

The problem is data transfer

Bob Clapp

SEP136 - Pg 125-130, 131-136, SEP138 - Pg 297



Relative Cost of Imaging Algorithms

Outline

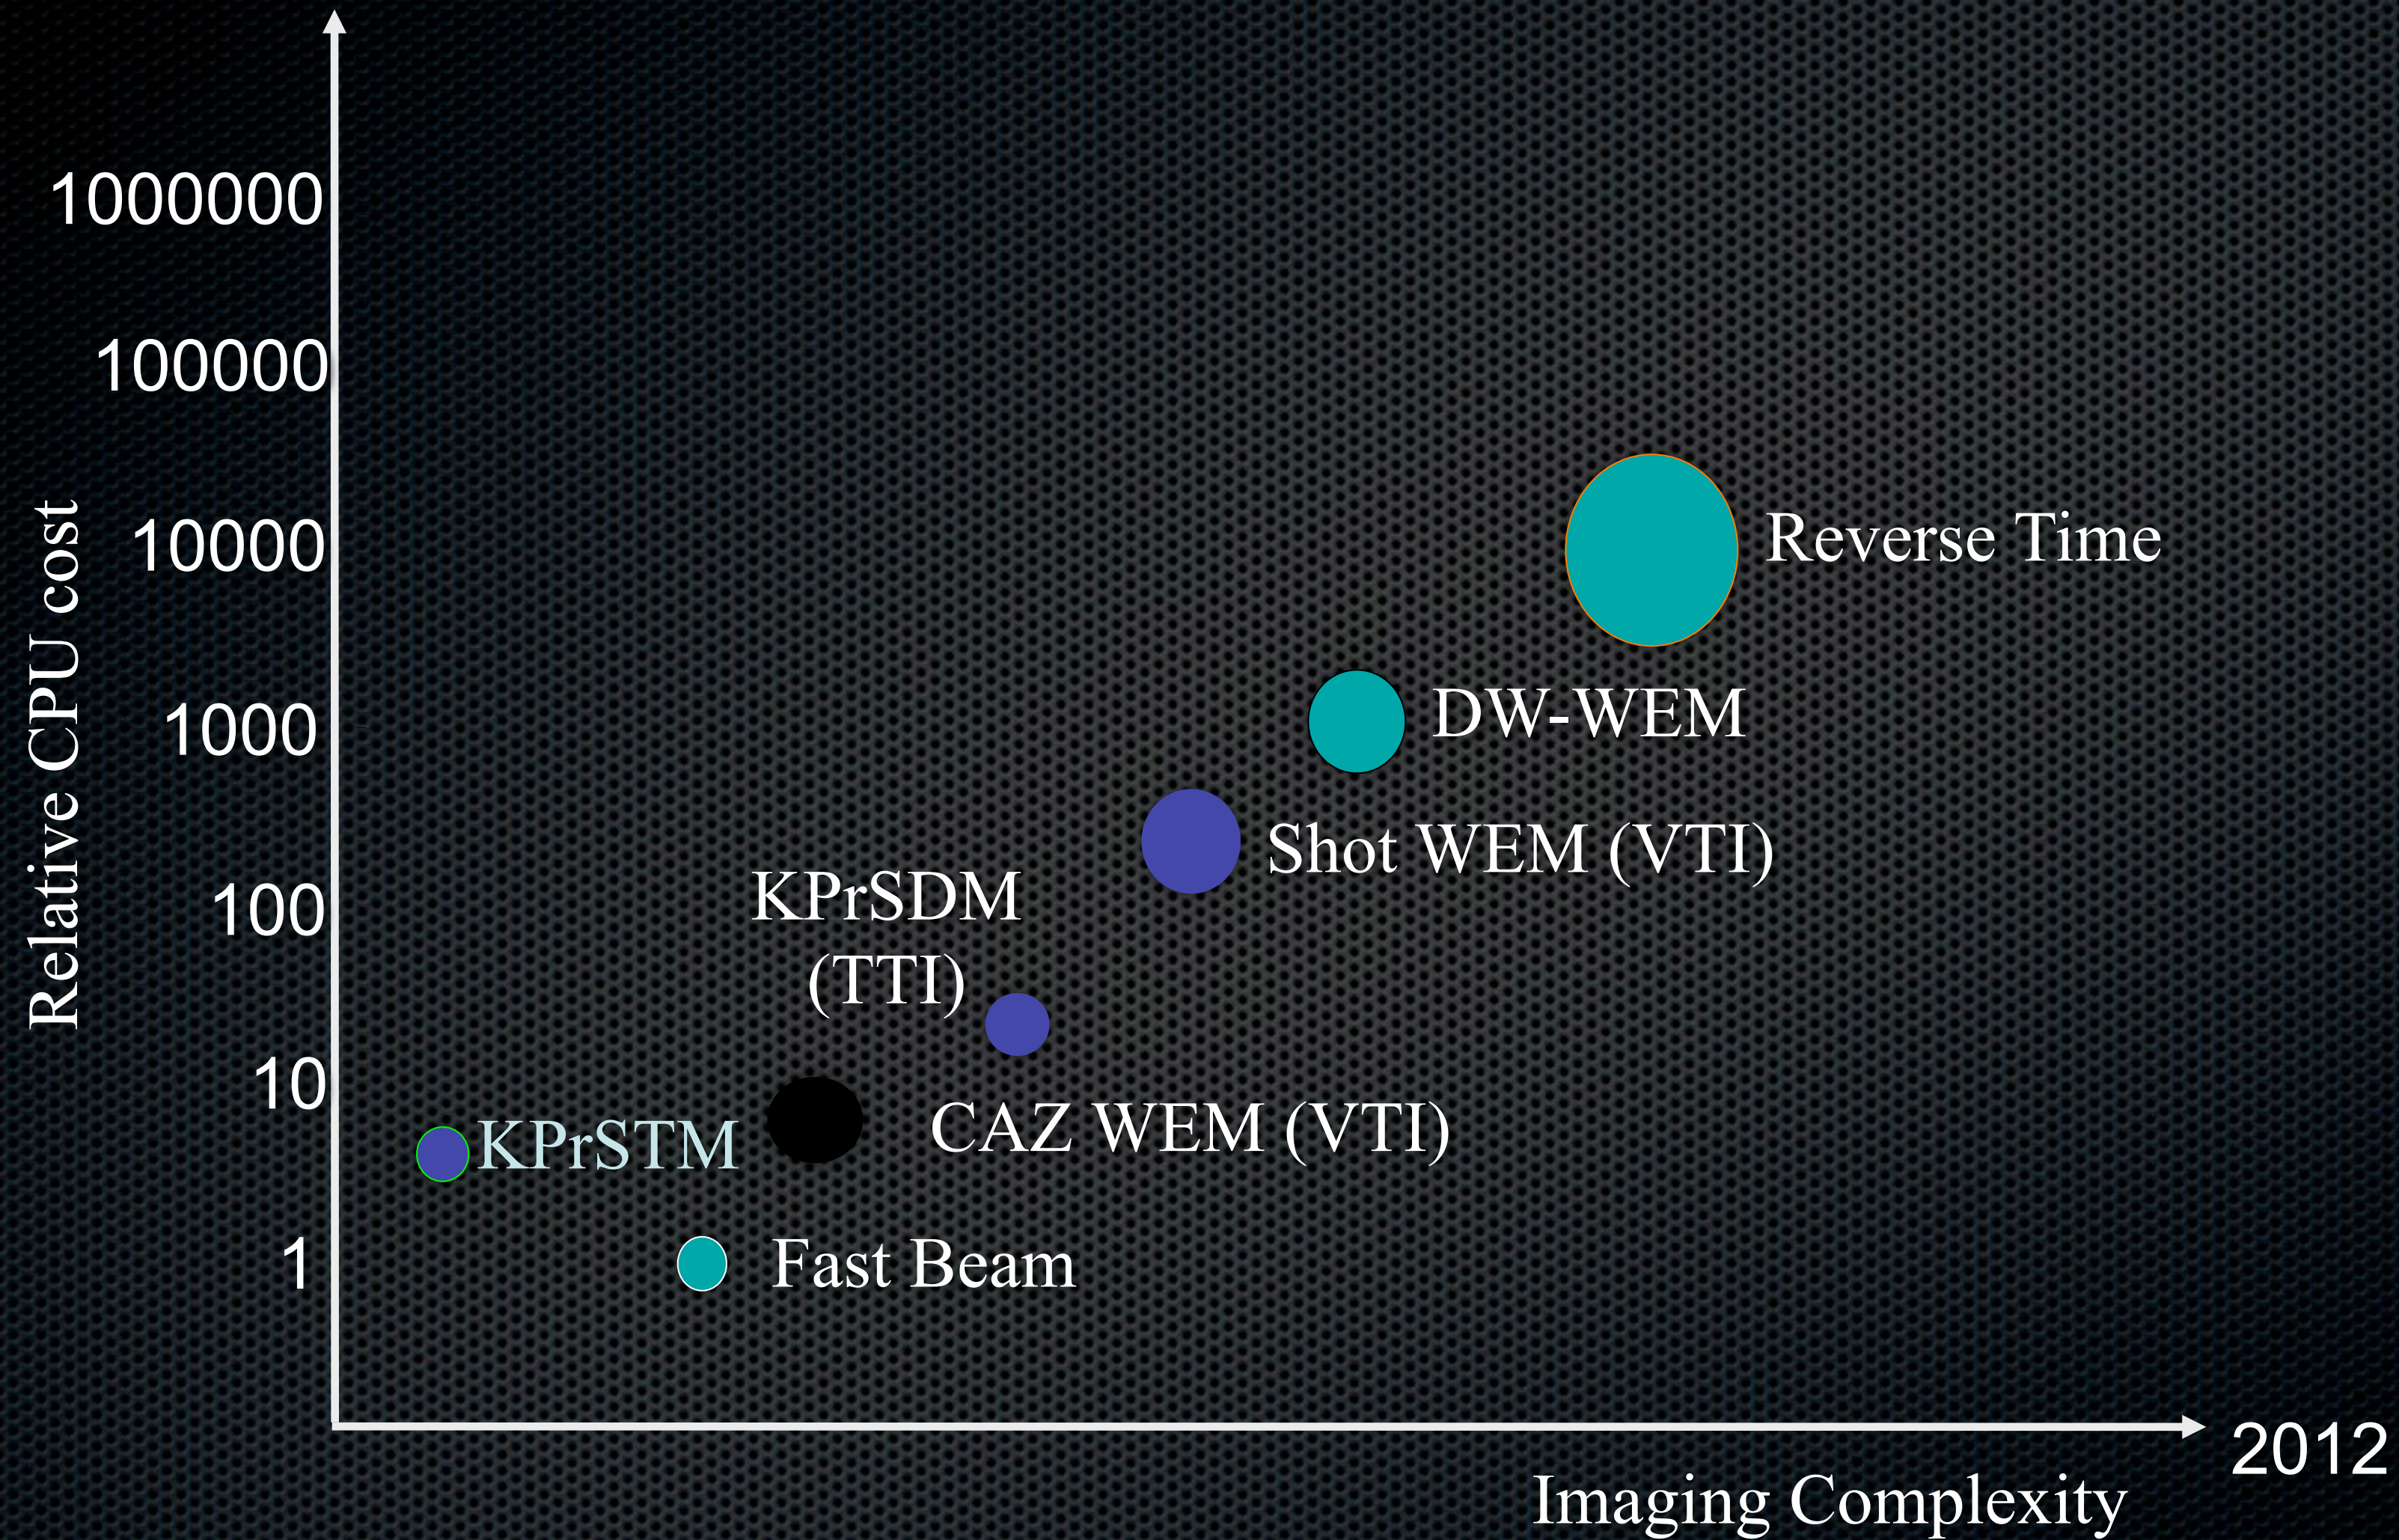
Motivation

Unique hardware

Algorithmic changes

Programming challenges

Demo



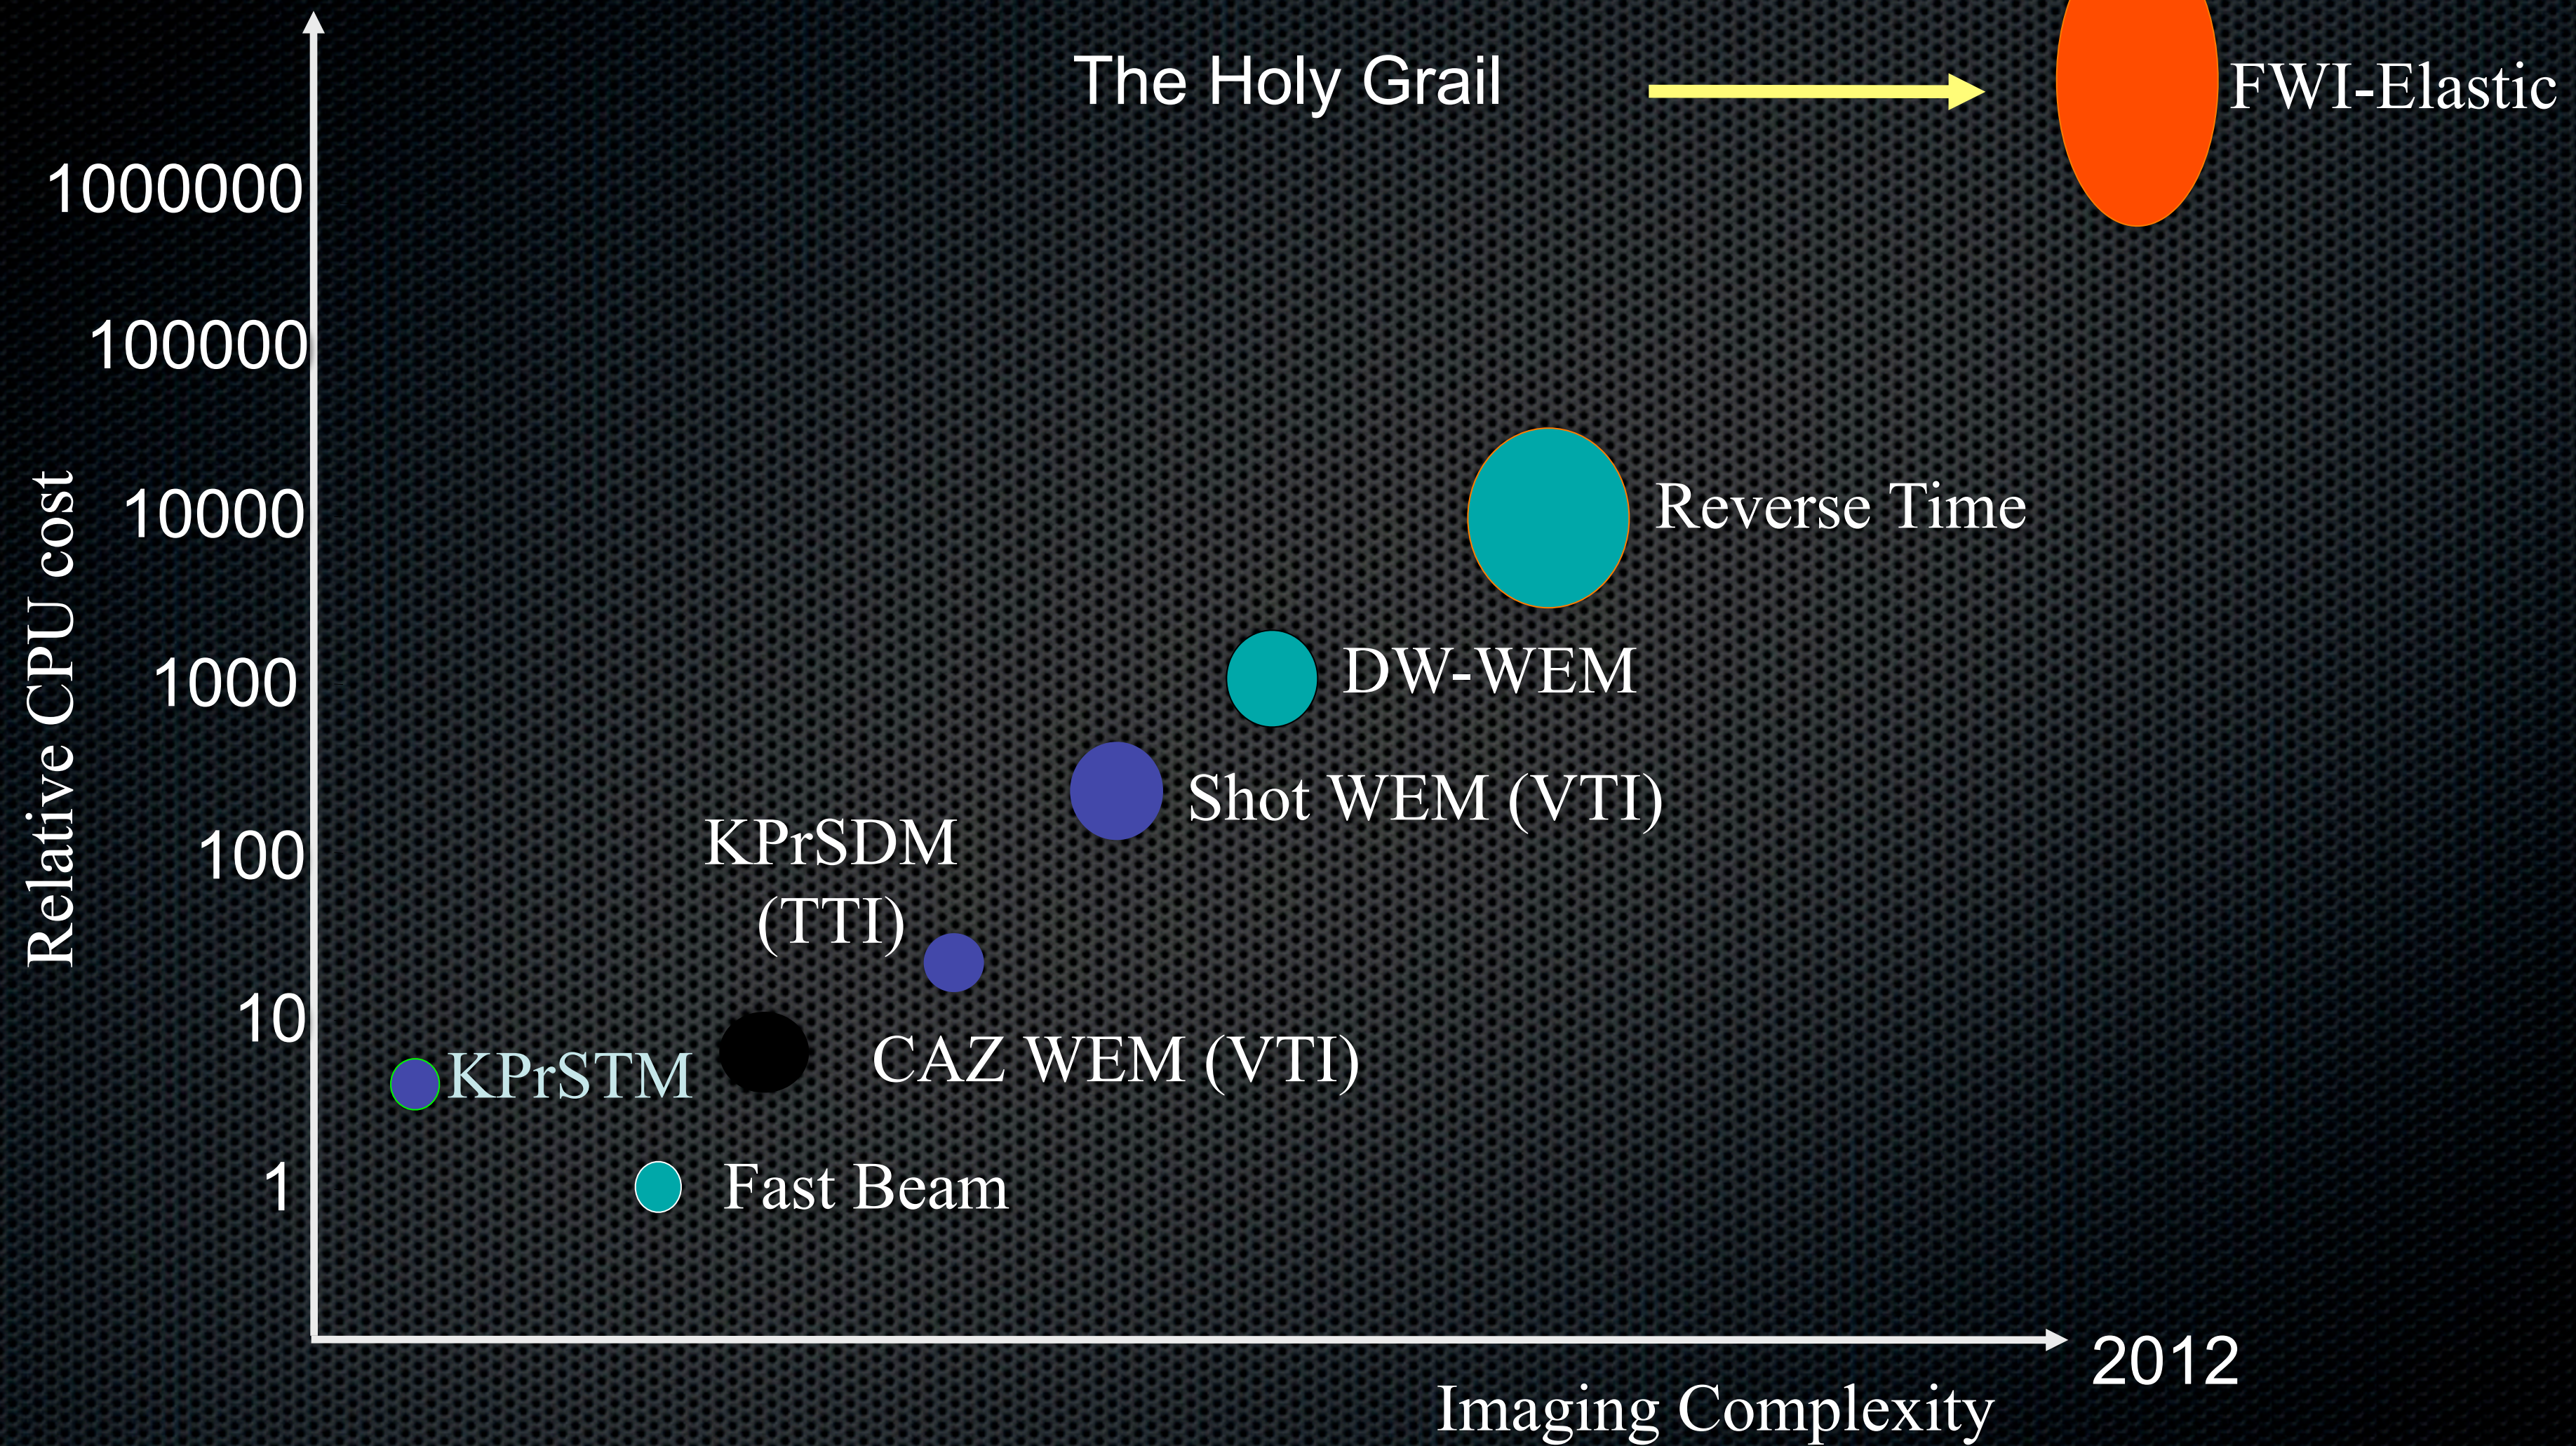
● : Memory and disk space cost



Bob Clapp

SEP Meeting 2009

Relative Cost of Imaging Algorithms



● : Memory and disk space cost

Outline

Motivation

Unique hardware

Algorithmic changes

Programming challenges

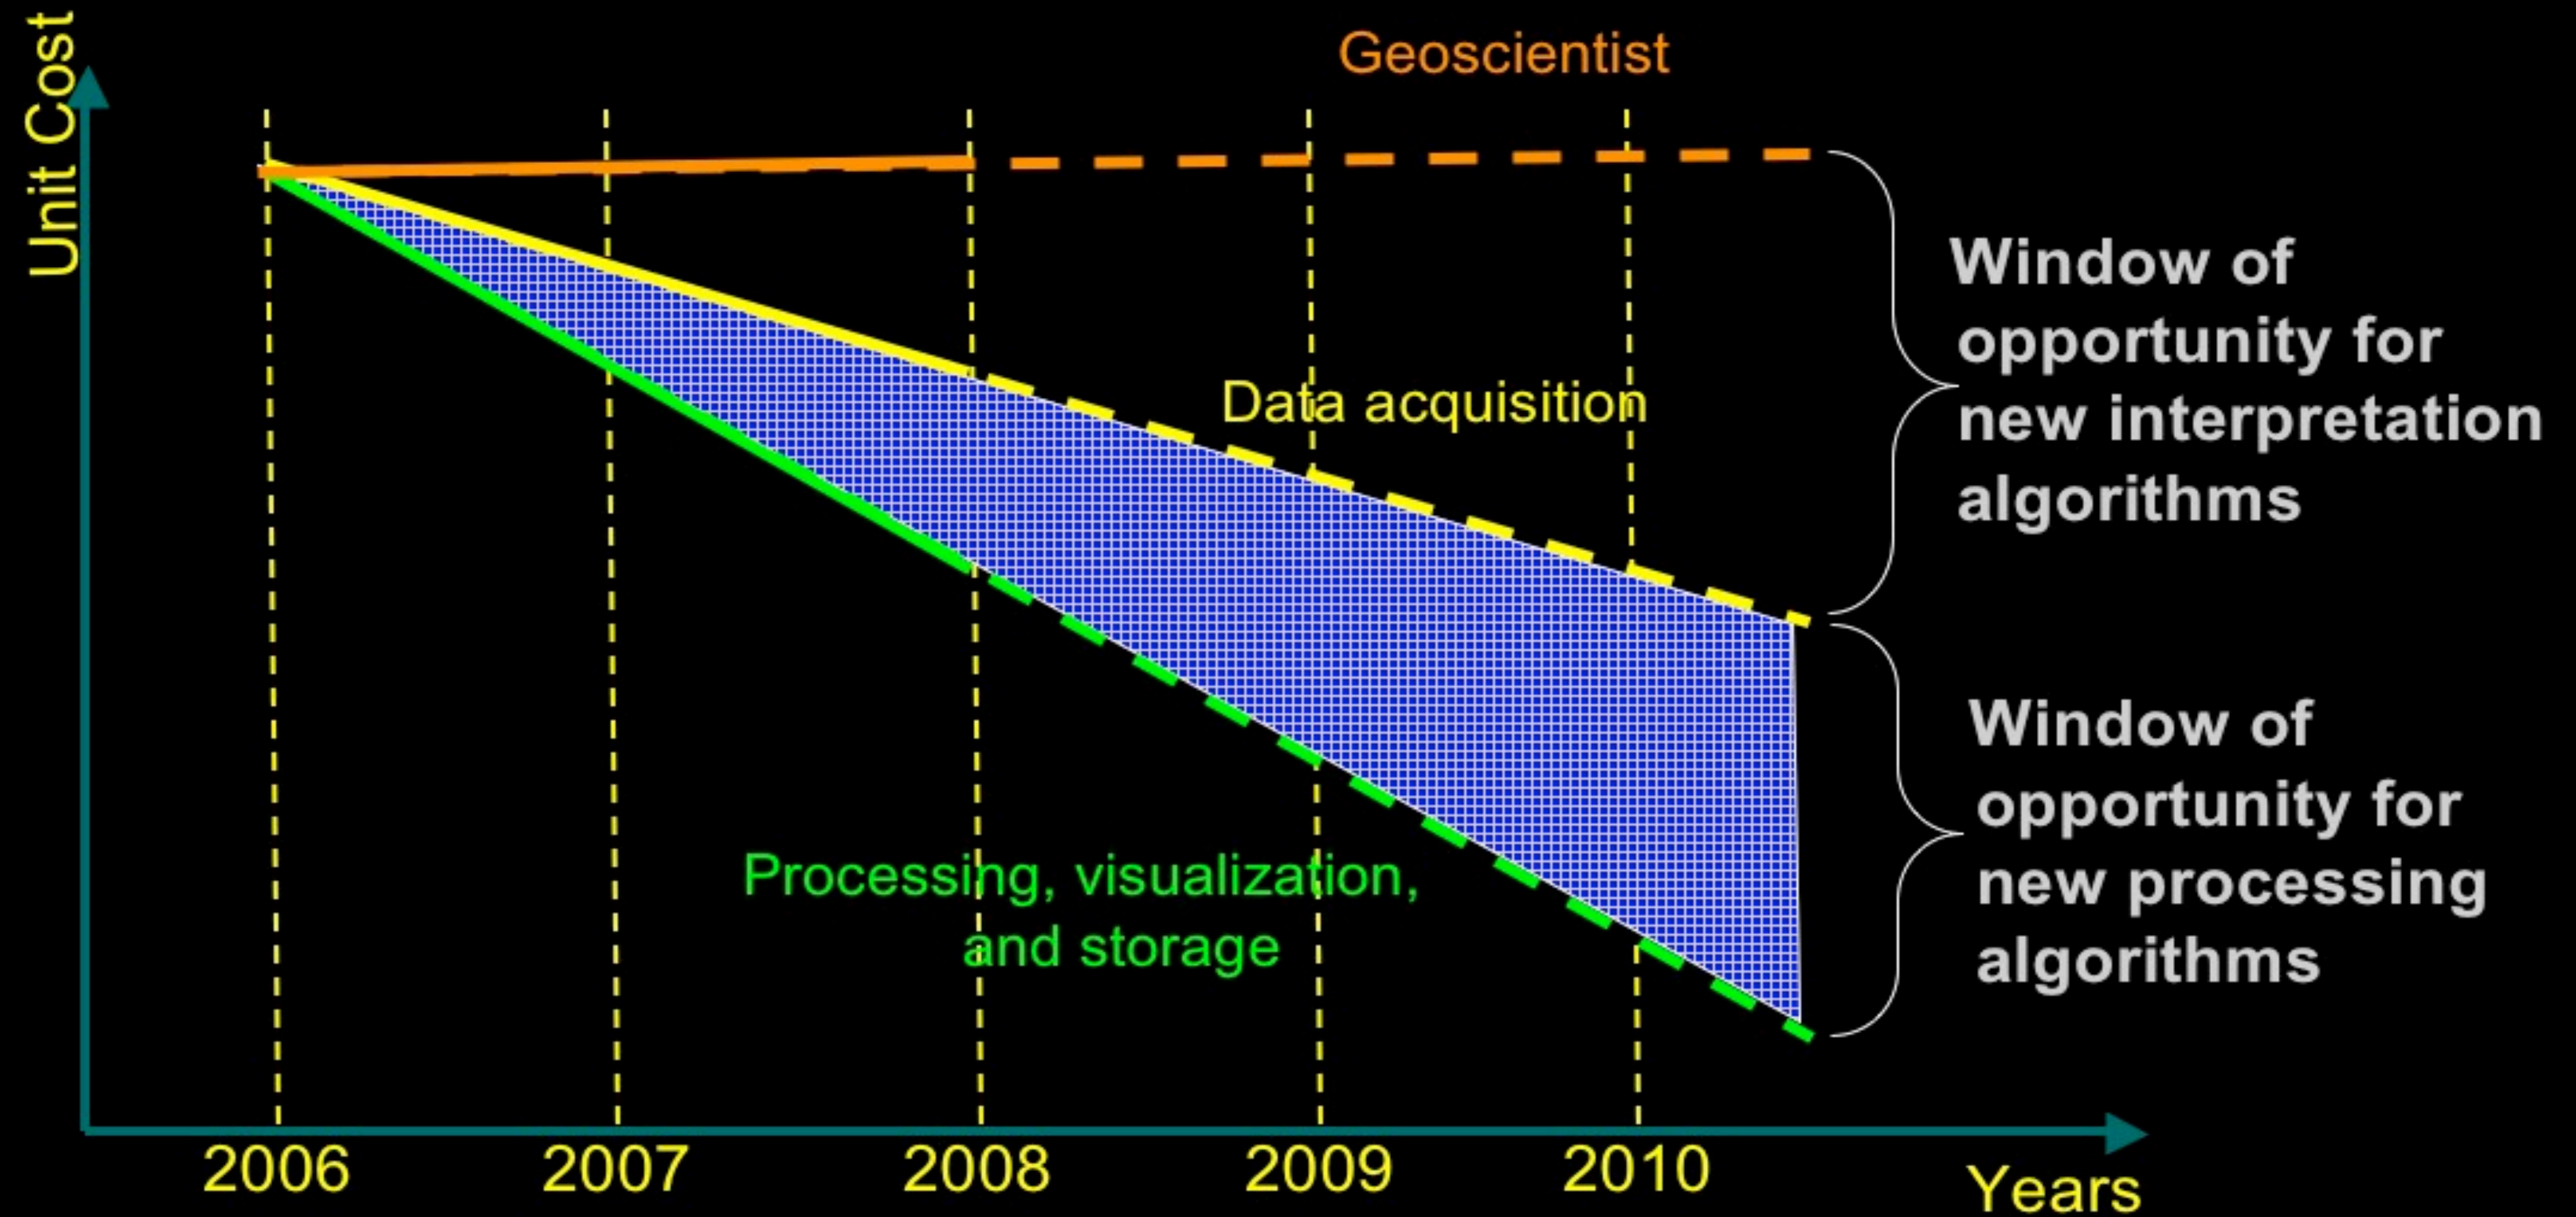
Demo



Bob Clapp

SEP Meeting 2009

Taking advantage of windows of opportunity



Outline

Motivation

Unique hardware

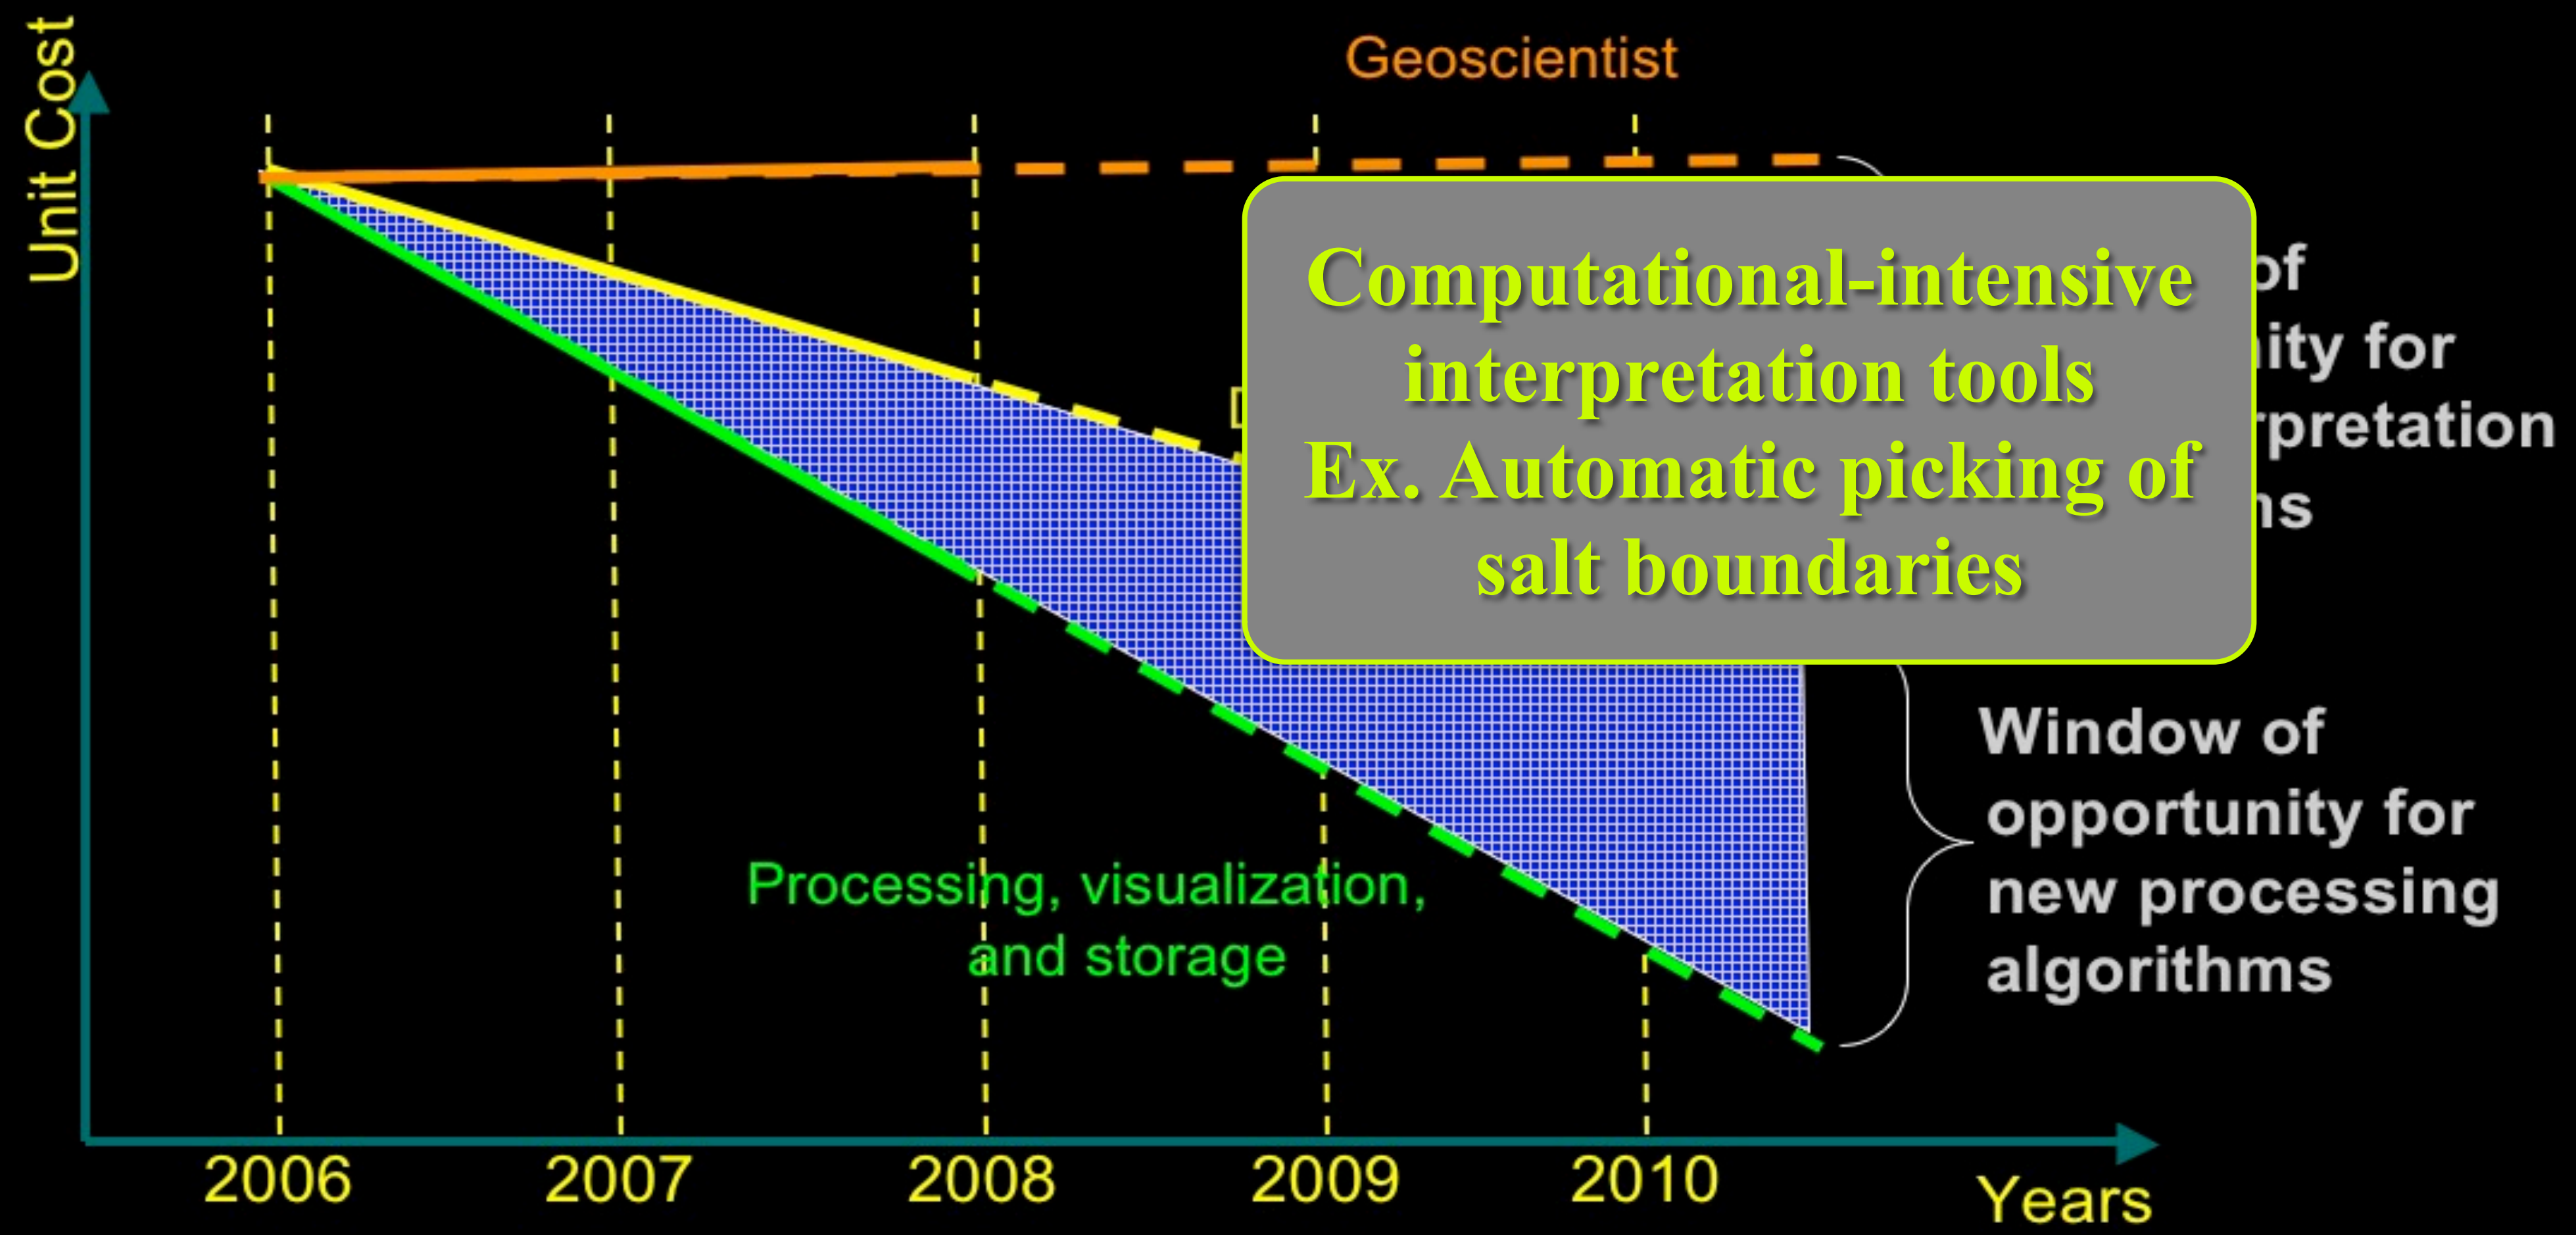
Algorithmic changes

Programming challenges

Demo



Taking advantage of windows of opportunities



Outline

Motivation

Unique hardware

Algorithmic changes

Programming challenges

Demo



Direct solution of least-squares imaging?

$$\hat{\mathbf{m}} = \left(\mathbf{L}'\mathbf{L} + \varepsilon\mathbf{R} \right)^{-1} \mathbf{L}'\mathbf{d}_{\text{obs}}$$



$$\hat{\mathbf{m}} = \left(\mathbf{H} + \varepsilon\mathbf{R} \right)^{-1} \mathbf{m}_{\text{mig}}$$

Outline

Motivation

Unique hardware

Algorithmic changes

Programming challenges

Demo



Target-oriented inversion in the image space

$$\hat{\mathbf{m}} = \left(\mathbf{L}'\mathbf{L} + \varepsilon\mathbf{R} \right)^{-1} \mathbf{L}'\mathbf{d}_{\text{obs}}$$



Precomputed

$$\hat{\mathbf{m}} = \left(\mathbf{H} + \varepsilon\mathbf{R} \right)^{-1} \mathbf{m}_{\text{mig}}$$

