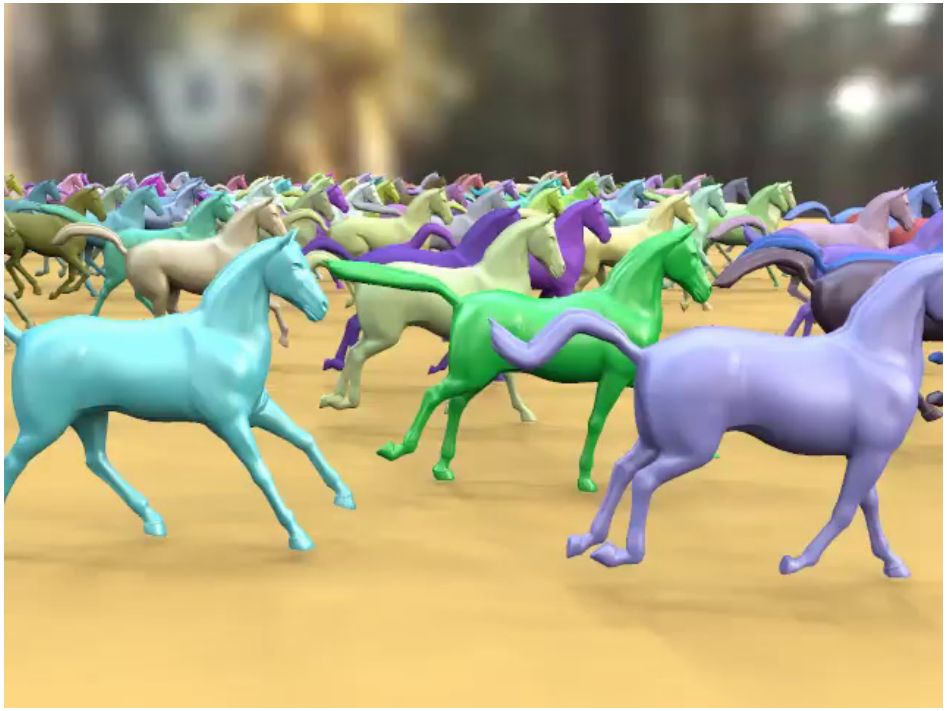

Our result: 49 sec/frame Reference: 3.5 hours/frame

EGSR 2008 - joint work with Edgar Velazquéz-Armendáriz, Fabio Pellacini and Kavita Bala

