Exceptional service in the national interest

Sandia National Laboratories

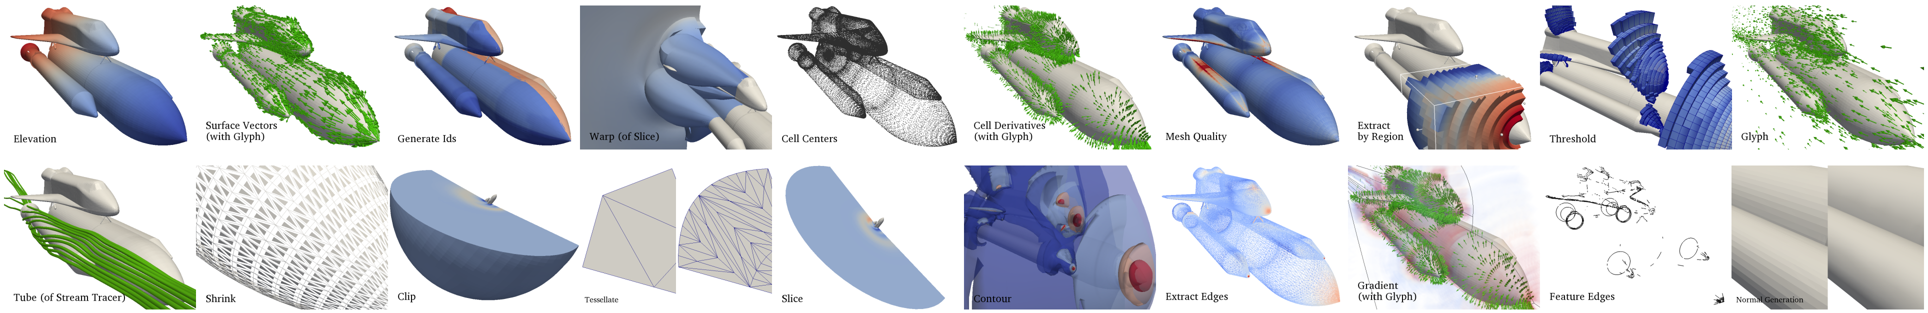

A Classification of Scientific Visualization Algorithms for Massive Threading

Kenneth Moreland Sandia National Laboratories
Berk Geveci Kitware, Inc.
Kwan-Liu Ma University of California at Davis
Robert Maynard Kitware, Inc.

November 17, 2013

U.S. DEPARTMENT OF ENERGY
NNSA National Nuclear Security Administration

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000. SAND NO. 2013-9677P

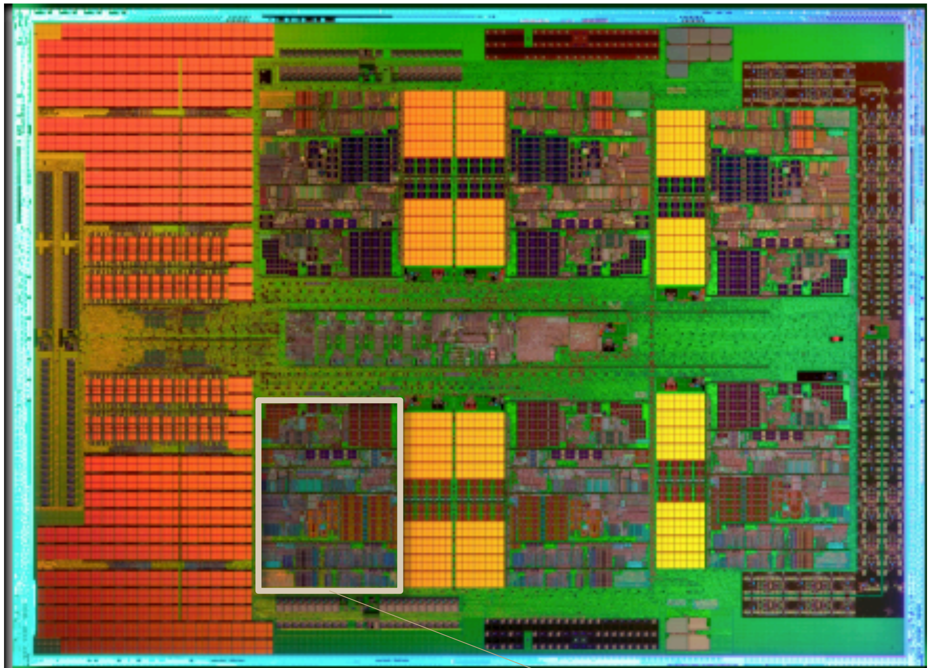

1mm

Cielo
AMD x86

Full x86 Core
+ Associated Cache
8 cores per die
MPI-Only feasible

1 x86
core

1 Kepler
core

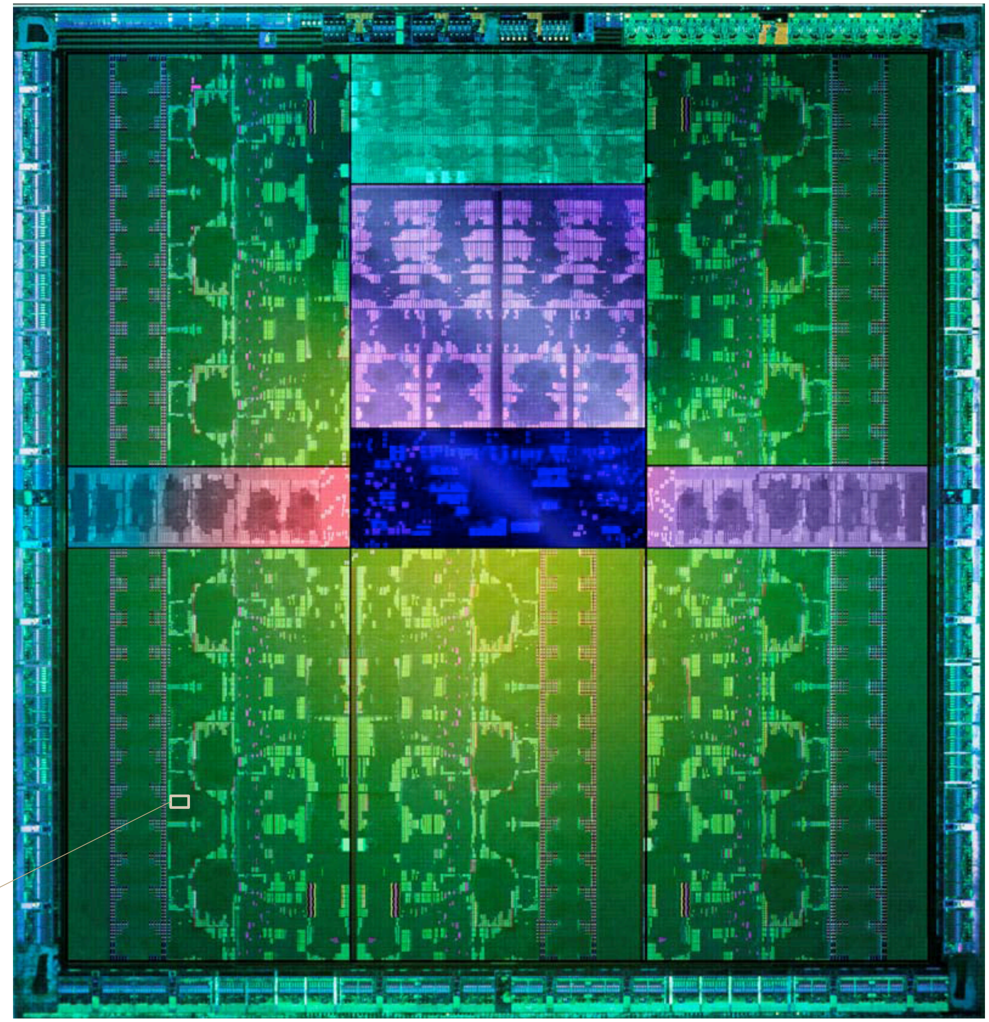

Trinity (option 1)
NVIDIA GPU

2,880 cores collected in 15 SMX
Shared PC, Cache, Mem Fetches
Reduced control logic
MPI-Only not feasible

Extreme Scale is Threads, Threads, Threads!

	Jaguar – XT5	Titan – XK7	Exascale*
Cores	224,256	299,008 and 18,688 gpu	1 billion
Concurrency	224,256 way	70 – 500 million way	10 – 100 billion way
Memory	300 Terabytes	700 Terabytes	128 Petabytes

- To succeed at extreme scale, you need to consider the finest possible level of concurrency
  - Expect each thread to process exactly one element
  - Disallow communication among threads

*Source: Scientific Discovery at the Exascale, Ahern, Shoshani, Ma, et al.

Data Parallel Visualization

- Duplicate processes run independently on different partitions of data.

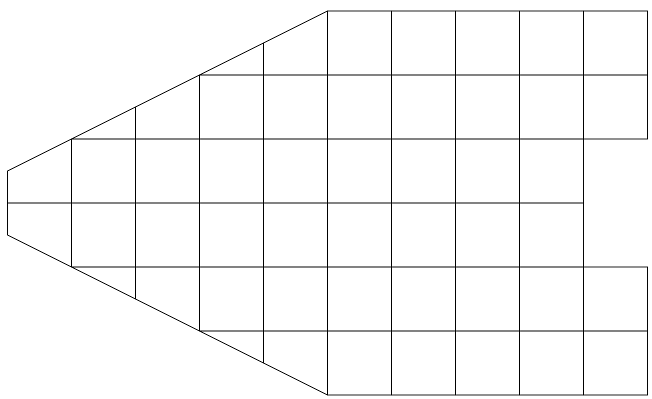

Data Parallel Visualization

- Duplicate processes run independently on different partitions of data.

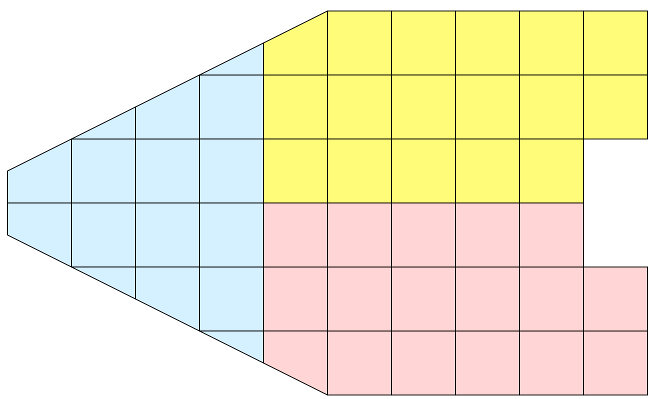

Data Parallel Visualization

- Duplicate processes run independently on different partitions of data.

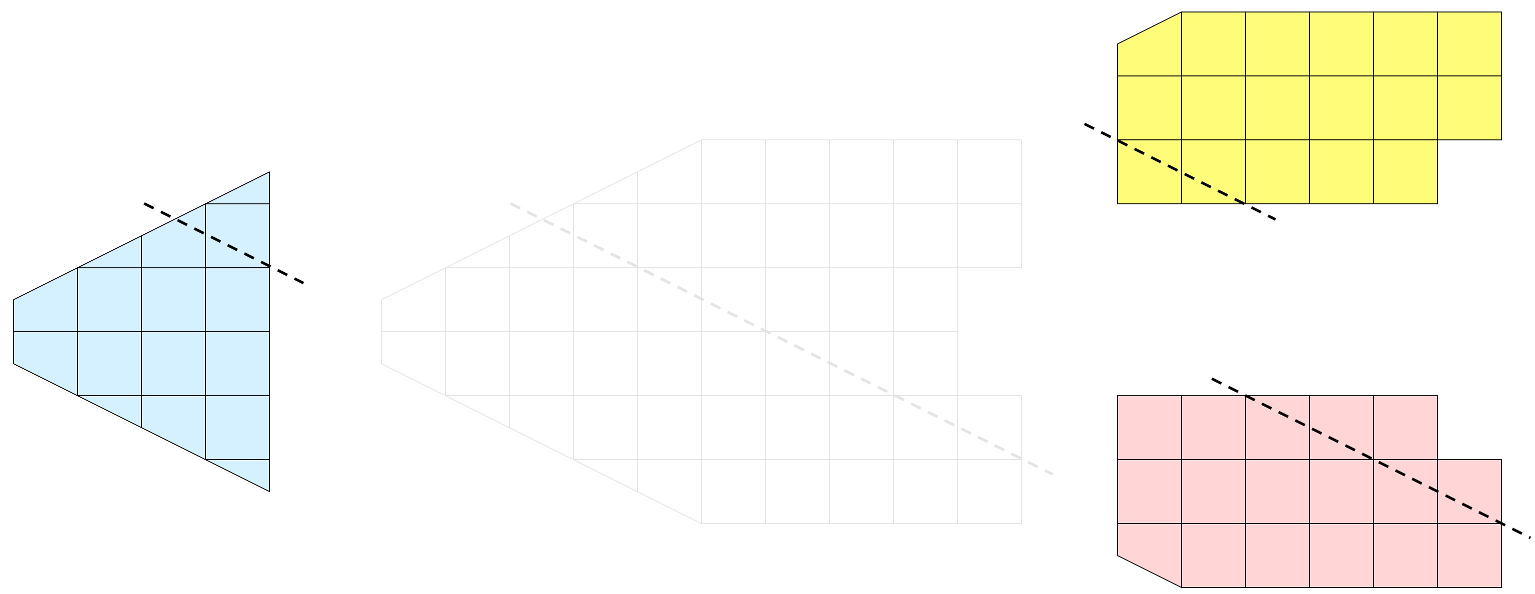

Data Parallel Visualization

- Duplicate processes run independently on different partitions of data.

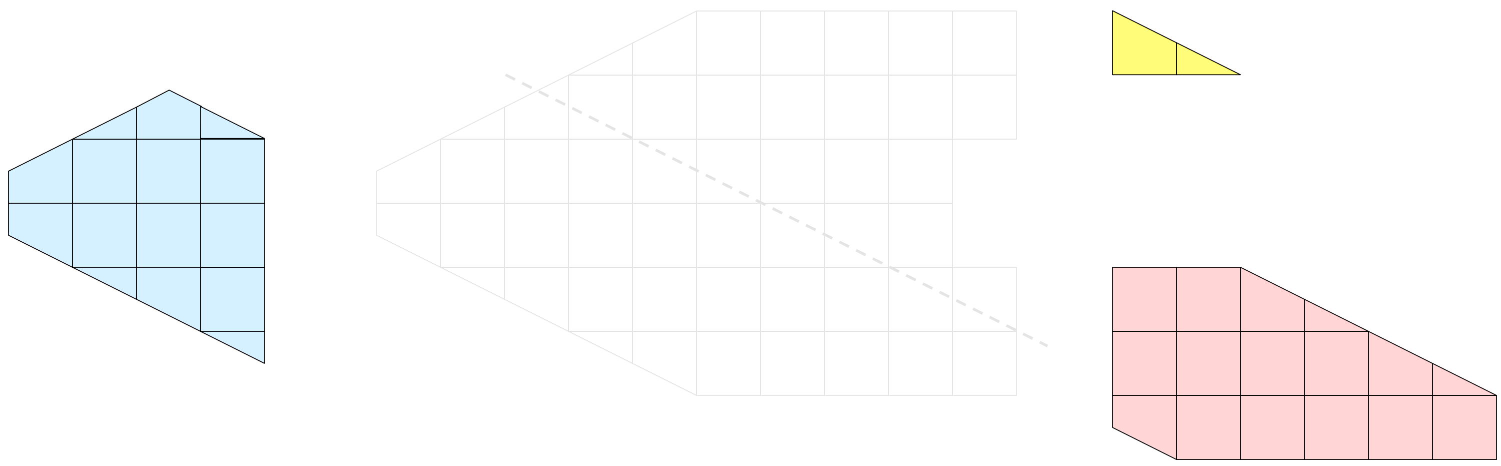

Data Parallel Visualization

- Duplicate processes run independently on different partitions of data.