Outline

Motivation

Unique hardware

Algorithmic changes

Programming challenges

Demo



3D Migration

Outline


Motivation

Unique hardware

Algorithmic changes

Programming challenges

Demo



Bob Clapp

SEP Meeting 2009

3D Inversion

Outline

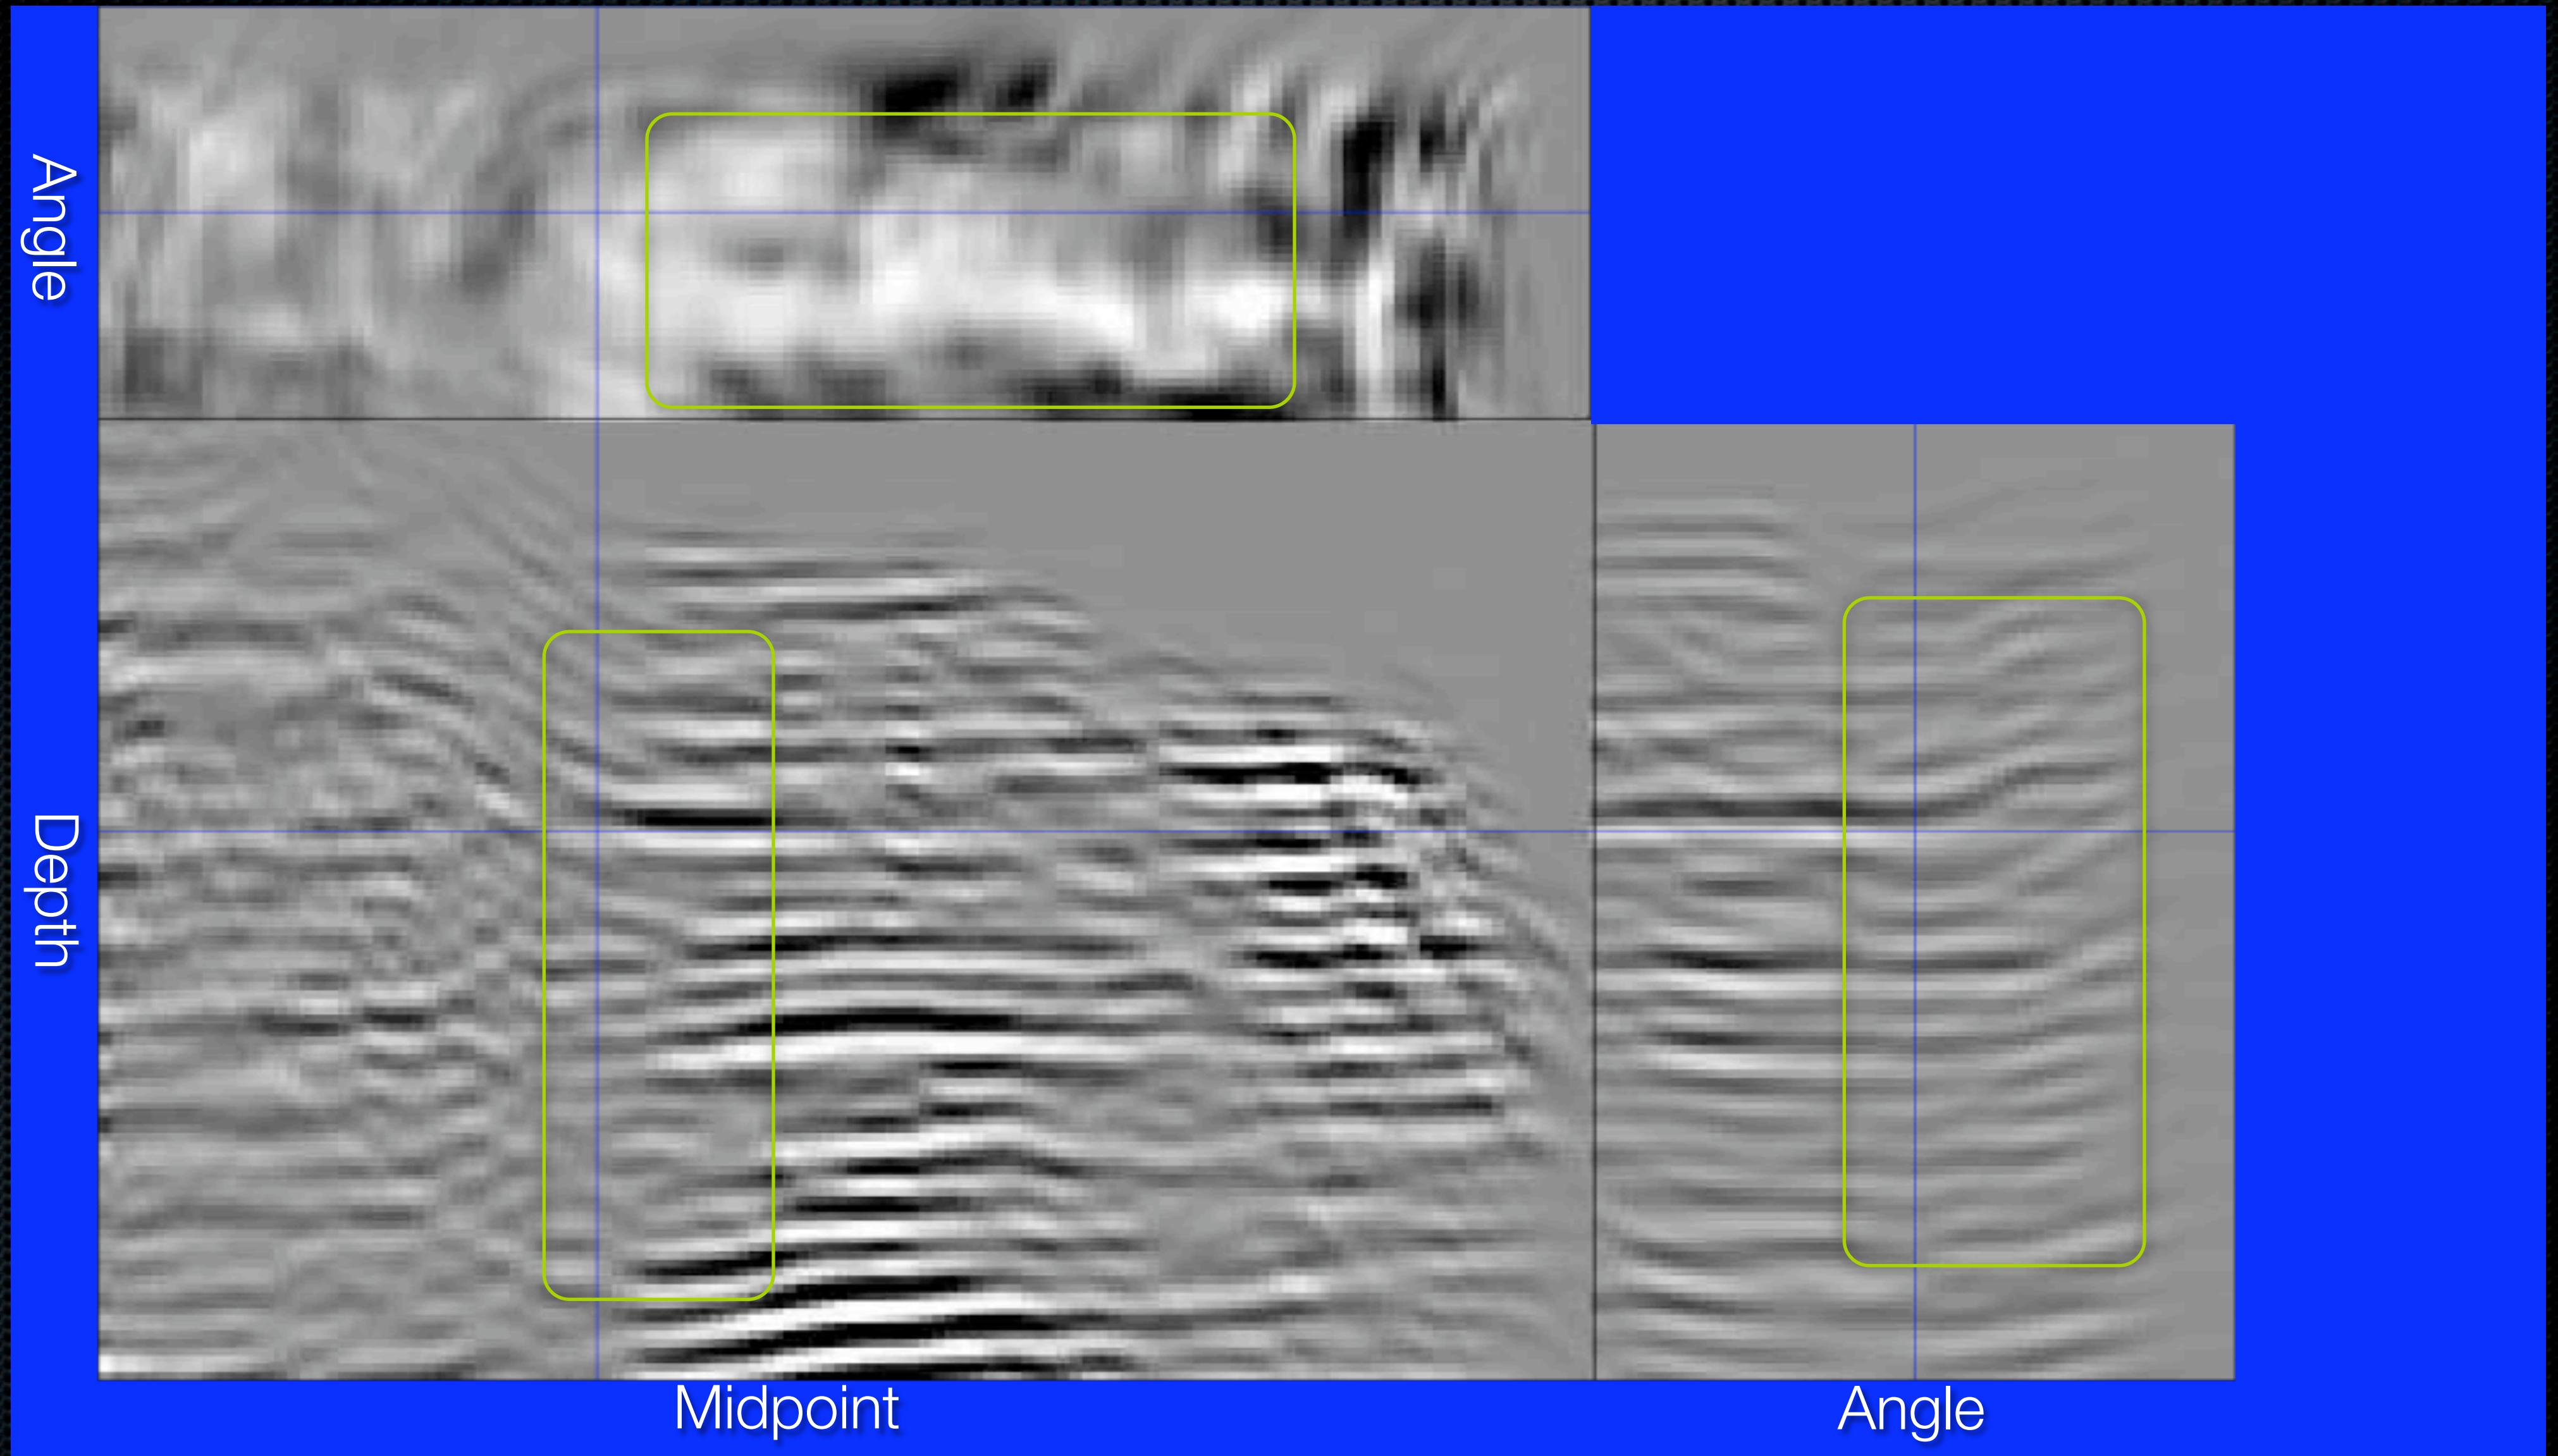
Motivation

Unique hardware

Algorithmic changes

Programming challenges

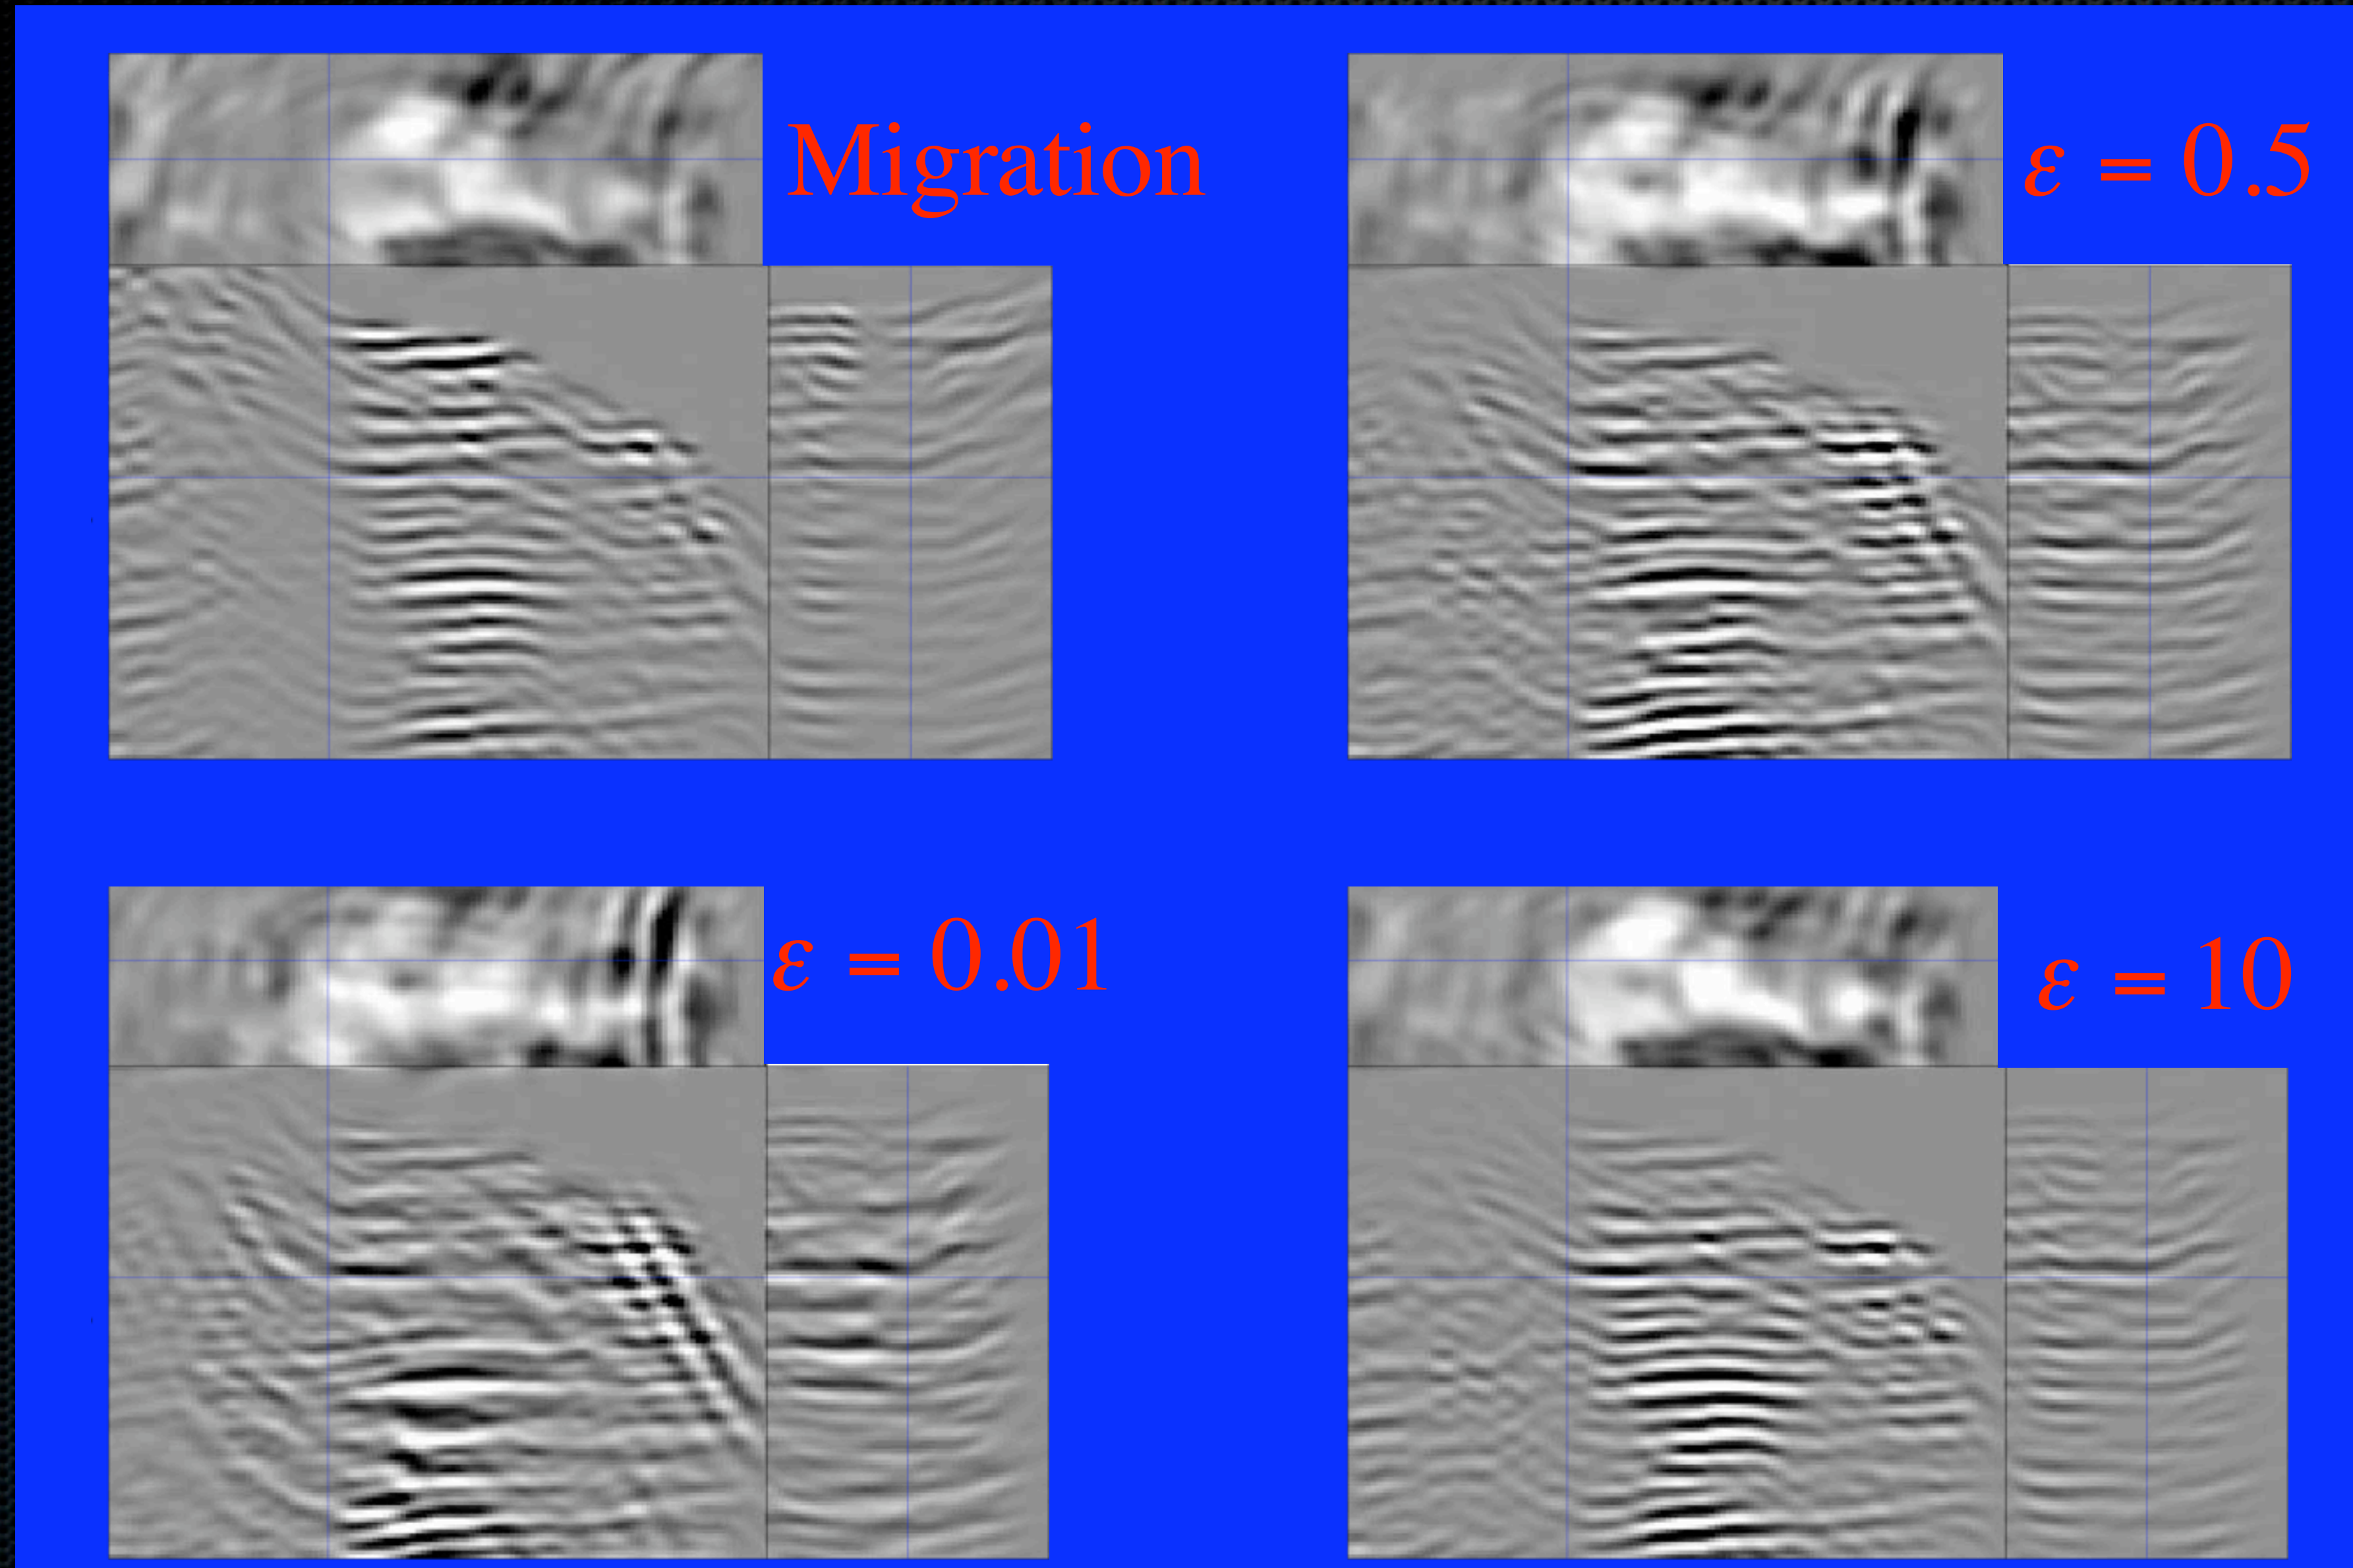
Demo



Bob Clapp

SEP Meeting 2009

Inversions with different ϵ



Outline

Motivation

Unique hardware

Algorithmic changes

Programming challenges

Demo



Bob Clapp

SEP Meeting 2009

Salt boundary picking by global image segmentation

Outline

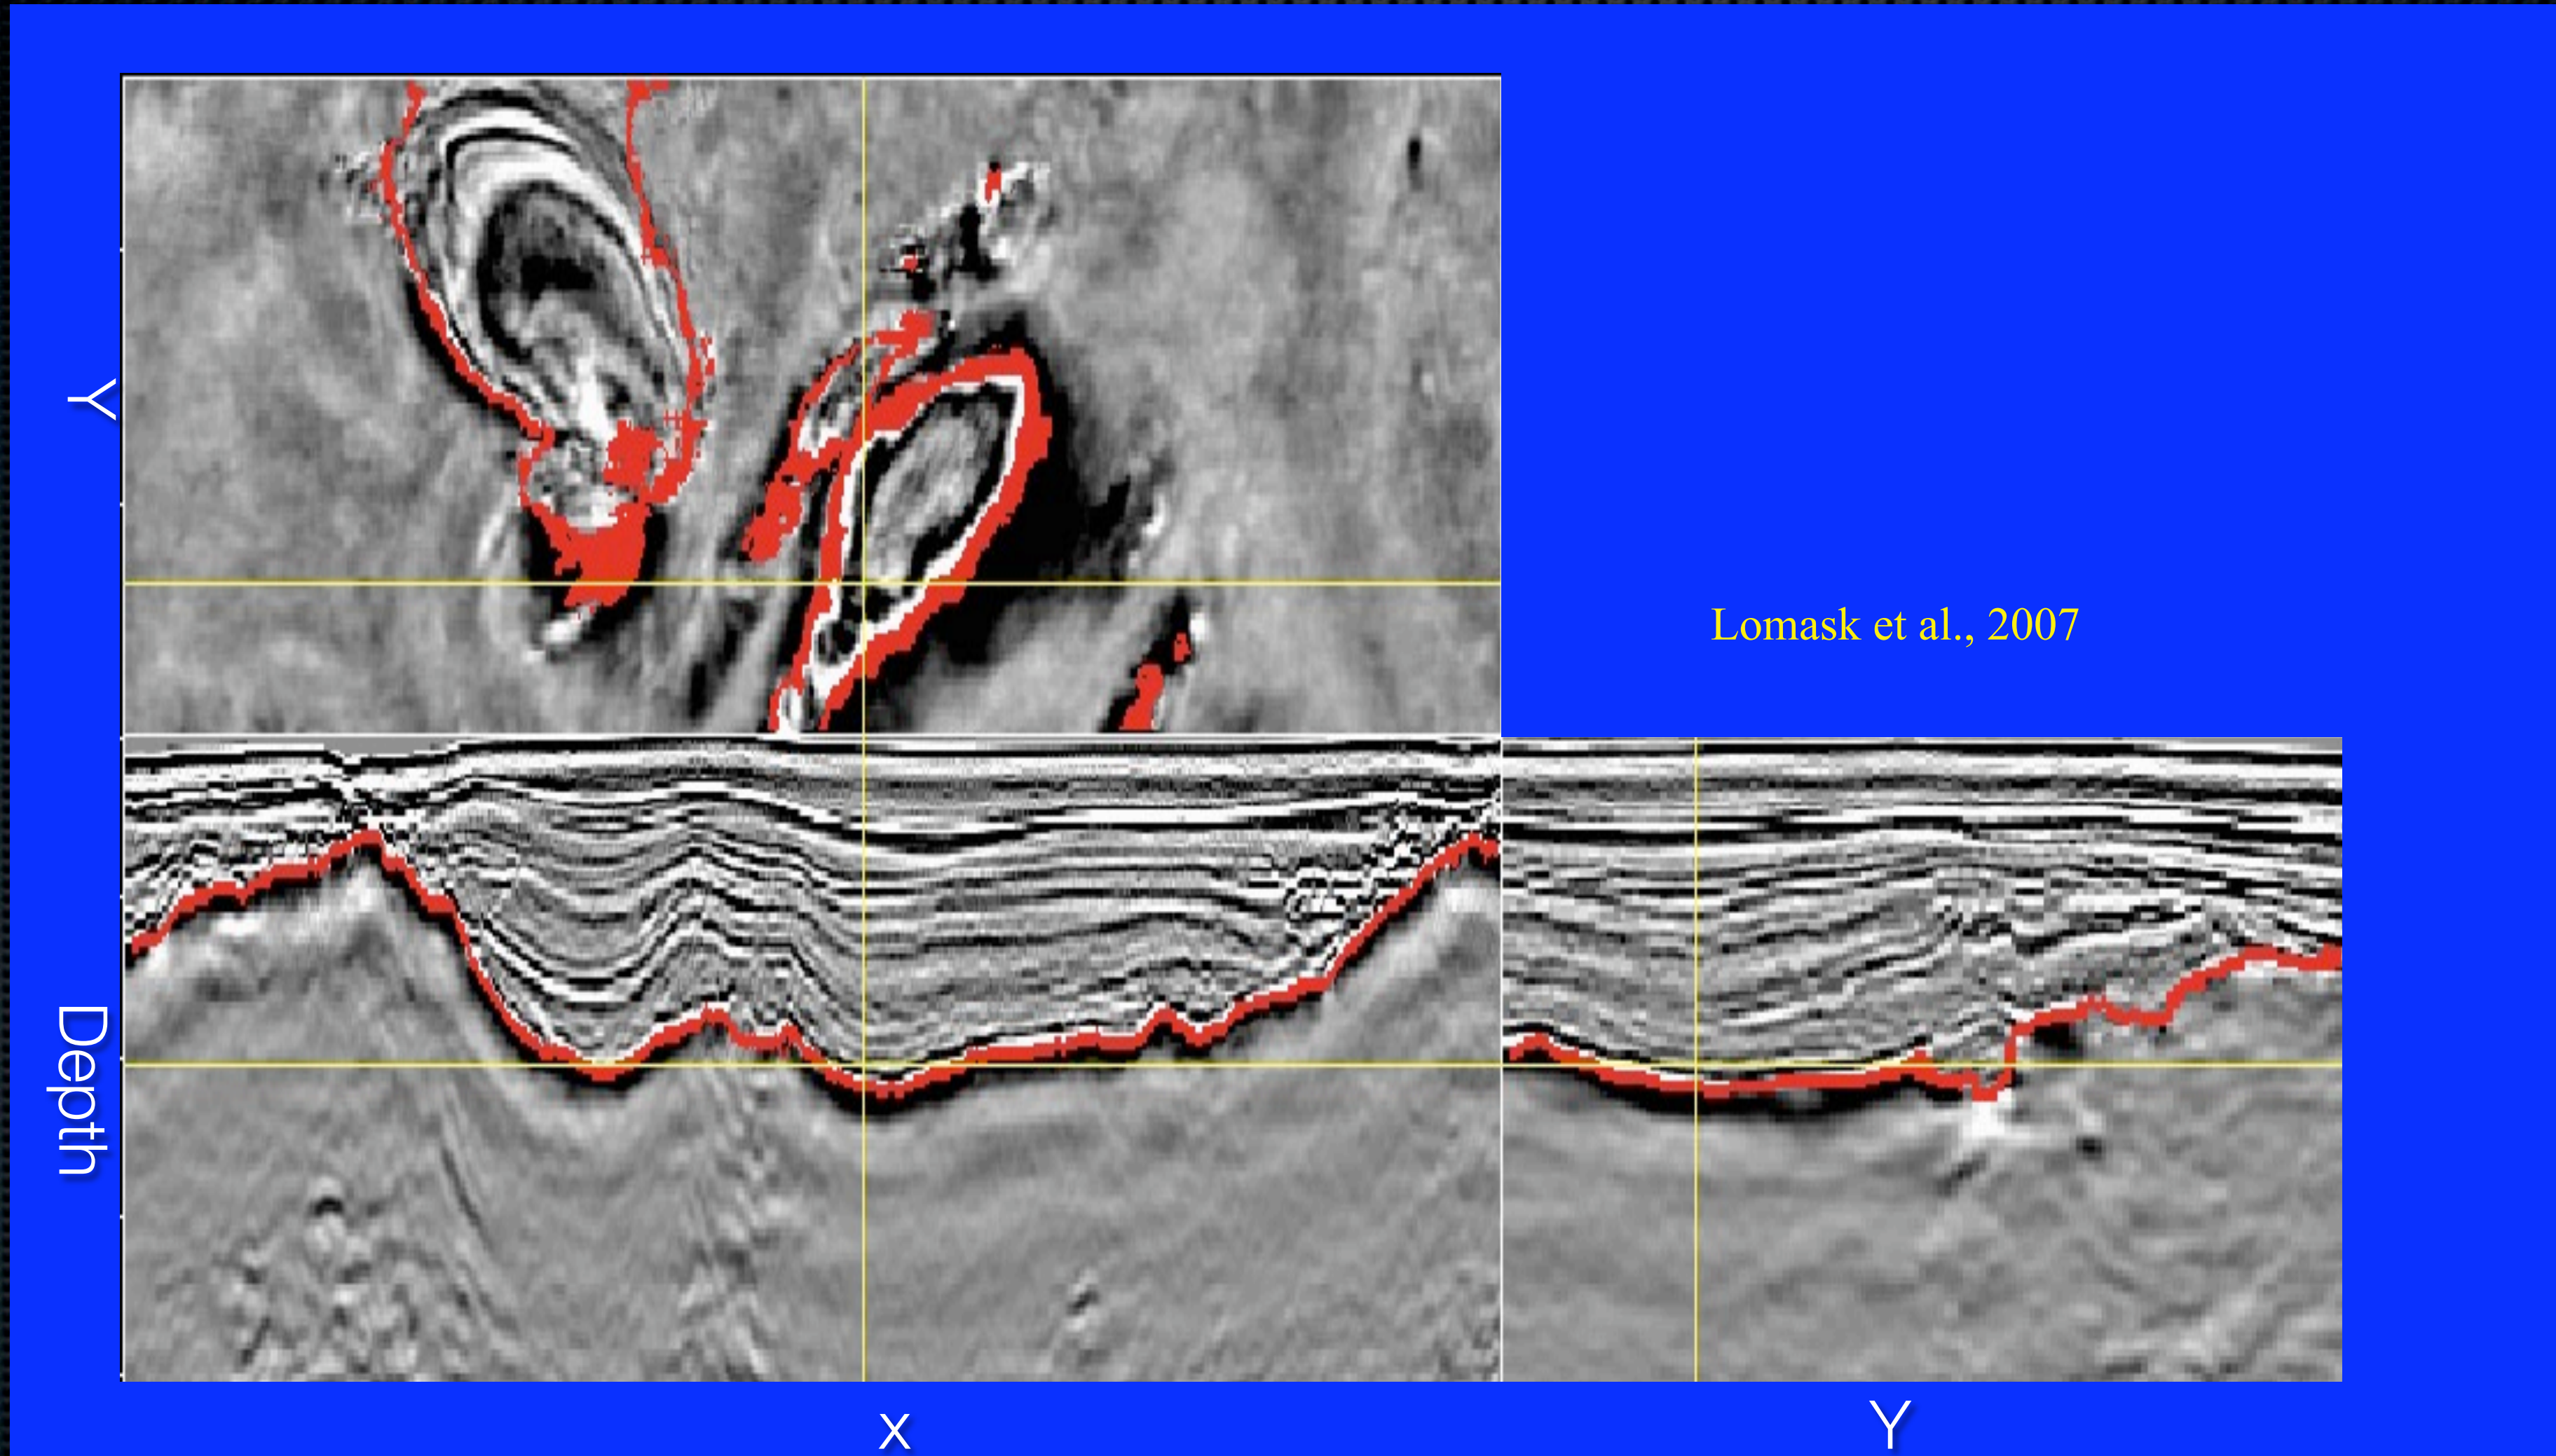
Motivation

Unique hardware

Algorithmic changes

Programming challenges

Demo



X
Bob Clapp

Y
SEP Meeting 2009



Everything is data transfer bound at some level



- Very small (16k)
- Latency: 1 nano second
- Bandwidth: 10s GB/s

Outline

Motivation

Unique hardware

Algorithmic changes


Programming challenges

Demo



Everything is data transfer bound at some level

- L1 Cache



- ✦ Megabytes in size
- ✦ Latency: 5-12 nanoseconds
- ✦ Bandwidth: 10s GB/s

Outline

Motivation

Unique hardware

Algorithmic changes

Programming challenges

Demo



Everything is data transfer bound at some level

L1 Cache

- L2/L3 Cache

- ✦ 10s of GBs in size
- ✦ Latency: 10s nanoseconds
- ✦ Bandwidth: 10s GB/s



Main memory

Outline

Motivation

Unique hardware

Algorithmic changes

Programming challenges

Demo




Everything is data transfer bound at some level

L1 Cache

L2/L3 Cache



Main memory



Network

- 10s of *fast* nodes
- Latency: 100,000s nanoseconds
- Bandwidth: 100s MB/s

Outline

Motivation

Unique hardware

Algorithmic changes

Programming challenges

Demo




Everything is data transfer bound at some level

L1 Cache

L2/L3 Cache



Main memory



Disk

- ✦ 100s-1000s GBs
- ✦ Latency: 3-10 million nanoseconds
- ✦ Total bandwidth .01 GB/s

Outline

Motivation

Unique hardware

Algorithmic changes

Programming challenges

Demo



Bob Clapp

SEP Meeting 2009

Everything is data transfer bound at some level

2-D acoustic modeling

Outline


Motivation

Unique hardware

Algorithmic changes

Programming challenges

Demo



L1 Cache

L2/L3 Cache

Main memory

Network

Disk



Everything is data transfer bound at some level

Outline


Motivation

Unique hardware

Algorithmic changes


Programming challenges

Demo



L1 Cache


L2/L3 Cache



Main memory



Network



Disk

2-D acoustic modeling

