

# Rethinking Visualization Frameworks for Extreme-Scale Computing

## Ultrascale Visualization Workshop

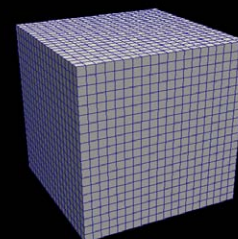
November 15, 2010

Kenneth Moreland  
Sandia National Laboratories

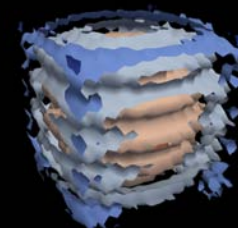
Stéphane Marchesin  
University of California at Davis

SAND 2010-8034P

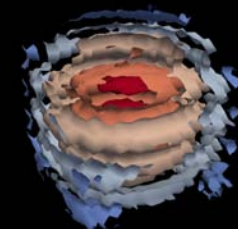
# Serial Visualization Pipeline



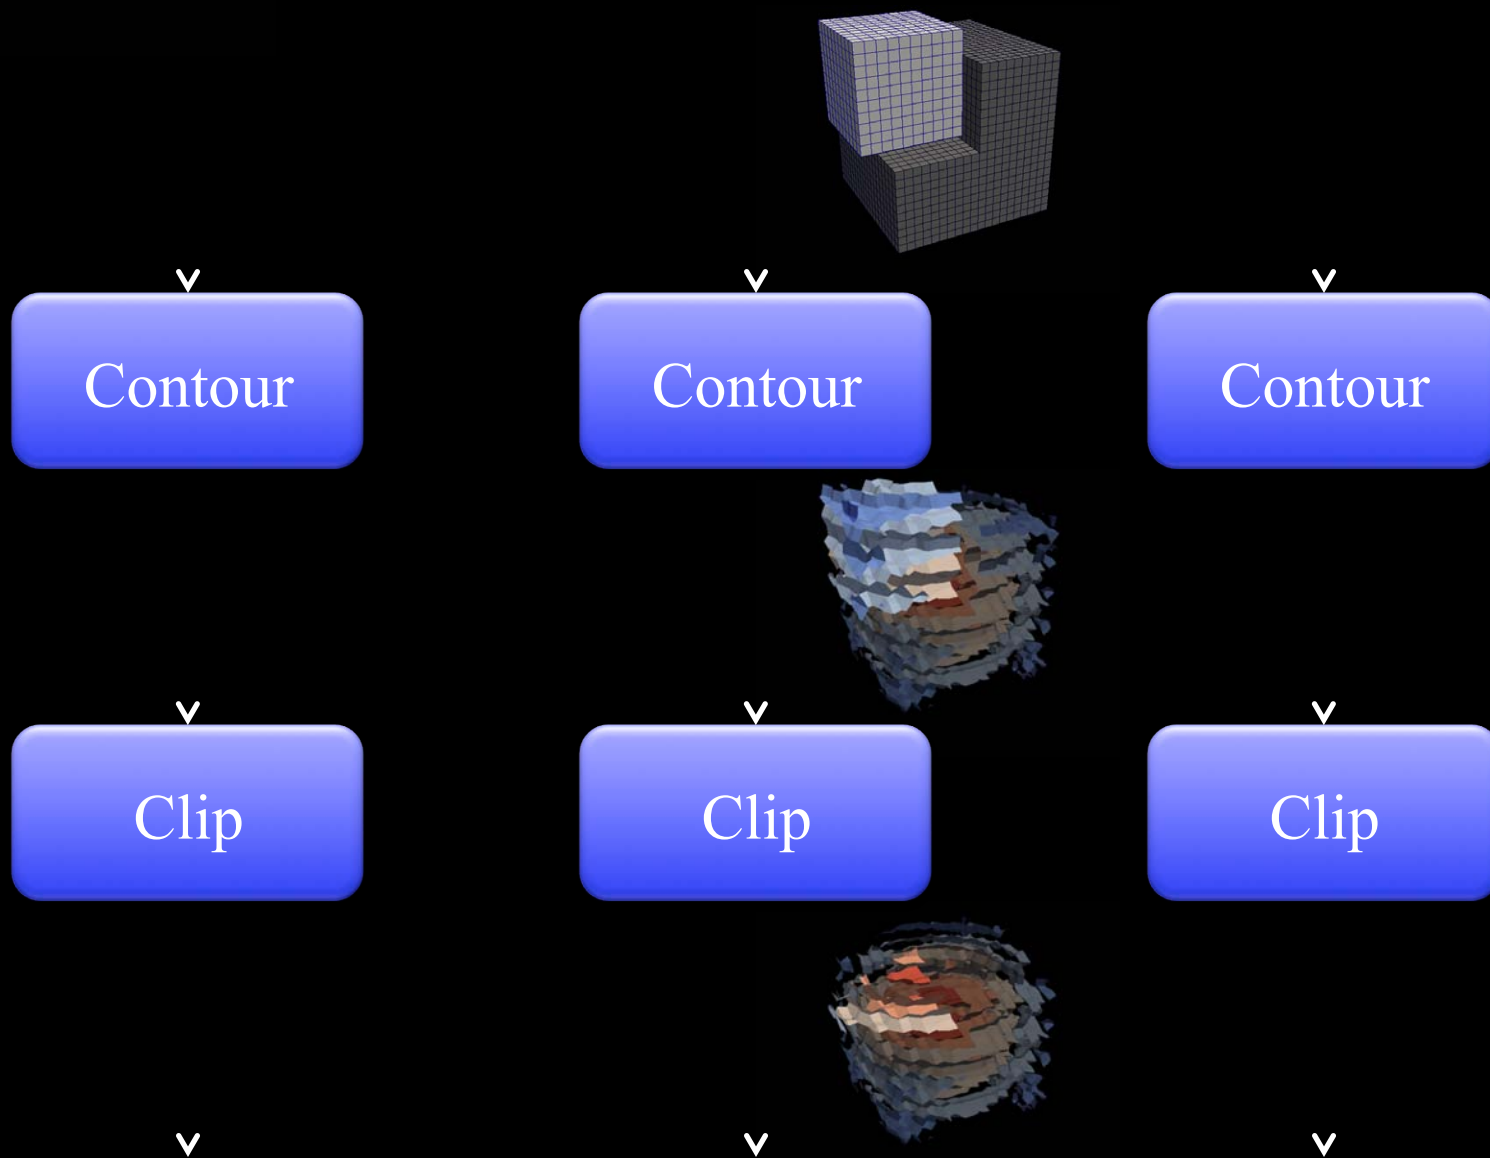
Contour



Clip



# Parallel Visualization Pipeline



# Exascale Projection

	<b>Jaguar – XT5</b>	<b>Exascale*</b>	<b>Increase</b>
Cores	224,256	100 million – 1 billion	~1,000×
Concurrency	224,256 way	10 billion way	~50,000×
Memory	300 Terabytes	128 Petabytes	~500×

# Exascale Projection

	Jaguar – XT5	Exascale*	Increase
Cores	224,256	100 million – 1 billion	~1,000×
Concurrency	224,256 way	10 billion way	~50,000×
Memory	300 Terabytes	128 Petabytes	~500×

## MPI Only?

Vis object code + state: 20MB

On Jaguar: 20MB × 200,000 processes = 4TB

On Exascale: 20MB × 10 billion processes = 200PB !

# Exascale Projection

	Jaguar – XT5	Exascale*	Increase
Cores	224,256	100 million – 1 billion	~1,000×
Concurrency	224,256 way	10 billion way	~50,000×
Memory	300 Terabytes	128 Petabytes	~500×

