SAND2023-07725C

SIGGRAPH 2023
LOS ANGELES+ 6-10 AUG

THE PREMIER CONFERENCE & EXHIBITION ON
COMPUTER GRAPHICS & INTERACTIVE TECHNIQUES

VTK-M

VISUALIZATION FOR THE EXASCALE ERA AND BEYOND

Sandia National Laboratories
U.S. DEPARTMENT OF ENERGY NNSA

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

ACKNOWLEDGEMENTS

SIGGRAPH 2023
LOS ANGELES+ 6-10 AUG

- This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, under Award Numbers 10-014707, 12-015215, and 14-017566.
- This research was supported by the Exascale Computing Project (17-SC-20-SC), a collaborative effort of two U.S. Department of Energy organizations (Office of Science and the National Nuclear Security Administration) responsible for the planning and preparation of a capable exascale ecosystem, including software, applications, hardware, advanced system engineering, and early testbed platforms, in support of the nation's exascale computing imperative.
- Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC., a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA-0003525.
- **Thanks to many, many partners in labs, universities, and industry.**

ECP EXASCALE COMPUTING PROJECT

ASC

Sandia National Laboratories

TERMINOLOGY

- FLOP – Floating Point Operations per Second
- teraflop – 10^{12} floating point operations per seconds
- petaflop – 10^{15} FLOPS
- exaflop – 10^{18} FLOPS
- Exascale – A system capable of achieving an exaflop (measured by the HPLinpack benchmark)

- For Comparison:
  - Current generation AMD/Intel CPUs – 100s of gigaflops
  - Current generation AMD/NVidia GPUs - ~100 teraflops

DOE EXASCALE CLASS COMPUTING SYSTEMS

FRONTIER (OAK RIDGE)

- 9472 AMD Epyc CPUs (>600,000 cores)
- 37,888 AMD MI250X GPUs (> 8.3 million cores)
- 74 Cabinets; Power consumption 21 MW
- World's first exascale computer
  - 1.1 exaflops sustained
  - 1.67 exaflops peak

AURORA (ARGONNE)

- 21,248 Intel Xeon CPU Max Series
- 63,744 Intel Data Center GPU Max Series
- Theoretical Peak: > 2 exaflops
- Fun facts:
  - 300 miles of optical cable
  - 44,000 gallons of water for cooling
  - Weighs > 600 tons

ORIGINS OF VTK-M

- Goals
  1. A single place for the visualization community to collaborate, contribute, and leverage massively threaded algorithms.
  2. Reduce the challenges of writing highly concurrent algorithms by using data parallel algorithms.
  3. Make it easier for simulation codes to take advantage these parallel visualization and analysis tasks on a wide range of current and next-generation hardware.

ORIGINS OF VTK-M

- Combined the strengths of three different projects
  - Extreme-scale Analysis and Visualization Library (EAVL), ORNL
    - New mesh layouts, memory efficiency, parallel algorithms, zero-copy for In-Situ support
  - Data Analysis Toolkit for Extreme Scale (Dax), SNL
    - ParaView plugin, large volumes through streaming to threaded filters
  - Piston, LANL
    - Data-parallel algorithms for many/multi-core, in-situ focused