3-D acoustic modeling



Everything is data transfer bound at some level

Outline


Motivation

Unique hardware

Algorithmic changes

Programming challenges

Demo



L1 Cache

L2/L3 Cache



Main memory



Network



Disk

2-D acoustic modeling

3-D acoustic modeling

Sparse matrix multiply



Everything is data transfer bound at some level

Outline


Motivation

Unique hardware

Algorithmic changes

Programming challenges

Demo



L1 Cache


L2/L3 Cache



Main memory



Network



Disk

2-D acoustic modeling

3-D acoustic modeling

Sparse matrix multiply

3-D RTM



Solutions????

Outline


Motivation

Unique hardware

Algorithmic changes

Programming challenges

Demo



Niagara 2



Sicortex



GPU



Convey



Ecoram



FPGA



Bob Clapp

SEP Meeting 2009

Niagara 2

Outline

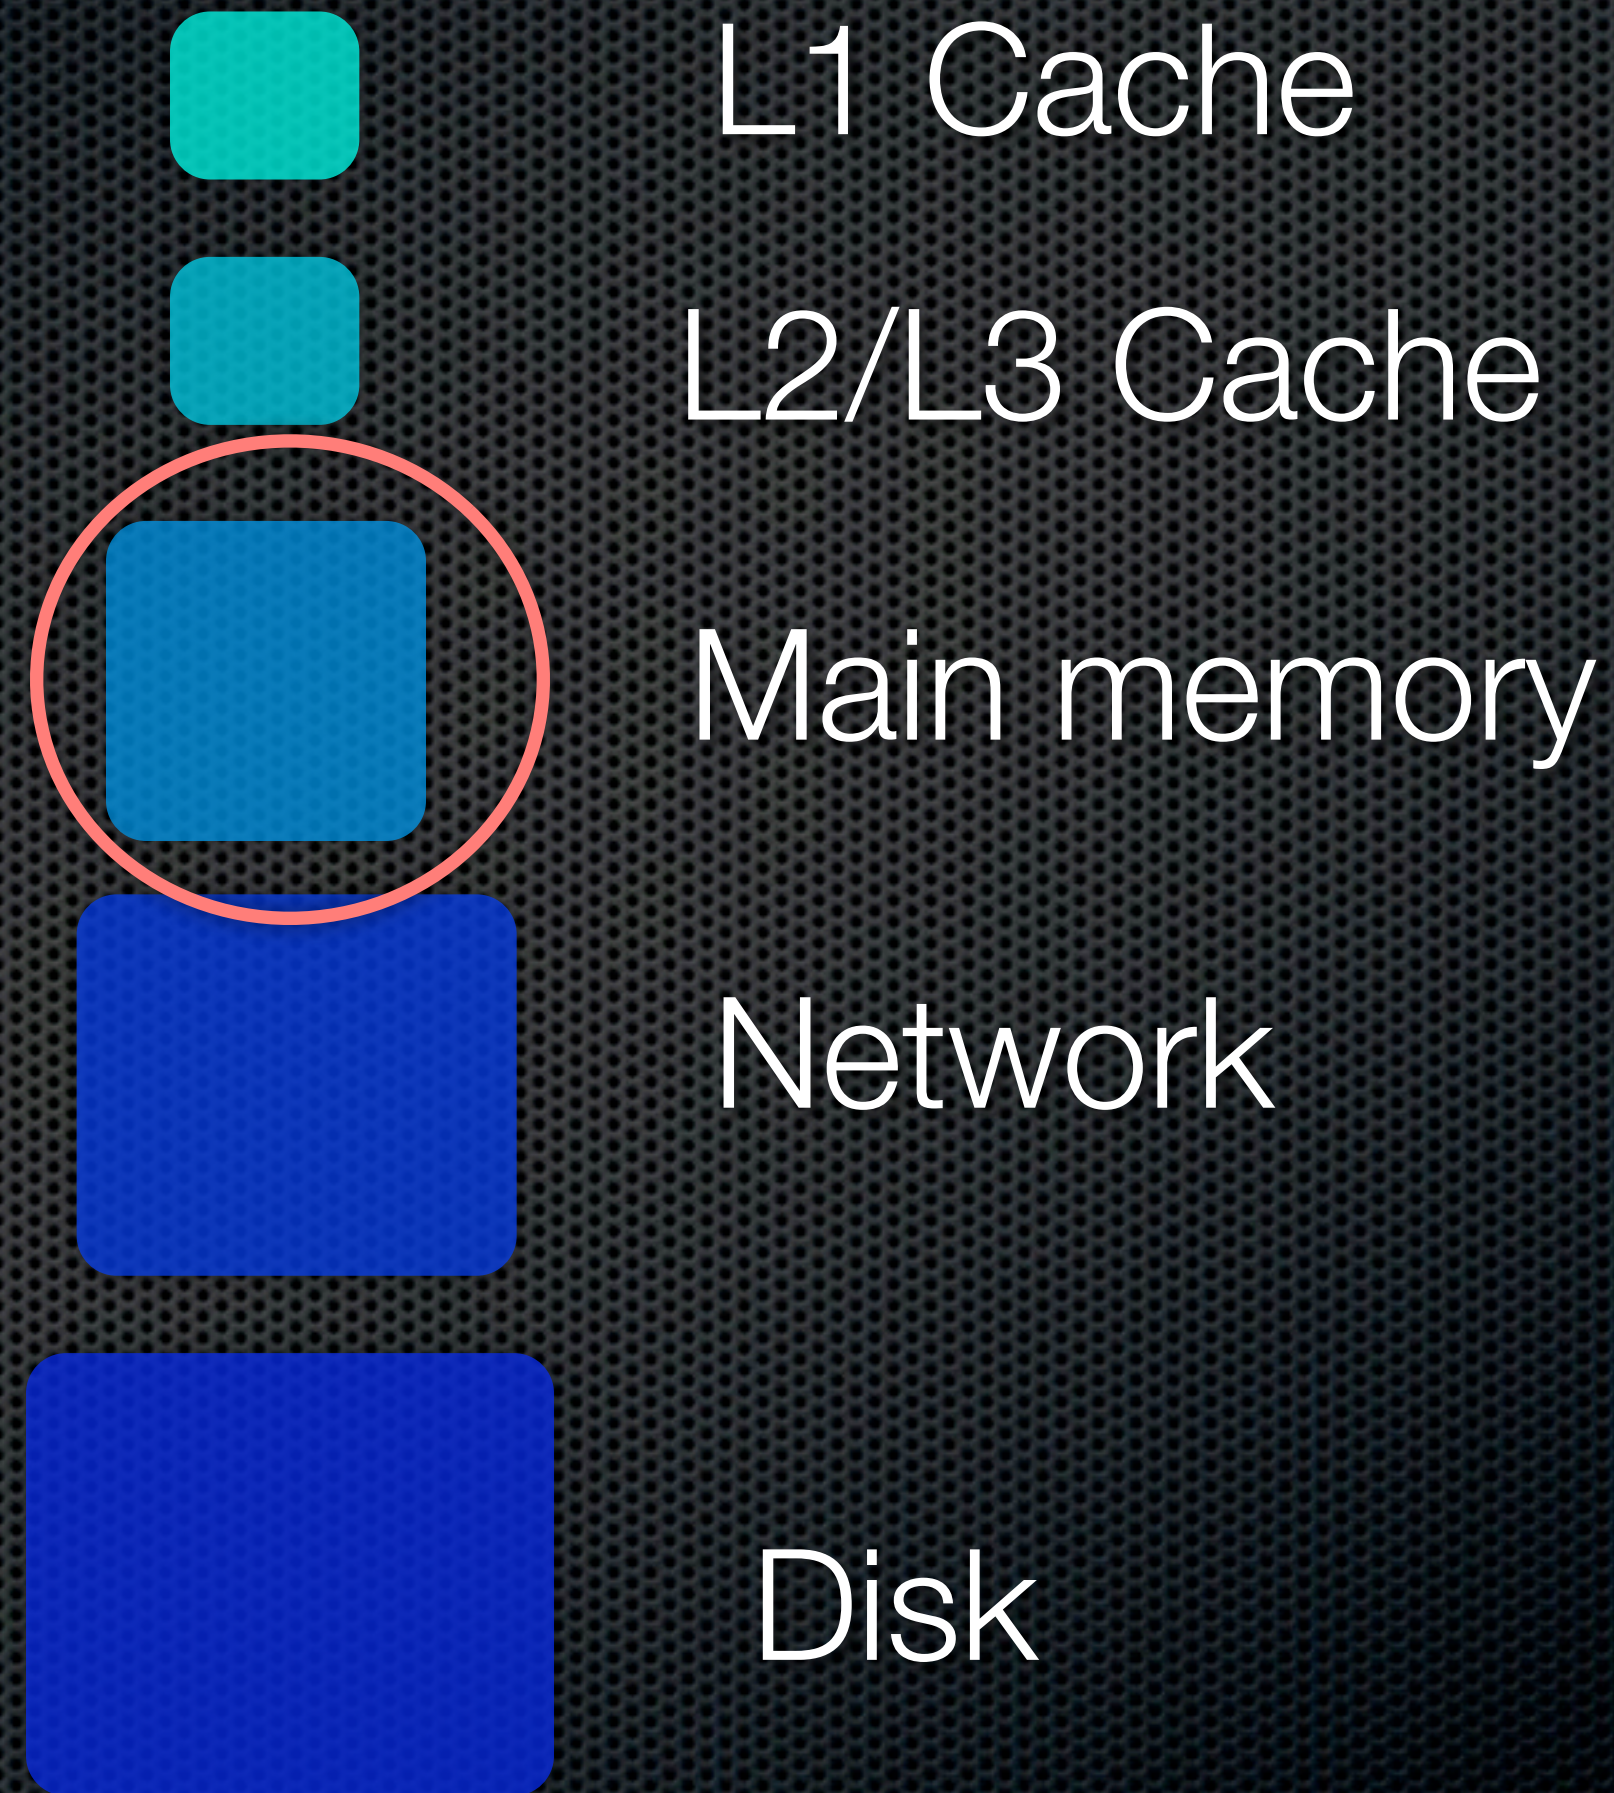

Motivation

Unique hardware

Algorithmic changes



Programming challenges

Demo



Niagara 2

Core



Outline

Motivation

Unique hardware

Algorithmic changes

Programming challenges

Demo





Bob Clapp

SEP Meeting 2009

Niagara 2

Core



Outline

Motivation

Unique hardware

Algorithmic changes

Programming challenges

Demo





Bob Clapp

SEP Meeting 2009

Niagara 2

Core



Outline

Motivation

Unique hardware

Algorithmic changes

Programming challenges

Demo





Bob Clapp

SEP Meeting 2009

Niagara 2

Core



Waiting on main
memory

Ready to compute

Outline

Motivation

Unique hardware

Algorithmic changes

Programming challenges

Demo





Bob Clapp

SEP Meeting 2009

Niagara 2

Core



Waiting on main
memory

Ready to compute

Outline

Motivation

**Unique
hardware**

Algorithmic
changes

Programming
challenges

Demo





Bob Clapp

SEP Meeting 2009

Niagara 2

Core



Waiting on main
memory

Ready to compute

Outline

Motivation

**Unique
hardware**

Algorithmic
changes

Programming
challenges

Demo





Bob Clapp

SEP Meeting 2009

Niagara 2

Core



Waiting on main
memory

Ready to compute

Outline

Motivation

**Unique
hardware**

Algorithmic
changes

Programming
challenges

Demo





Bob Clapp

SEP Meeting 2009

Niagara 2

Core



Waiting on main
memory

Ready to compute

Outline

Motivation

**Unique
hardware**

Algorithmic
changes

Programming
challenges

Demo



Bob Clapp

SEP Meeting 2009

Niagara 2

Outline



Motivation

Unique hardware

Algorithmic changes

Programming challenges

Demo



- ✦ Latency and bandwidth to main memory is the problem
- ✦ Multiple threads per core to hide memory latency
- ✦ Large pipe to memory