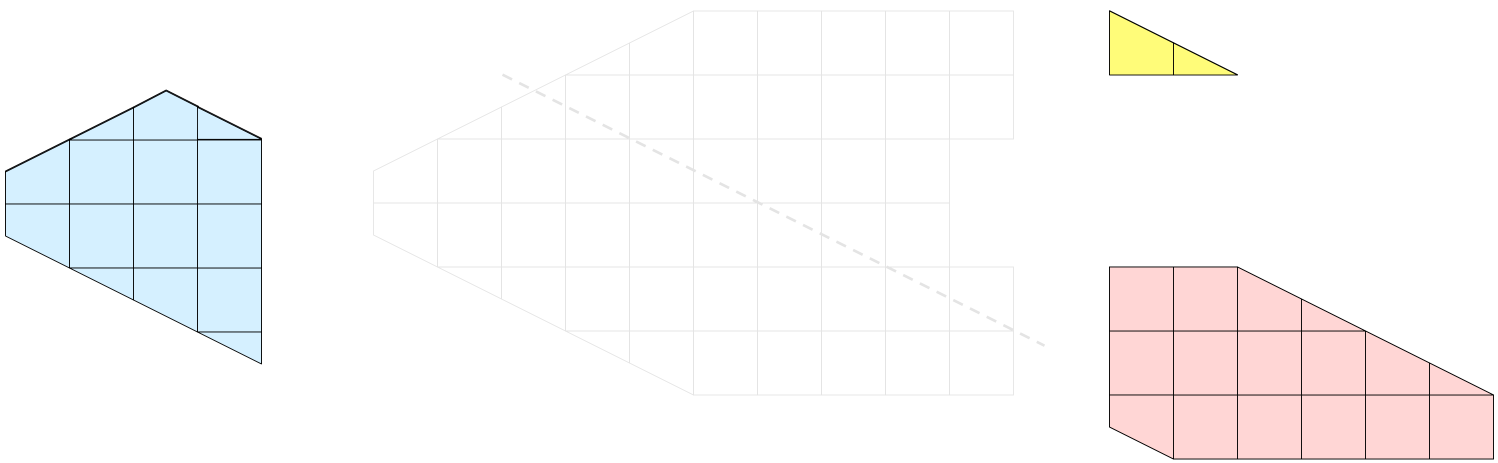

Key Principles

[Law et al., 1999]

Sandia National Laboratories

- Data Separability

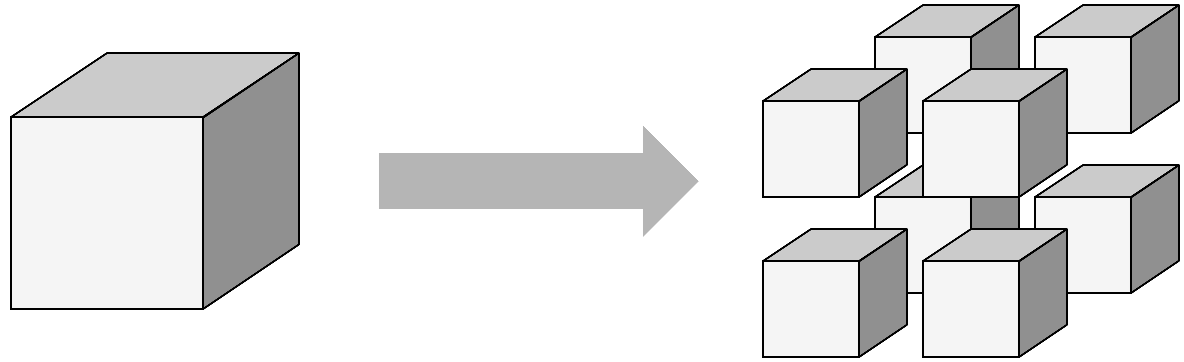

- Mappable Input

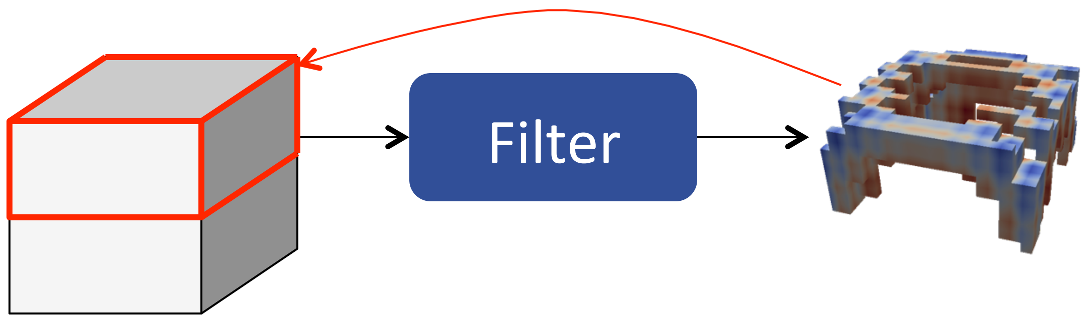

- Result Invariant

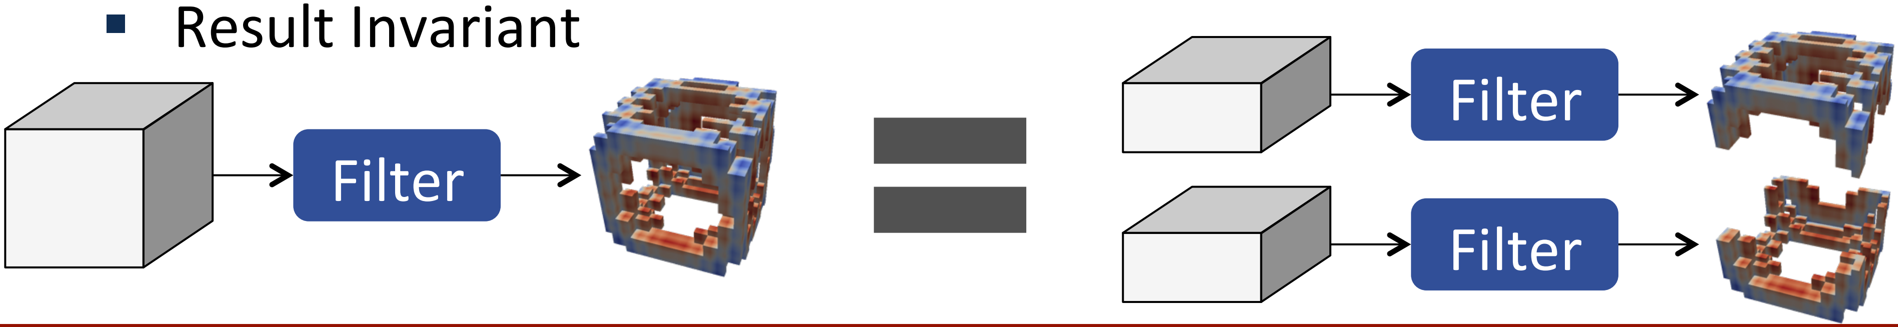

Mappable Input Can Break Down

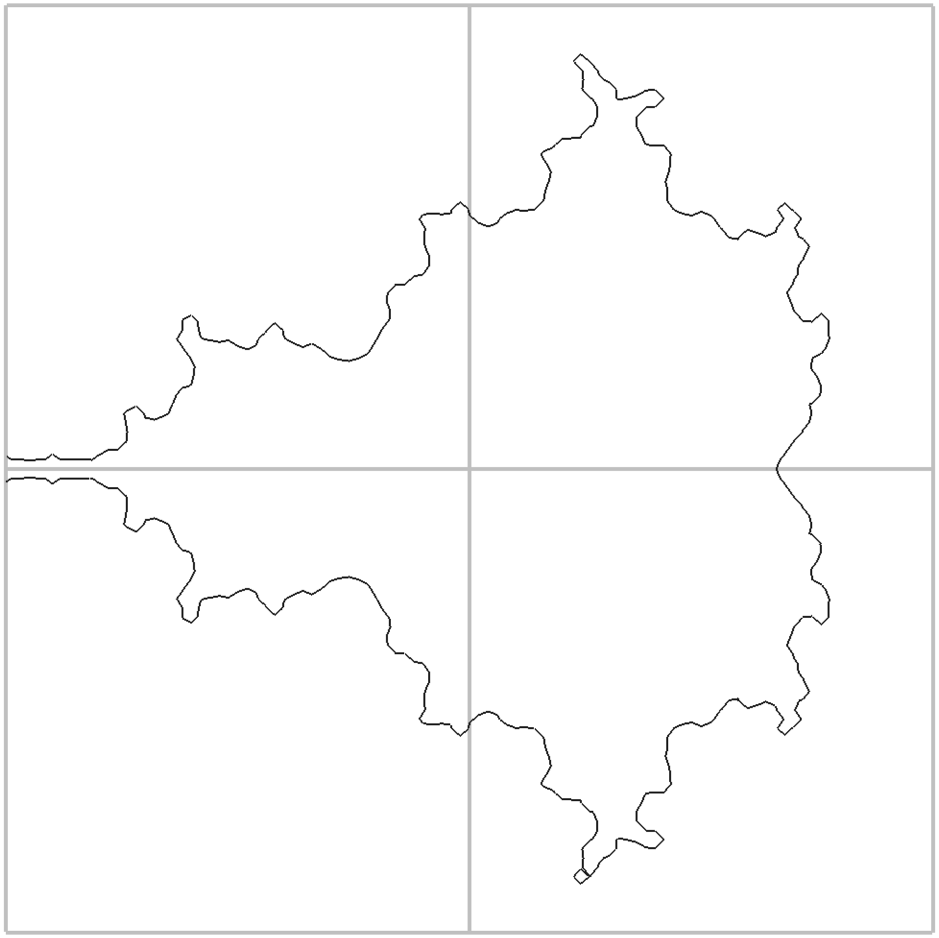

Coarse Parallelism

Result reasonably divided among partitions.

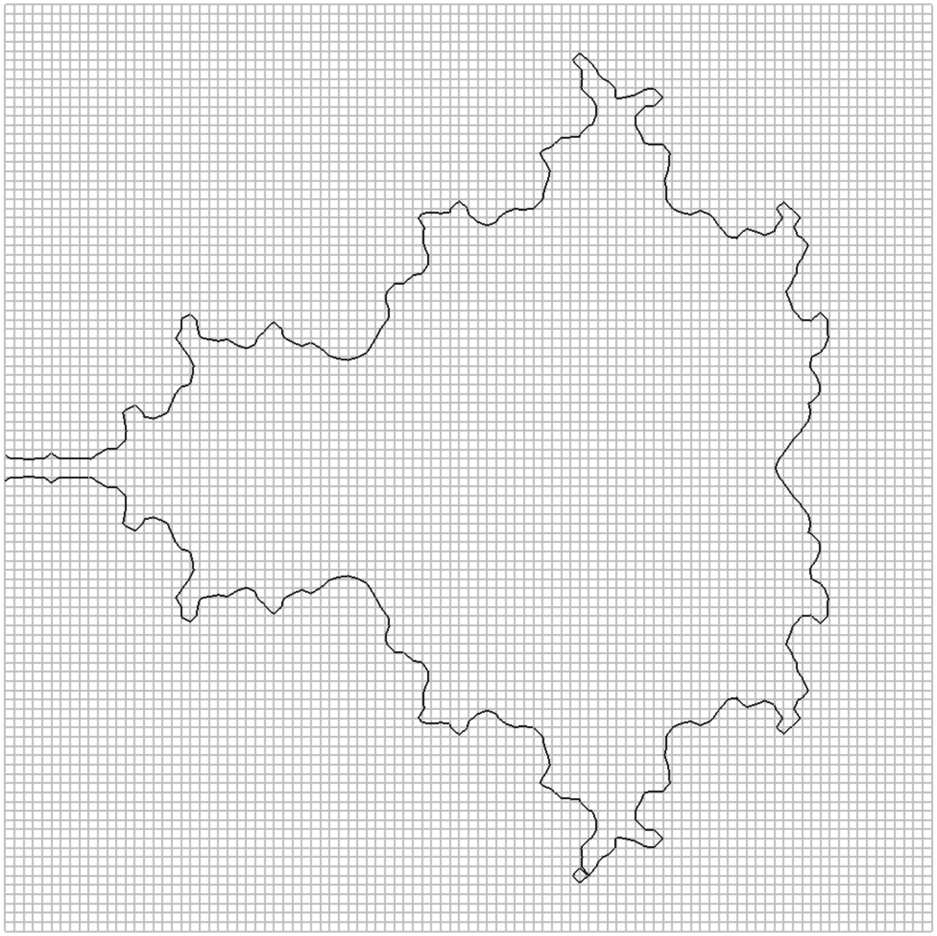

Massive Parallelism

Most partitions are empty.

Result Invariant Can Break Down

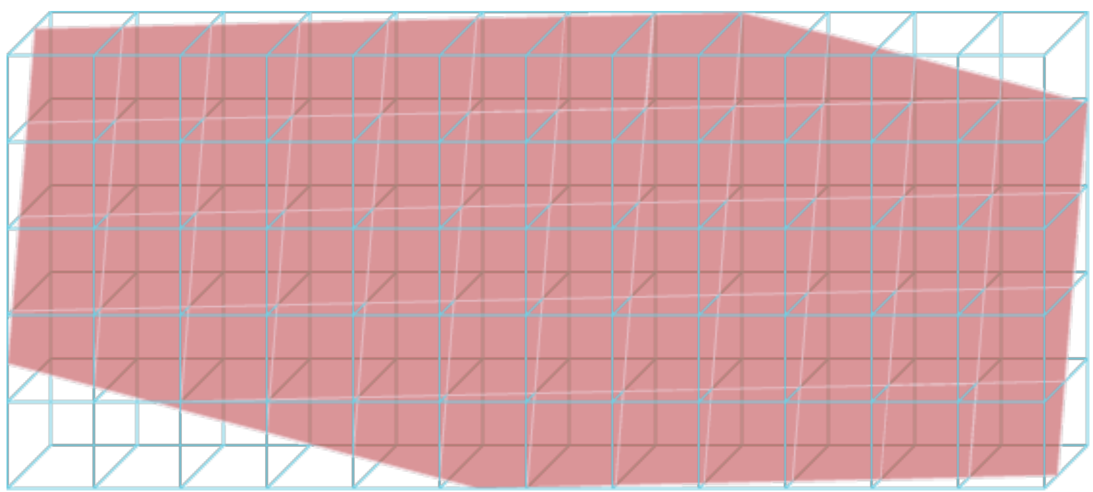

Result Invariant Can Break Down

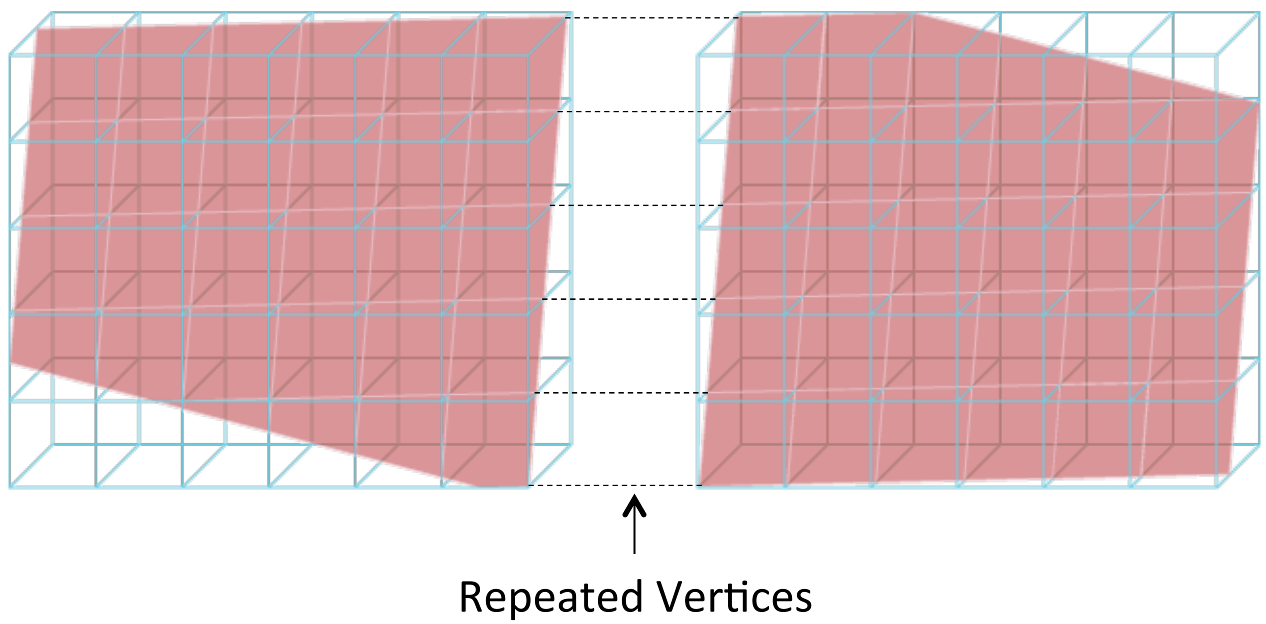

Result Invariant Can Break Down

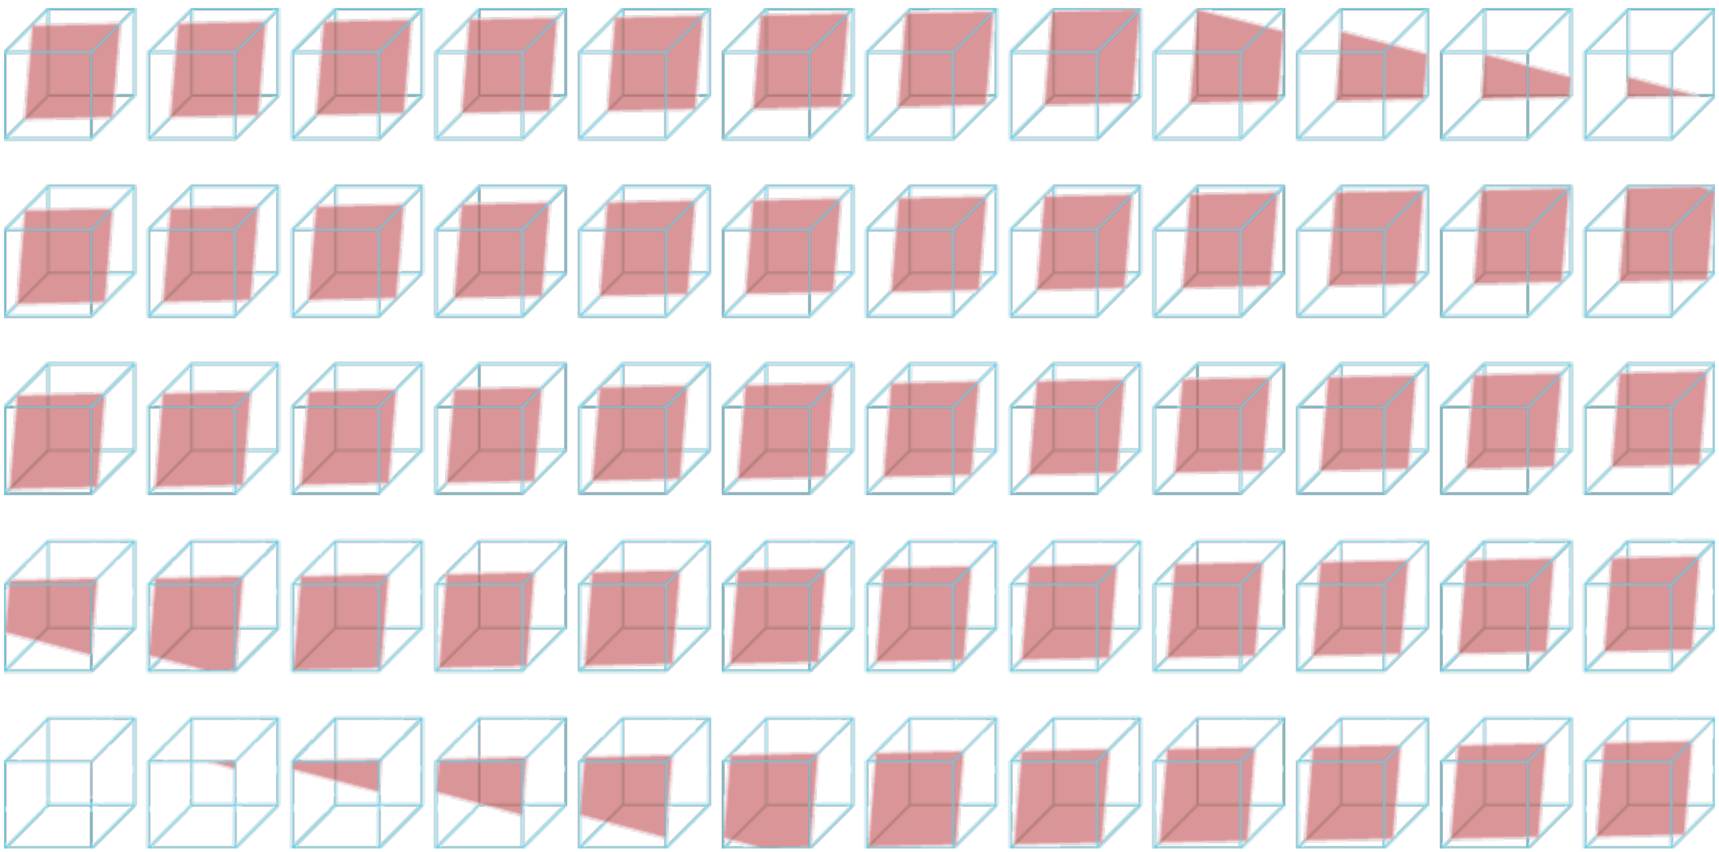

New Key Principles

Coarse Parallel

- Data Separability
- Mappable Input
- Result Invariant

Massive Parallel

- Data Separability
- Discoverable Input Mapping
- Collective Work

Characterization and Classification

Sandia National Laboratories

Key Principles

- Data Separability
- Discoverable Input Mapping
- Collective Work

Characterization Parameters

- Separable Element
- Point Mapping
- Cell Mapping
- Field Mapping
- Collective Work