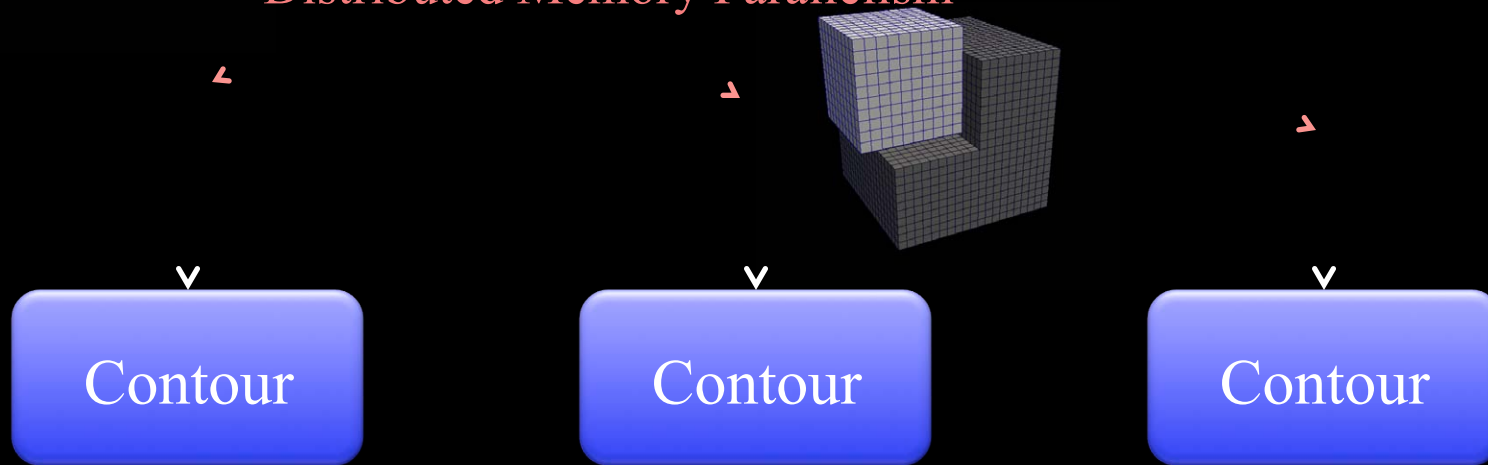
Visualization pipeline too heavyweight?

On Jaguar: 1 trillion cells → 5 million cells/core

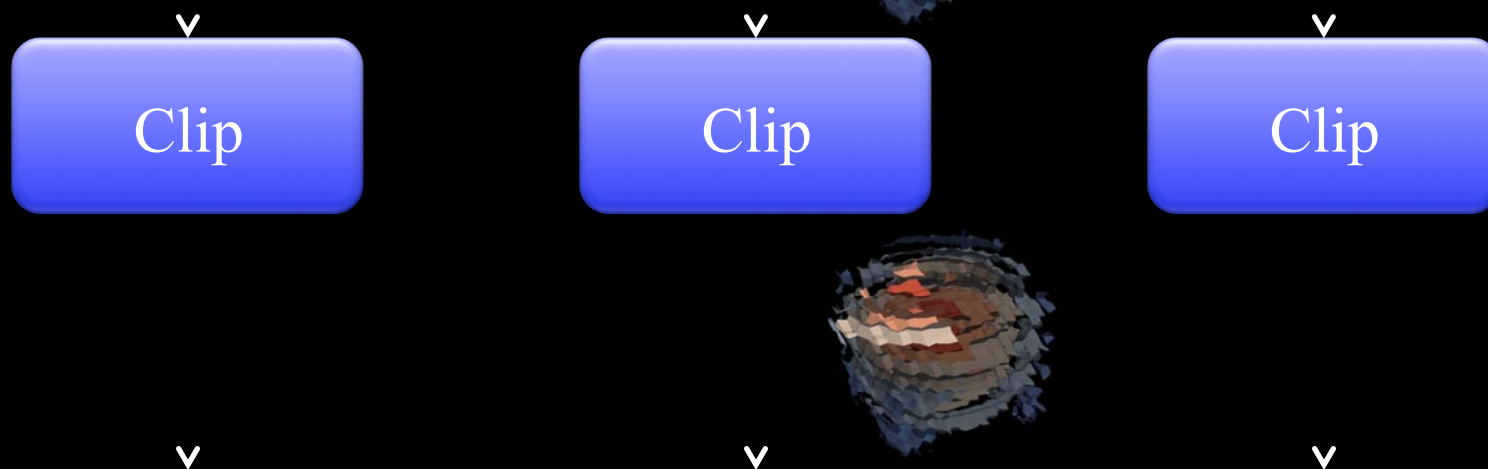
On Exascale: 500 trillion cells → 50K cells/core

