LA-UR-12-26127

Approved for public release; distribution is unlimited.

The PISTON Software Framework for Visualization and Analysis on Next-Generation Multi-core Architectures
Christopher Sewell Li-ta Lo James Ahrens
Visualization Frameworks for Multi-Core and Many-Core Architectures Panel, The International Conference for High Performance Computing, Networking, Storage, and Analysis, November 14, 2012



Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by the Los Alamos National Security, LLC for the National Nuclear Security Administration of the U.S. Department of Energy under contract DE-AC52-06NA25396. By acceptance of this article, the publisher recognizes that the U.S. Government retains a nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or to allow others to do so, for U.S. Government purposes. Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy. Los Alamos National Laboratory does not endorse the viewpoint of a publication or guarantee its technical correctness.

The PISTON Software Framework for Visualization and Analysis on Next-Generation Multi-core Architectures

### Chris Sewell, Li-Ta Lo, and James Ahrens Los Alamos National Laboratory













### Overview

Goal: Portability and performance for visualization and analysis operators on current and next-generation supercomputers

 Main idea: Write operators using only data-parallel primitives (scan, reduce, etc.)

 Requires architecture-specific optimizations for only for the small set of primitives

PISTON is built on top of NVIDIA's Thrust Library



### Motivation and Background

- Current production visualization software does not take full advantage of acceleration hardware and/or multi-core architecture
- Research on accelerating visualization operations are mostly hardware-specific; few were integrated in visualization software
- Standards such as OpenCL may allow program to run cross-platform, but usually still requires many architecture specific optimizations to run well
- Data parallelism: independent processors performs the same task on different pieces of data (see Blelloch, "Vector Models for Data Parallel Computing")
- Due to the massive data sizes we expect to be simulating we expect data parallelism to be a good way to exploit parallelism on current and next generation architectures
- Thrust is a NVidia C++ template library for CUDA. It can also target other backends such as OpenMP, and allows you to program using an interface similar the C++ Standard Template Library (STL)



### Questions for the Panel

- What fundamental problem are you trying to solve?
  - Portable performance for visualization and analysis operators
- What are your plans to deal with exascale-specific issues (massive concurrency, distributed memory, memory overhead, fault tolerance)?
  - Data-parallel programming model; requires architecture-specific optimizations for a limited set of "embarrassingly parallel" primitive operators
- What is your philosophy for dealing with ambiguity of the exascale architecture (multiple swim lanes, heterogeneous architectures)?
  - Data-parallel programming model; requires architecture-specific optimizations for a limited set of "embarrassingly parallel" primitive operators
- How is your technology implemented?
  - Extension of NVIDIA's Thrust library
- What is the long-term result for this effort? (Production software? Research prototype?)
  - Open-source repository (available on-line now); integrated into VTK/ParaView (available on-line now)



### Videos of PISTON in Action







# Brief Introduction to Data-Parallel Programming and Thrust

### What algorithms does Thrust provide?

• Sorts	input	4	5	2	1	3
• Transforms	 transform(+1)				2	 4
Reductions	inclusive_scan(+)	4	9	11	12	15
	exclusive_scan(+)	0	4	9	11	12
• Scans	exclusive_scan(max)	0	4	5	5	5
	transform_inscan(*2,+)	8	18	22	24	30
<ul> <li>Binary searches</li> </ul>	for_each(-1)	3	4	1	0	2
	sort	1	2	3	4	5
Stream compactions	copy_if(n % 2 == 1)	5	1	3		
	reduce(+)					15
Scatters / aathers						
• Scallers / gallers	input1	0	0	2	4	8
	input2	3	4	1	0	2
Challenge: Write operators in terms of						
these primitives only	upper_bound	3	4	2	2	3
	permutation_iterator	4	8	0	0	2

#### **Reward: Efficient, portable code**



### Isosurface with Marching Cube – the Naive Way

- Classify all cells by transform
- Use copy\_if to compact valid cells.
- For each valid cell, generate same number of geometries with flags.
- Use copy\_if to do stream compaction on vertices.
- This approach is too slow, more than 50% of time was spent moving huge amount of data in global memory.
- Can we avoid calling copy\_if and eliminate global memory movement?