VTK-M ARCHITECTURE

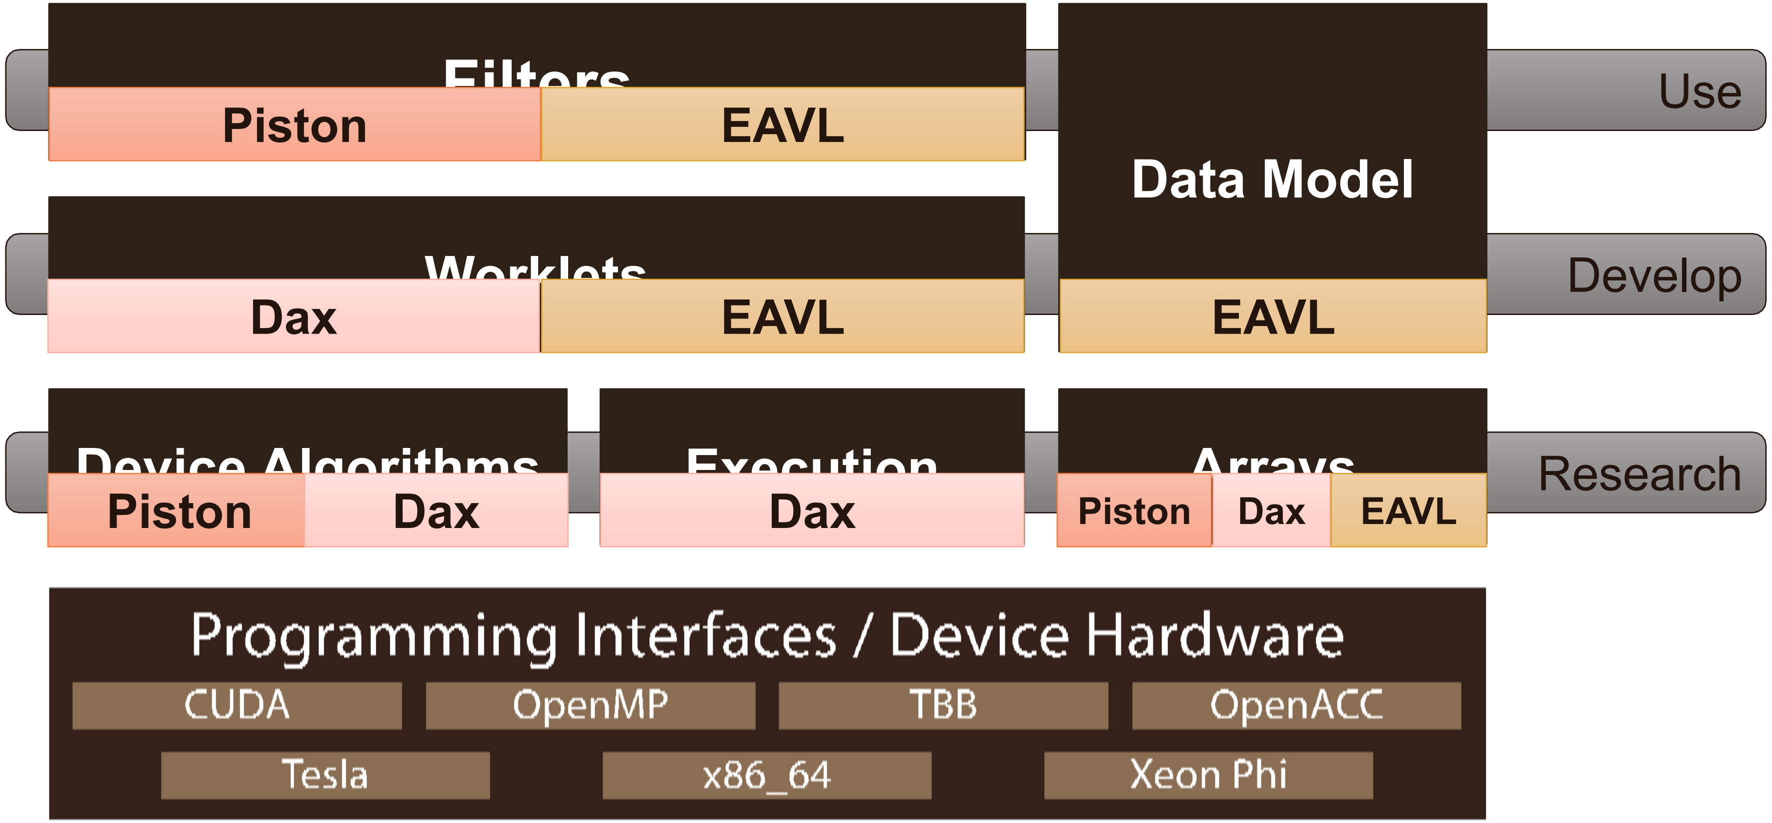

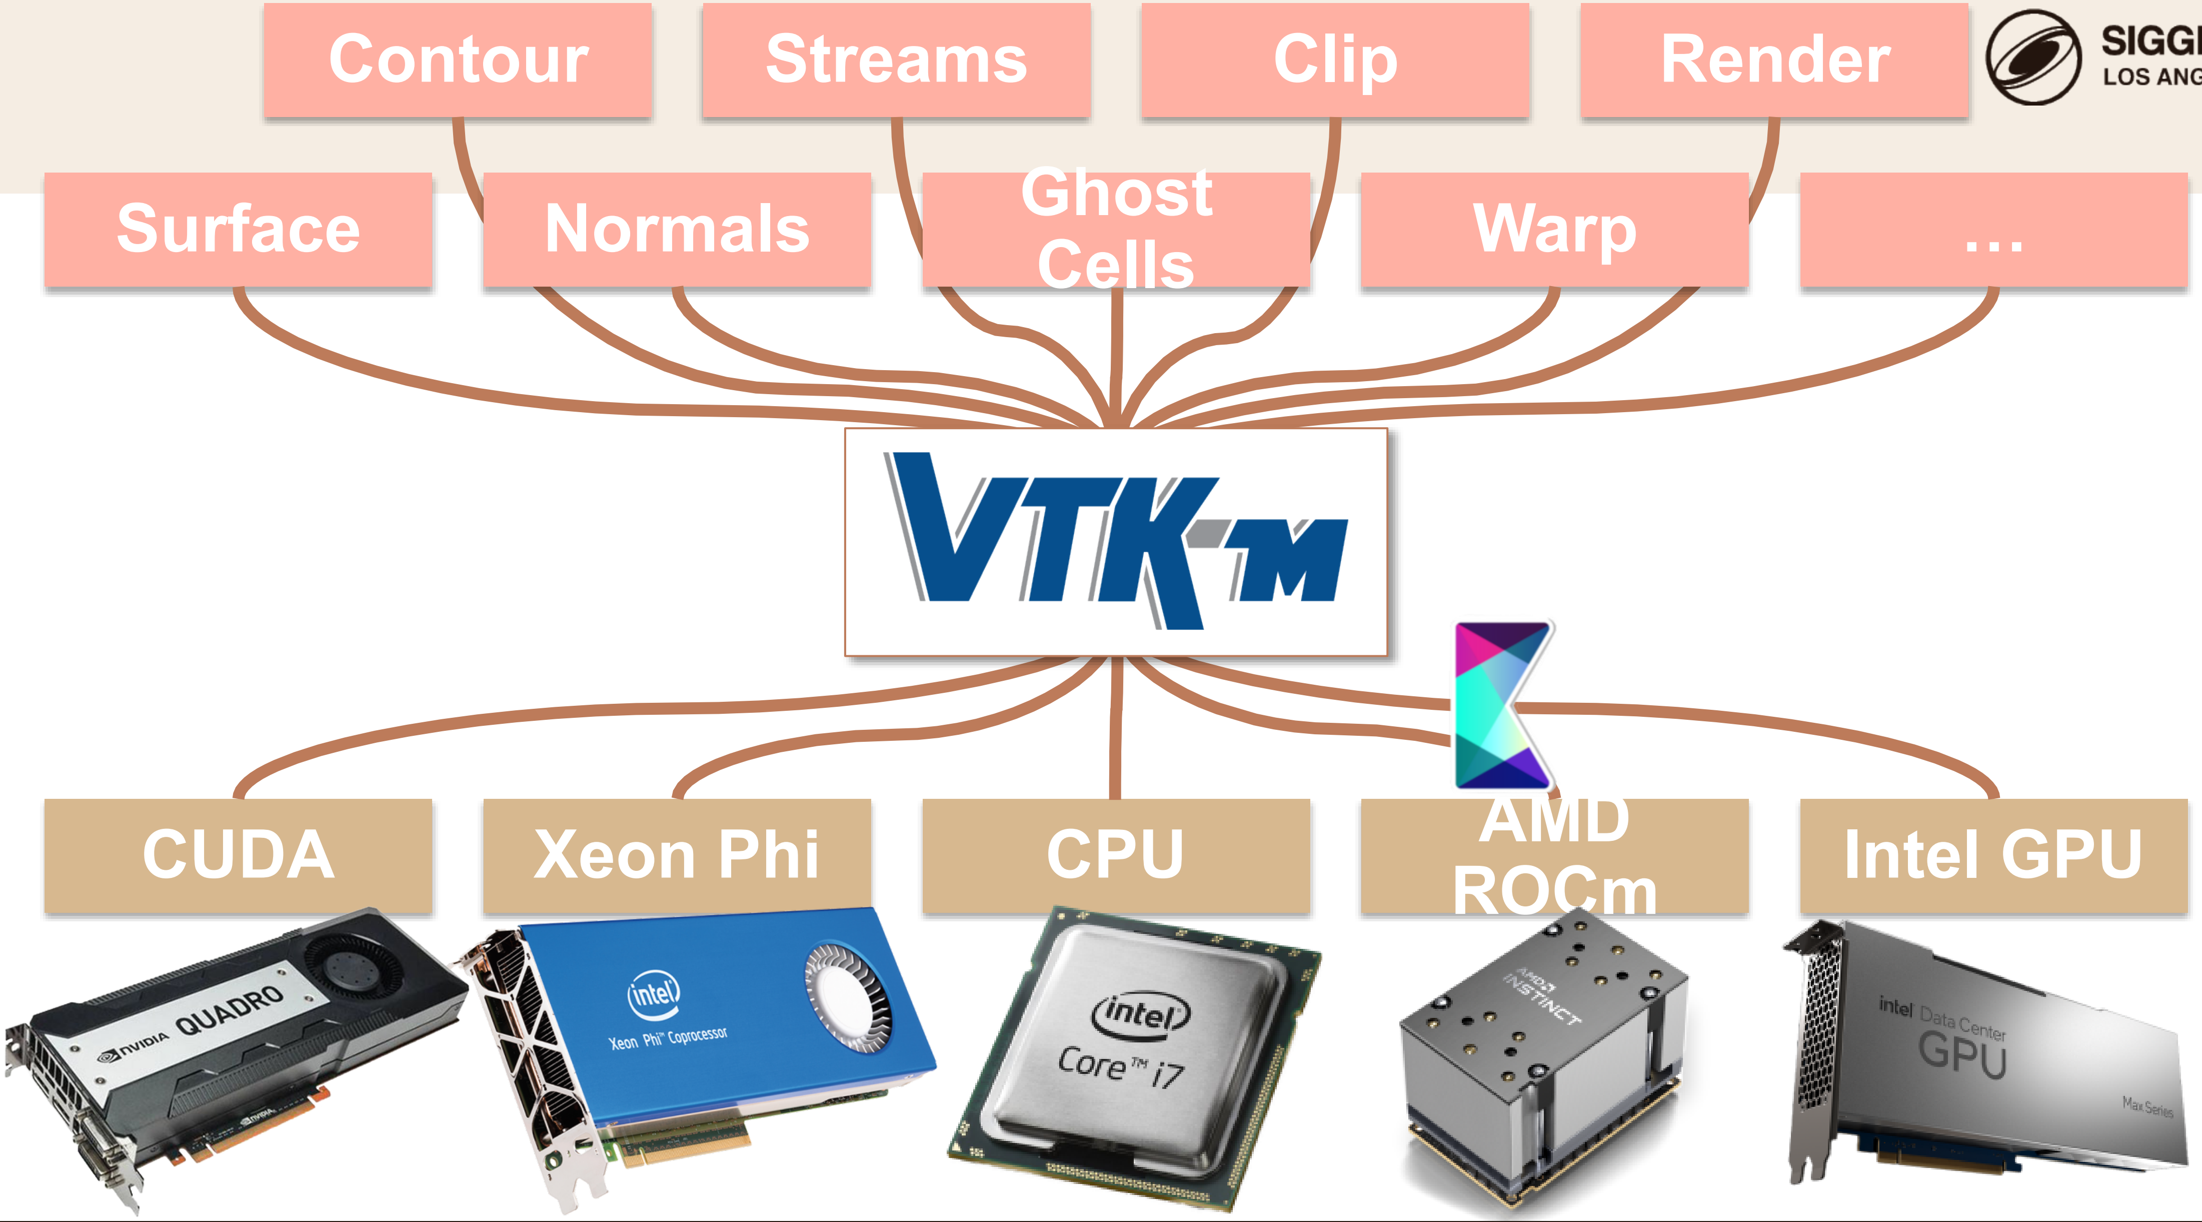
Contour
Streams
Clip
Render
Surface
Normals
Ghost Cells
Warp
…
VTK-m
CUDA
Xeon Phi
CPU
AMD ROCm
Intel GPU
QUADRO
intel
Xeon Phi Coprocessor
intel
Core™ i7
AMD INSTINCT
intel Data Center GPU
Max Series