Basic Mapping

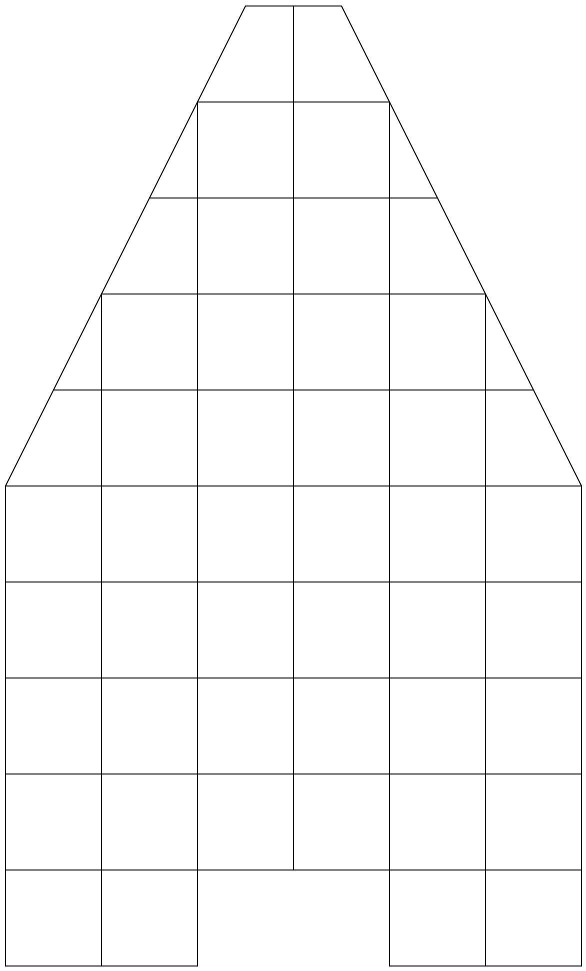

Input

Output

Basic Mapping

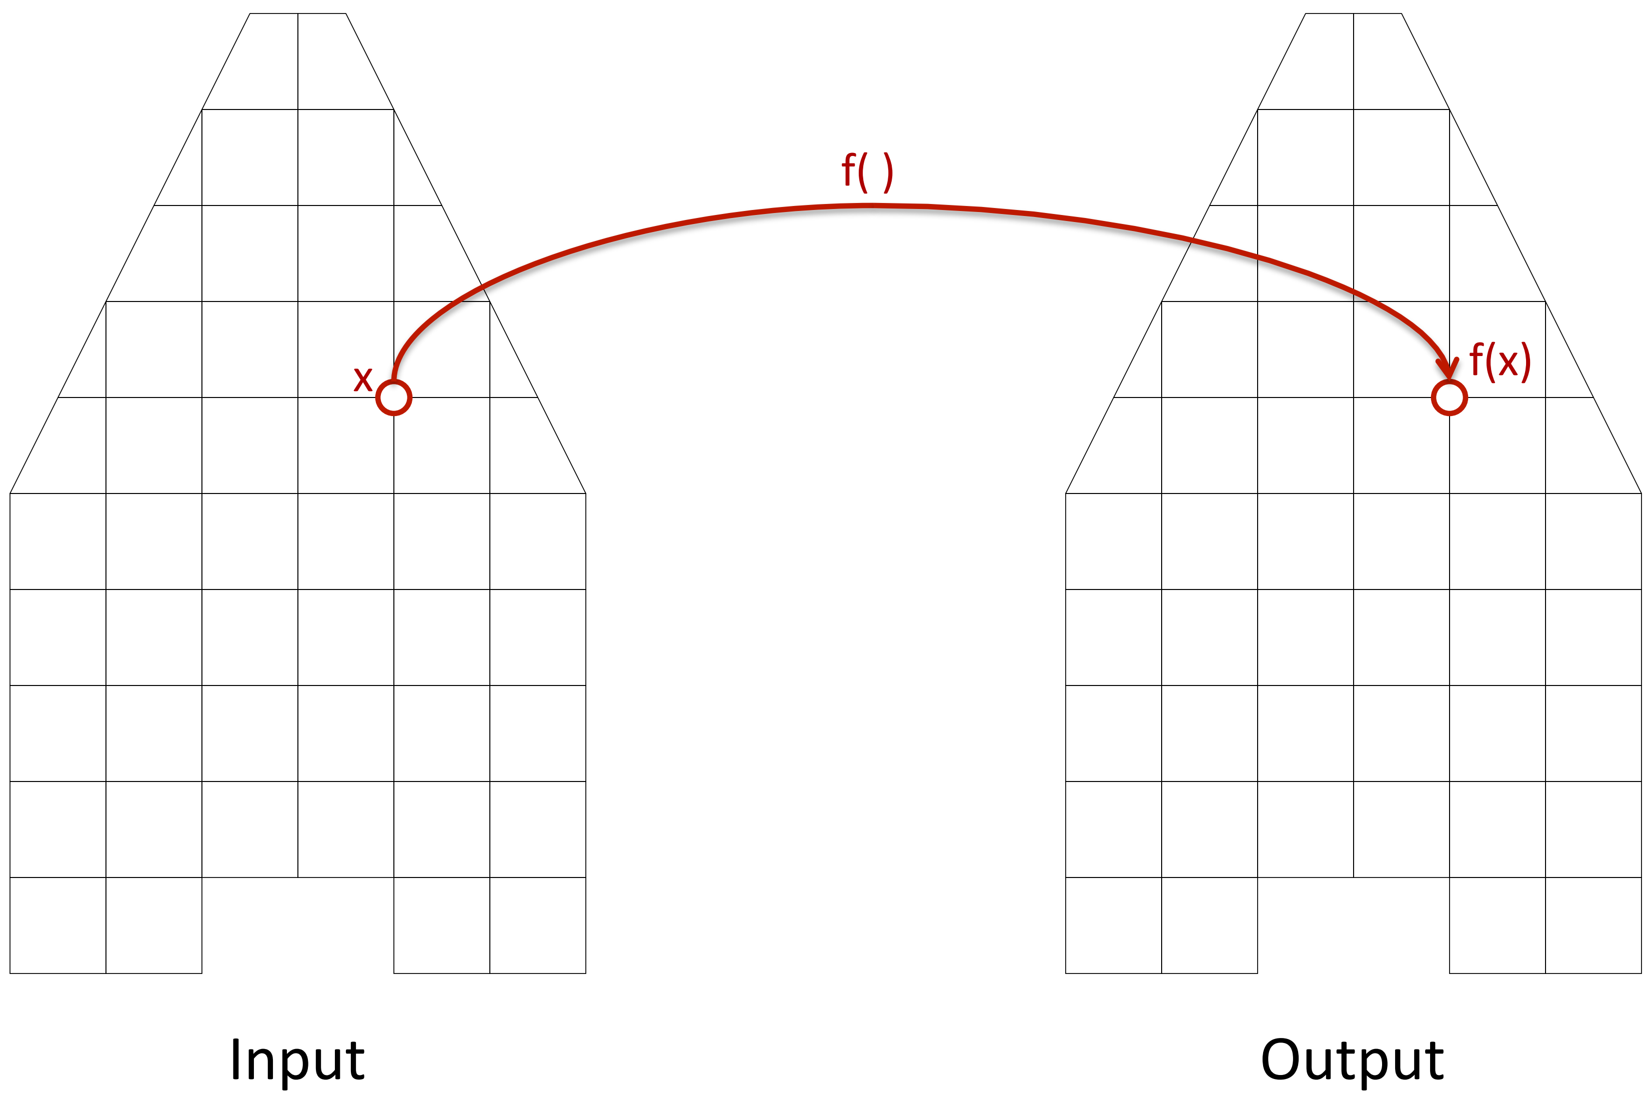

Sandia National Laboratories

Basic Mapping

Separable Element	Any
Point Mapping	Identity
Cell Mapping	Identity
Field Mapping	1 to 1
Collective Work	None

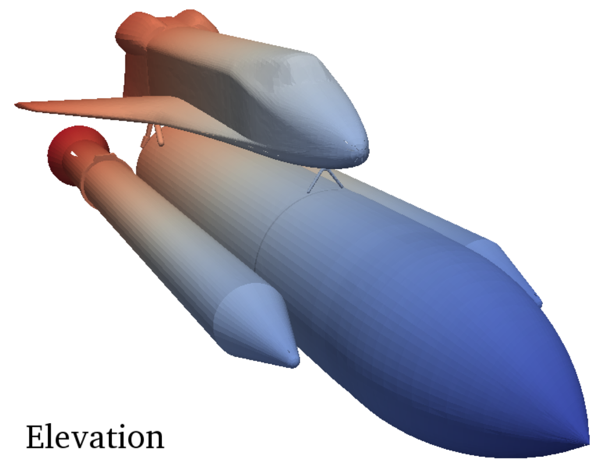
Elevation

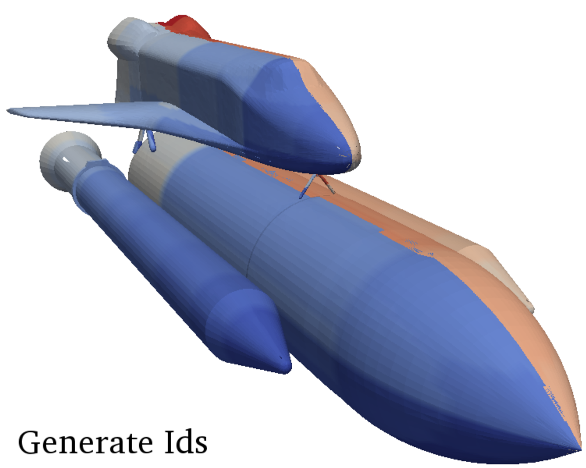
Generate Ids

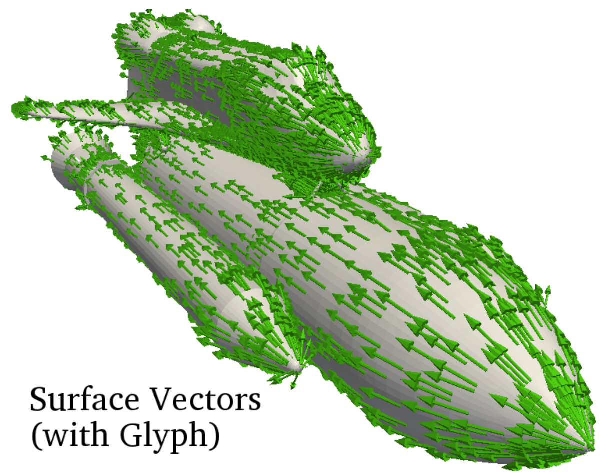
Surface Vectors (with Glyph)

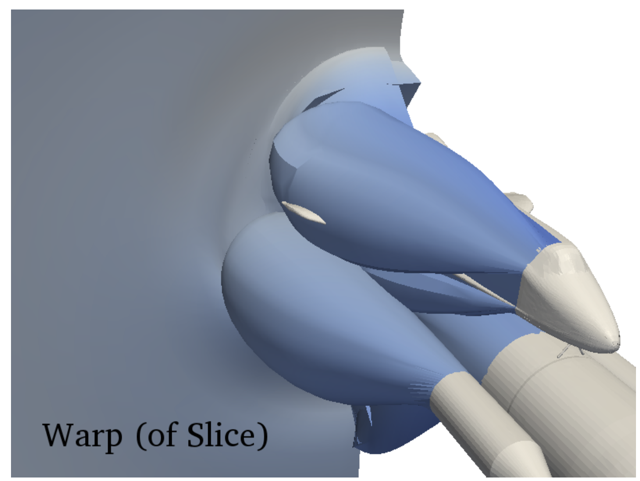
Warp (of Slice)