# Hybrid Parallel Pipeline

Distributed Memory Parallelism



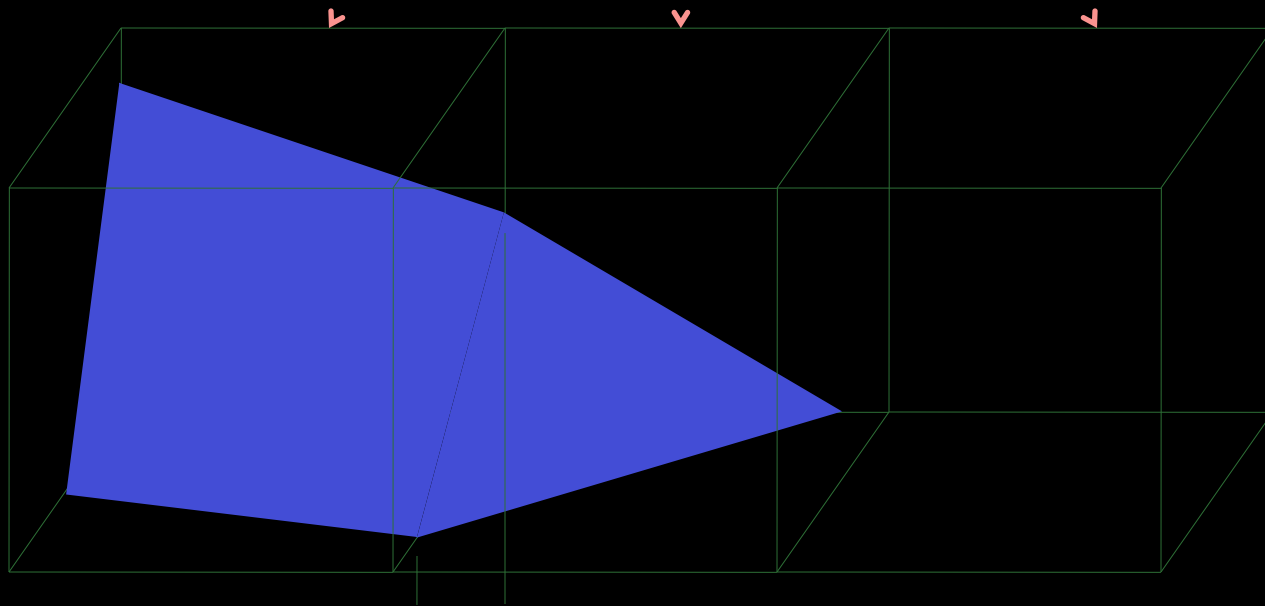
Shared Memory  
Parallel Processing



# Threaded Programming is Hard

## Example: Marching Cubes

Easy because cubes can be processed in parallel, right?

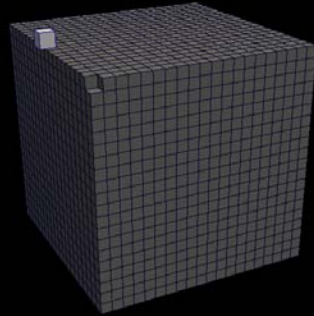


How do you resolve coincident points?  
How do you capture topological connections?

How do you pack the results?



# Revisiting the Pipeline



Filter



- Lightweight Object
- Serial Execution
- No explicit partitioning
- No access to larger structures
- No state