MCD³ ARCHITECTURE

- **Meta-DPPs**: which are parallel processing patterns that involve one or more DPPs. The word choice of "meta" is meant to evoke its definition of "denoting something of a higher or second-order kind."
- **Convenience routines:** which encapsulate common operations for scientific visualization.
- **DPPs:** which provide parallel processing patterns.
- **Data management:** which insulates algorithms from data layout complexities. These complexities range from how data is organized (e.g., structure-of-arrays vs array-of-structures) to different types of meshes (e.g., unstructured, rectilinear, etc.) to different memory spaces (e.g., host memory, device memory, or unified managed memory)
- **Devices:** which enable code to run on a given hardware architecture.

K. Moreland, R. Maynard, D. Pugmire, A. Yenpure, A. Vacanti, M. Larsen, and H. Childs. Minimizing Development Costs for Efficient Many-Core Visualization Using MCD3. Parallel Computing, 108:102834, Dec. 2021.

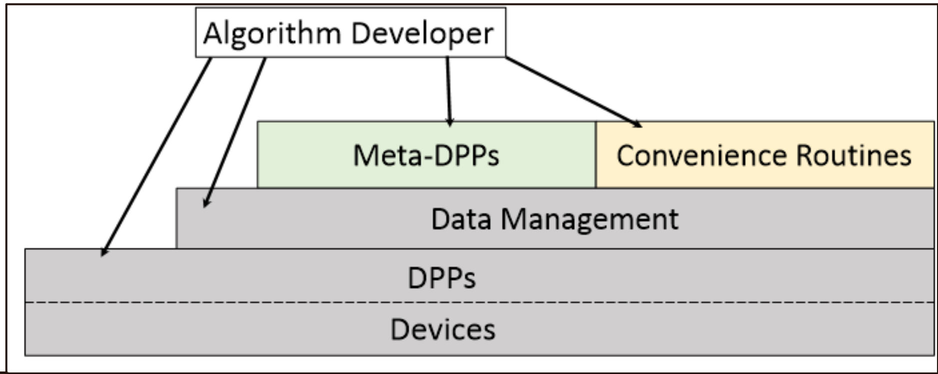

MAP FIELD

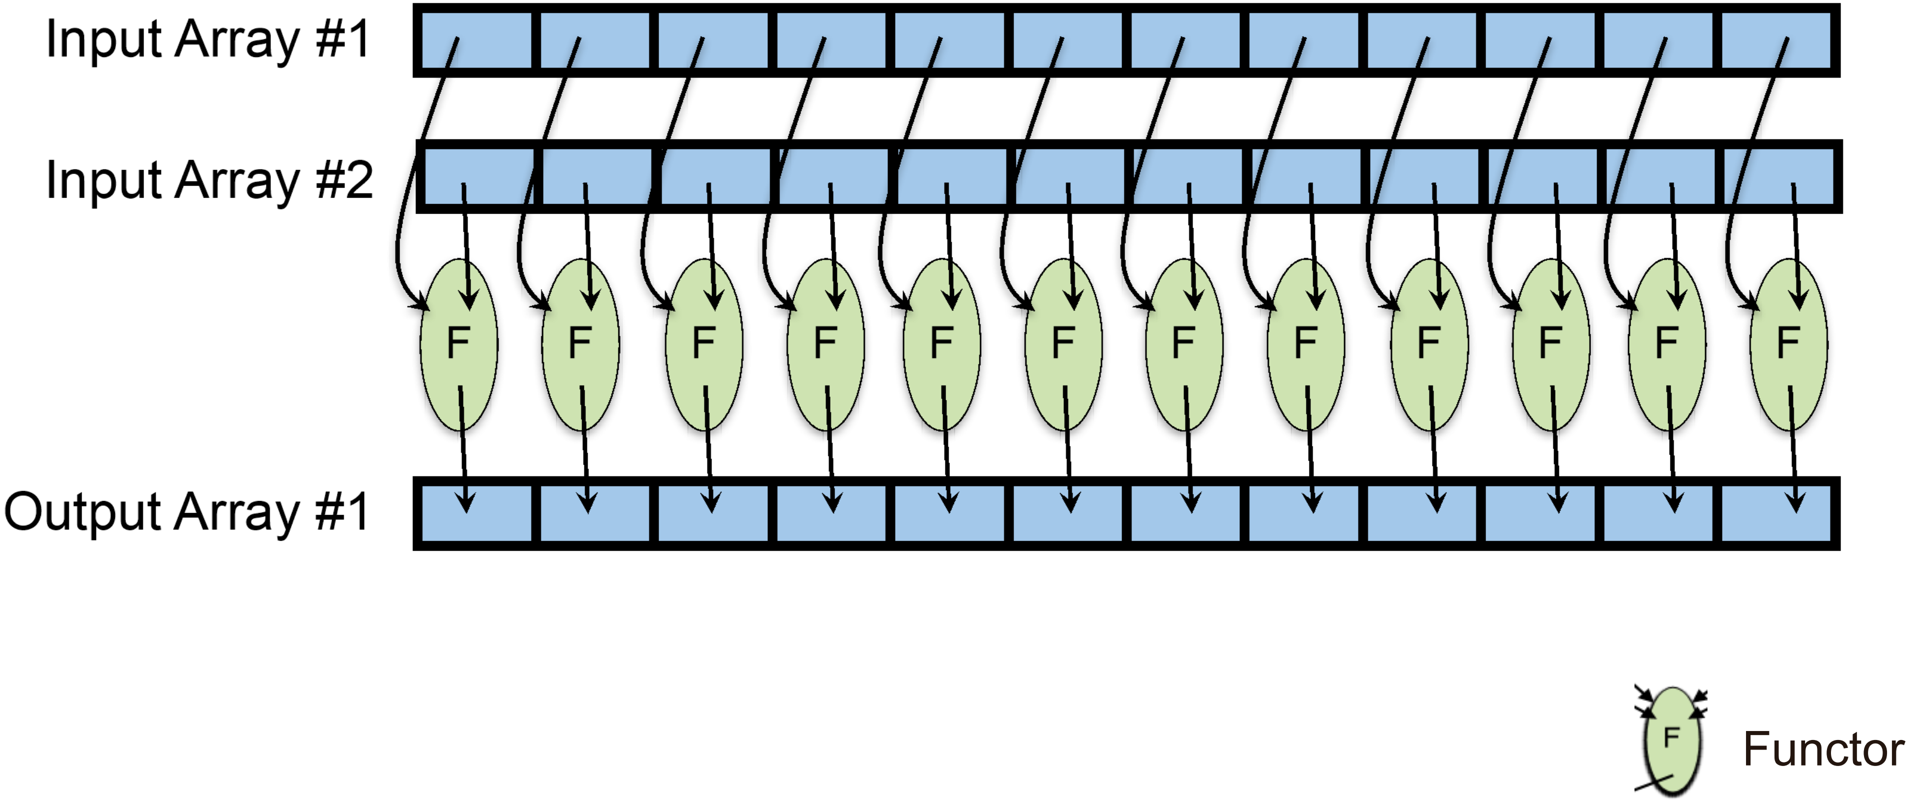
Input Array #1
Input Array #2
F F F F F F F F F F F F
Output Array #1
F
Functor