Basic Mapping

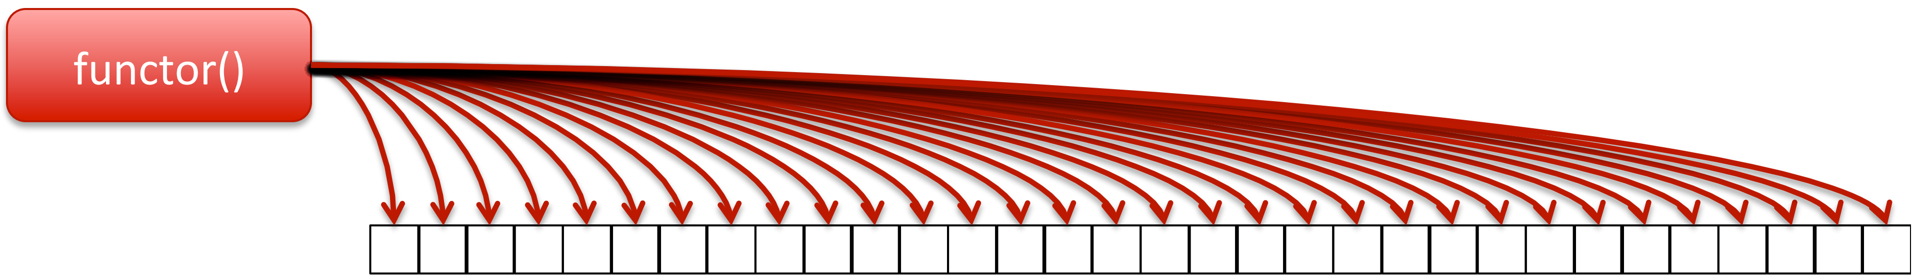

Map by Cell

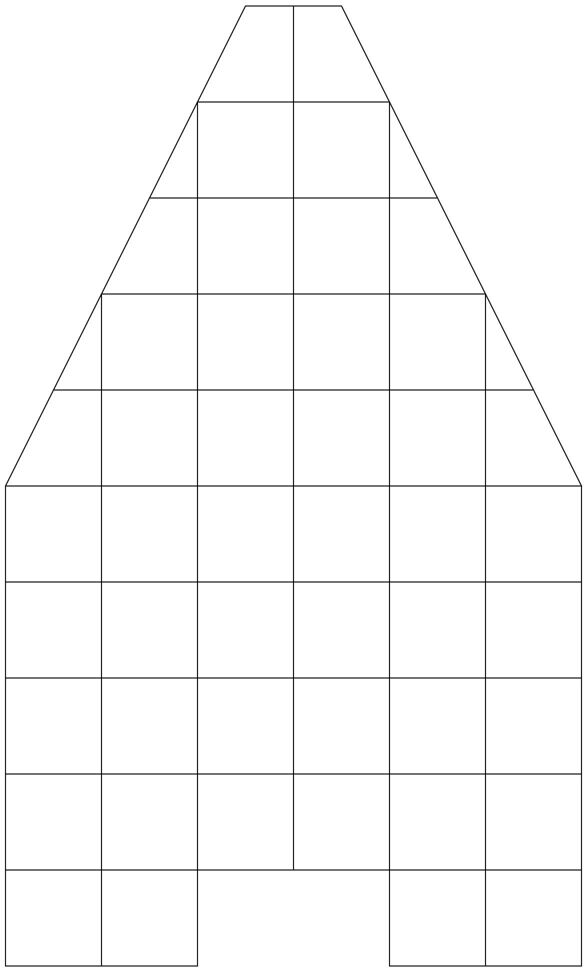

Input

Output

Map by Cell

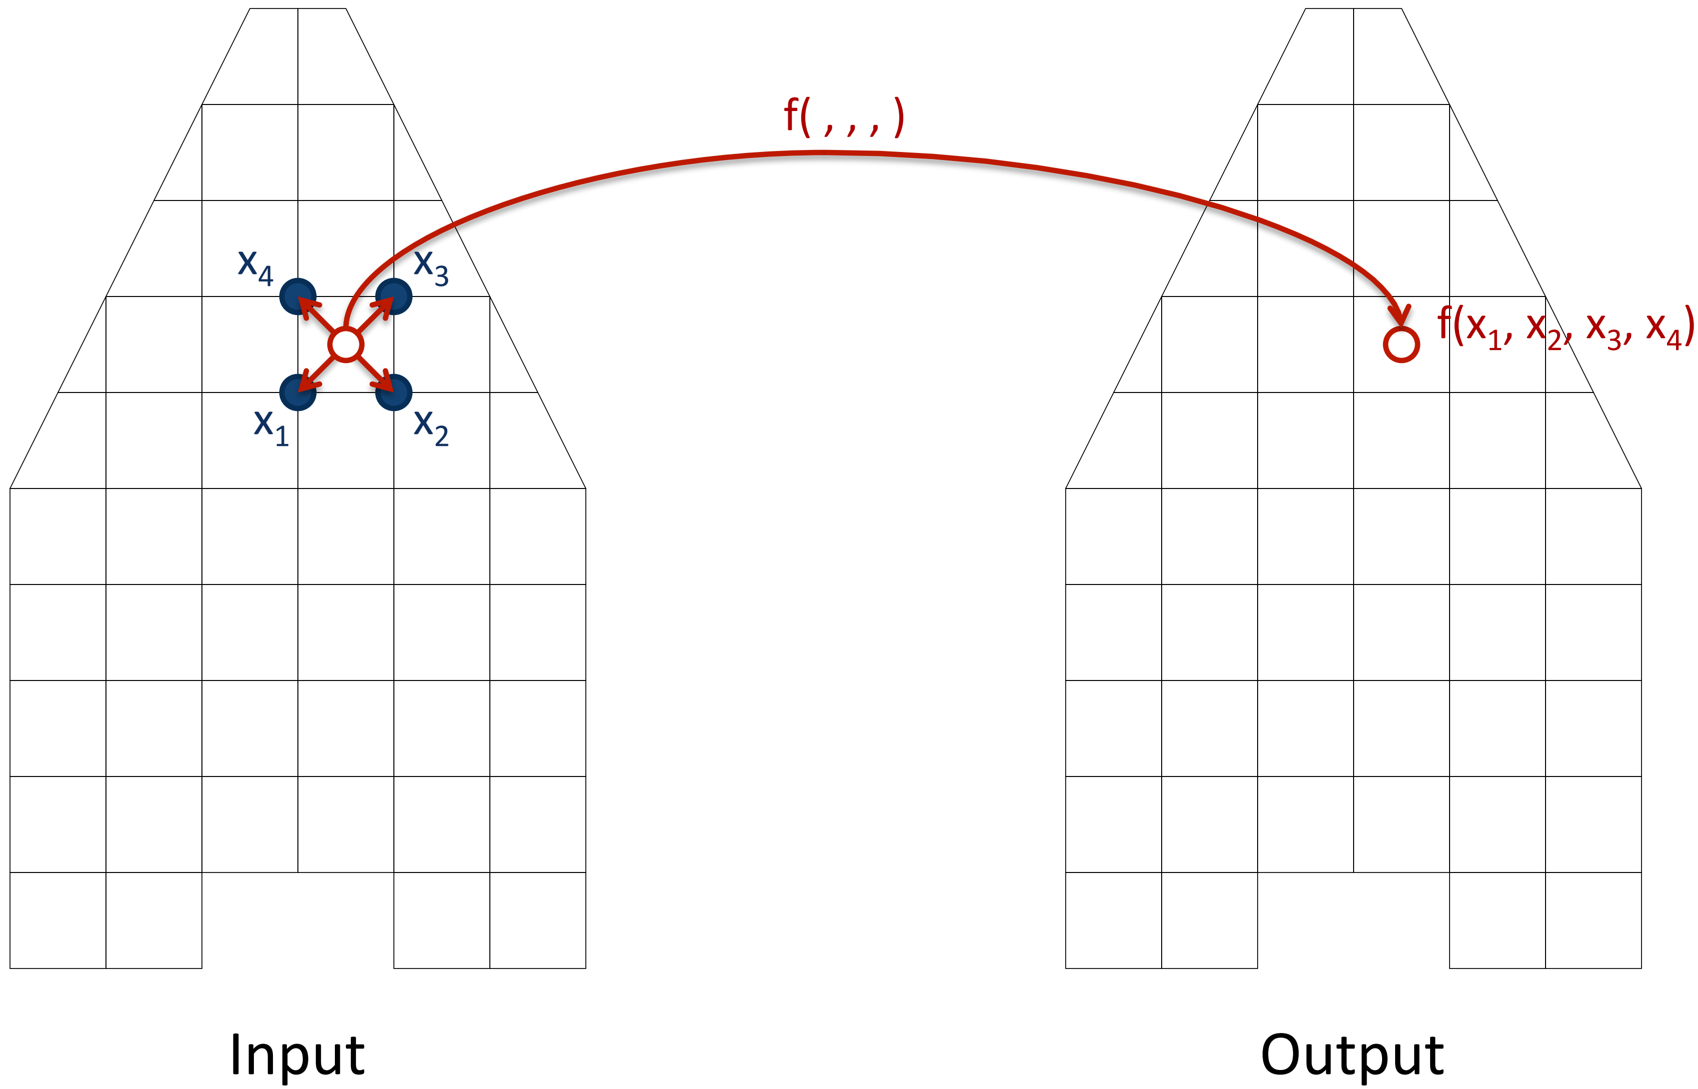

Map by Cell

Separable Element	Cell
Point Mapping	Identity
Cell Mapping	Identity
Field Mapping	Points on cell to cell
Collective Work	None

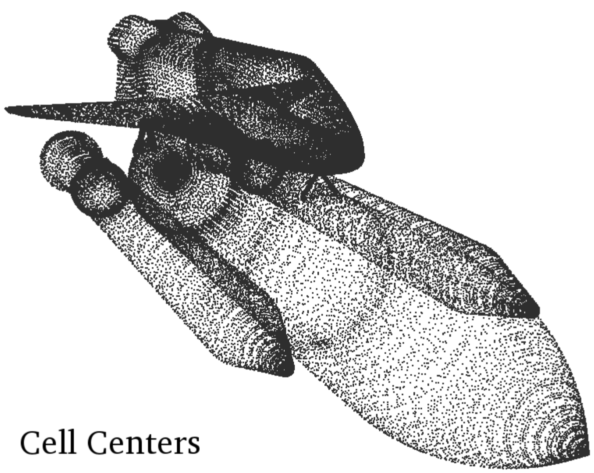

Cell Centers

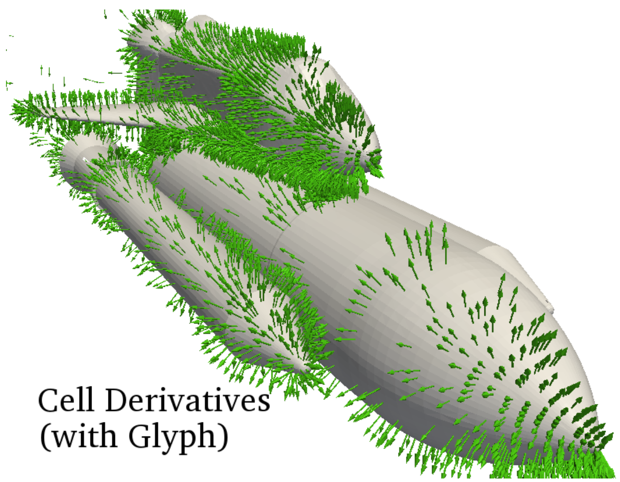

Cell Derivatives
(with Glyph)