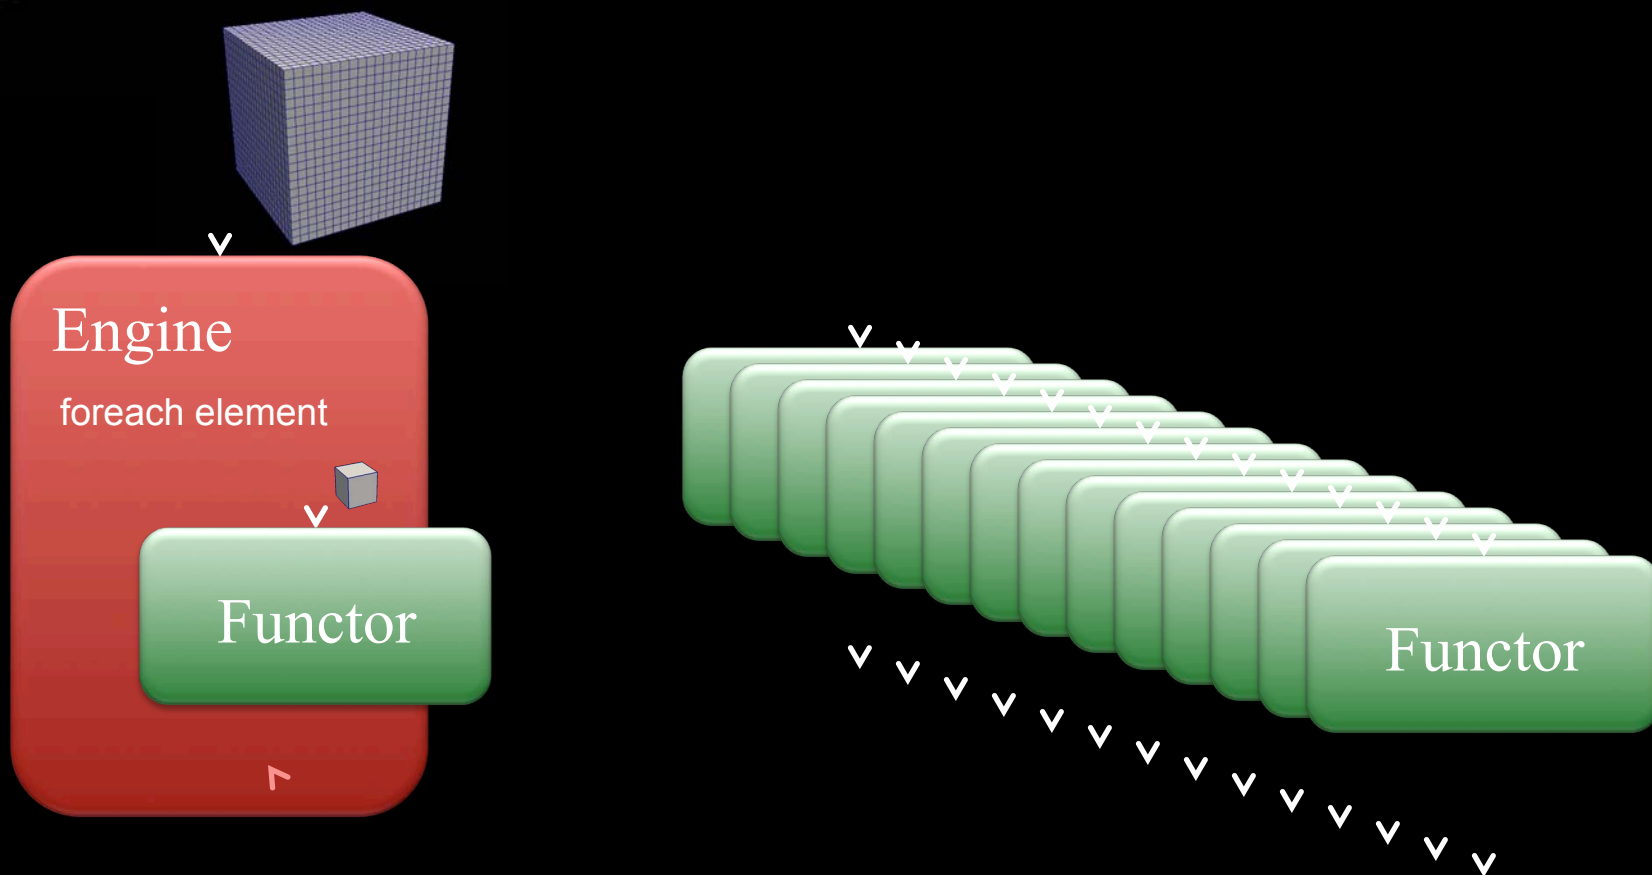
function ( in , out  )

# Functor

function ( in  , out  )