POINT NEIGHBORHOOD

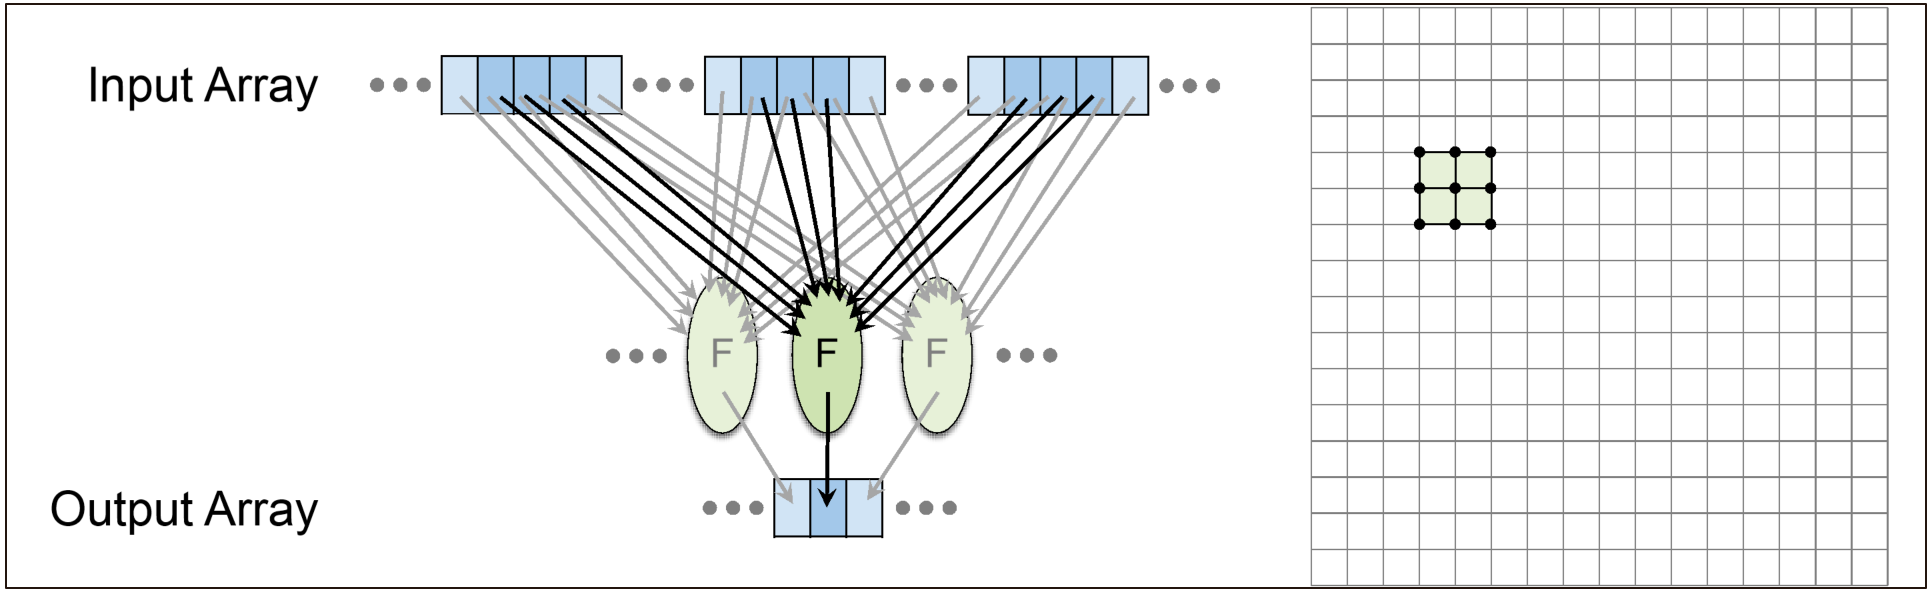

Functor

REDUCE BY KEY

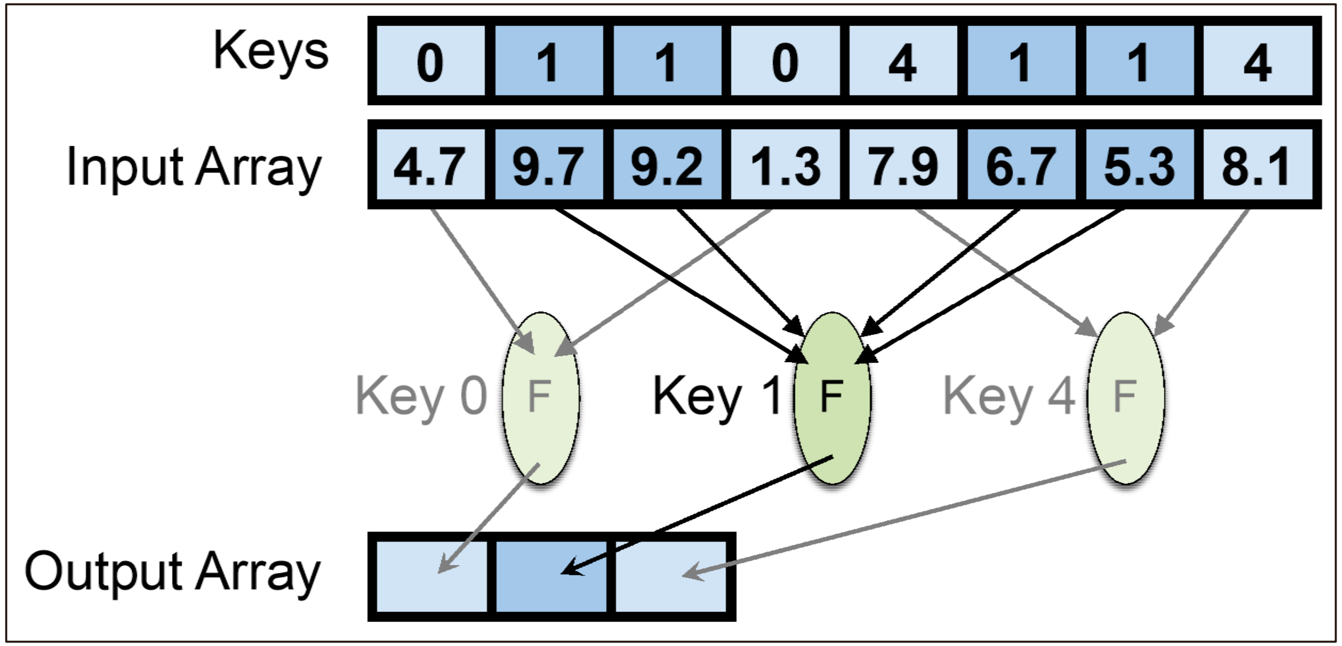

VISIT POINT WITH CELLS