Niagara 2

Outline

Motivation



Unique hardware

Algorithmic changes

Programming challenges

Demo

- ✦ Shot profile migration
- ✦ 3-D modeling
- ✦ Sparse matrix inverse



- ✦ Latency and bandwidth to main memory is the problem
- ✦ Multiple threads per core to hide memory latency
- ✦ Large pipe to memory



Niagara 2

Outline

Motivation



Unique hardware

Algorithmic changes

Programming challenges

Demo


- ✦ Shot profile migration
- ✦ 3-D modeling
- ✦ Sparse matrix inverse



Niagara 2 is actually FPU bound for many of our applications



Sicortex




L1 Cache



L2/L3 Cache



Main memory



Network



Disk

Outline

Motivation

Unique hardware


Algorithmic changes

Programming challenges

Demo




High Performance Interconnect

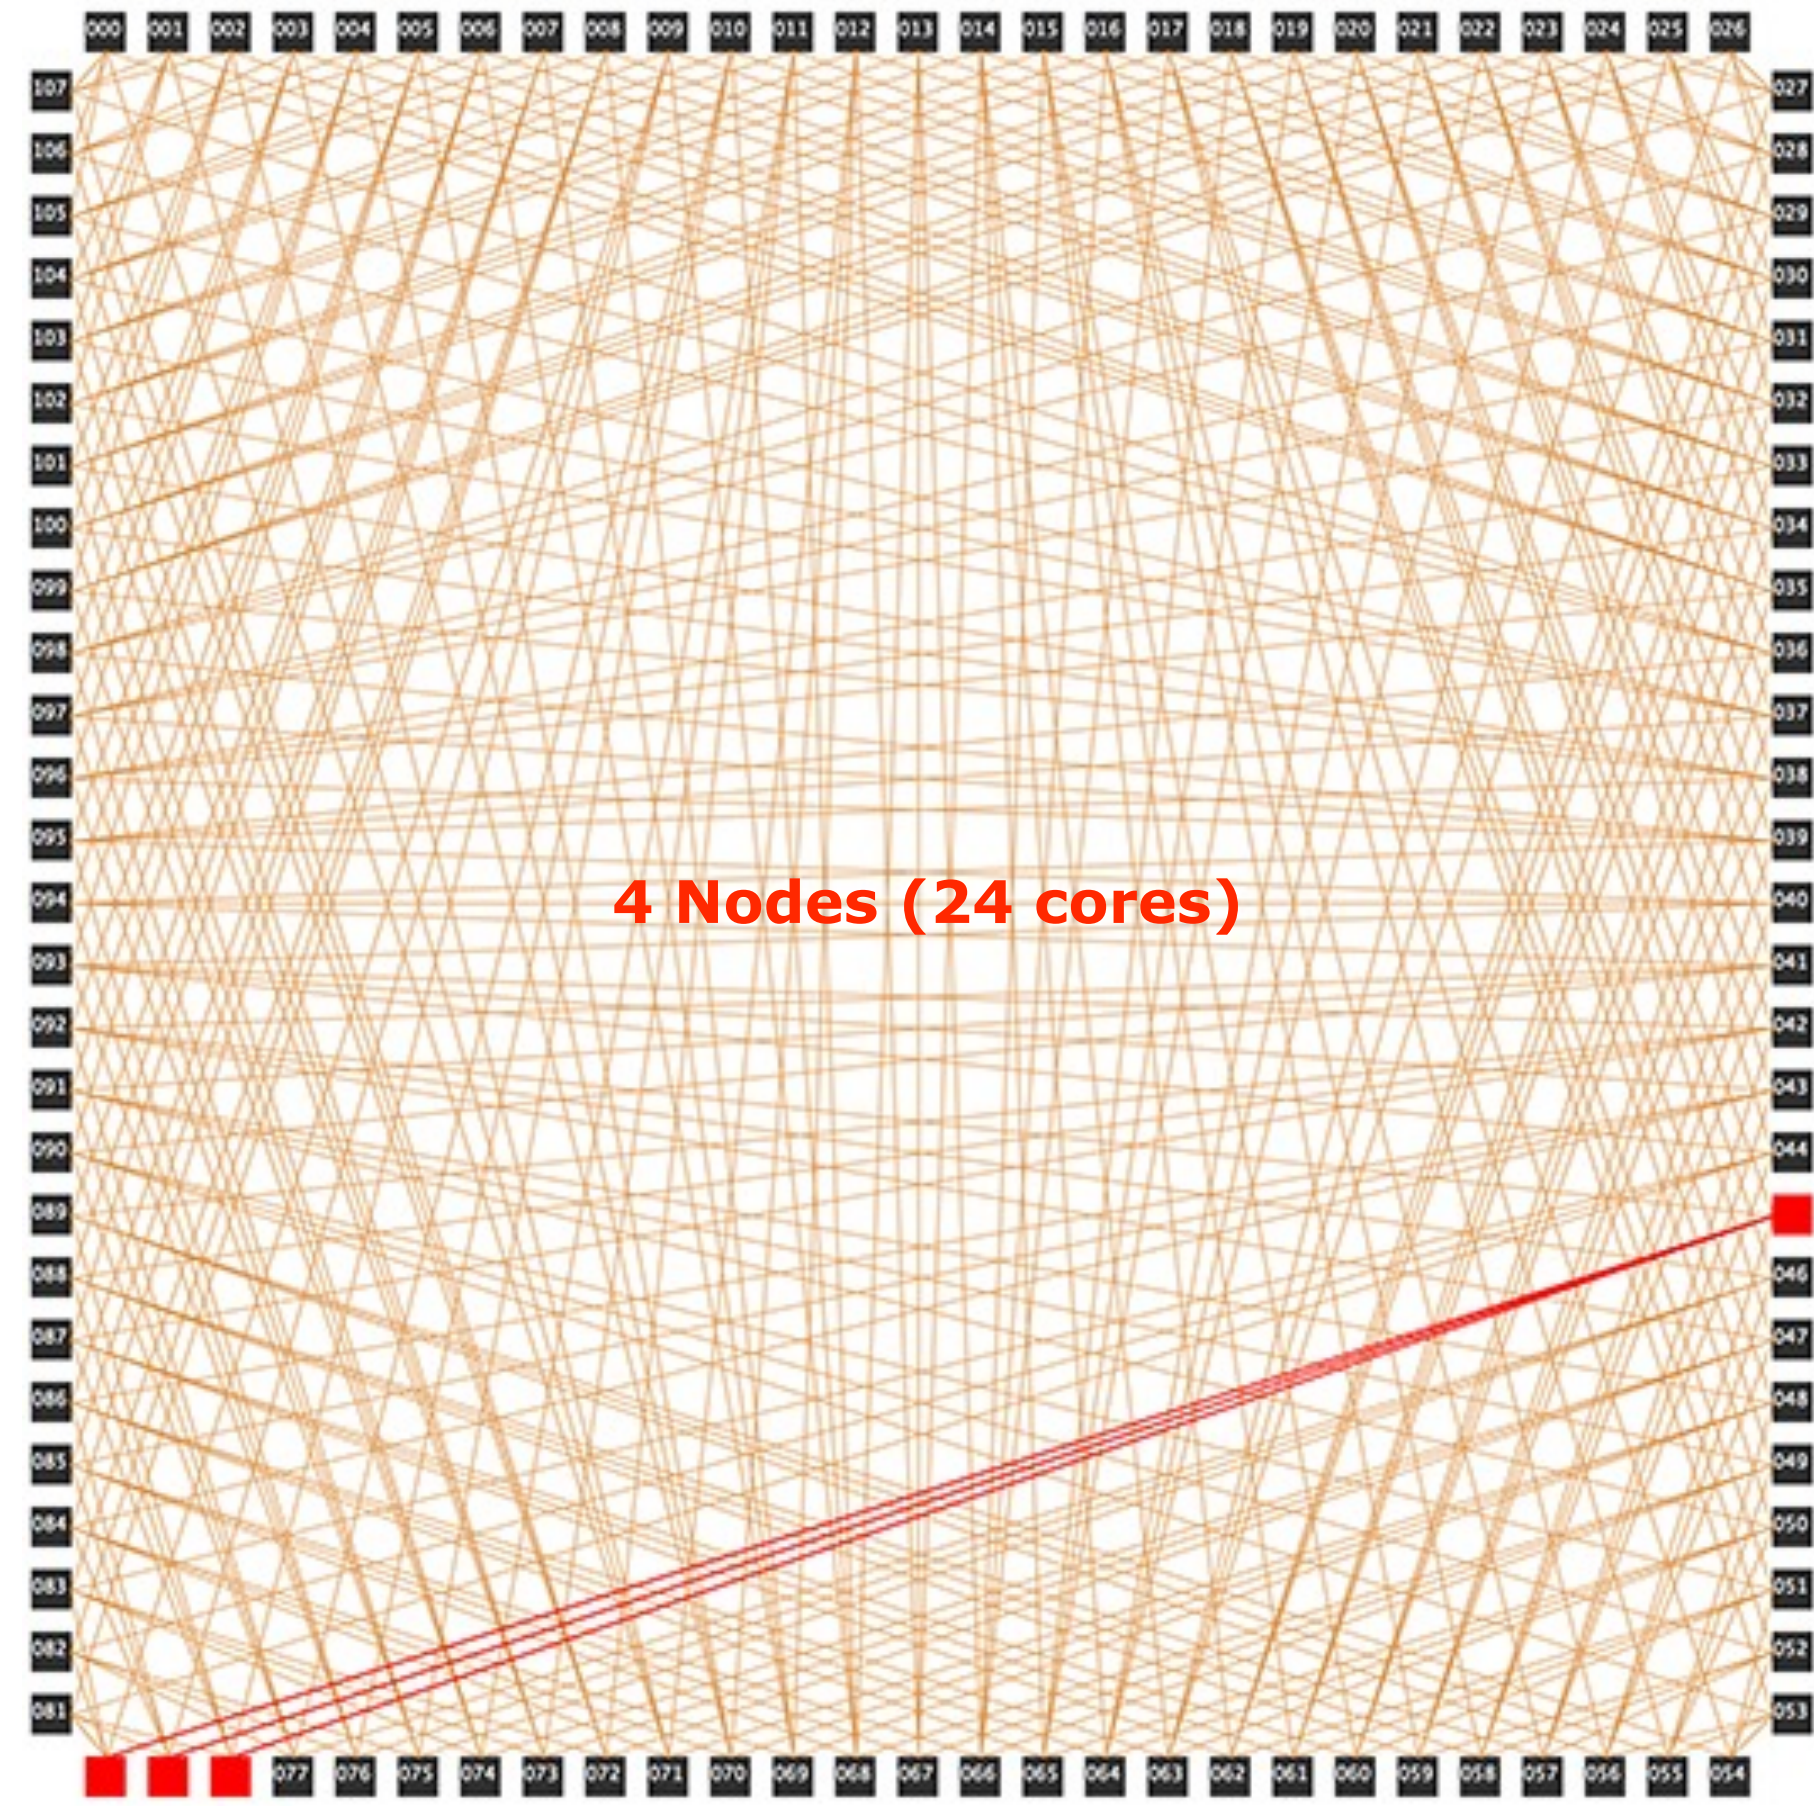


108 nodes (648 cores)