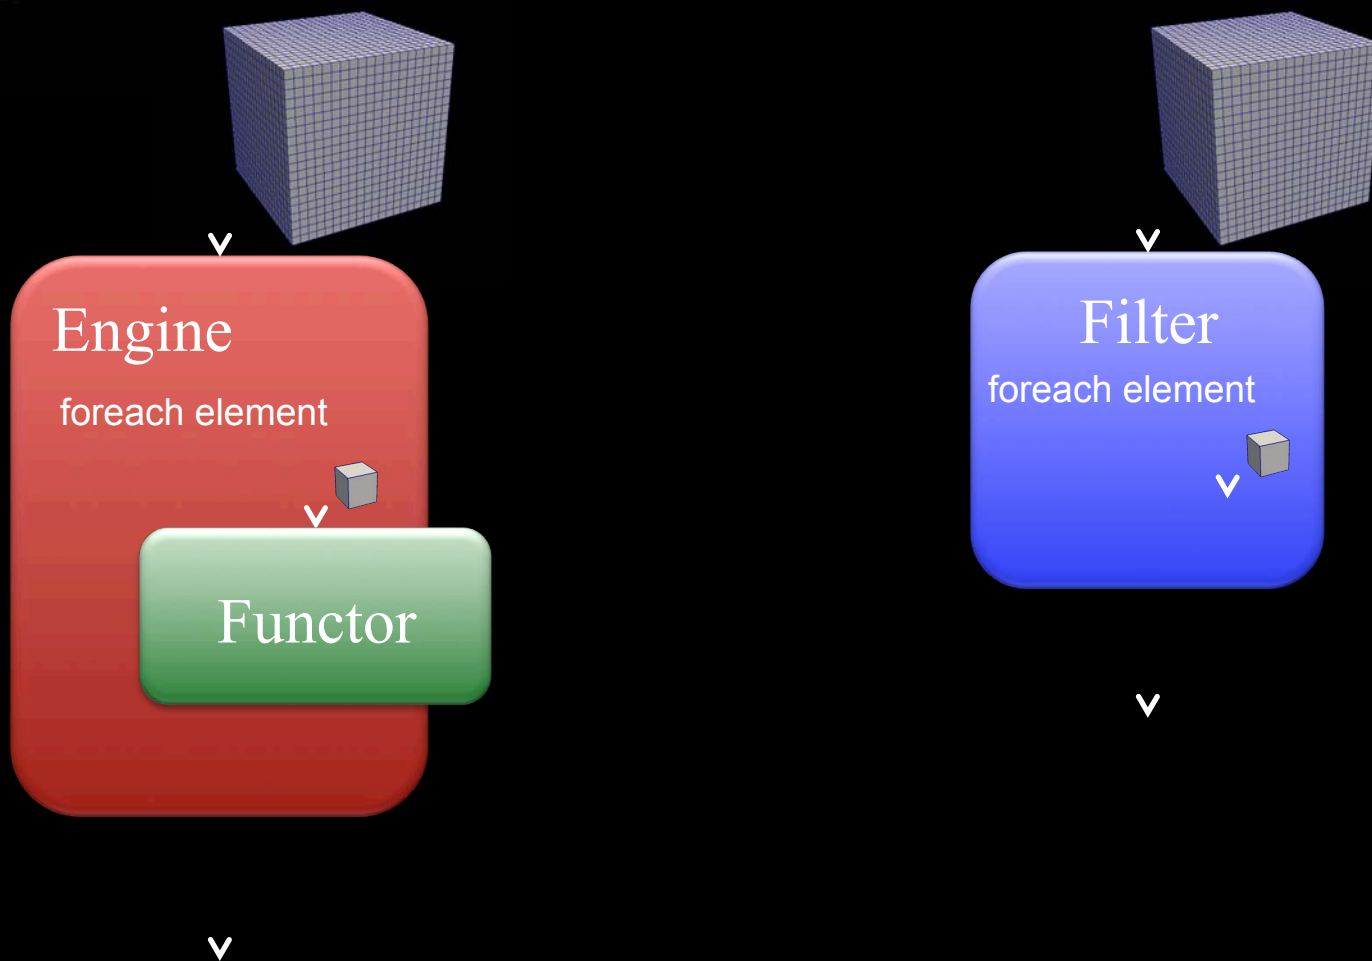
# Iteration Mechanism



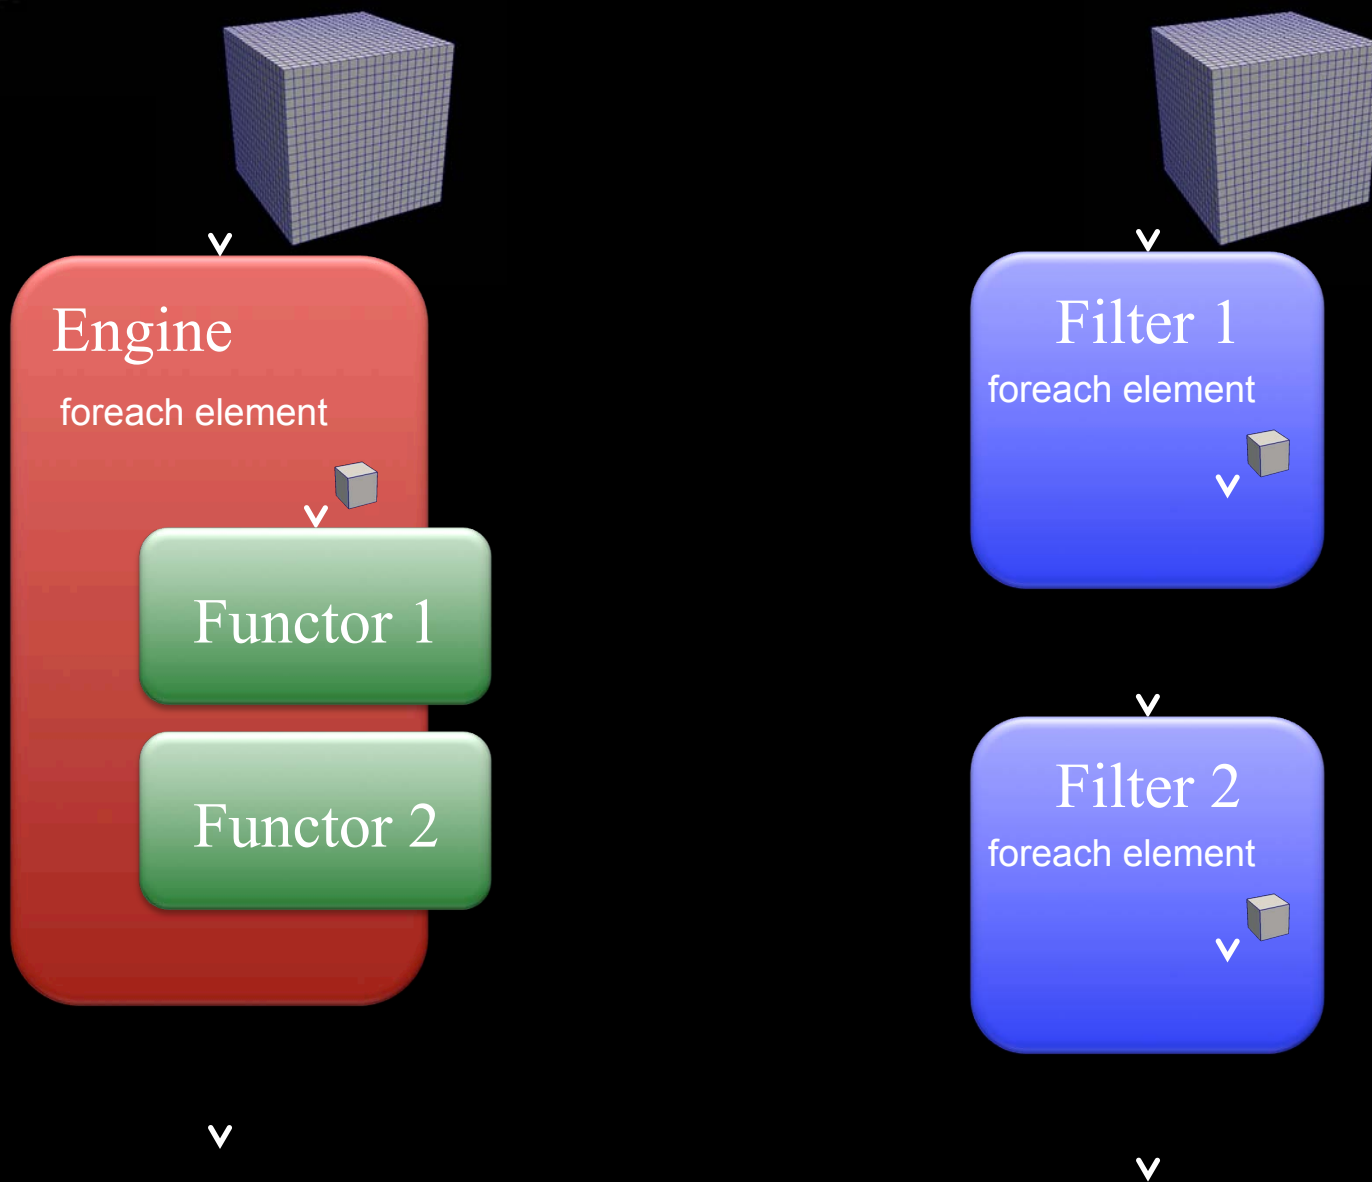
Conceptual  
Iteration

Reality: Iterations can be  
scheduled in parallel.

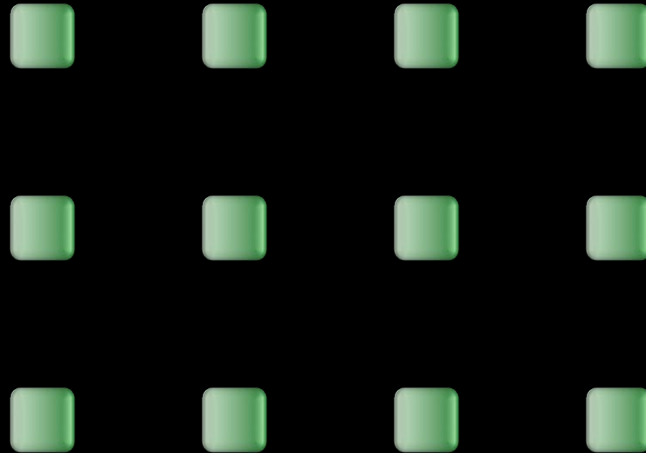
# Comparison



# Comparison

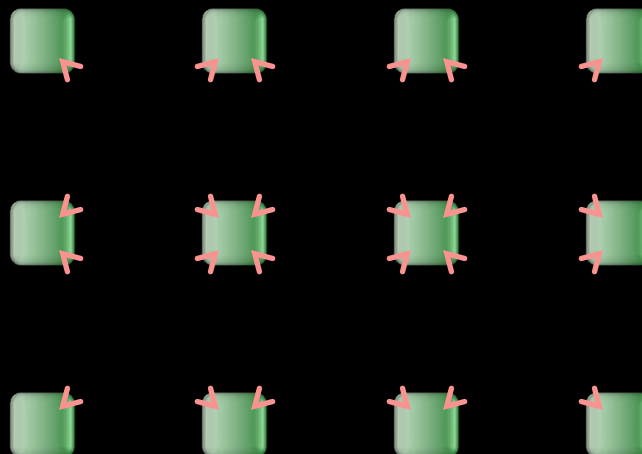