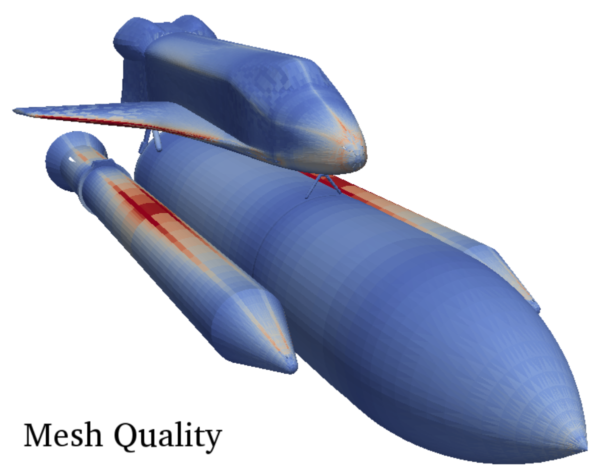

Mesh Quality

Map by Cell

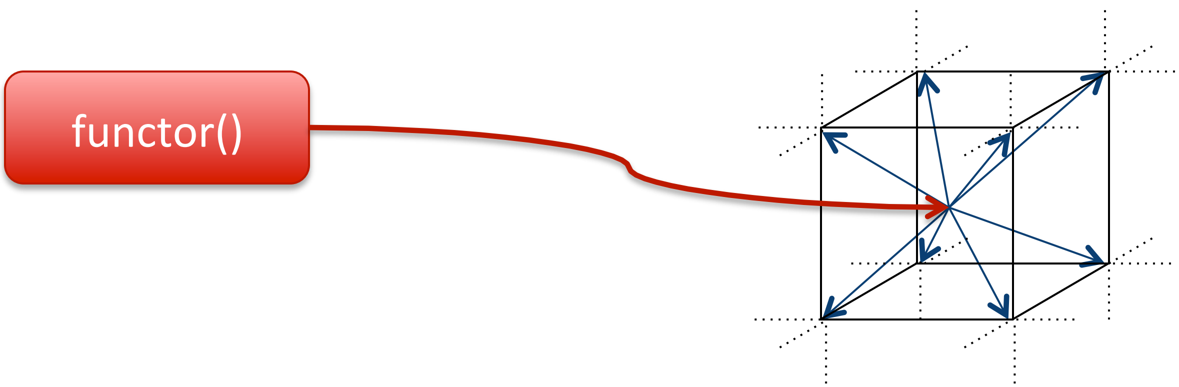

Reconnect Cell

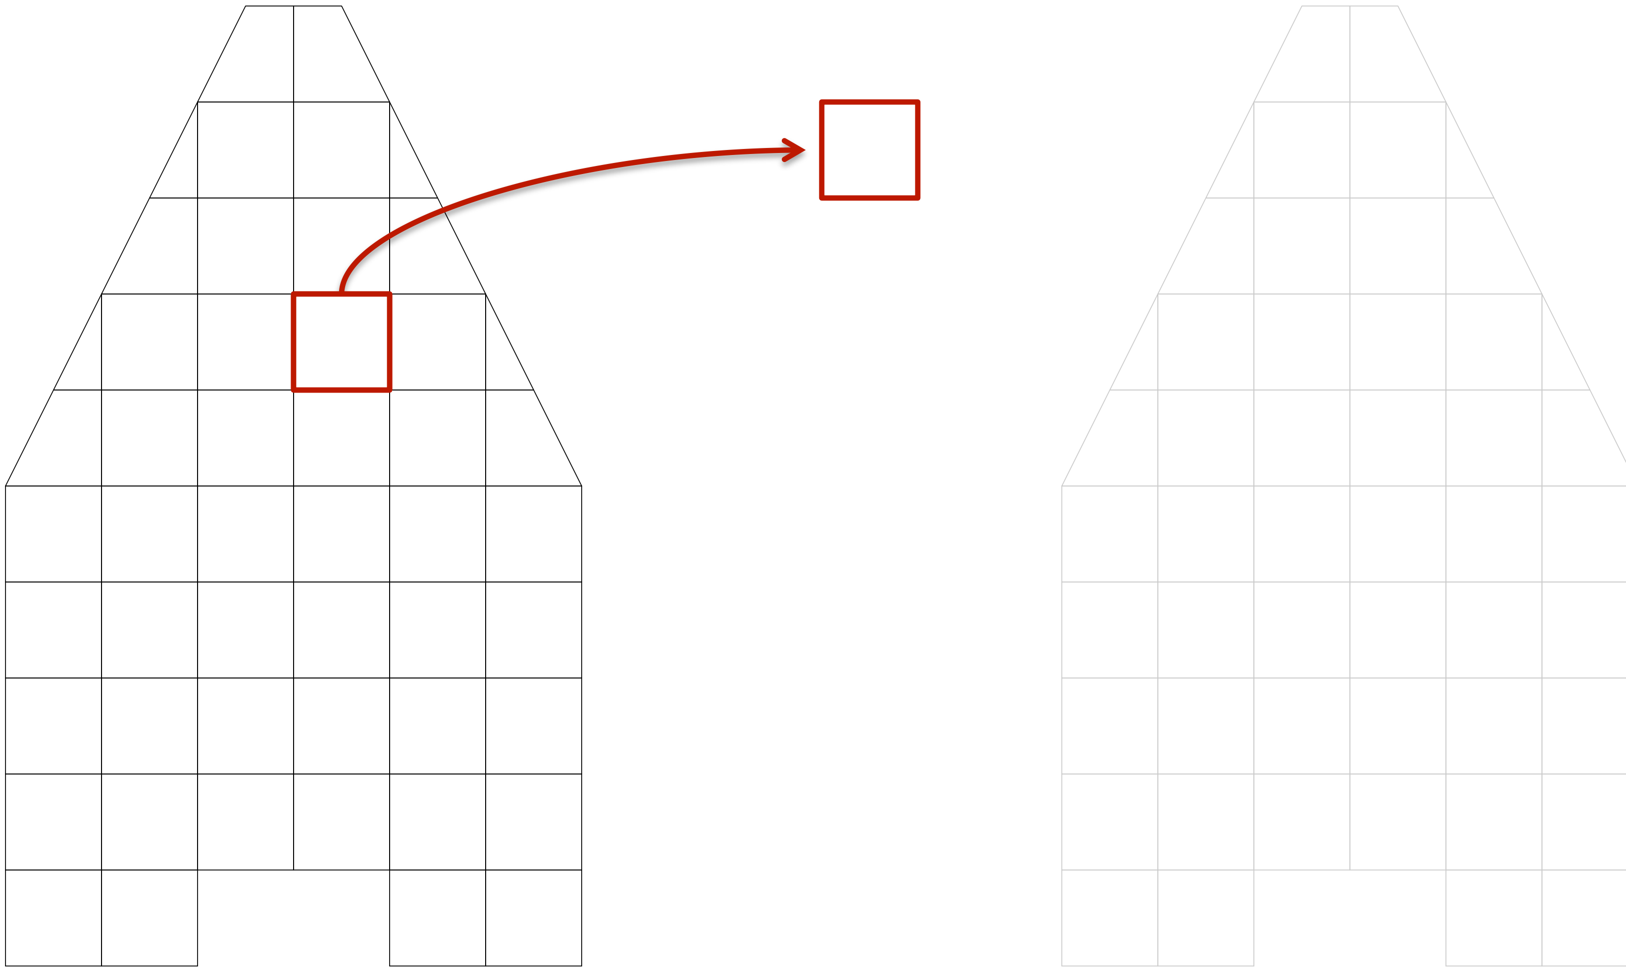

Input Output

Reconnect Cell

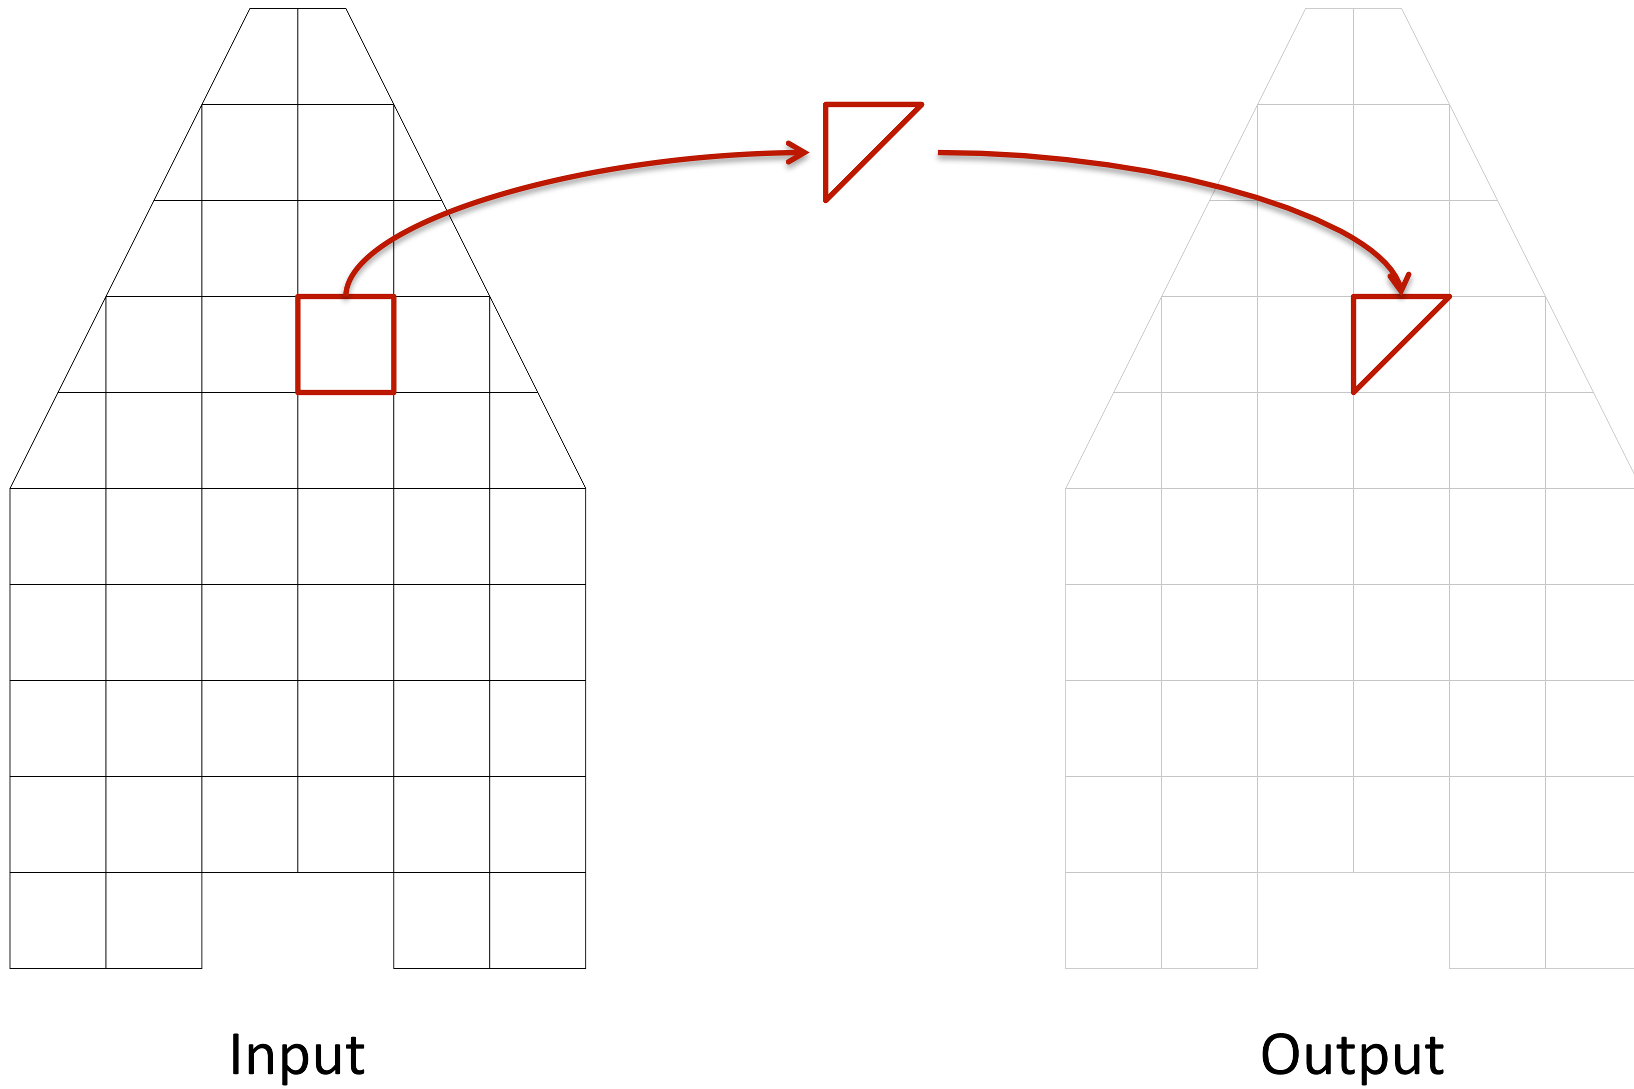

Reconnect Cell

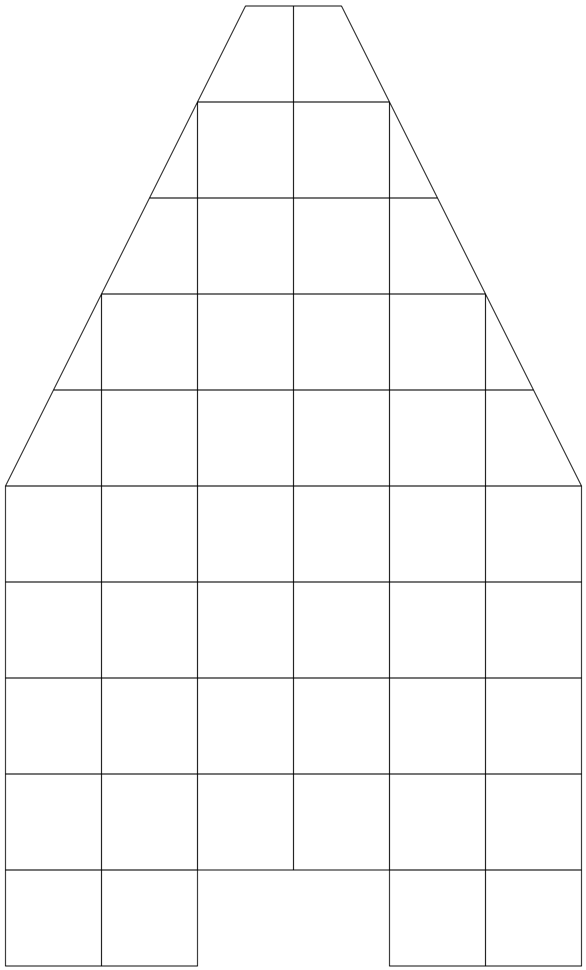

Input

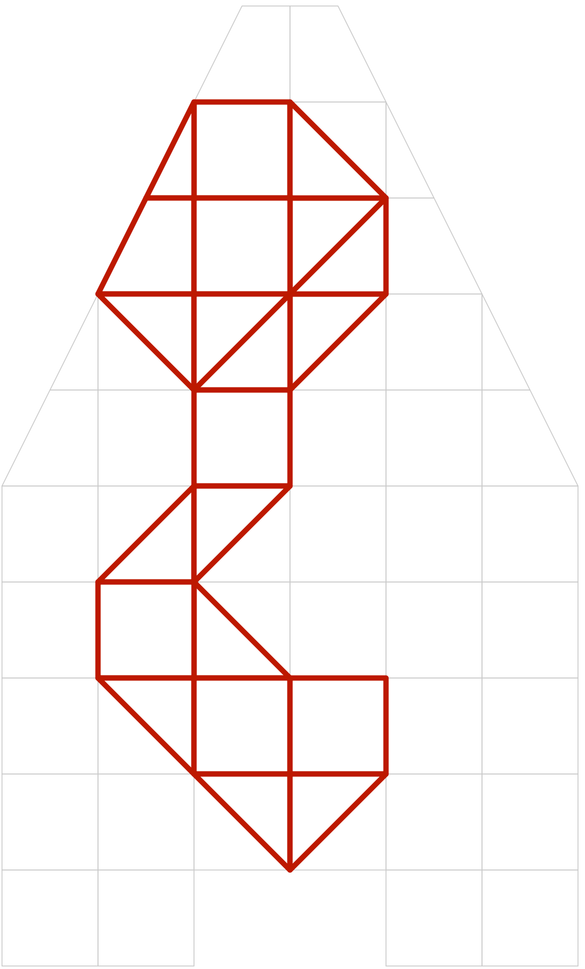

Output

Reconnect Cell

Separable Element	Cell
Point Mapping	1 to 0 or 1
Cell Mapping	1 to 0 or more
Field Mapping	Identity
Collective Work	None (or find unused)

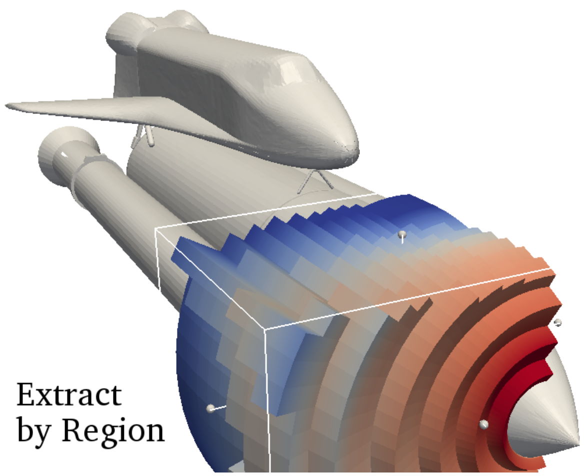

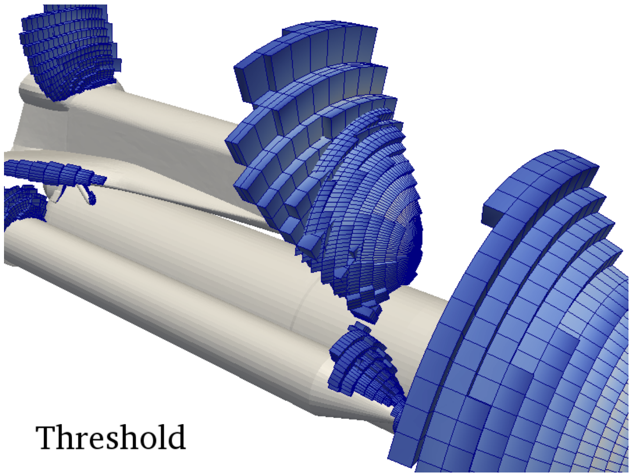

Reconnect Cell

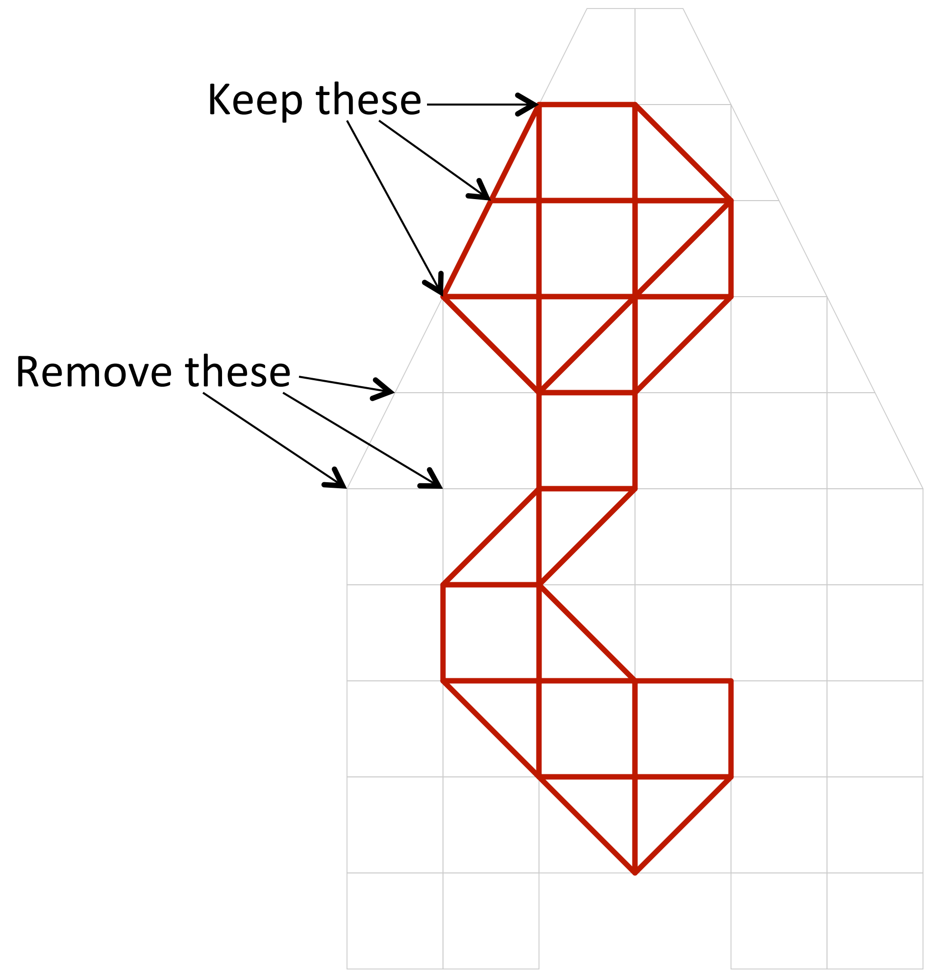

Output

Build Independent Topology

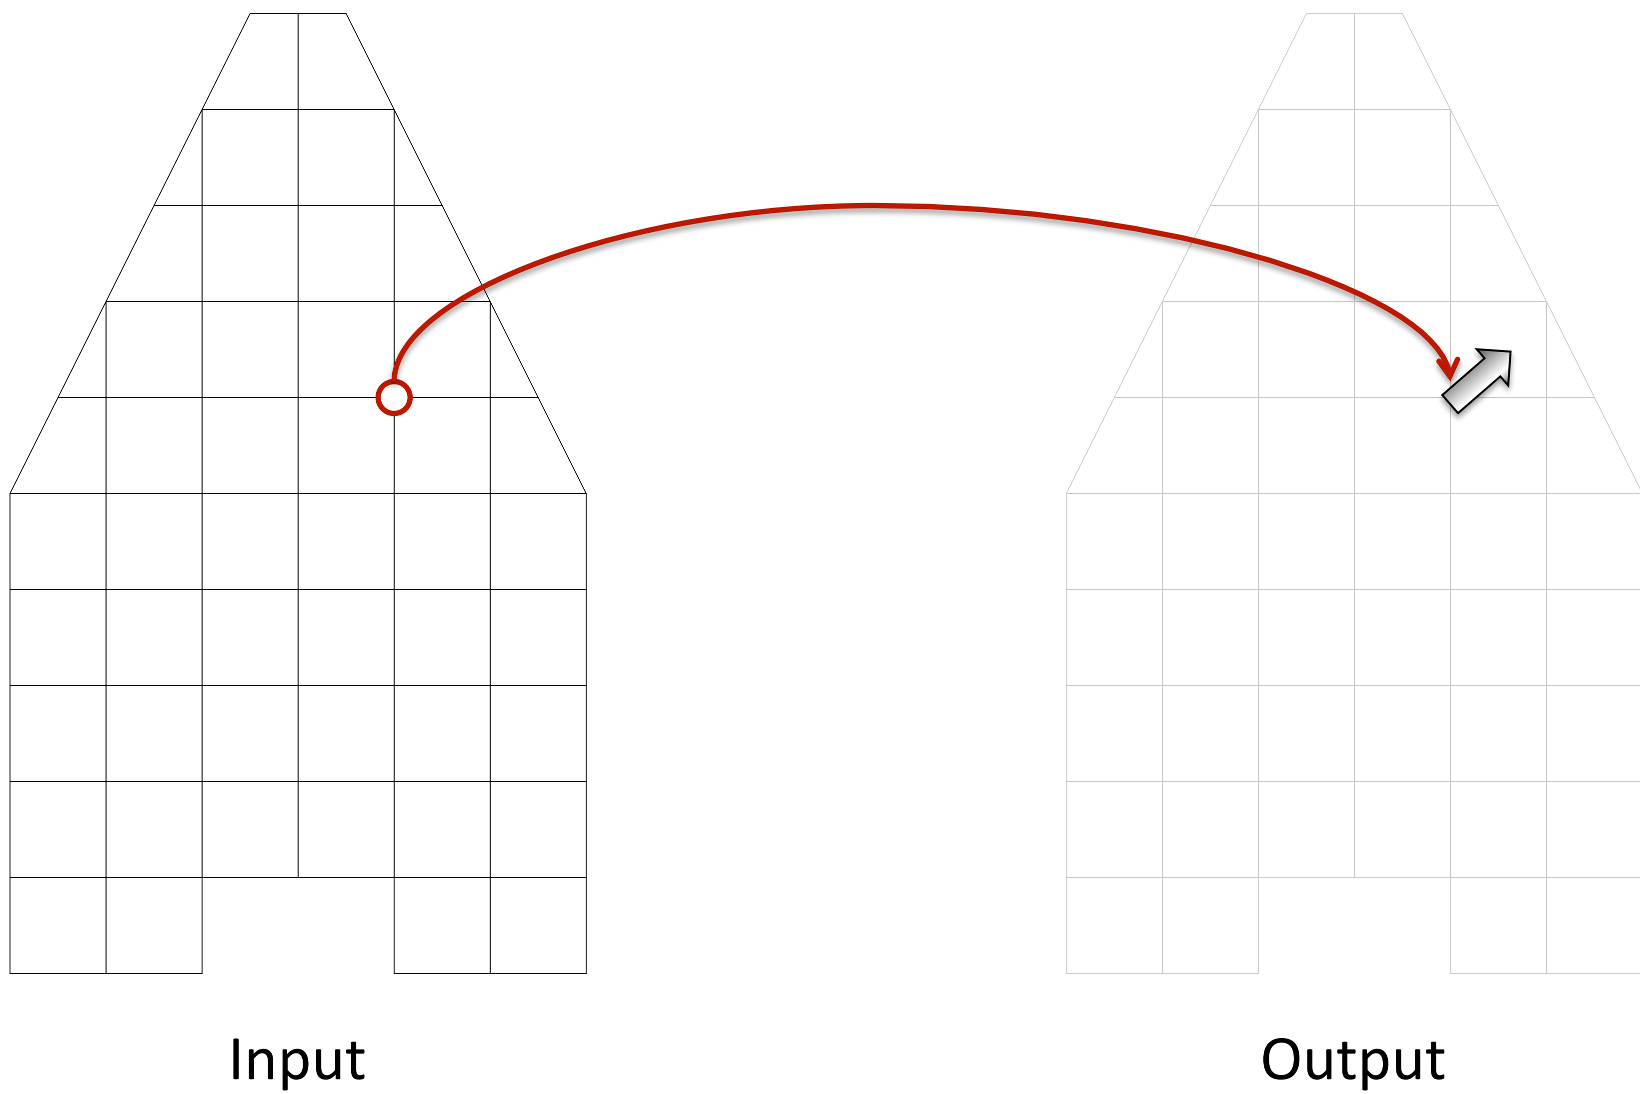

Build Independent Topology

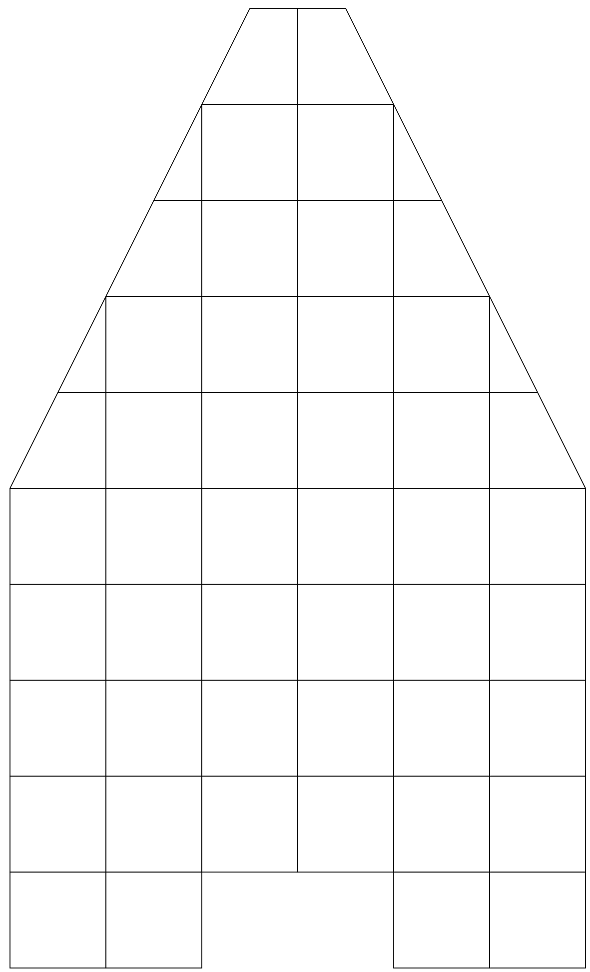

Input

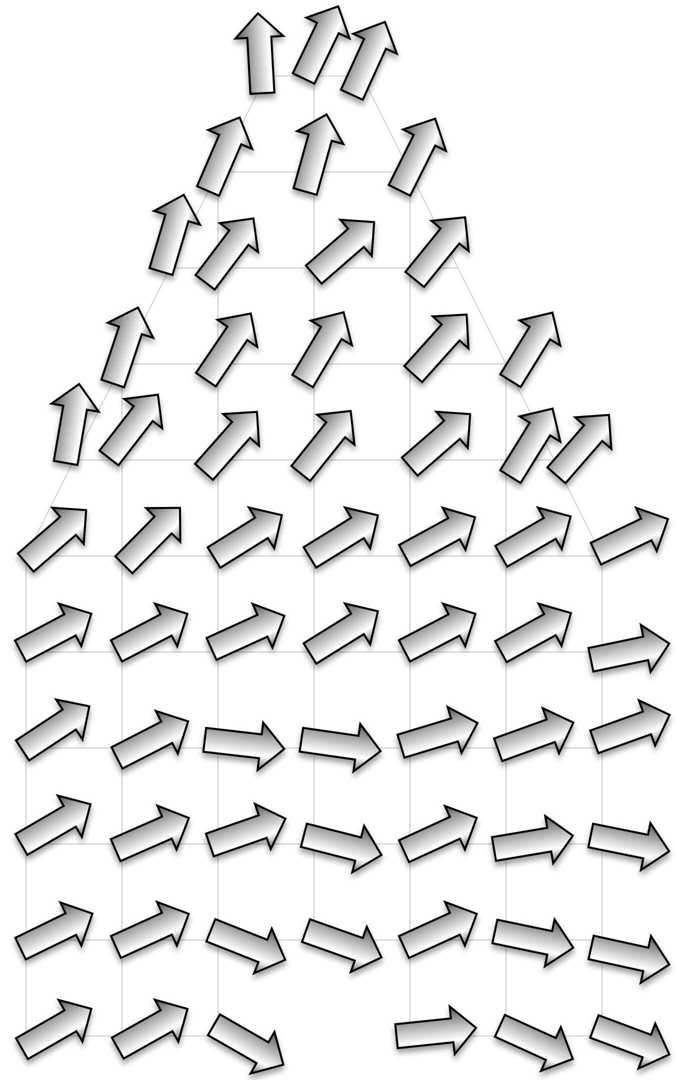

Output

Build Independent Topology

Sandia National Laboratories

Separable Element	Any
Point Mapping	1 element to many points
Cell Mapping	1 element to many cells (constant number)
Field Mapping	Identity
Collective Work	None