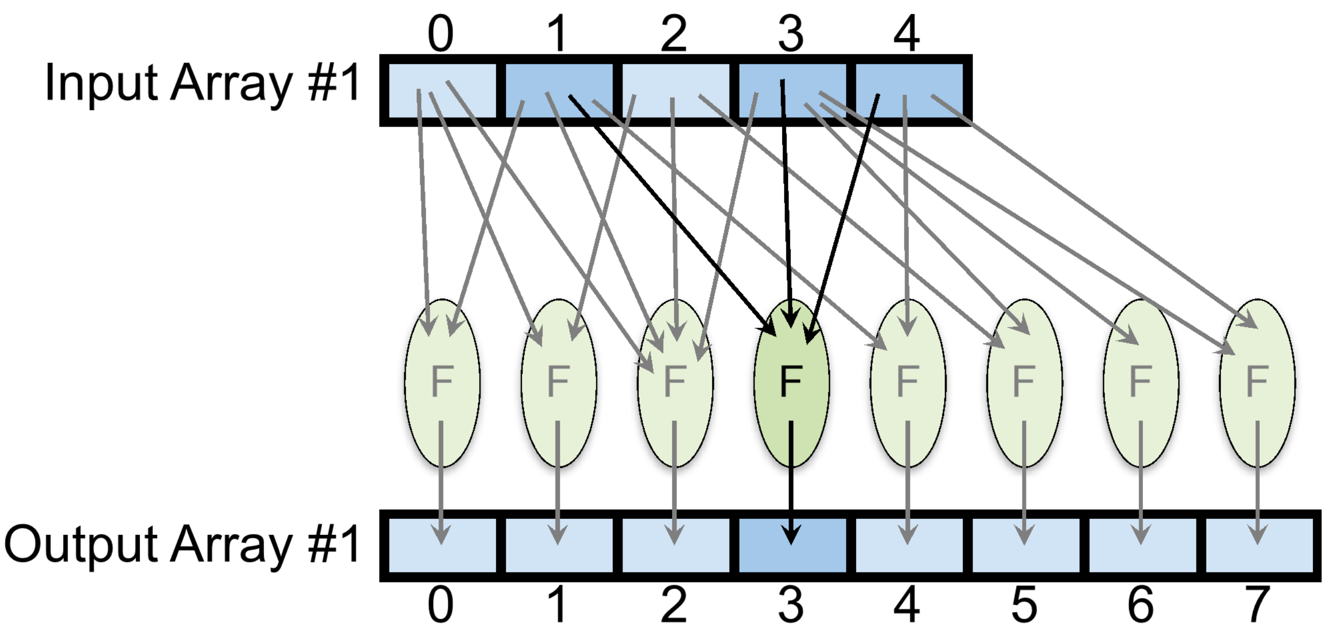

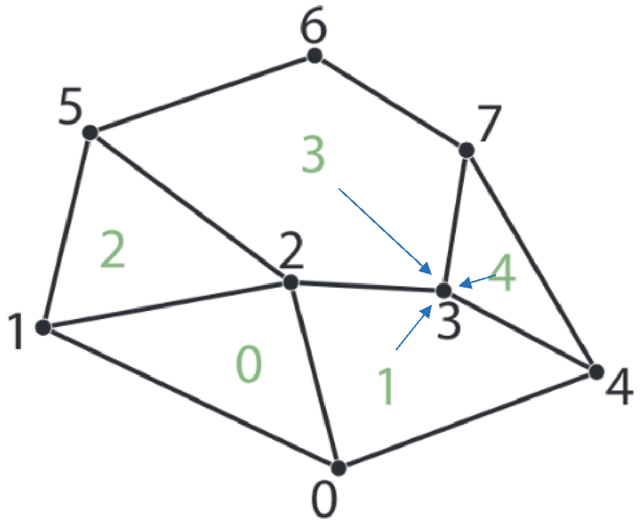

VISIT CELL WITH POINTS

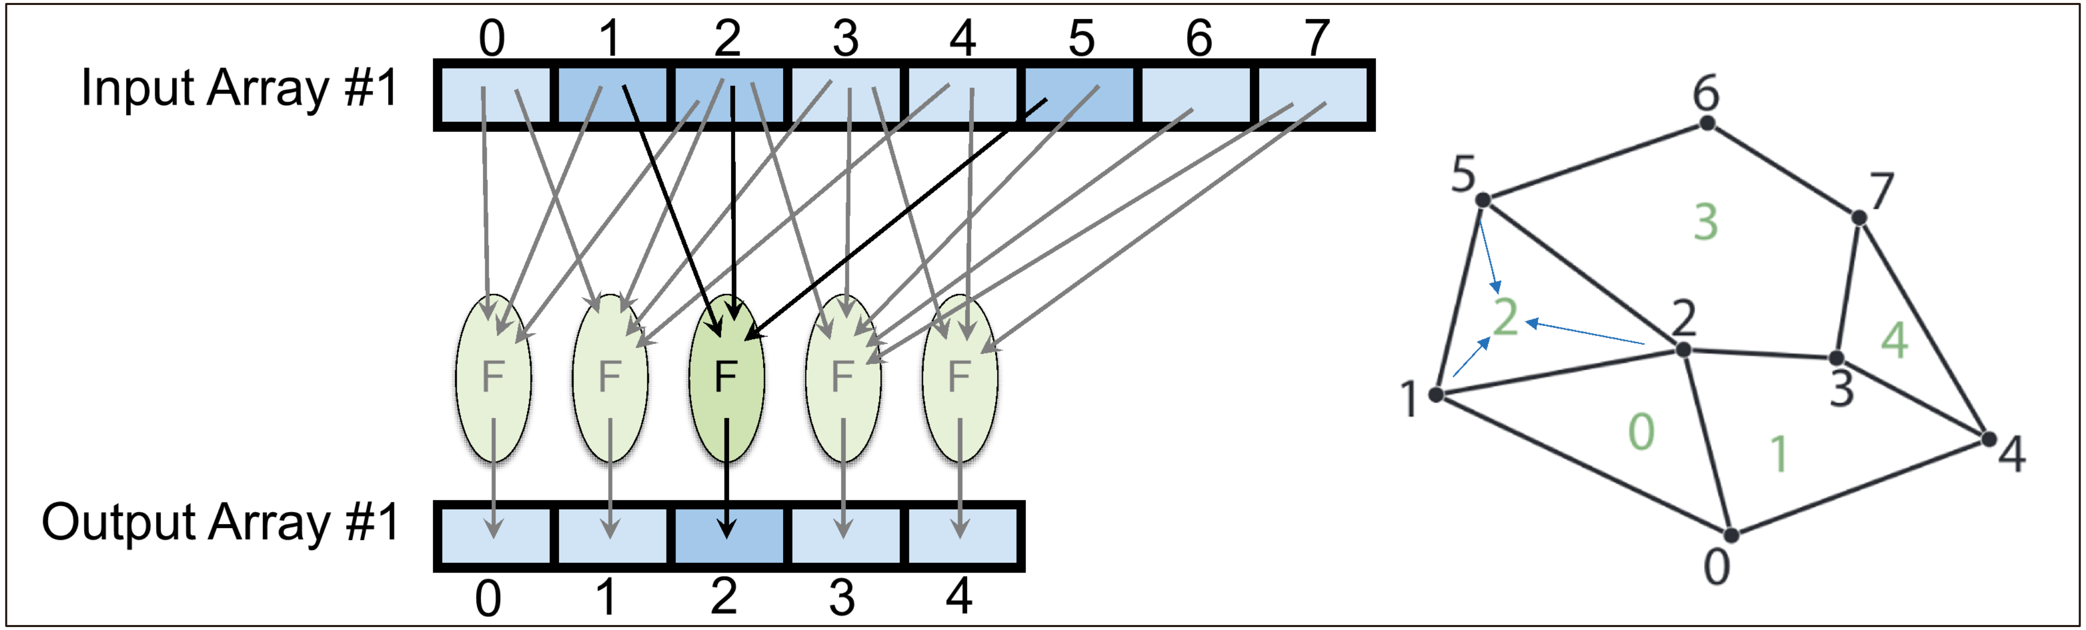

Functor

DOES IT SCALE?