Derek Juba, University of Maryland, College Park

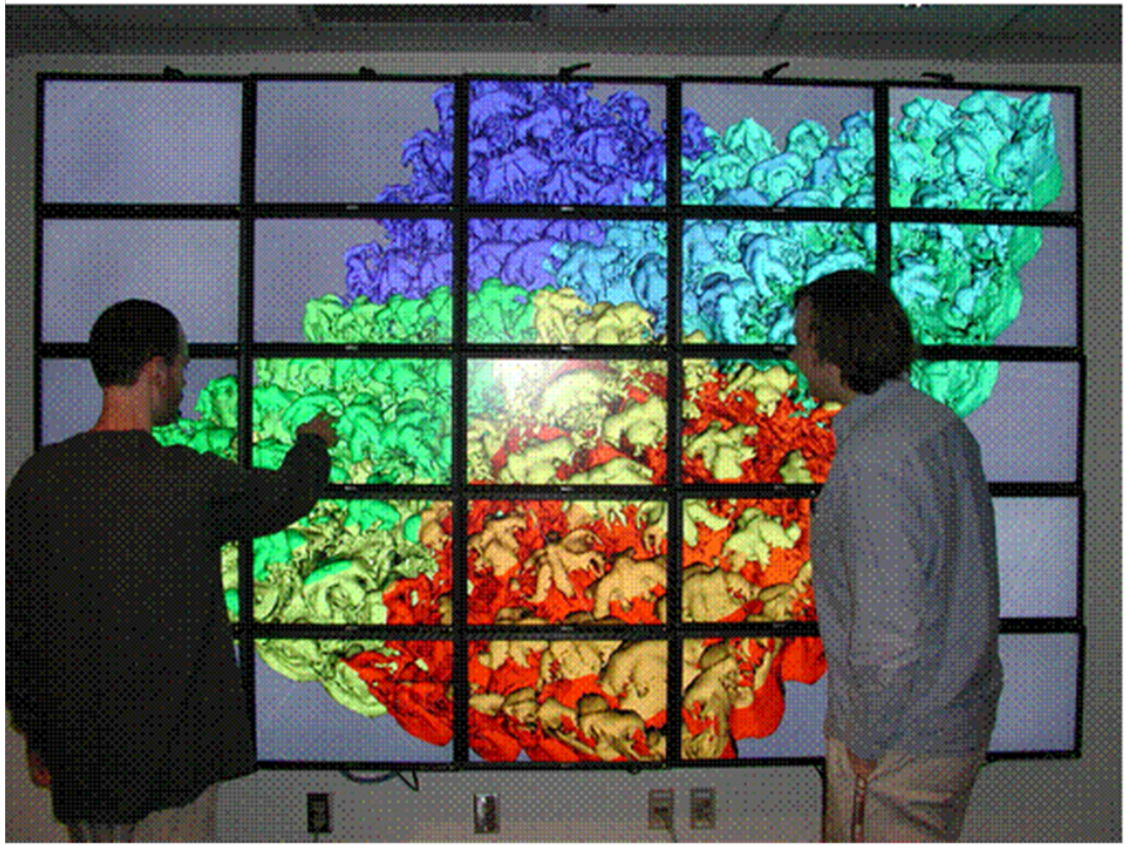

NVISION 08
THE WORLD OF VISUAL COMPUTING

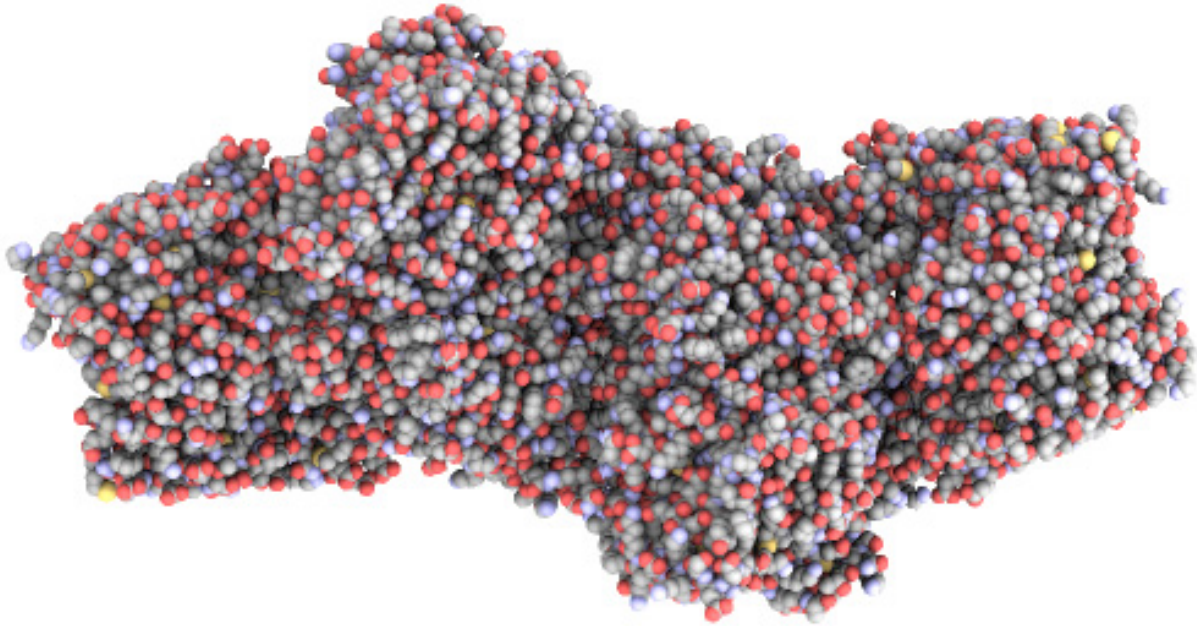

NVIDIA.