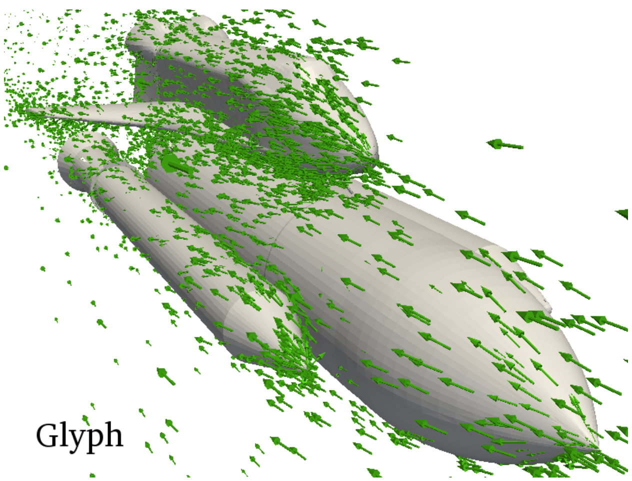
Glyph


Shrink

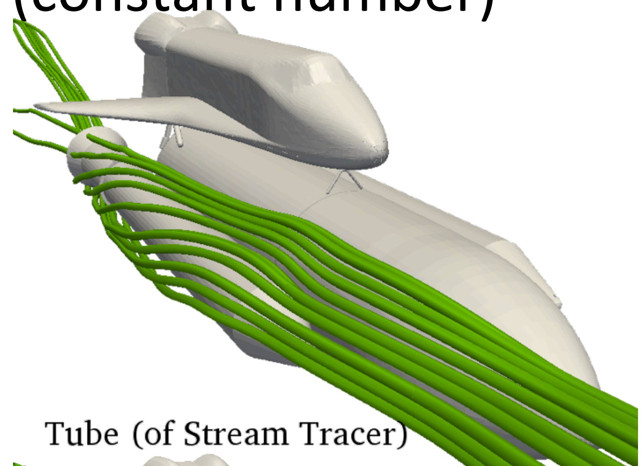
Tube (of Stream Tracer)

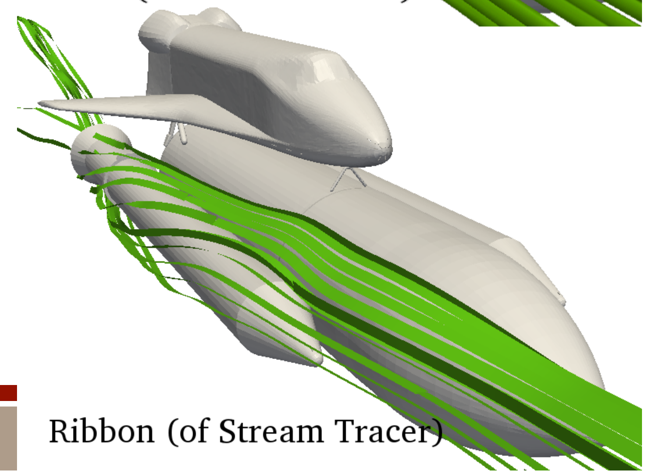
Ribbon (of Stream Tracer)