Kautz Graph Topology
Low Latency
High Throughput
Fault Tolerant



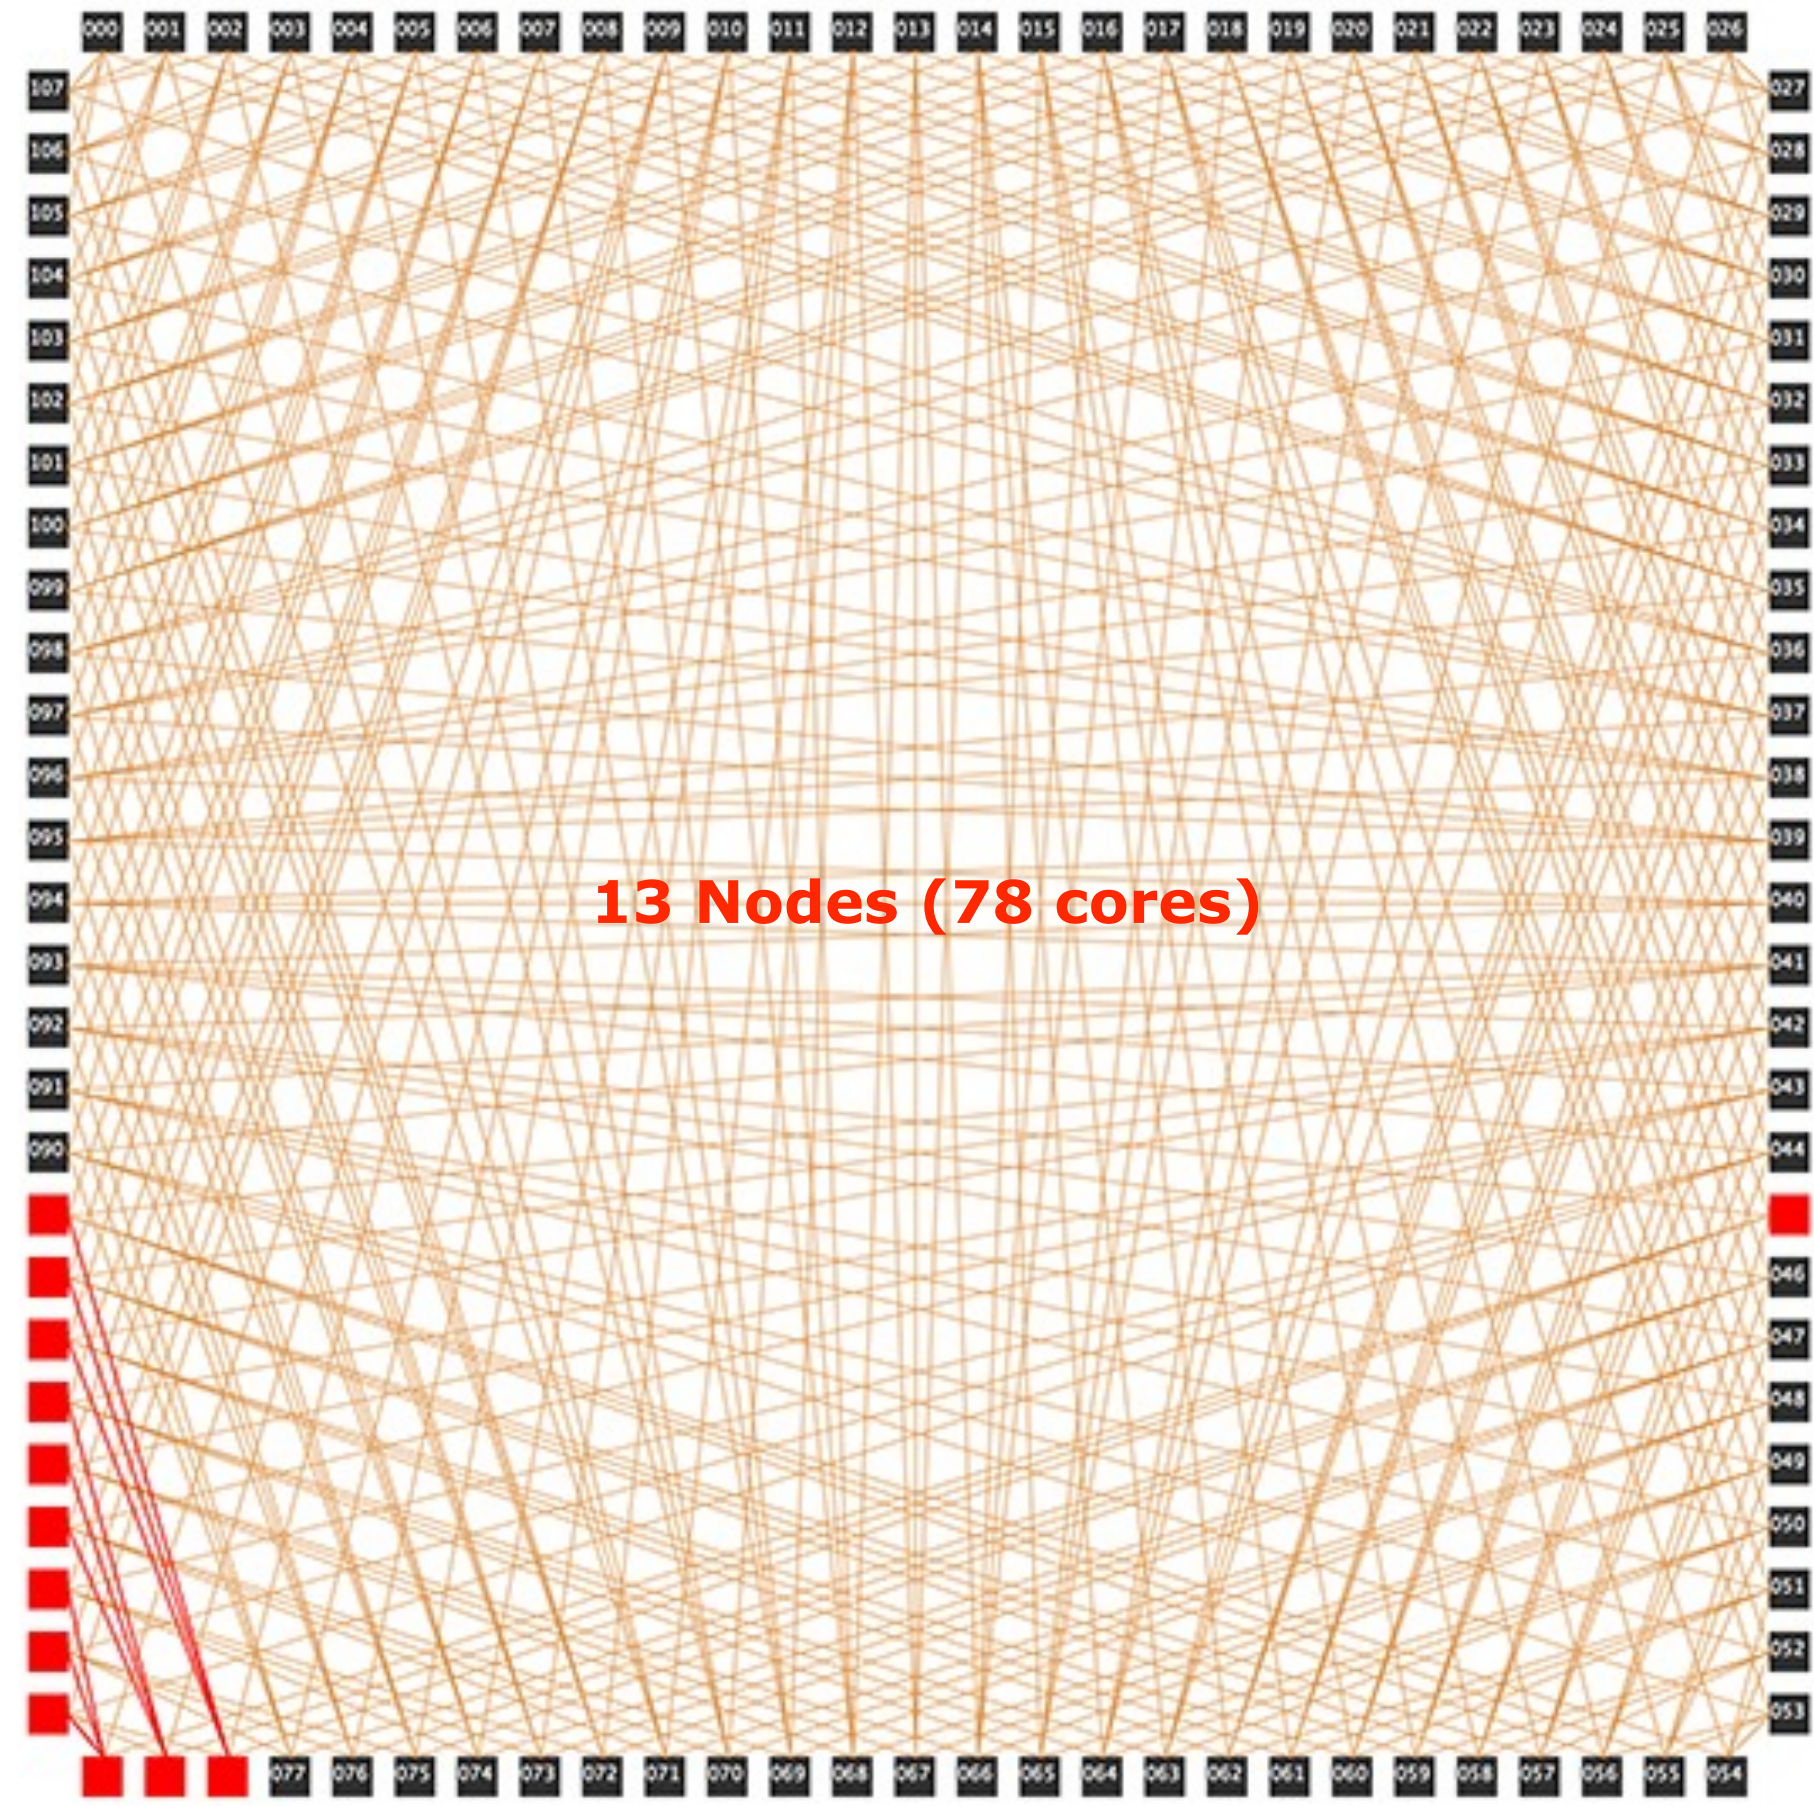
High Performance Interconnect




Kautz Graph Topology
Low Latency
High Throughput
Fault Tolerant



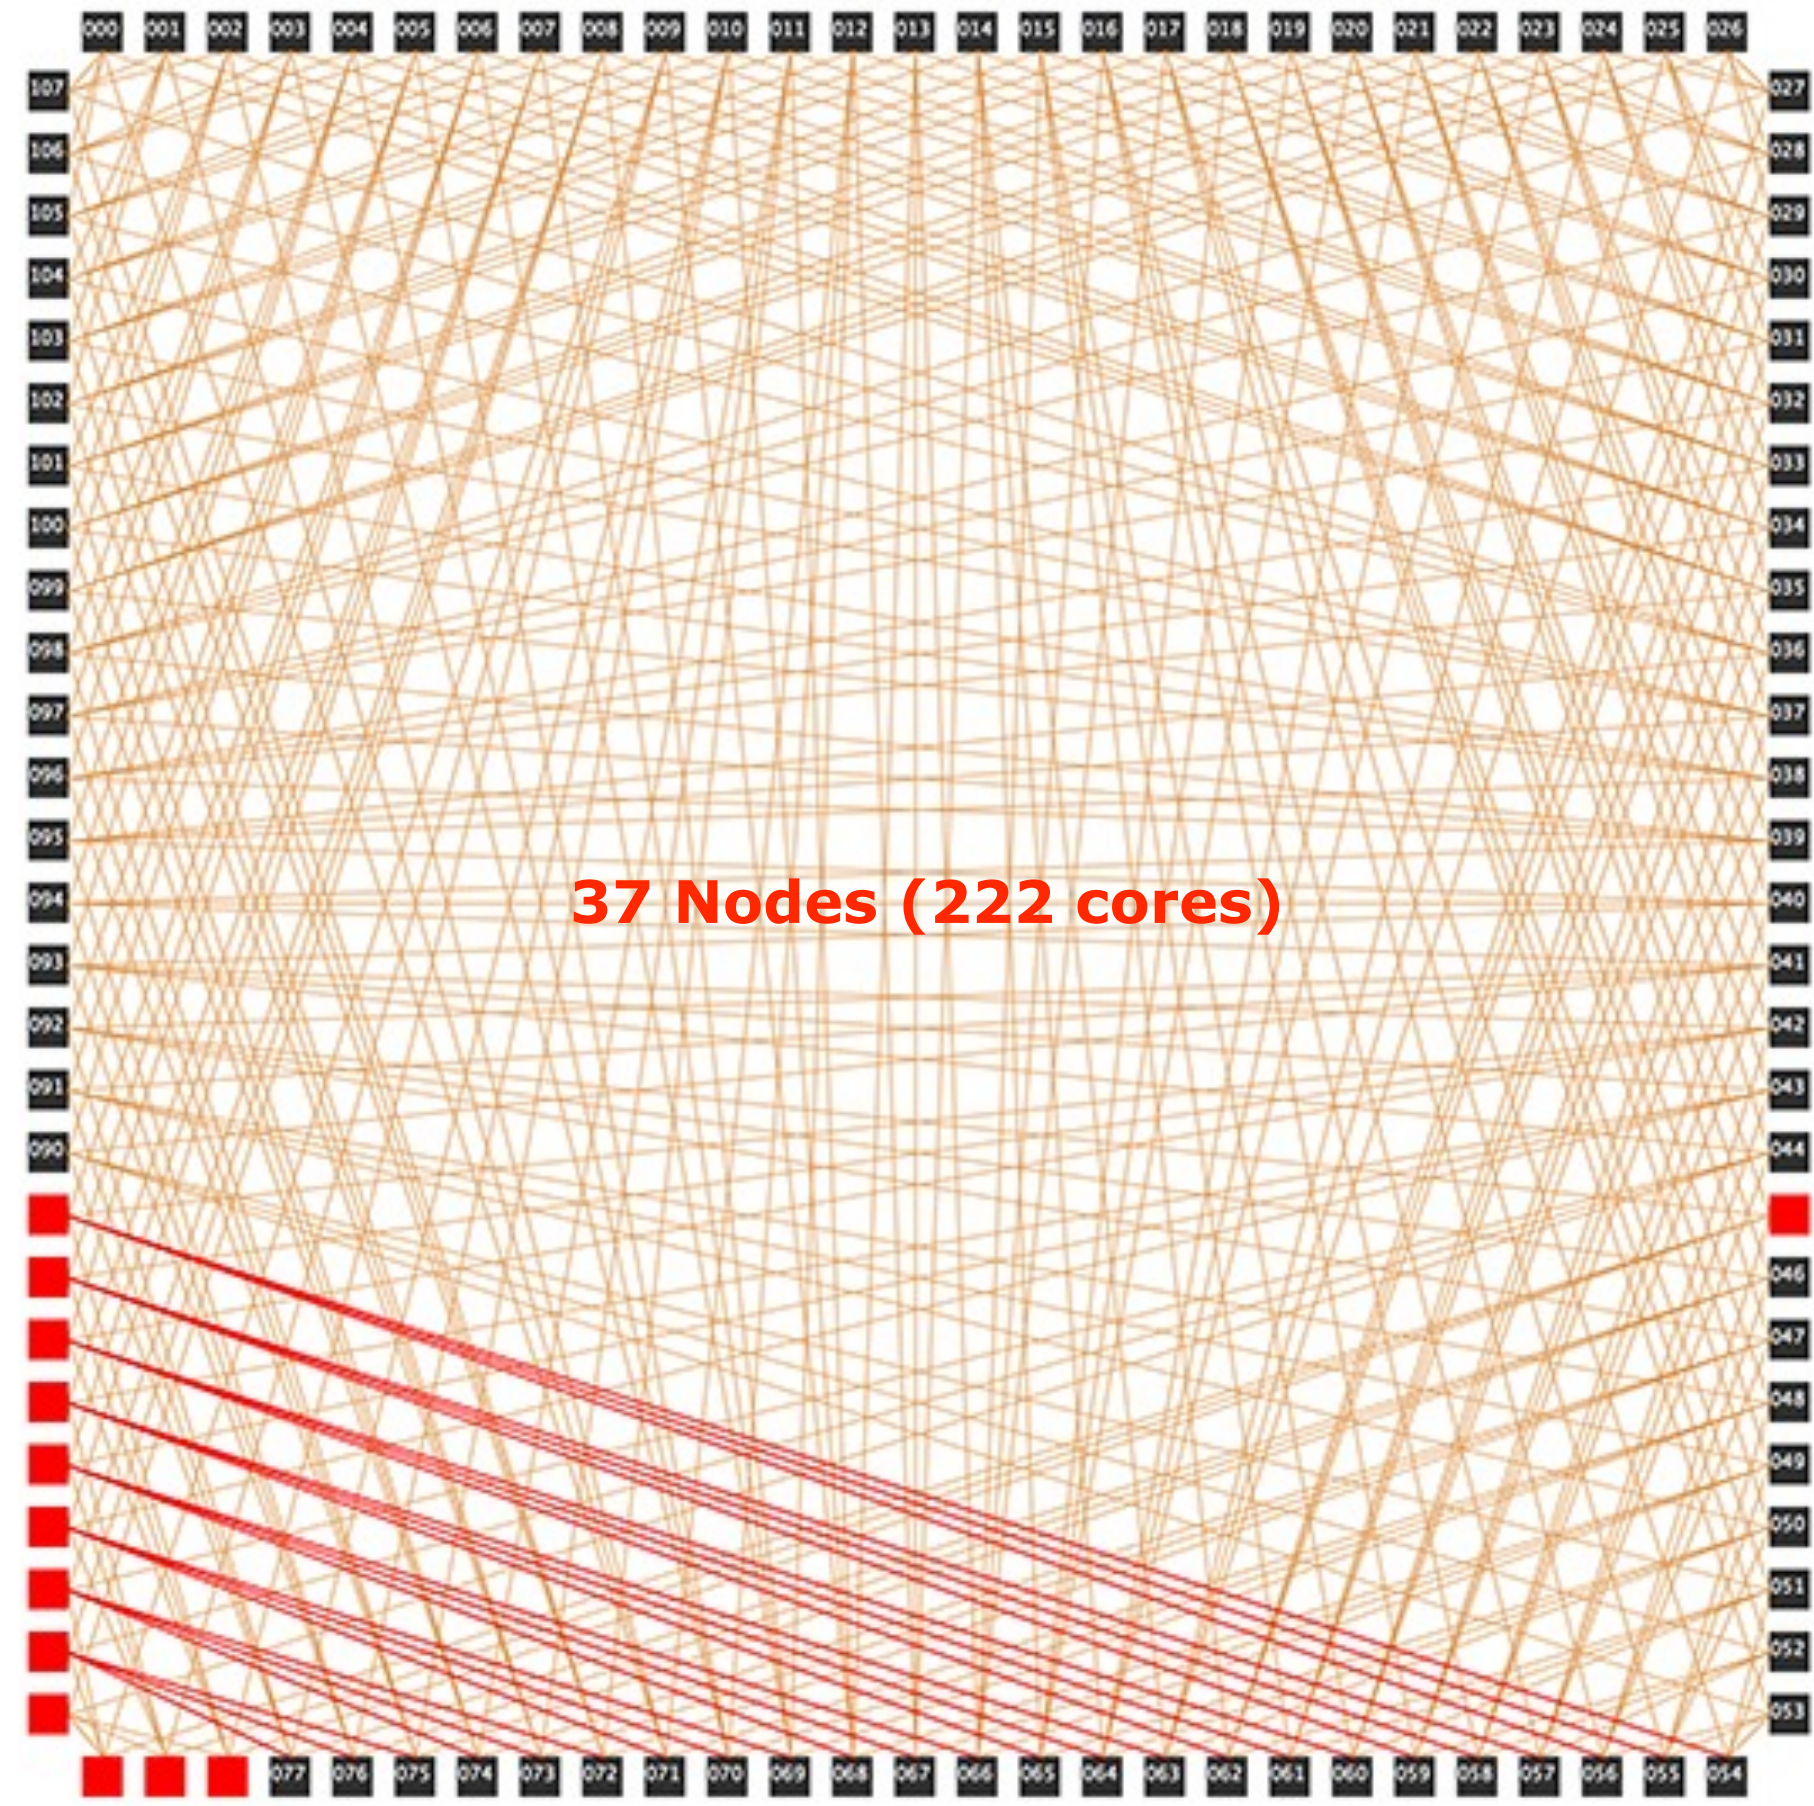
High Performance Interconnect




Kautz Graph Topology
Low Latency
High Throughput
Fault Tolerant




High Performance Interconnect




Kautz Graph Topology
Low Latency
High Throughput
Fault Tolerant




High Performance Interconnect



Kautz Graph Topology
Low Latency
High Throughput
Fault Tolerant




High Performance Interconnect

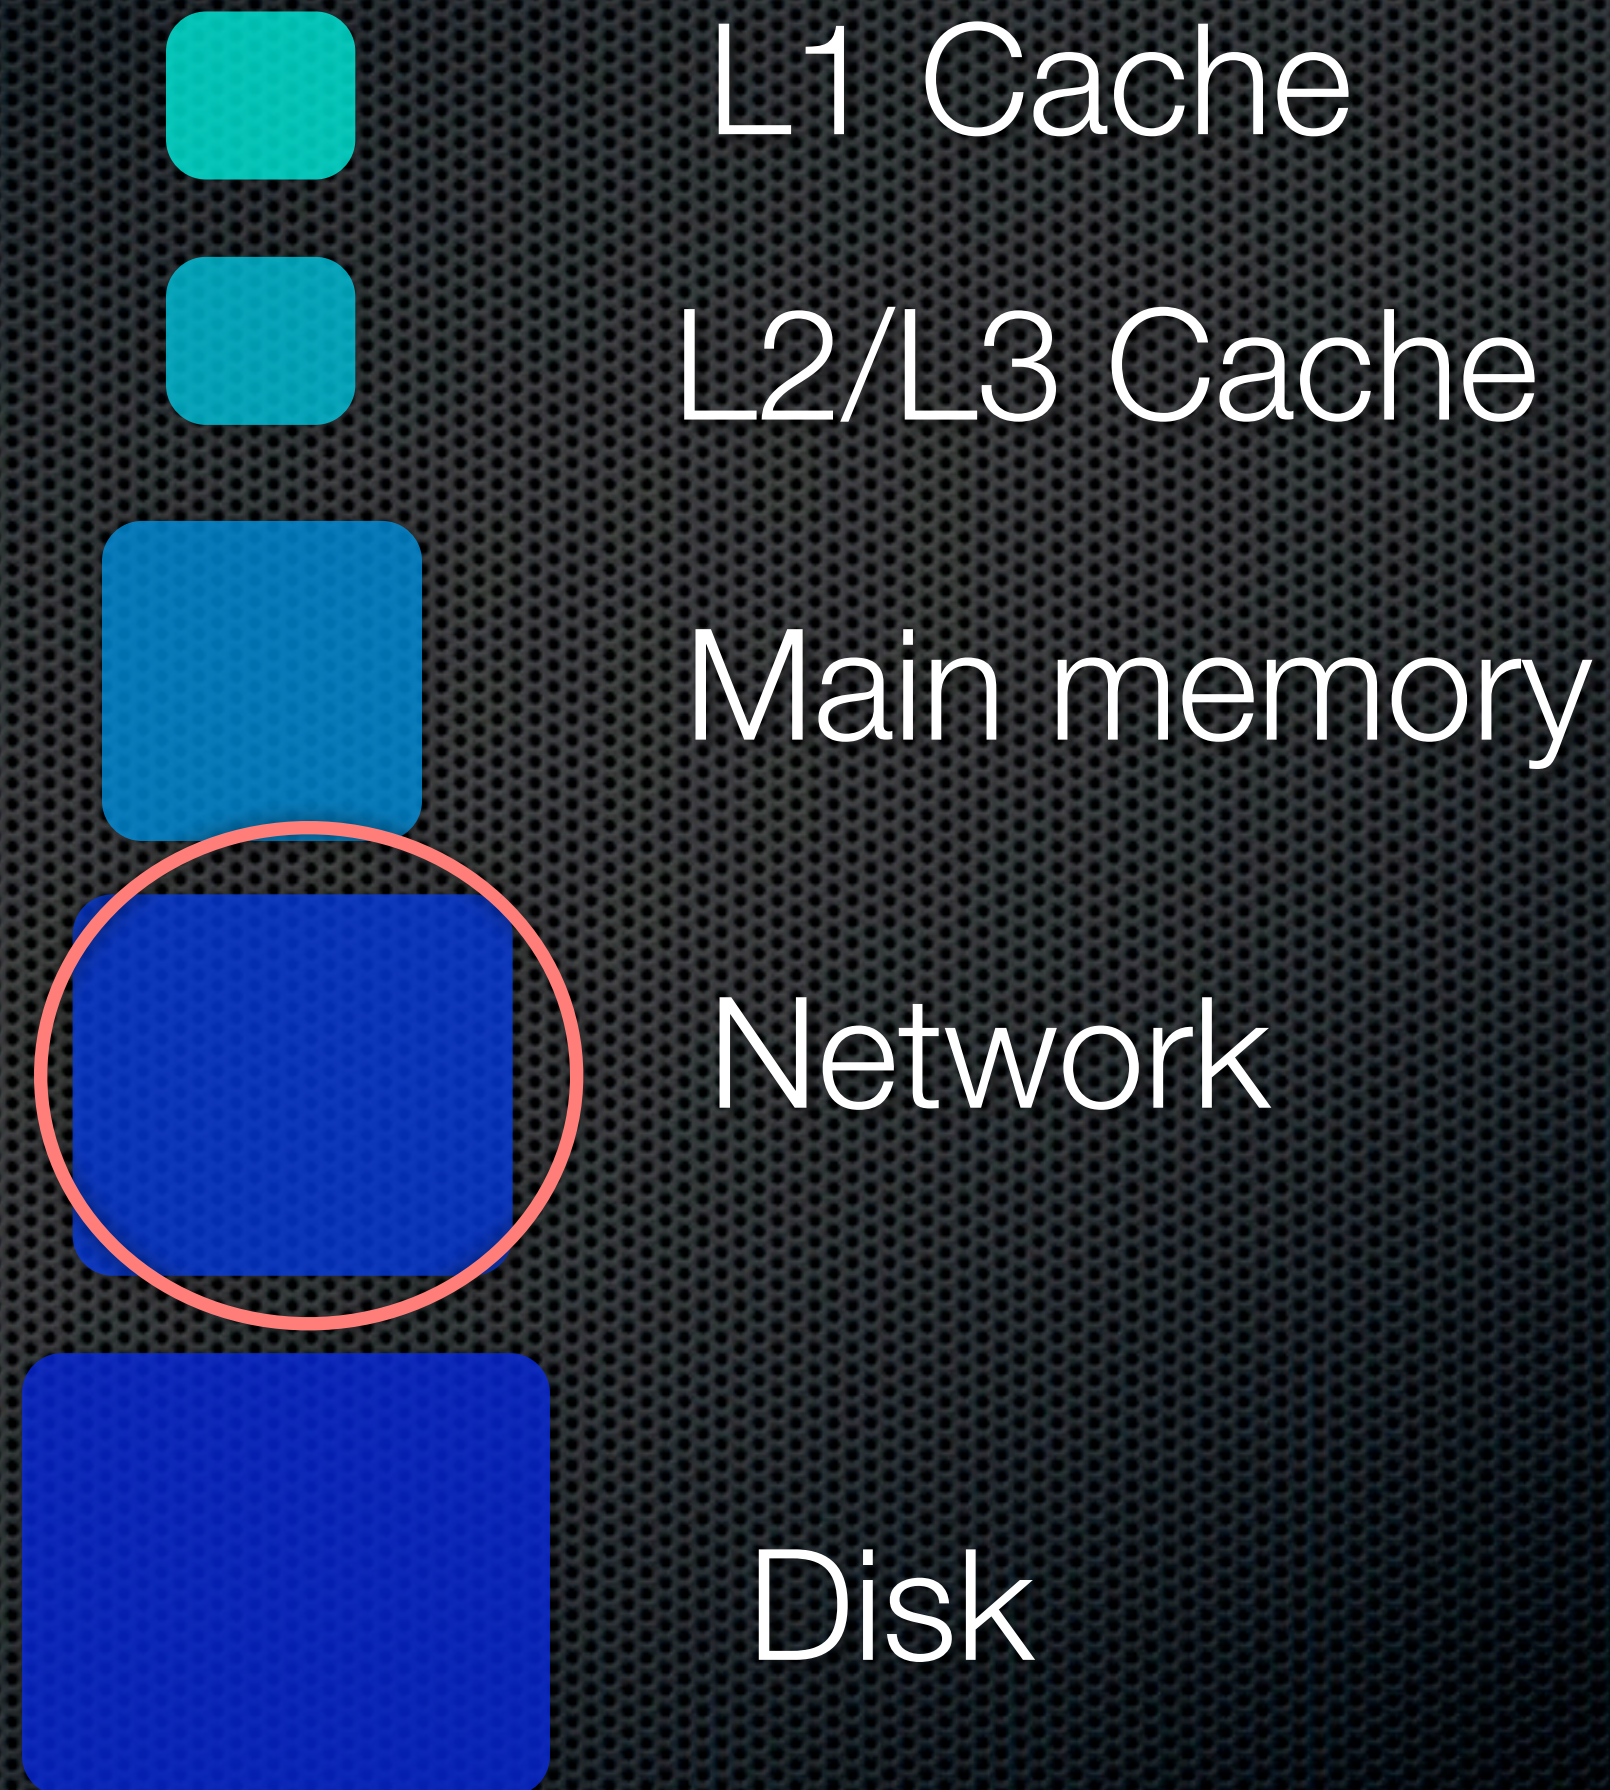


Kautz Graph Topology

Low Latency
High Throughput
Fault Tolerant



Sicortex



- Everything is communication bound
 - Latency
 - Bandwidth for large number of nodes
- Shrink problems to fit better into L1, L2 cache