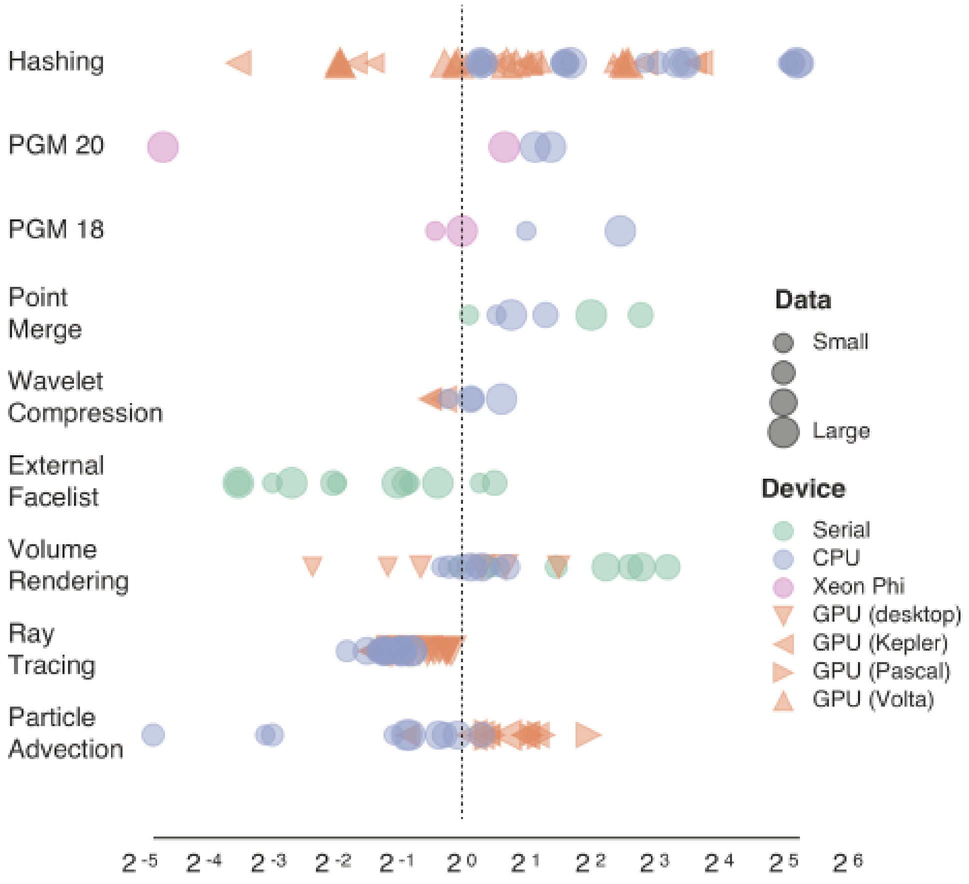

Algorithm	CPUs	GPUs	X. Phi	Serial	Total
External facelist	-	-	-	0.34	0.34
PGM 18	3.32	-	0.87	-	1.69
PGM 20	2.39	-	0.25	-	0.78
Particle advection	0.38	1.53	-	-	0.76
Point merge	1.82	-	-	3.10	2.38
Ray tracing	0.47	0.55	-	-	0.51
Volume rendering	1.13	0.83	-	3.10	1.43
Wavelet compression	1.13	0.75	-	-	0.92
Hashing	5.97	1.45	-	-	2.94
Total	1.45	0.95	0.47	1.48	**1.14**