Build Connected Topology

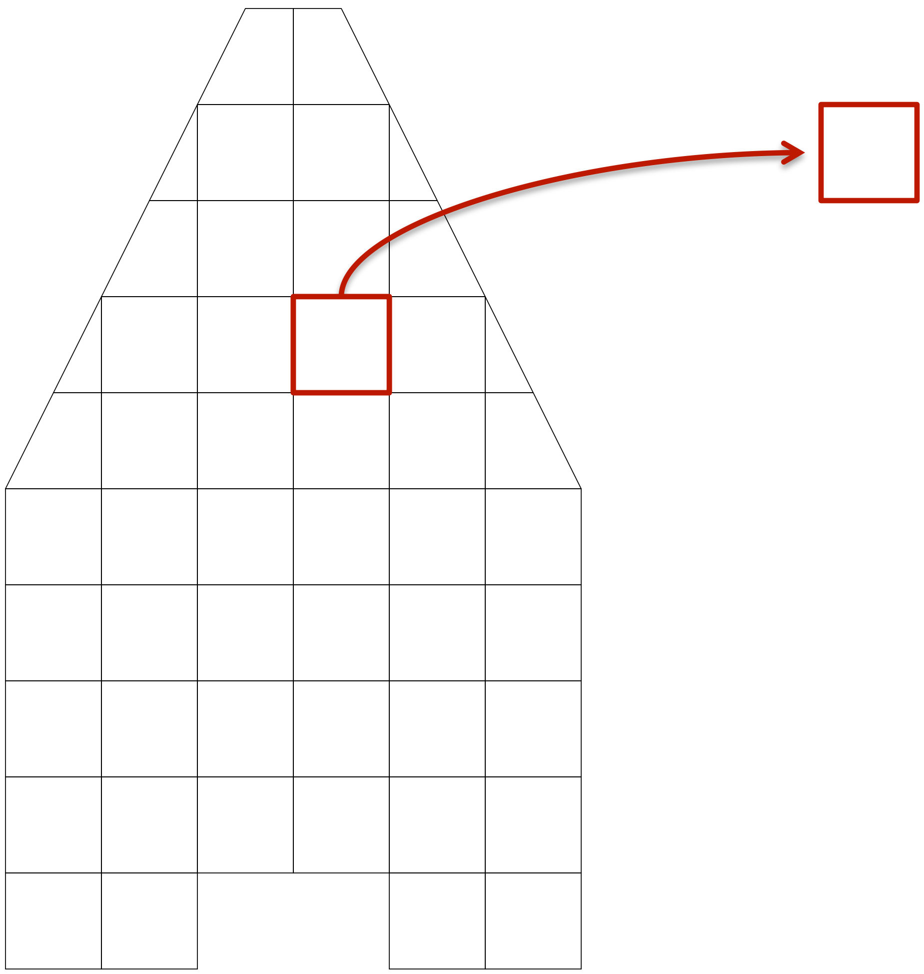

Input

Output

Build Connected Topology

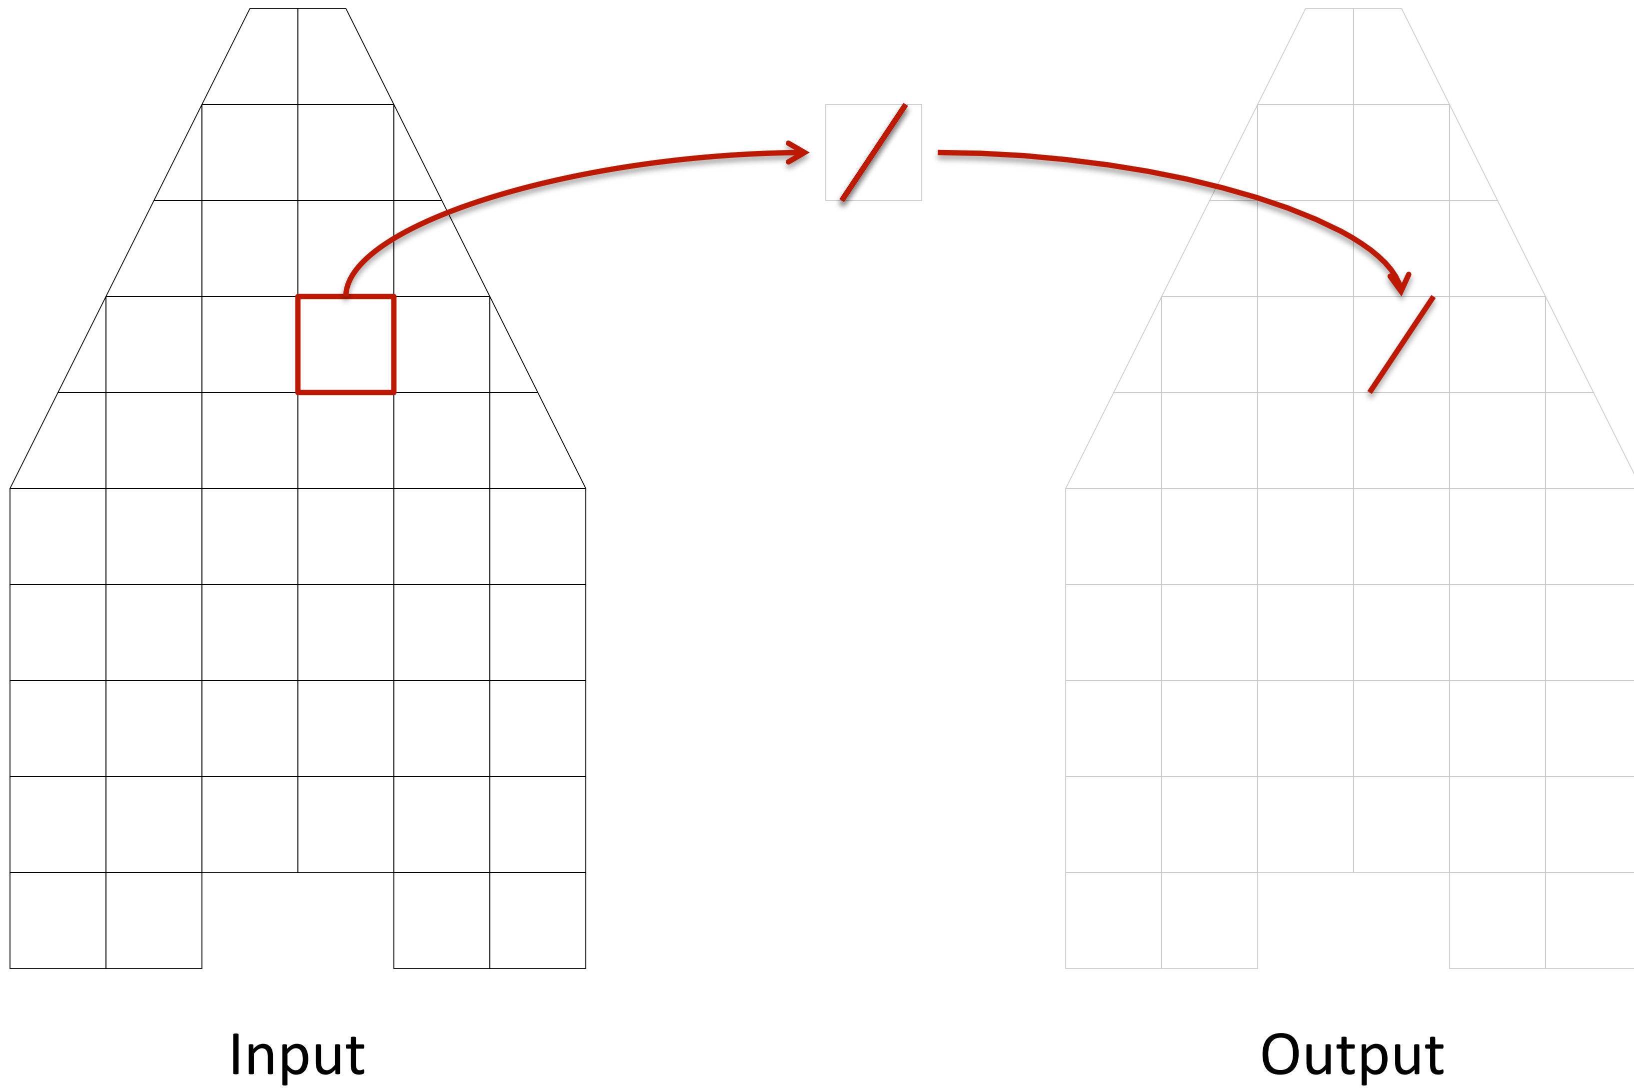

Build Connected Topology

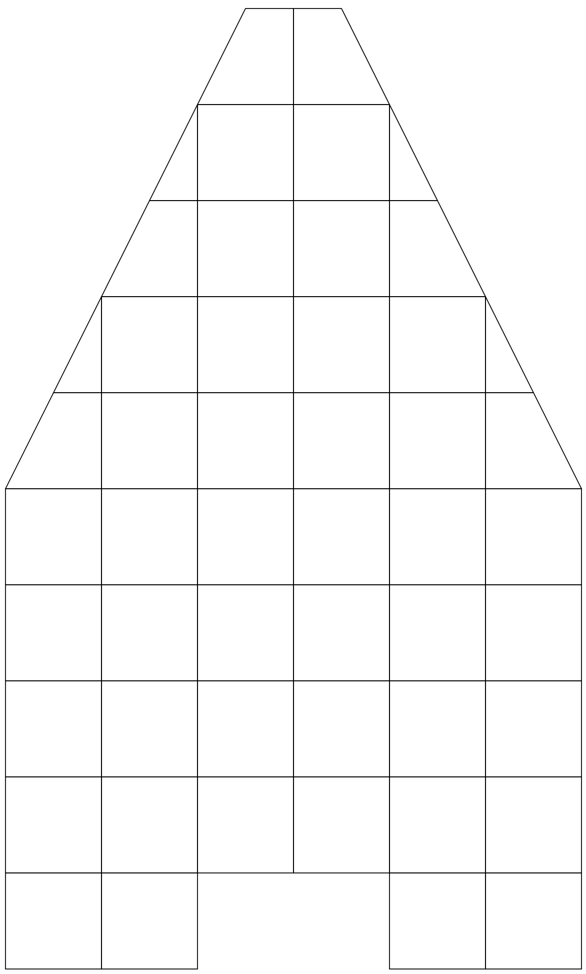

Input

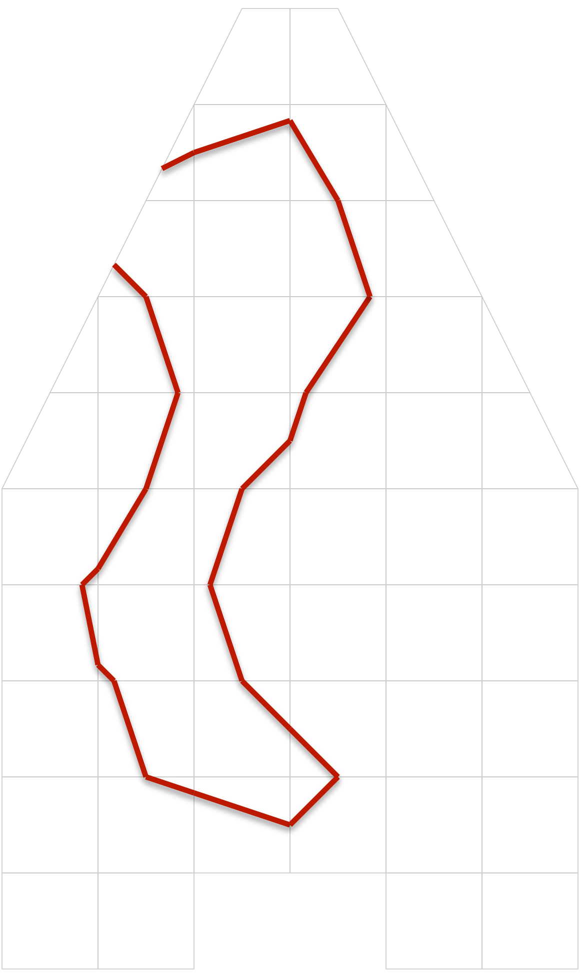

Output

Build ConnectedTopology

Separable Element Cell

Point Mapping 1 cell to 0 or more points

Cell Mapping 1 to 0 or more

Field Mapping Interpolated points

Collective Work Resolve duplicate points

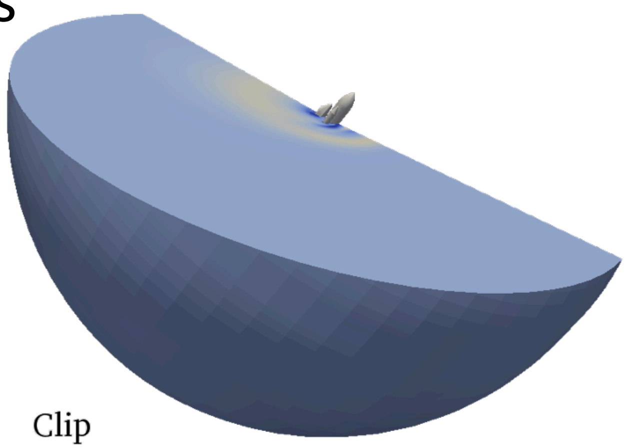

Clip

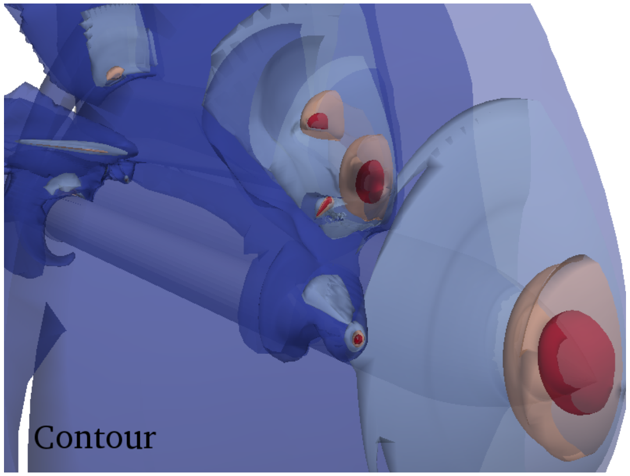

Contour

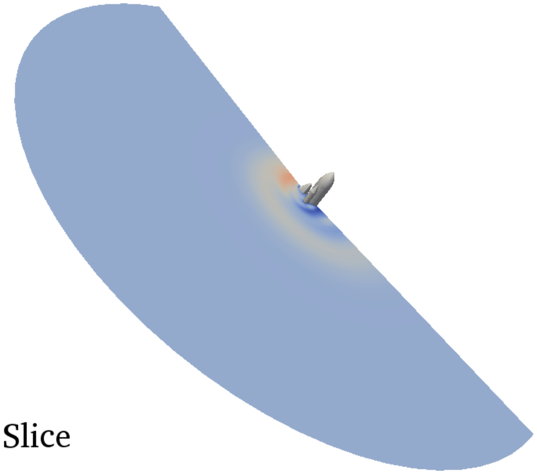

Slice

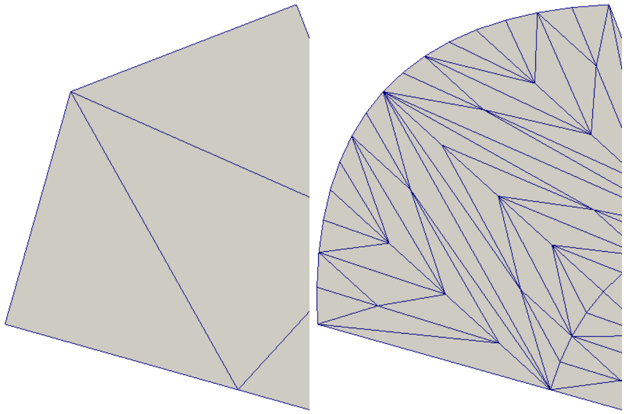

Tessellate

Build Connected Topology

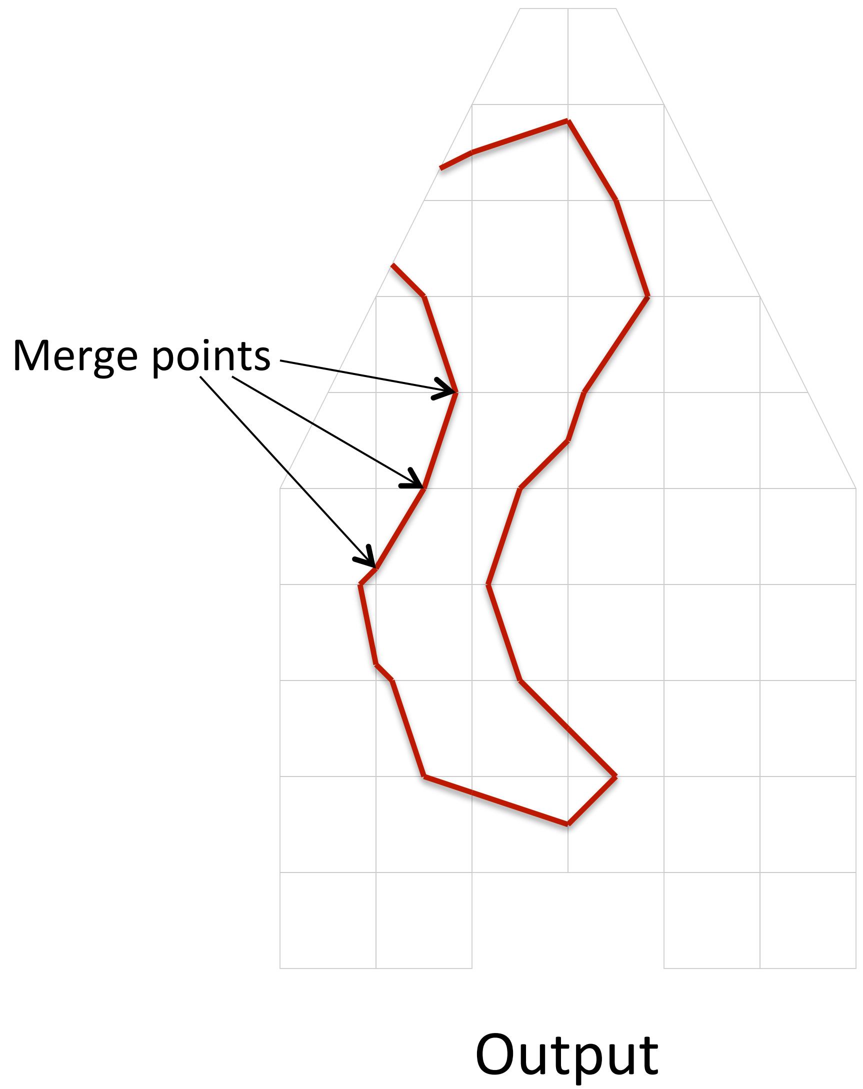

Capture Cell Adjacencies

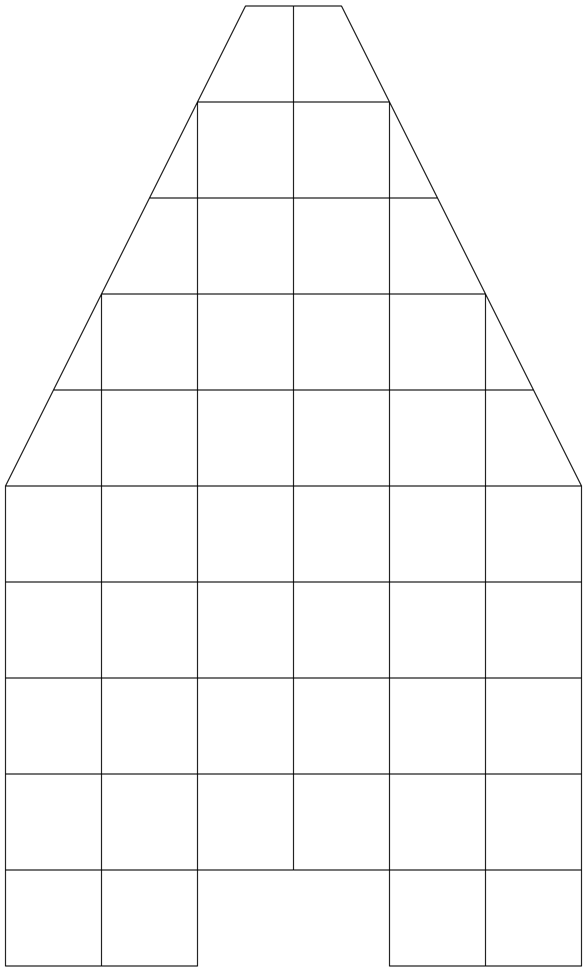

Input

Output

Capture Cell Adjacencies

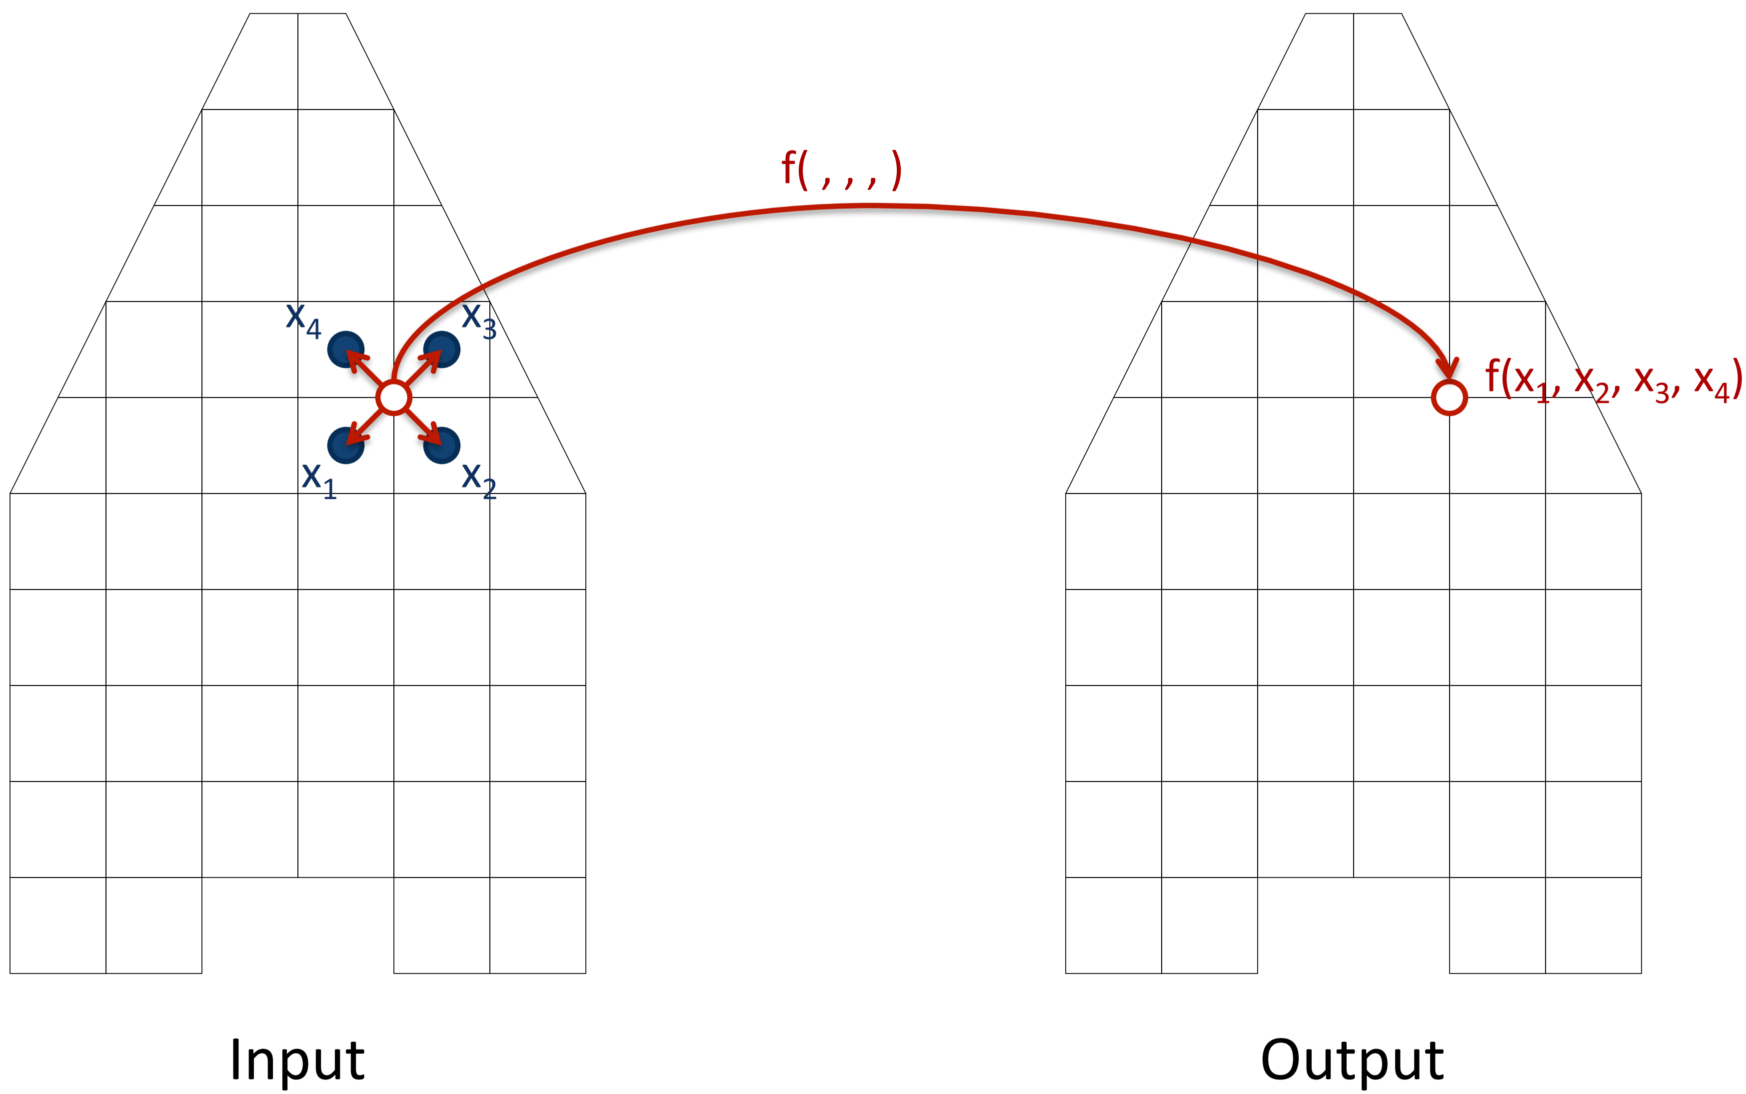

Capture Cell Adjacencies

Separable Element	Point, edge, or face
Point Mapping	Identity
Cell Mapping	Identity
Field Mapping	Interpolated incident fields
Collective Work	Find incidence relationships

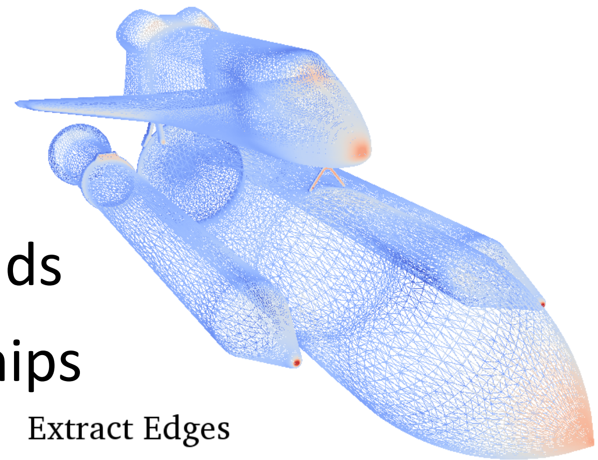

Extract Edges

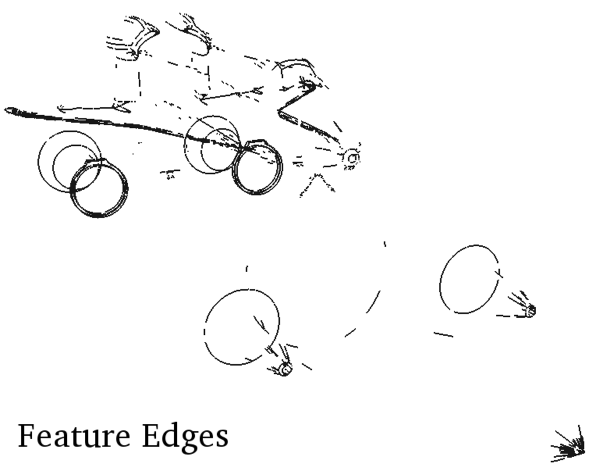

Feature Edges

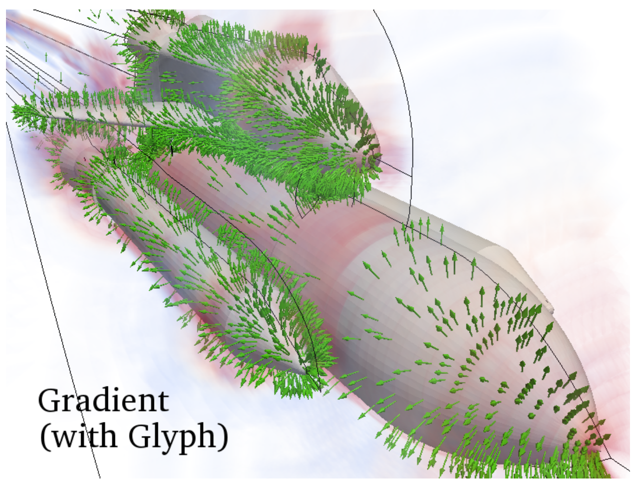

Gradient (with Glyph)

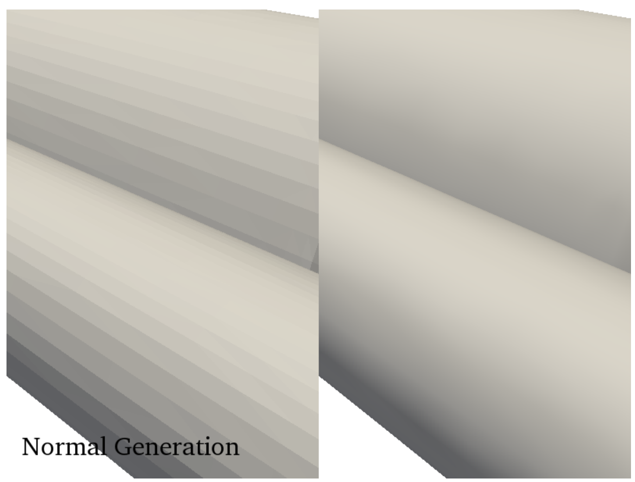

Normal Generation

Capture Cell Adjacencies

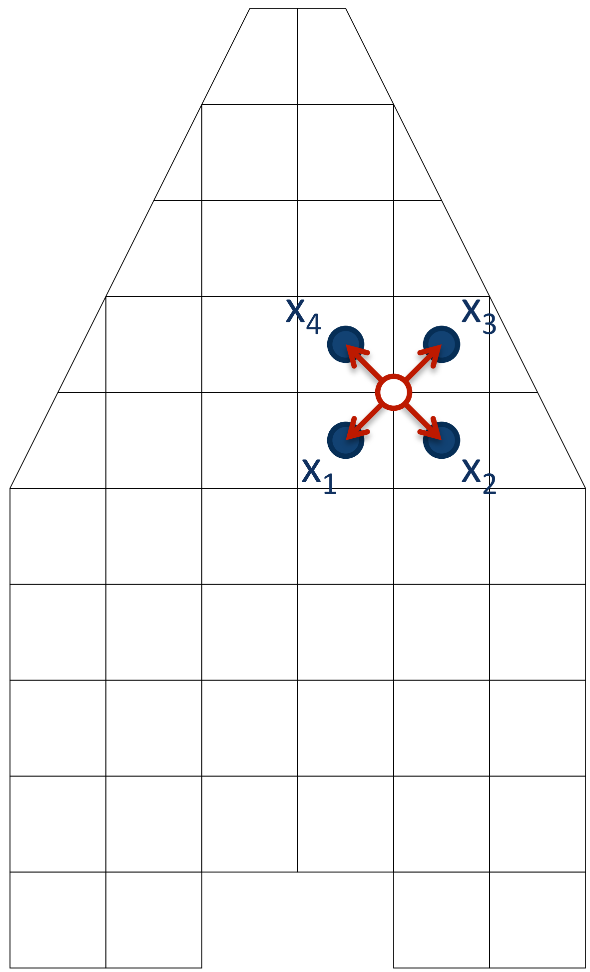

Incidence relationships may not be explicitly captured by topology structure.