Outline

Motivation

Unique hardware

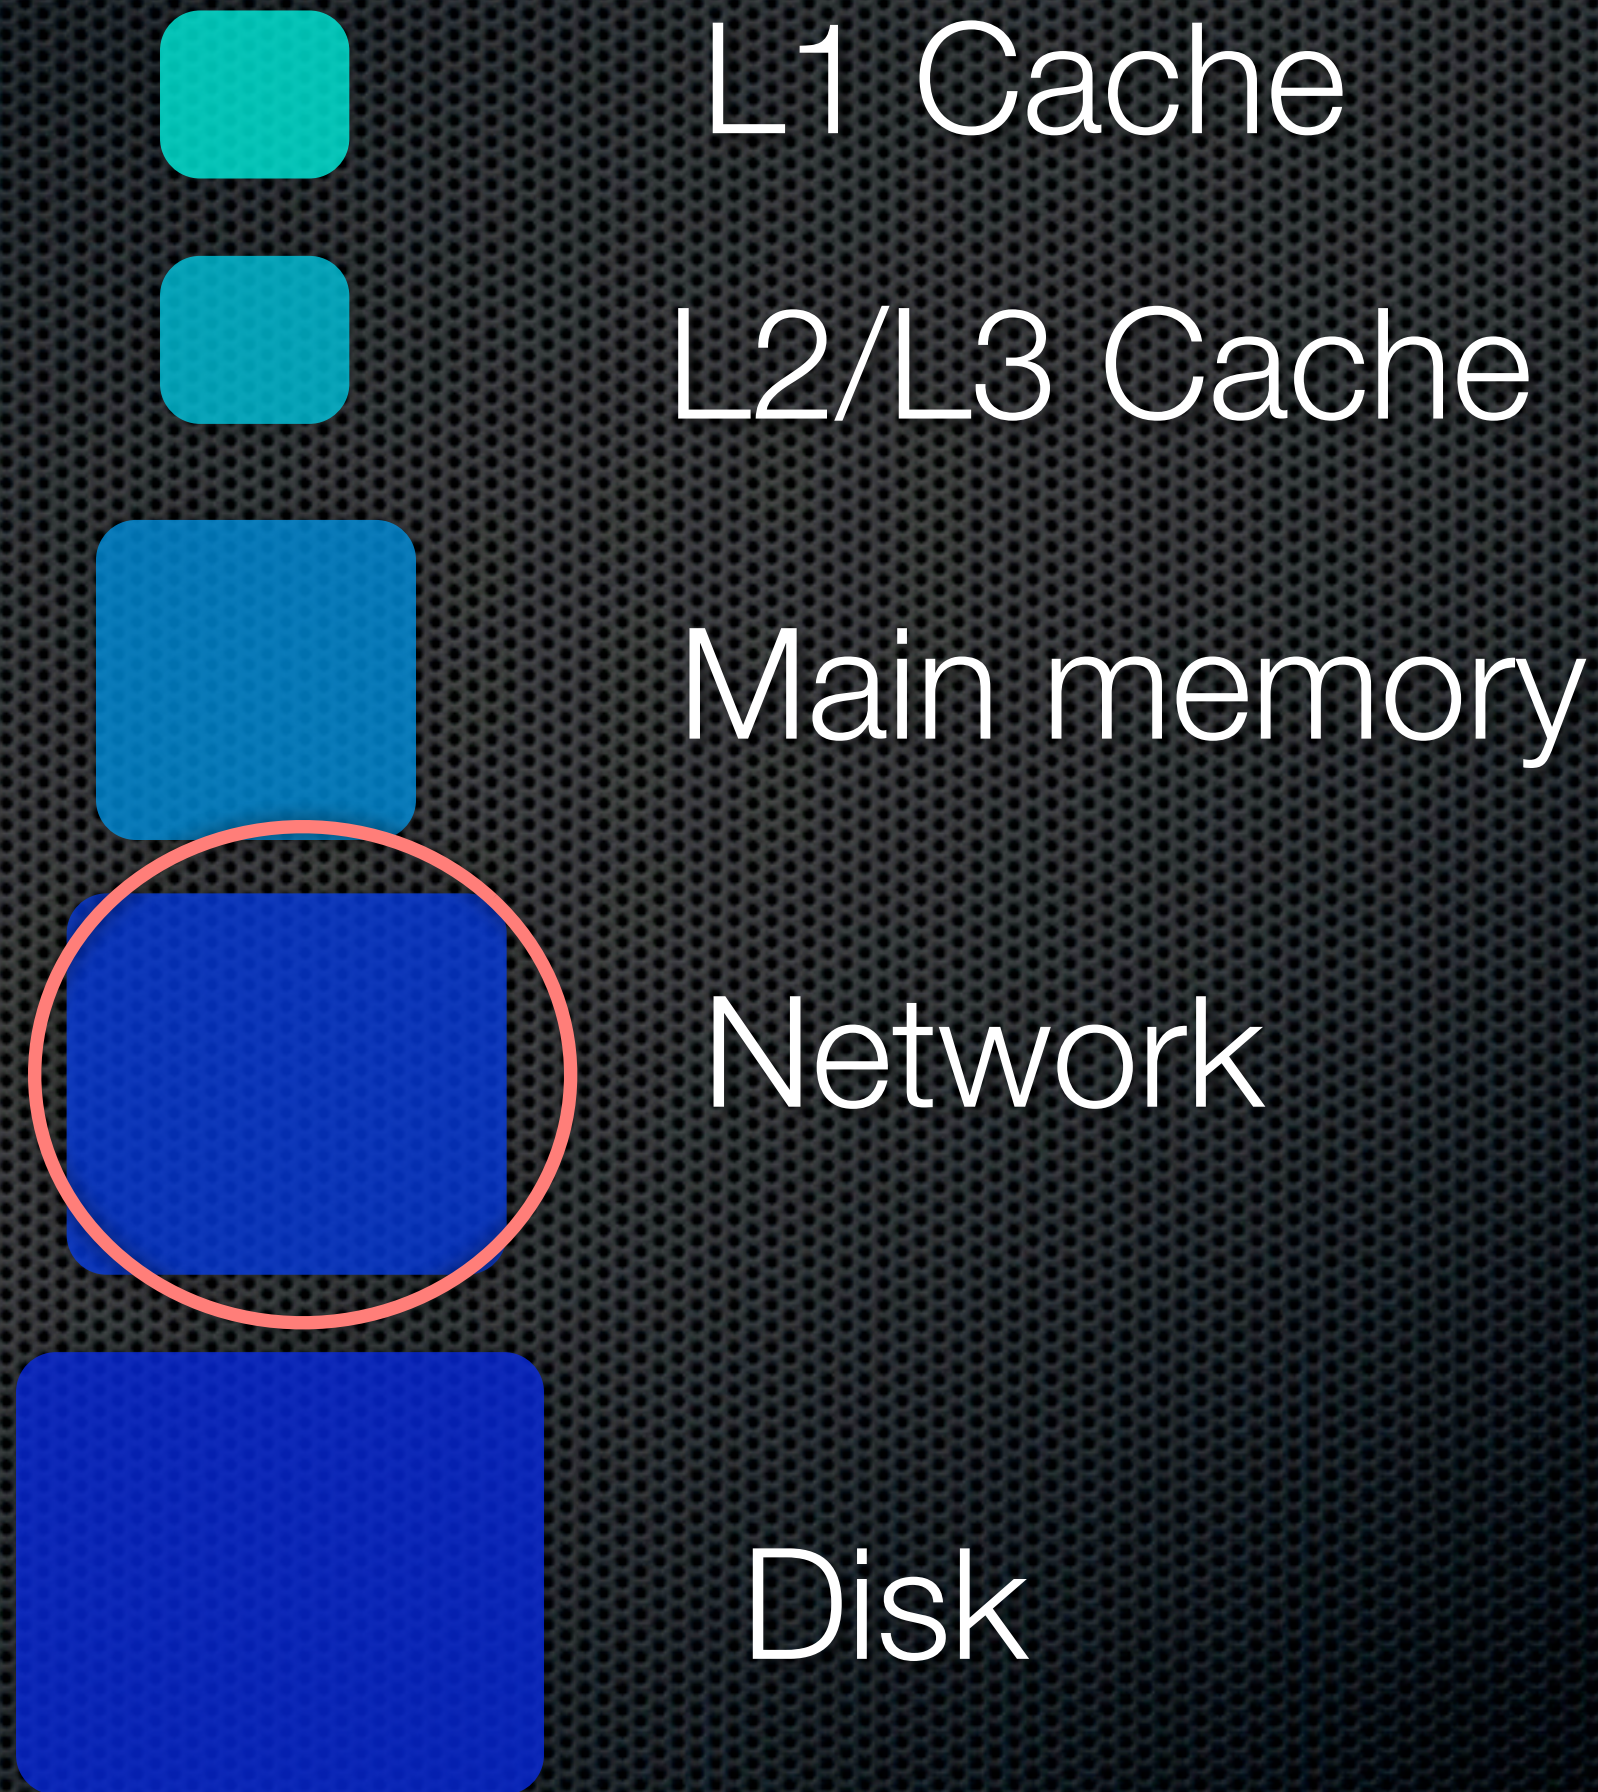
Algorithmic changes

Programming challenges

Demo



Sicortex



- ✦ Everything is communication bound
 - ✦ Latency
 - ✦ Bandwidth for large number of nodes
- ✦ Shrink problems to fit better into L1, L2 cache

Outline

Motivation

Unique hardware

Algorithmic changes

Programming challenges

Demo

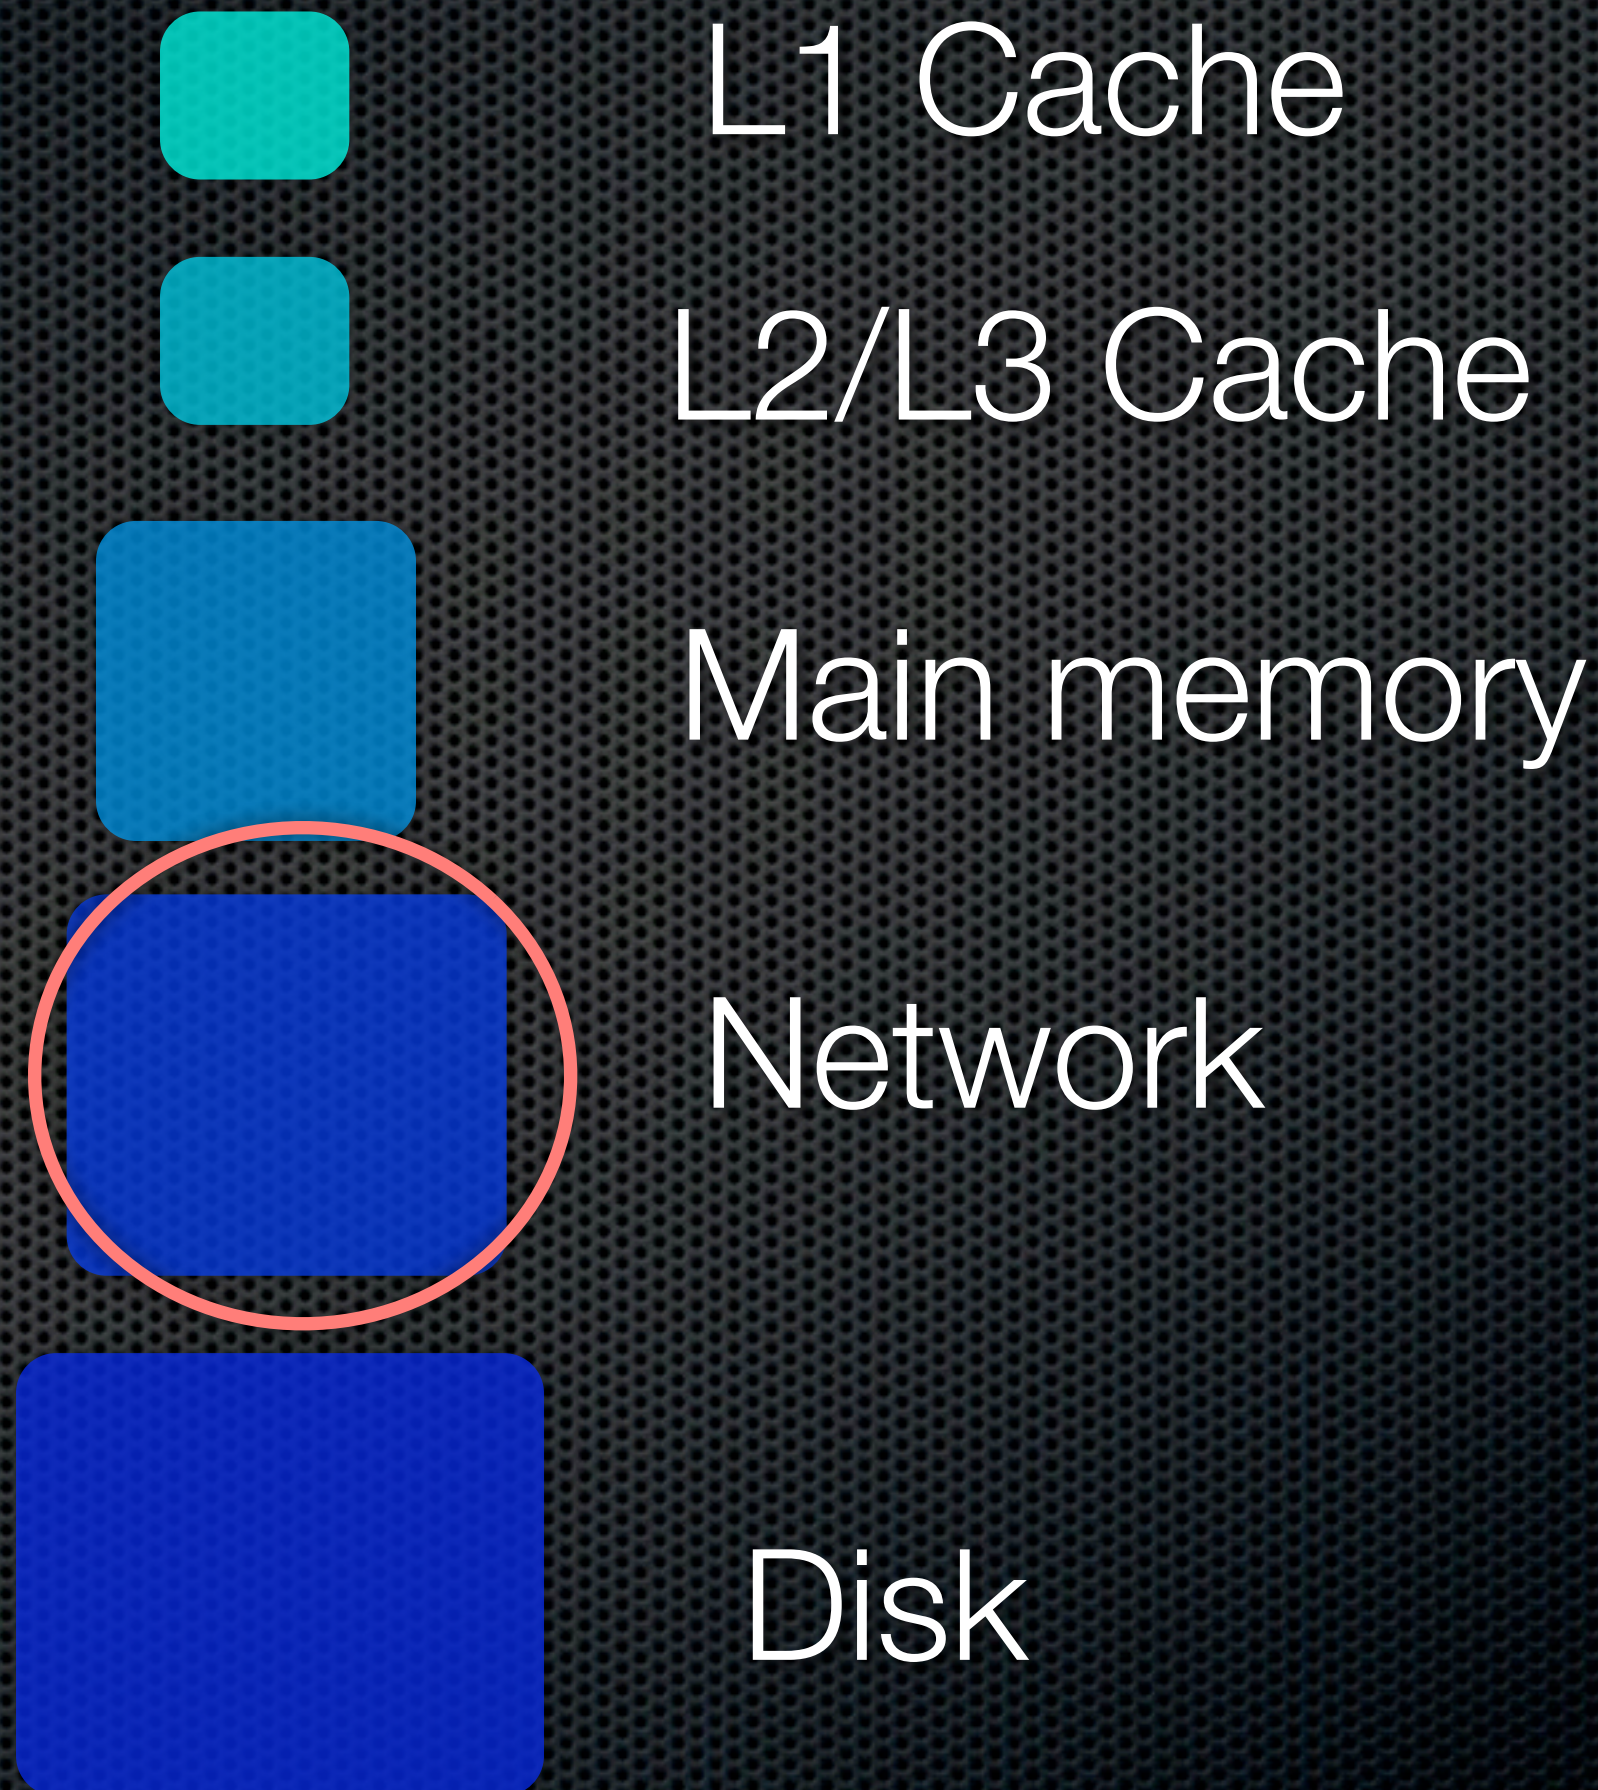
- ✦ Source-receiver migration (FFT)
- ✦ Acoustic modeling
- ✦ Eigenvector calculation



Bob Clapp

SEP Meeting 2009

Sicortex



- ✦ Source-receiver migration (FFT)
- ✦ Acoustic modeling
- ✦ Eigenvector calculation

- ✦ Compiler has problems
- ✦ Problems become L1/L2/Ram bound
- ✦ IO system needs work