## Isosurface with Marching Cube – Optimization

- Inspired by HistoPyramid
- The filter is essentially a mapping from input cell id to output vertex id
- Is there a "reverse" mapping?
- If there is a reverse mapping, the filter can be very "lazy"
- Given an output vertex id, we only apply operations on the cell that would generate the vertex
- Actually for a range of output vertex ids





## Isosurface with Marching Cubes Algorithm





## Variations on Isosurface: Cut Surfaces and Threshold

- Cut surface
  - Two scalar fields, one for generating geometry (cut surface) the other for scalar interpolation
  - Less than 10 LOC change, negligible
     performance impact to isosurface
  - One 1D interpolation per triangle vertex
- **Threshold** 
  - Classify cells, this time based on whether value at each vertex falls within threshold range, then stream compact valid cells and generate geometry for valid cells
  - Additional pass of cell classification and stream compaction to remove interior cells









## **Additional Operators**

### Blelloch's "Vector Models for Data-Parallel Computing"

Data Structures

Graphs: Neighbor reducing, distributing excess across edges Trees: Leaffix and rootfix operations, tree manipulations Multidimensional arrays Computational Geometry Generalized binary search k-D tree Closest pair Quickhull Merge Hull Graph Algorithms Minimum spanning tree Maximum flow Maximal independent set Numerical Algorithms Matrix-vector multiplication Linear-systems solver Simplex Outer product Sparse-matrix multiplication

### Current prototypes

- Glyphs
- Halo finder for cosmology simulations
- "Boid" simulation (flocking birds)



LOS AIAMOS NATIONAL LABORATORY EST 1943







### **PISTON** Performance



EST. 1943

## Integration with VTK and ParaView

- Filters that use PISTON data types and algorithms integrated into VTK and ParaView
- Utility filters interconvert between standard VTK data format and PISTON data format (thrust device vectors)
- Supports interop for on-card rendering



## Extending PISTON's Portability: Architectures

#### Prototype OpenCL backend

- Successfully implemented isosurface and cut plane operators in OpenCL with code almost identical to that used for the Thrust-based CUDA and OpenMP backends
- With interop on AMD FirePro V7800, we can run at about 6 fps for 256<sup>3</sup> data set (2 fps without interop)

#### Renderer

- Allows generation of images on systems without OpenGL
- Rasterizing and ray-casting versions (using K-D Tree)
- Inter-node parallelism
  - VTK Integration

FST 1943

- Domain partitioned by VTK's MPI libraries
- Each node uses PISTON filters to compute results for its portion of domain
- Results combined by VTK's compositors
- Distributed implementations of Thrust primitives using MPI (in progress)









LA-UR-12-26127

# Extending PISTON's Portability: Data Types

- Curvilinear coordinates
  - Multiple layers of coordinate transformations
  - Due to kernel fusion, very little performance impact
- Unstructured / AMR data
  - Tetrahedralize uniform grid or unstructured grid (e.g., AMR mesh)
  - Generate isosurface geometry based on look-up table for tetrahedral cells
  - Next step: Develop PISTON operator to tetrahedralize grids, and/or to compute isosurface directly on AMR grid







### LA-UR-12-26127

## **PISTON In-Situ**

- VPIC (Vector Particle in Cell) Kinetic Plasma Simulation Code
  - Implemented first version of an in-situ adapter based on Paraview CoProcessing Library (Catalyst)
  - Three pipelines: vtkDataSetMapper, vtkContourFilter, vtkPistonContour
  - CoGL
    - Stand-alone meso-scale simulation code developed as part of the Exascale Co-Design Center for Materials in Extreme Environments
    - Studies pattern formation in ferroelastic materials using the Ginzburg–Landau approach
    - Models cubic-to-tetragonal transitions under dynamic strain loading
    - Simulation code and in-situ viz implemented using PISTON



Output of vtkDataSetMapper and vtkPistonContour filters on Hhydro charge density at one timestep of VPIC simulation







## **PISTON's New Companion Project: PINION**

A portable, data-parallel software framework for physics simulations

- Data structures that allow scientists to program in a way that maps easily to the problem domain rather than dealing directly with 1D host/device vectors
- Operators that provide data-parallel implementations of analysis and computational functions often used in physics simulations
- Backends that optimize implementations of data parallel primitives for one or two emerging supercomputer hardware architectures



FST 1943

## **PISTON Open-Source Release**

- Open-source release
  - Stable tarball: <u>http://viz.lanl.gov/projects/PISTON.html</u>
  - Current repository: <a href="https://github.com/losalamos/PISTON">https://github.com/losalamos/PISTON</a>



### Acknowledgments

- The SciDAC Institute of Scalable Data Management, Analysis and Visualization (SDAV), funded by the DOE Office of Science through the Office of Advanced Scientific Computing Research.
- NNSA ASC CCSE Program
- The Exascale Co-Design Center for Materials in Extreme Environments, funded by the DOE Office of Advanced Scientific Computing Research (ASCR)
- Los Alamos Laboratory Directed Research and Development Program