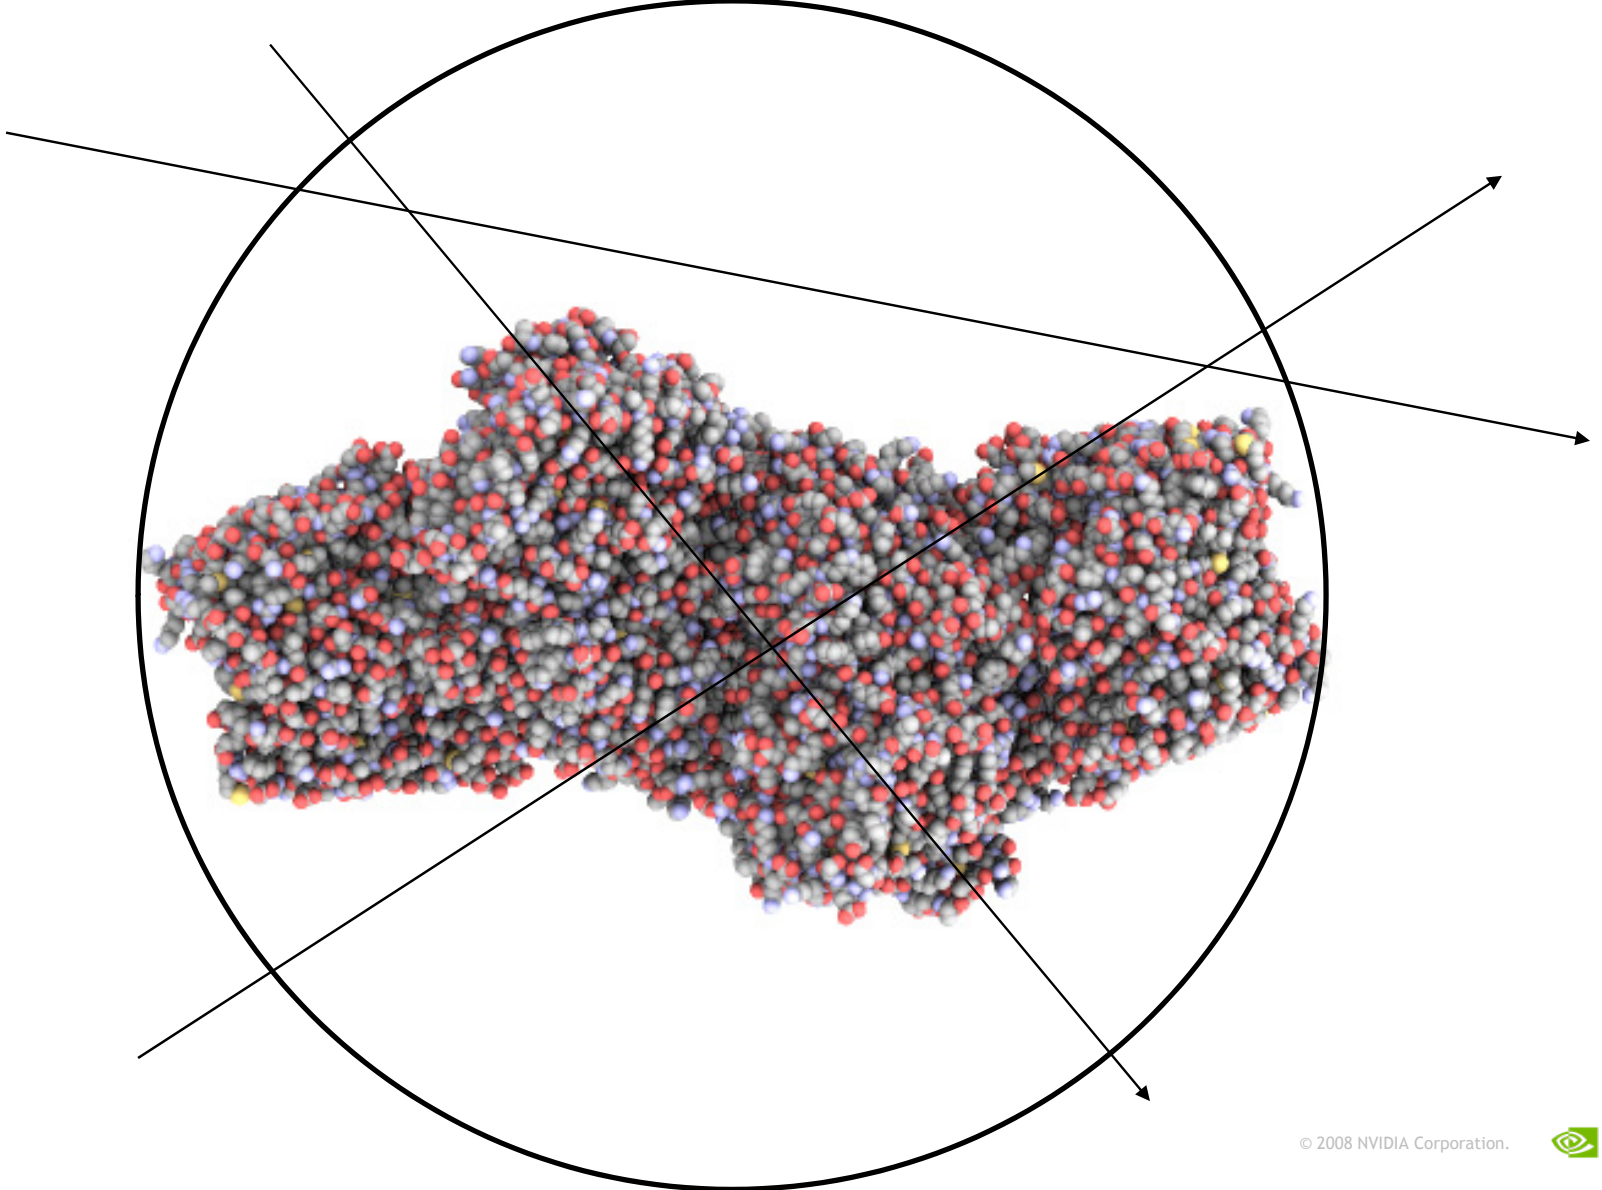

Shubhabrata Sengupta, University of California at Davis

Goal

- Build library of primitives to enable *interesting* algorithms on graphics processors

Why?

- Graphics Processors are powerful general purpose computing platforms
- But lack primitives to build high-level algorithms

But…

- GPGPU has been around for a while…
- But usually deals with problems with regular data access patterns which map well and easily to graphics processors

However…

- Interesting problems require *irregular* data access
- Sorting, building spatial data structures, sparse linear systems…

Primitives to the rescue

- Prefix-sum and associated primitives enable such algorithms

Primitives, Implemented

- CUDPP – CUDA data parallel primitives library
  http://www.gpgpu.org/developer/cudpp/
- Efficient implementations of scan, compact and sort

Spatial Data Structures

- Building them using data parallel primitives
- Linear BVH with Christian Lauterbach and Michael Garland
- Focus of this summer's work at Microsoft Research with Hugues Hoppe

Future Work

- Other spatial data structures
- Primitives for sparse linear systems

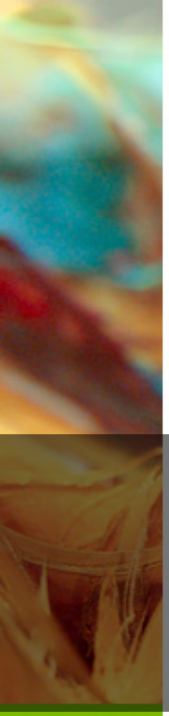

NVISION 08
THE WORLD OF VISUAL COMPUTING

Jay Steele, Clemson Univeristy

NVIDIA.

GPU Clusters

- Purple Hardware : 48 Quadro FX 5600
- Purple Software : eases distribution of CUDA kernels across GPUs in a cluster
  - Lattice-Boltzmann methods
  - Ray tracing
- New work : high level language for scientists
  - Auto parallelization for specific domains
  - Power of GPU clusters to non-CS scientists