Input

Globally Reduce

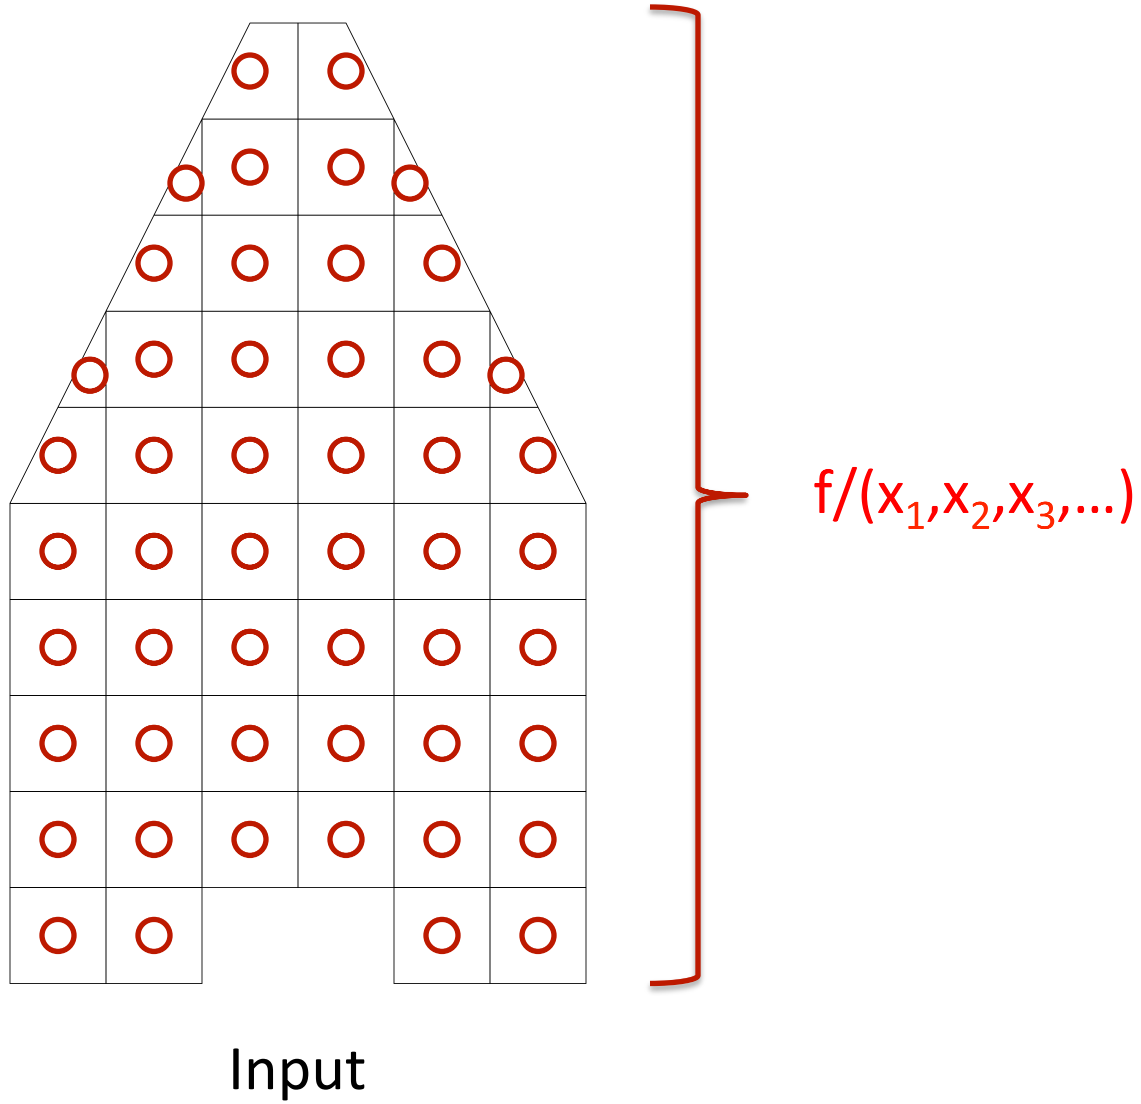

Globally Reduce

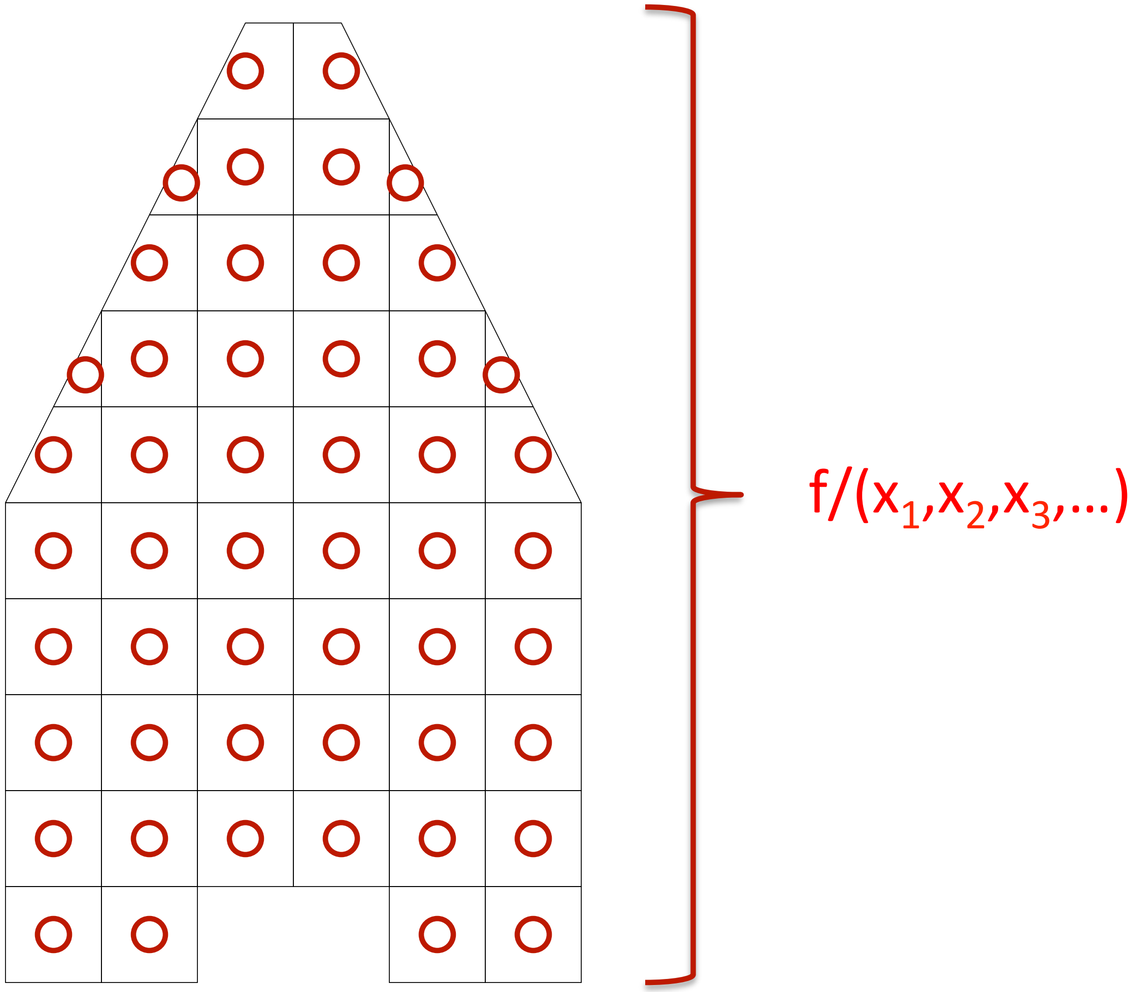

Input

Globally Reduce

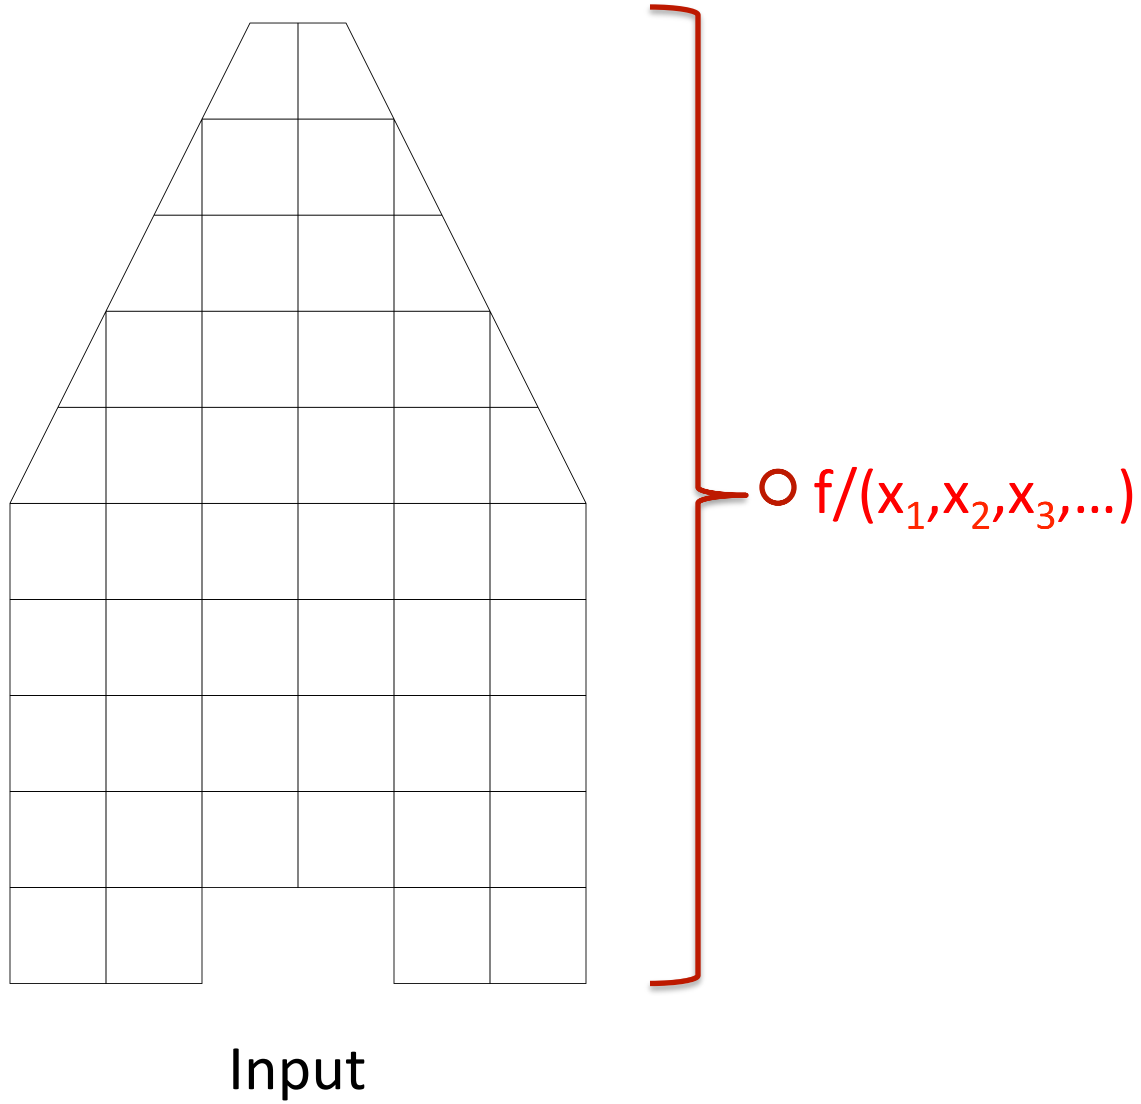

Globally Reduce

Separable Element	Any
Point Mapping	None in output
Cell Mapping	None in output
Field Mapping	All to 1
Collective Work	Global Reduction
Algorithms	Histogram Integrate Outline (find bounds) Statistics

Query Data

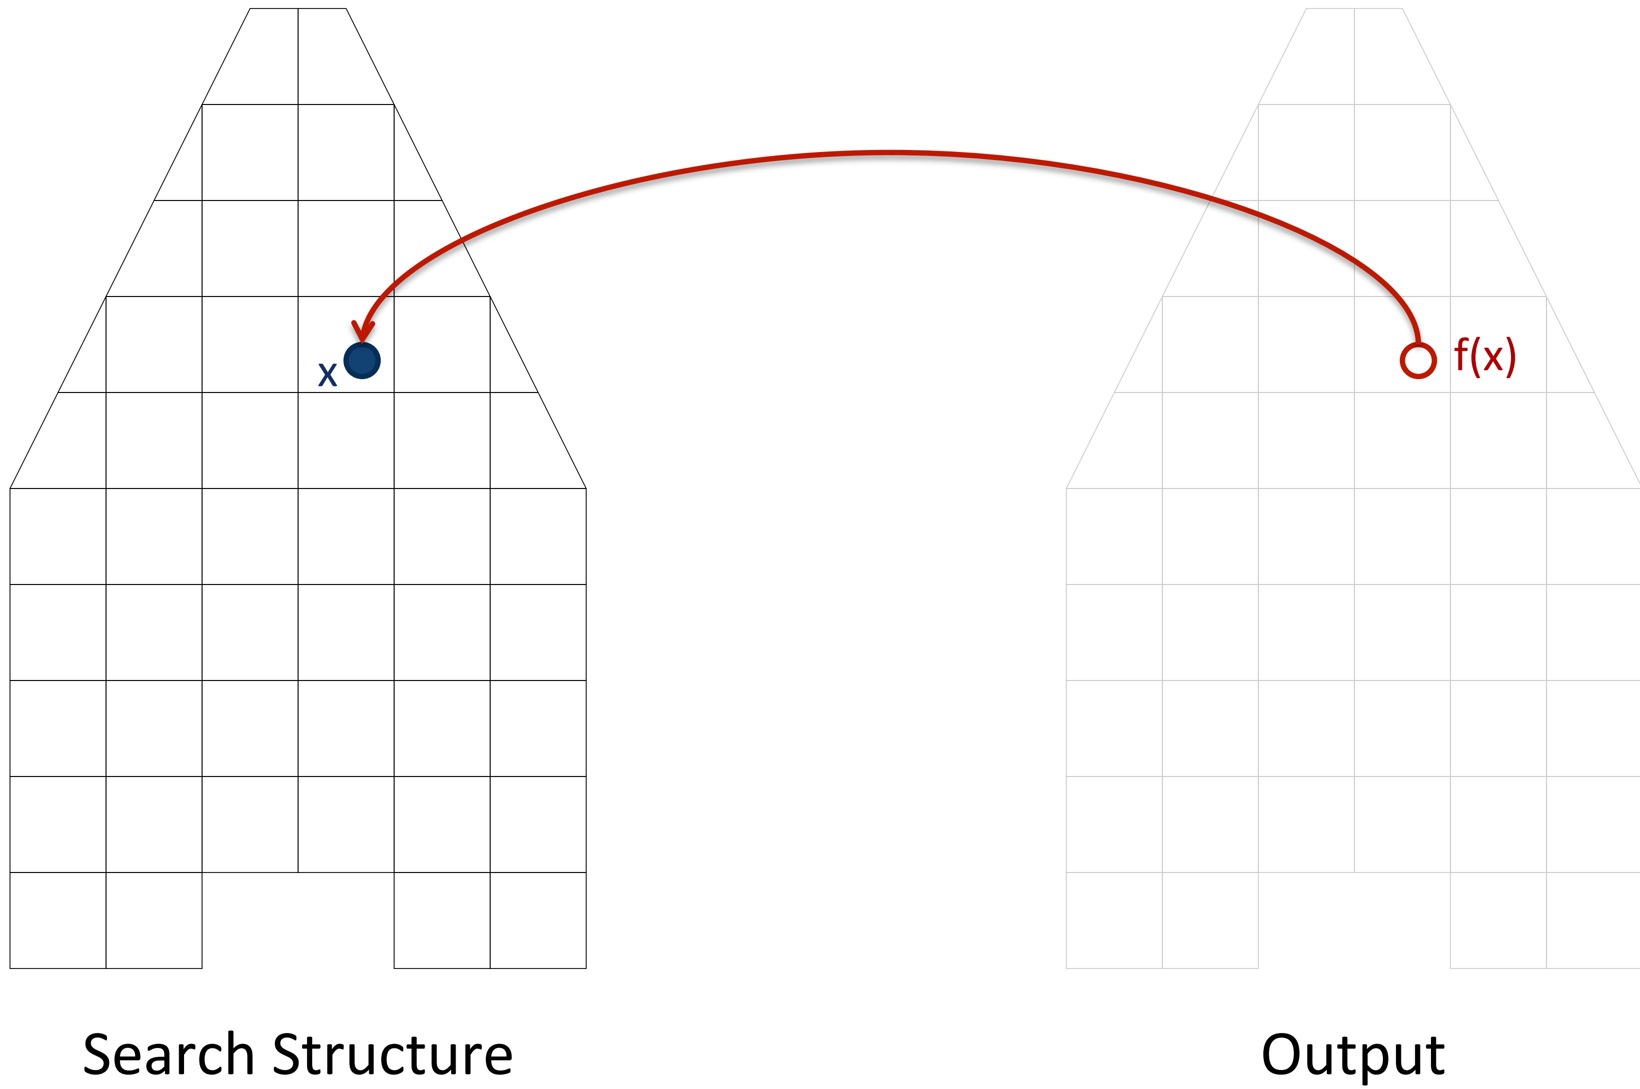

Query Data

Separable Element	Point, key, or query
Point Mapping	Identity
Cell Mapping	Identity
Field Mapping	1 query to 1 output
Collective Work	Building query structure

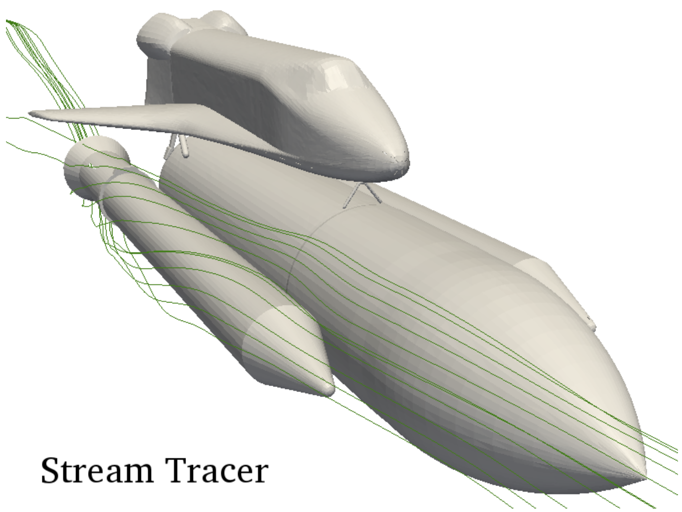

Stream Tracer

Summary

Sandia National Laboratories

Name	Separable Element	Point Mapping	Cell Mapping	Field Mapping	Collective Work	Example Algorithm
Basic Mapping	Any	Identity	Identity	1 to 1	None	Generate Ids
Map by Cell	Cell	Identity	Identity	Points on cell to cell	None	Cell Centers
Reconnect Cell	Cell	1 to 0 or 1	1 to 0 or more	Identity	None	Threshold
Build Independent Topology	Any	1 element to many points	1 element to many cells	Identity	None	Glyph
Build Connected Topology	Cell	1 cell to 0 or more points	1 to 0 or more	Interpolated points	Resolve duplicate points	Contour
Capture Cell Adjacencies	Point, edge, or face	Identity	Identity	Interpolated incident fields	Find incidence relationships	Normal Generation
Globally Reduce	Any	None in output	None in output	All to 1	Global reduction	Histogram
Query Data	Point, key, or query	Identity	Identity	1 query to 1 output	Building query structure	Stream Tracer

Conclusion

- Accelerators and other emerging architectures use massive threading that takes parallelism to a whole different scale
- Visualization algorithms can be characterized by:
  - How data is separated
  - How input elements are mapped to output elements
  - What collective work is performed
- Algorithms that share all three characterizations can be implemented with similar mechanisms
  - We can exploit this observation to implement new algorithms