SIGGRAPH 2023
LOS ANGELES+ 6-10 AUG

RECENT RESULTS




ACCELERATING PARAVIEW

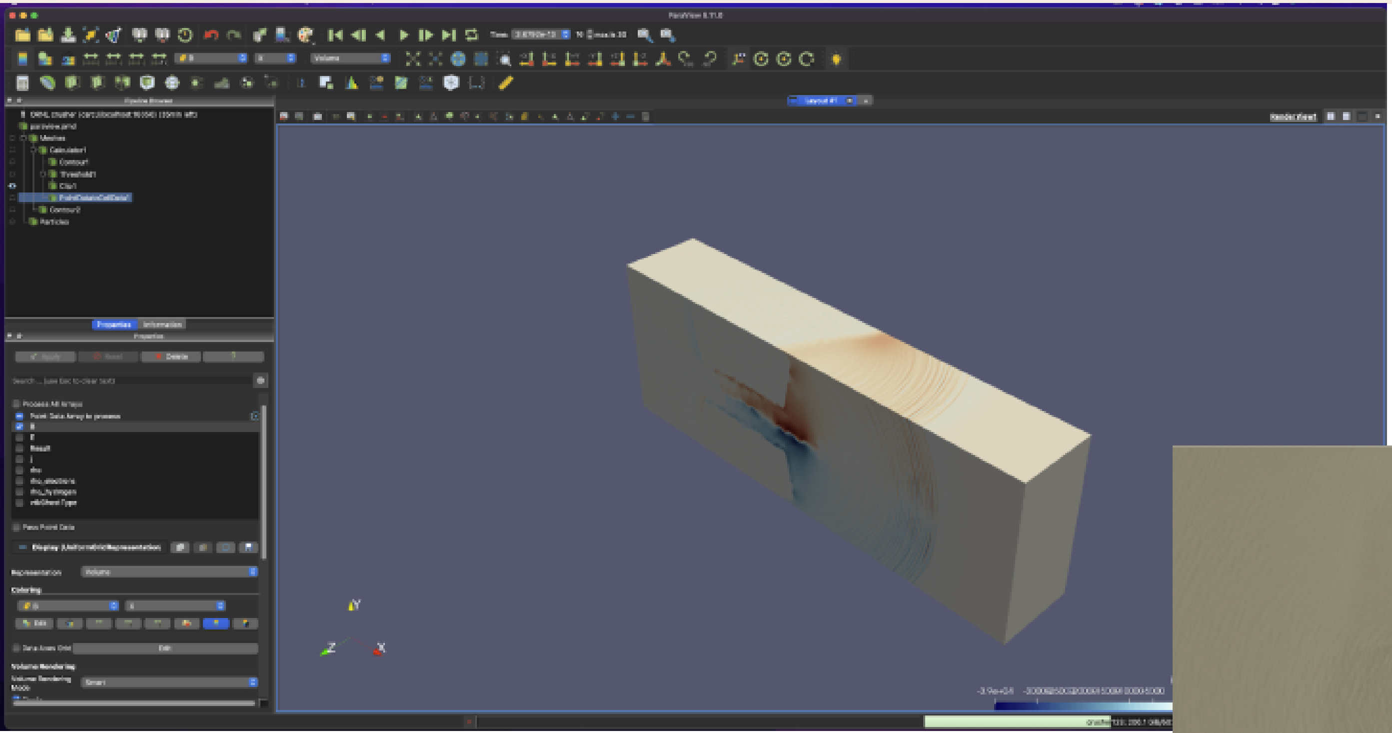

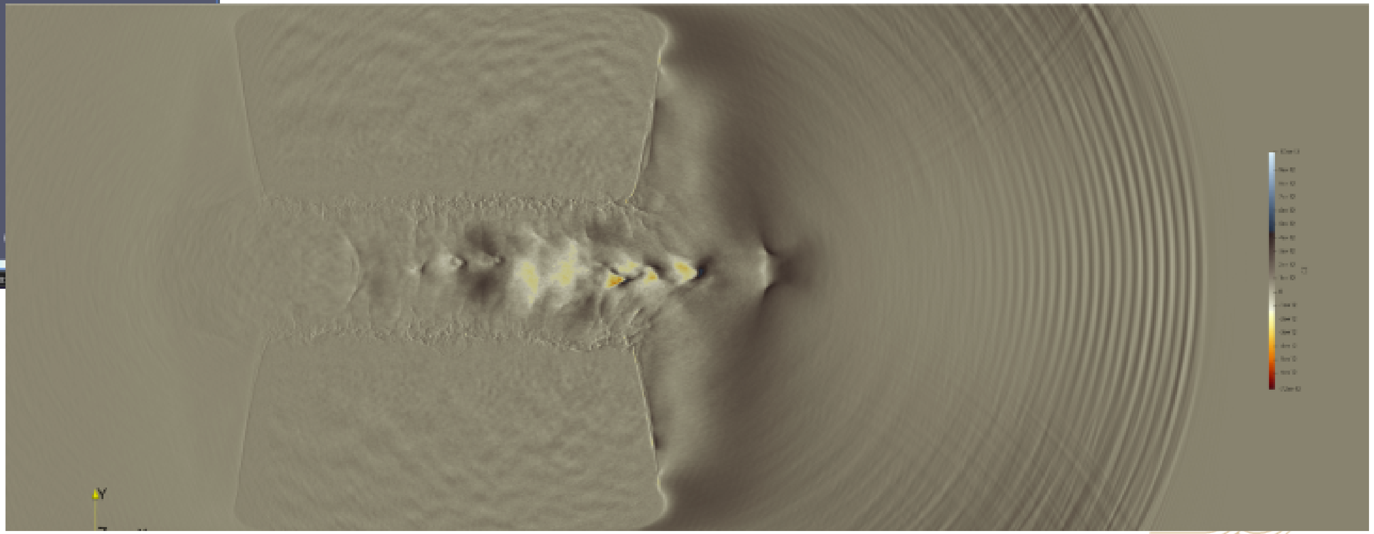

ACCELERATING POINCARÉ PLOTS

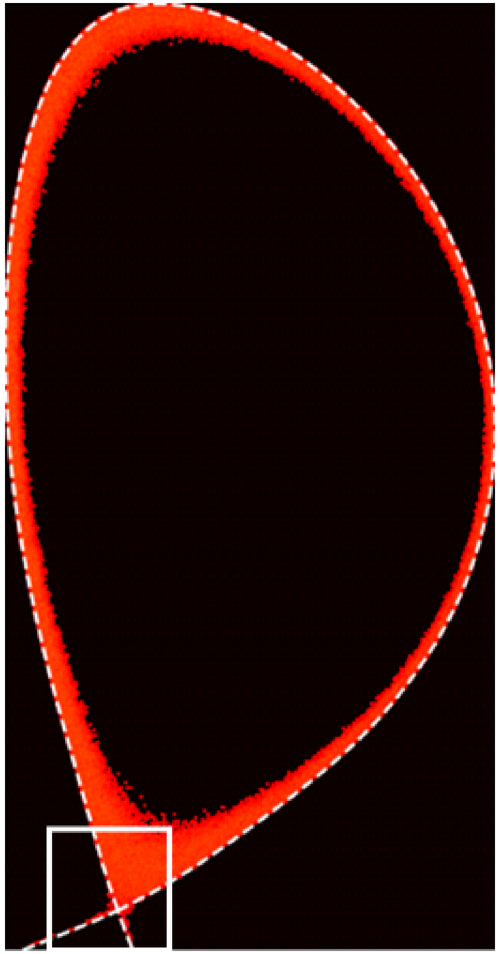

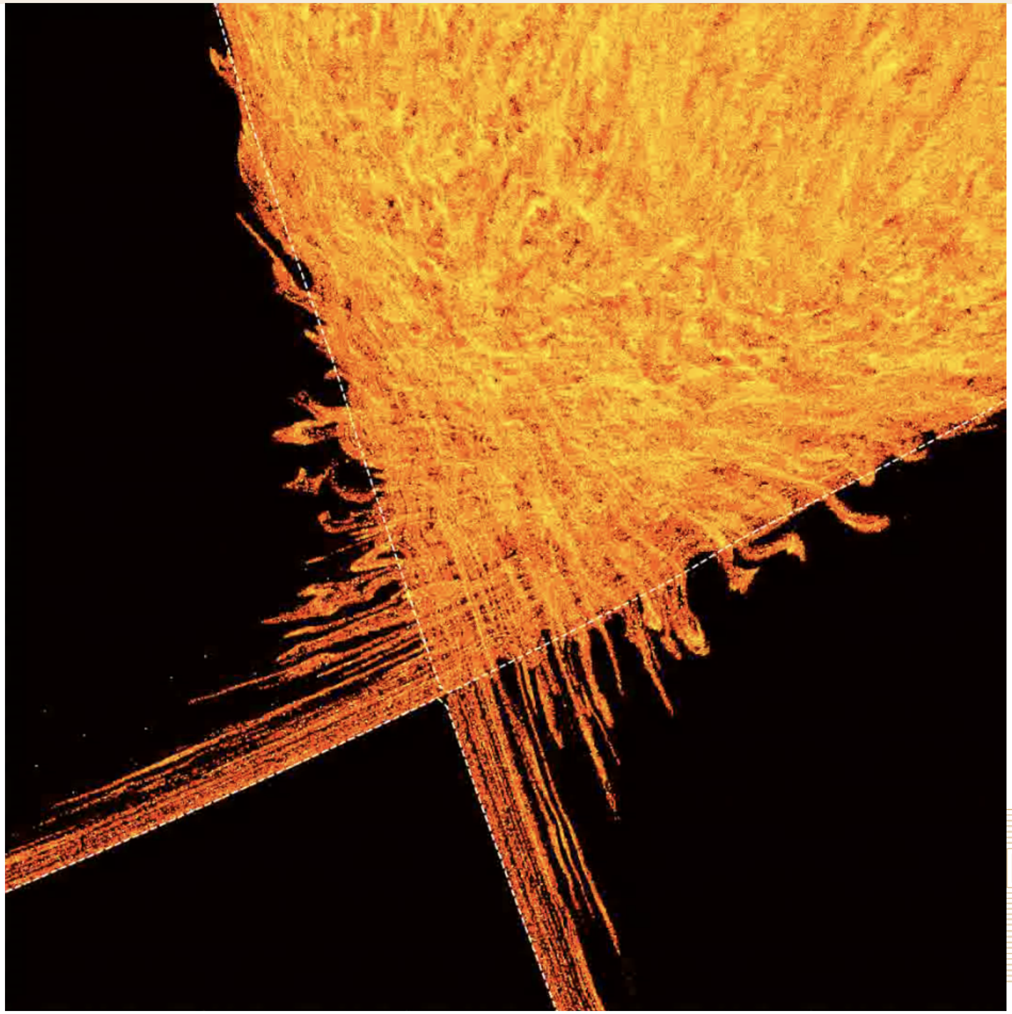

VOLUME RENDERING WITH SHADOWS

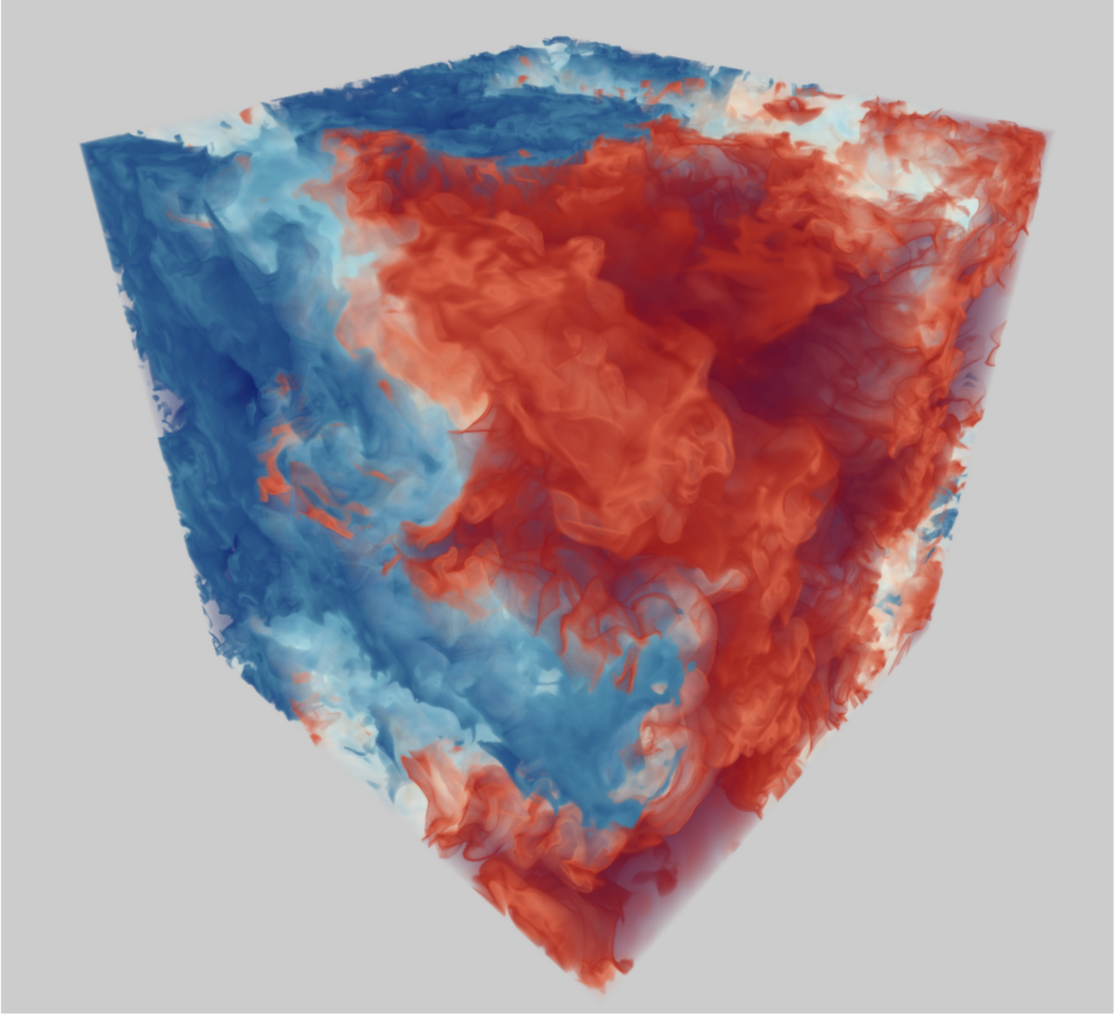

- M. Mathai, M. Larsen, and H. Childs. A Distributed-Memory Parallel Approach for Volume