# Engine Types: Map



Example Usage: Vector Magnitude

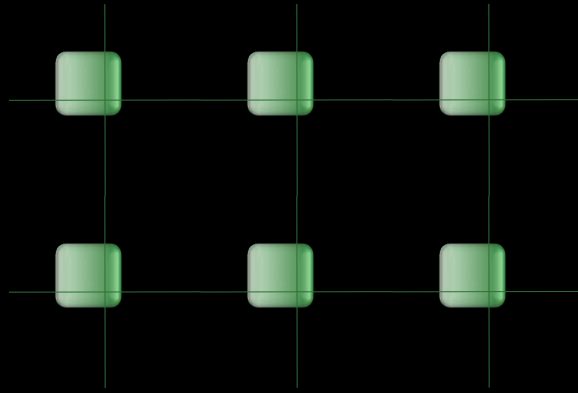
# Engine Type: Topological Reduce



Example Usages: Cell to Point, Normal Generation

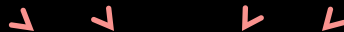


# Engine Types: Generate Geometry



Example Usages: Subdivide, Marching Cubes

# Engine Types: Compress



Example Usage: Marching Cubes

# Conclusion

- Why now? Why not before?
  - Rules of efficiency have changed.
- Concurrency: Coarse → Fine
- Execution cycles become free
- Minimizing DRAM I/O critical
- The current approach is unworkable
  - The incremental approach is unmanageable
- Designing for exascale requires lateral thinking