Rendering with Shadows. To appear at the IEEE Symposium for Large Data Analysis and Visualization (LDAV), October 2023.

VOLUME RENDERING WITH SHADOWS

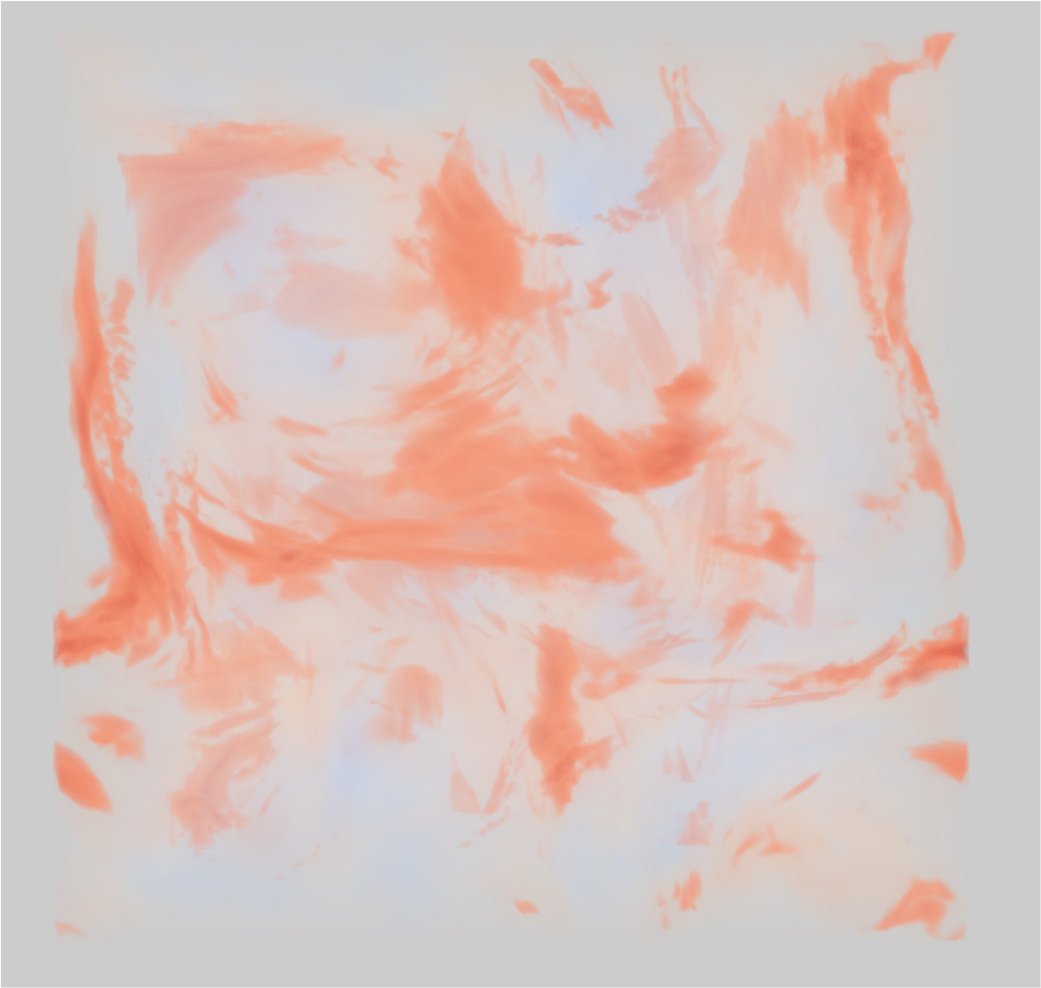

- M. Mathai, M. Larsen, and H. Childs. A Distributed-Memory Parallel Approach for Volume Rendering with Shadows. To appear at the IEEE Symposium for Large Data Analysis and Visualization (LDAV), October 2023.

RENDERING AT SCALE ON FRONTIER