Outline

Motivation

Unique hardware

Algorithmic changes

Programming challenges

Demo



Bob Clapp

SEP Meeting 2009

Ecoram

Outline


Motivation

Unique hardware

Algorithmic changes

Programming challenges

Demo



L1 Cache




L2/L3 Cache



Main memory



Network



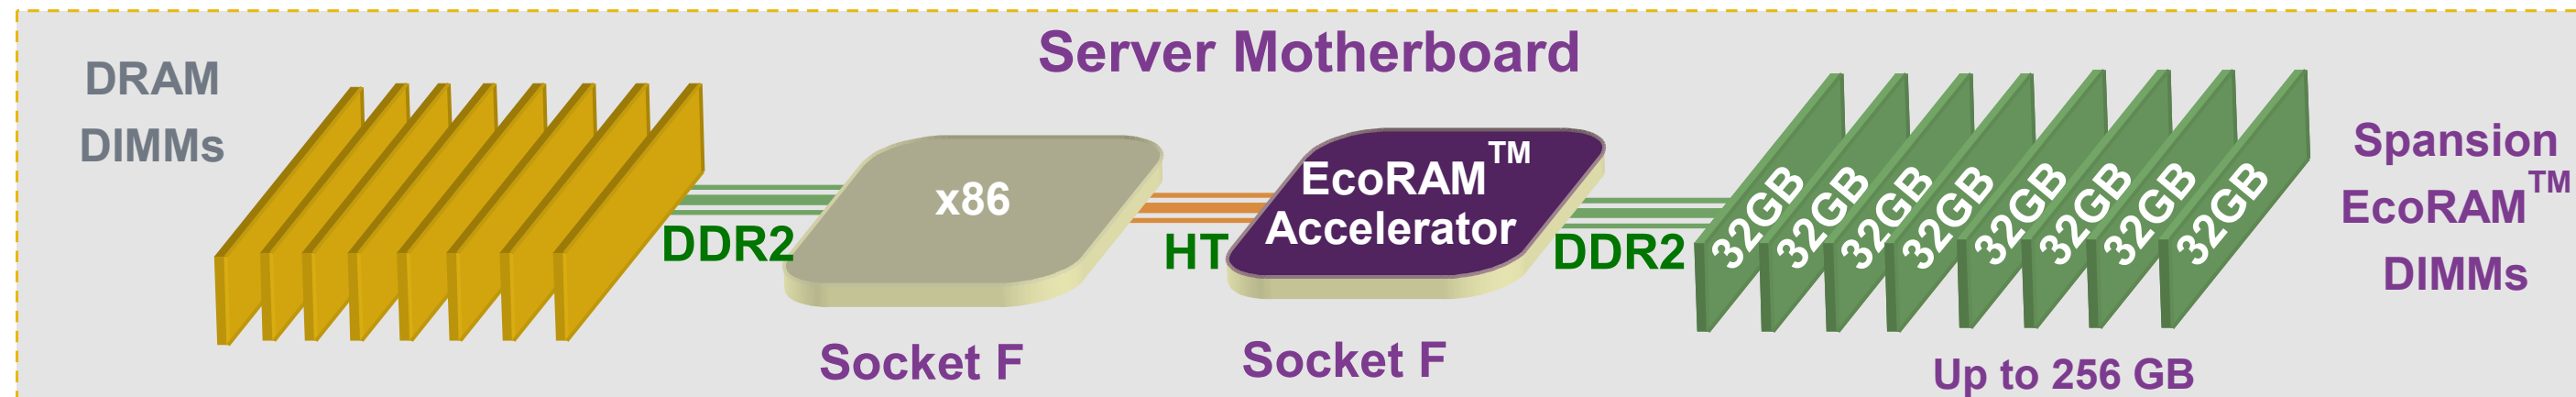
Disk



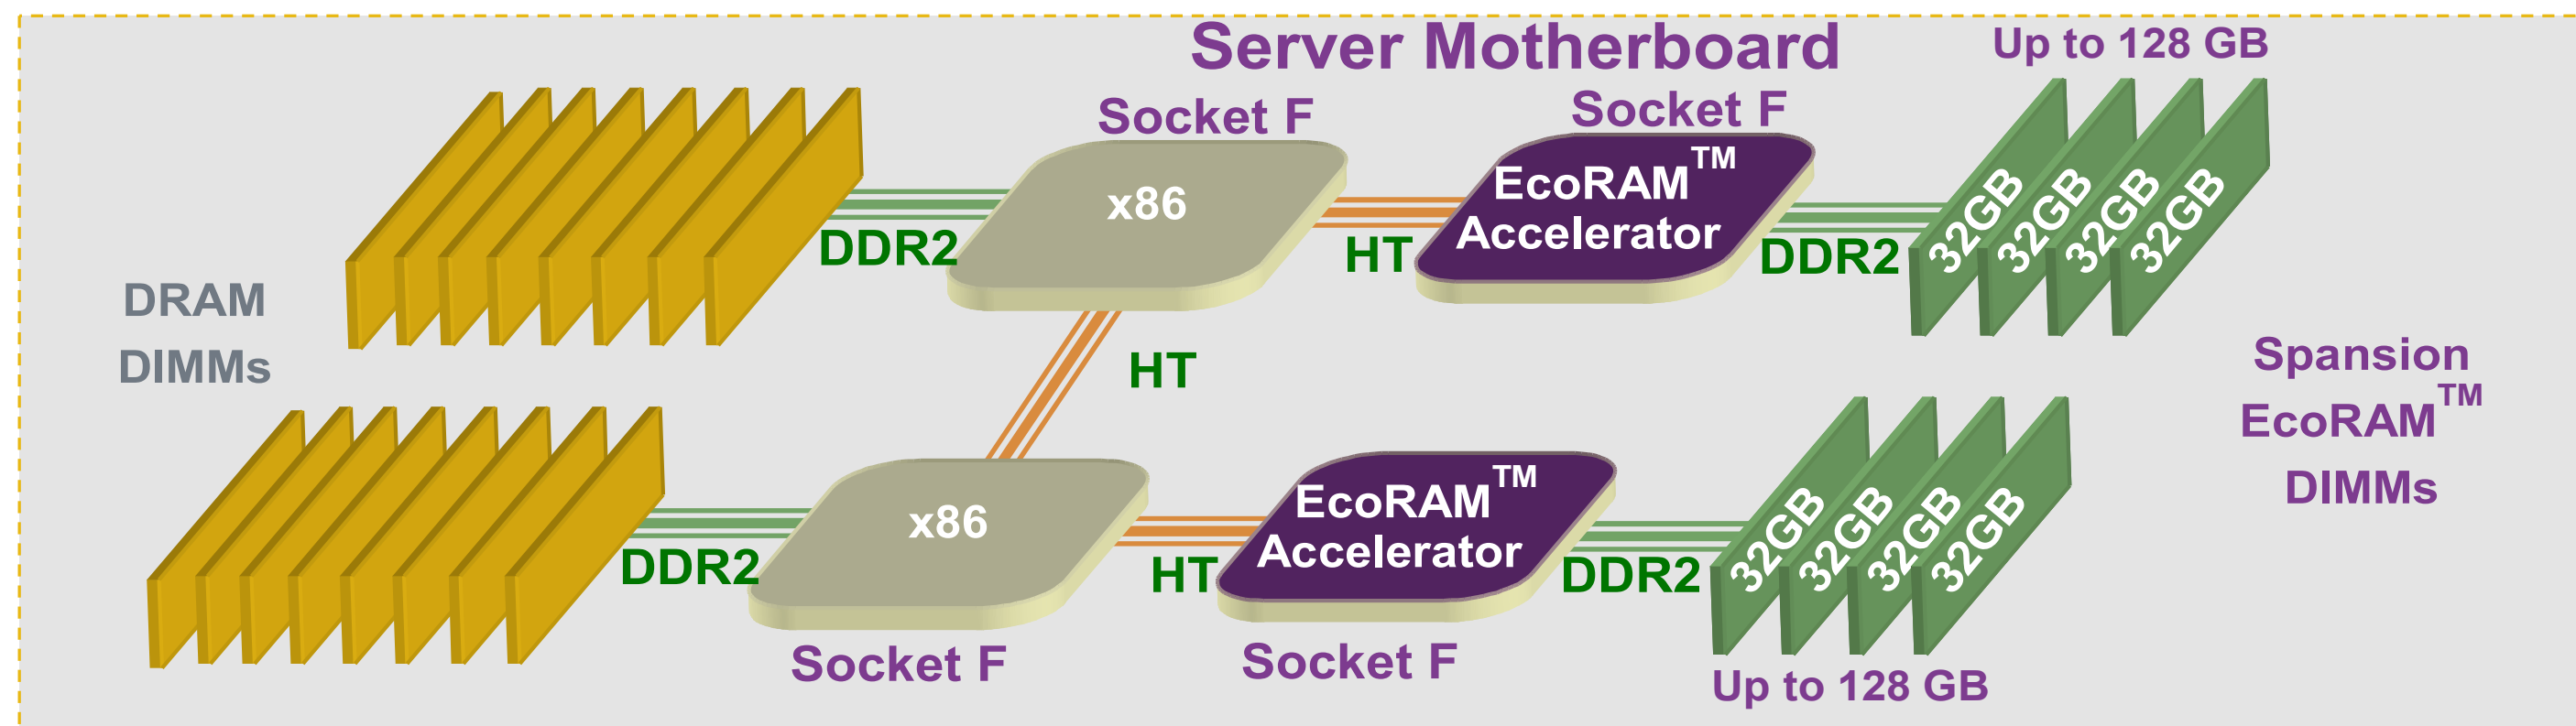
Bob Clapp

SEP Meeting 2009

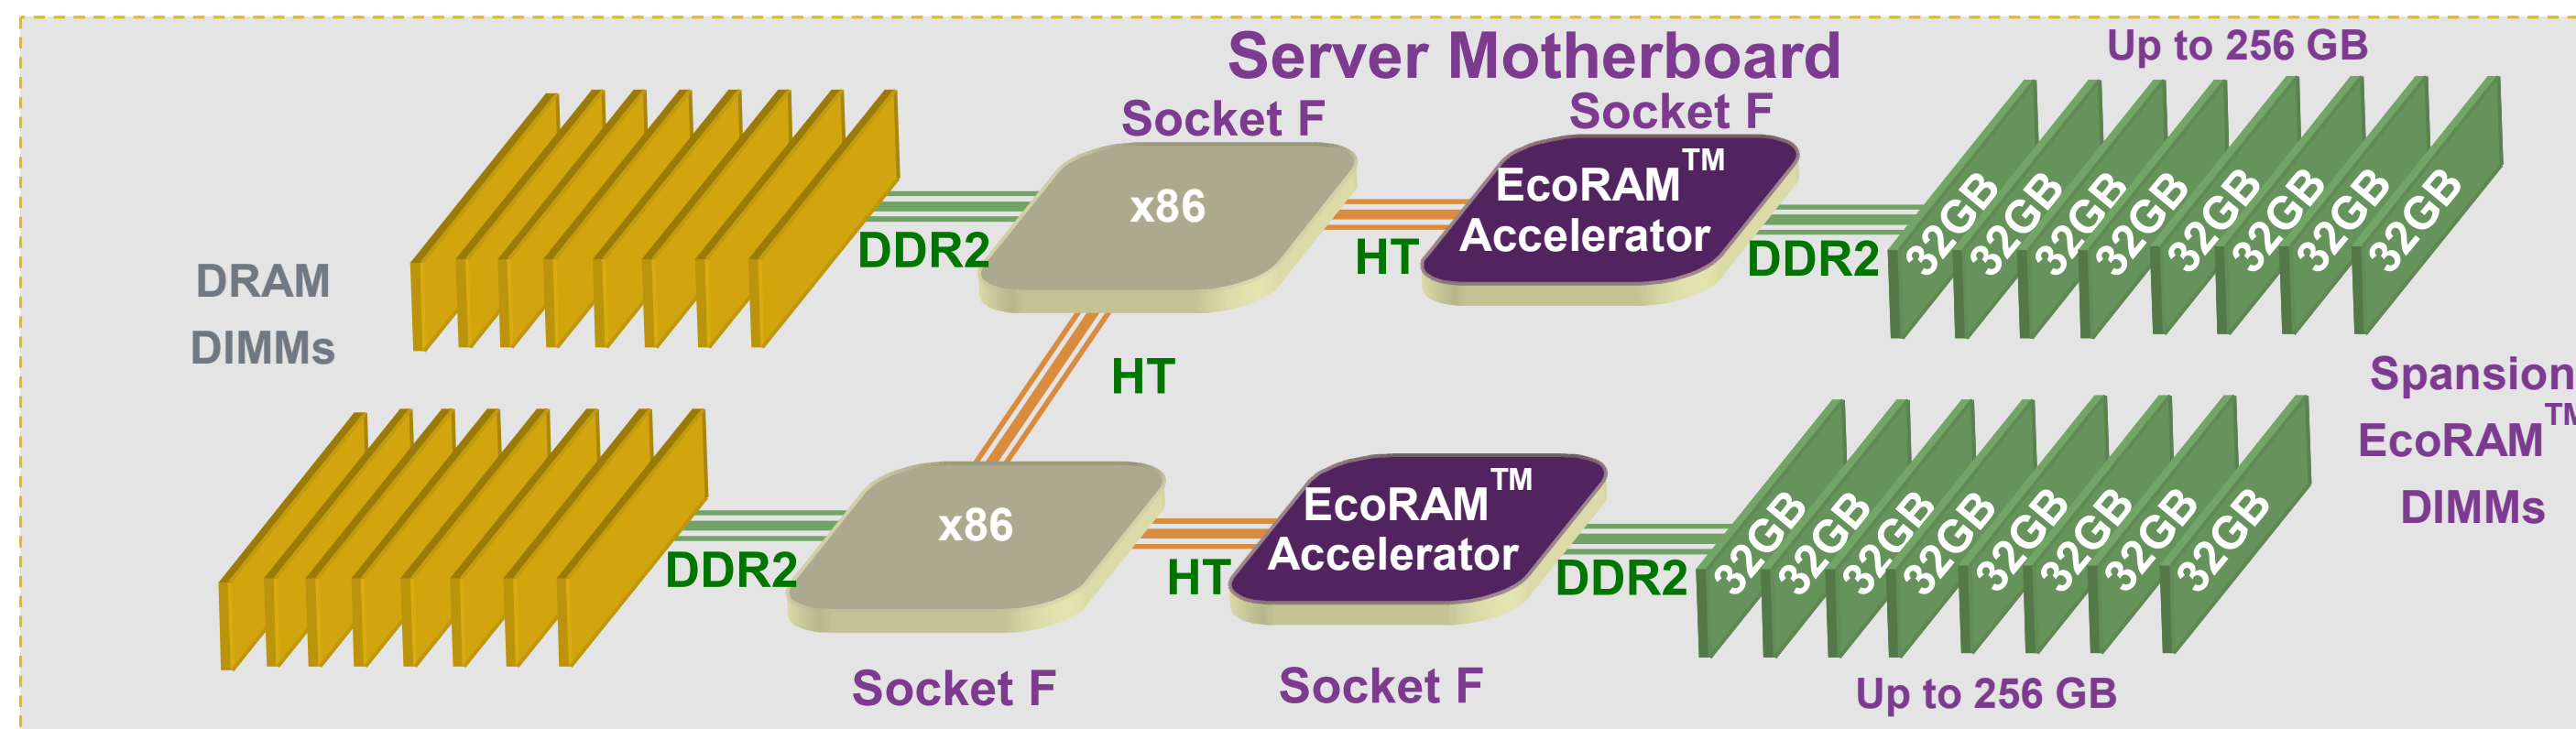
EcoRAM Server Configurations



256GB
Single-Accelerator

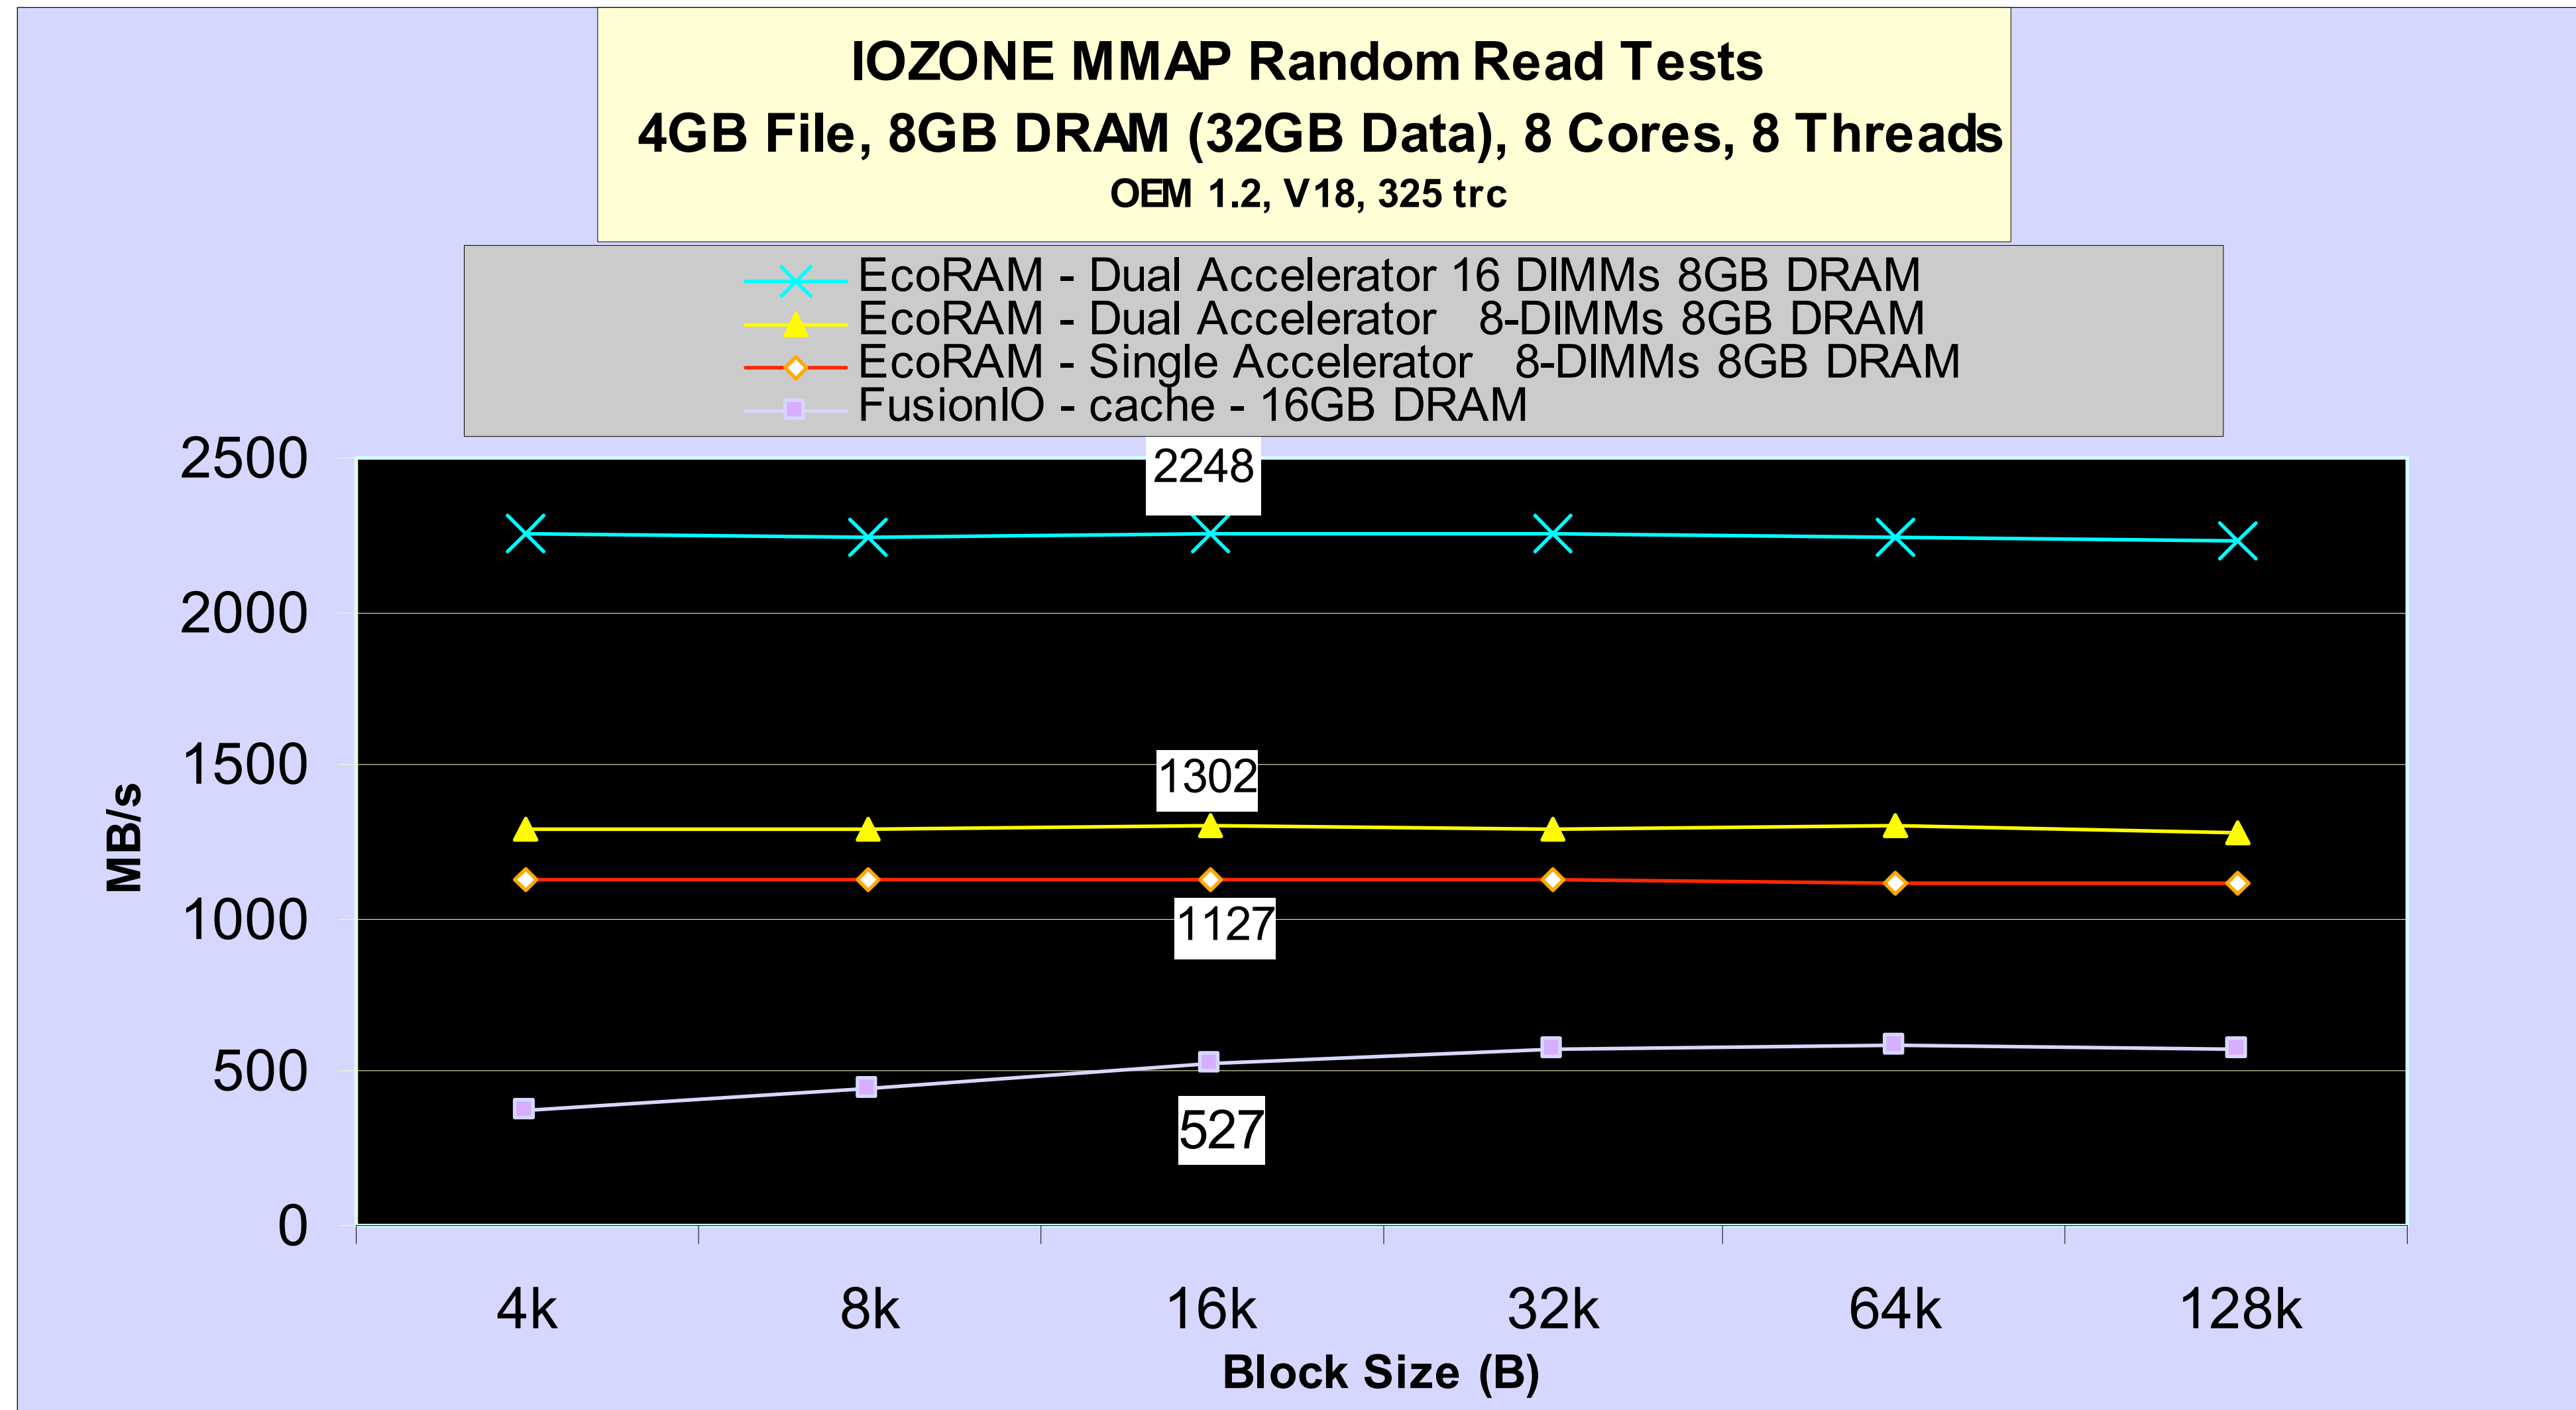


256GB
Dual-Accelerator



512GB
Dual-Accelerator

EcoRAM Raw Read Throughput Review



Ecoram

Outline



Motivation

Unique hardware

Algorithmic changes

Programming challenges

Demo



- Latency and bandwidth to disk is the limiting factor
- Replace RAM with solid state at 8x the density
- Speed up reads with FPGA/ASIC
- Make it easy to program



Bob Clapp

SEP Meeting 2009

Ecoram

Outline


Motivation

Unique hardware


Algorithmic changes

Programming challenges

Demo



- ✦ Visualization
- ✦ Very large transposes



- ✦ Latency and bandwidth to disk is the limiting factor
- ✦ Replace RAM with solid state at 8x the density
- ✦ Speed up reads with FPGA/ASIC
- ✦ Make it easy to program



Bob Clapp

SEP Meeting 2009

Ecoram

Outline


Motivation

Unique hardware


Algorithmic changes

Programming challenges

Demo



- ✦ Visualization
- ✦ Very large transposes



- ✦ Cost
- ✦ Can do the same thing with a networked system
- ✦ Not big enough memory options yet



Bob Clapp

SEP Meeting 2009

Accelerators

Outline

Motivation

Unique hardware

Algorithmic changes

Programming challenges

Demo



Accelerator memory system




Accelerator



Main memory



Network



Disk

- ✦ Accelerator bandwidth GB/s
- ✦ Internal memory systems vary dramatically
- ✦ Require different, sometimes significantly, programming paradigms



GPU

Outline

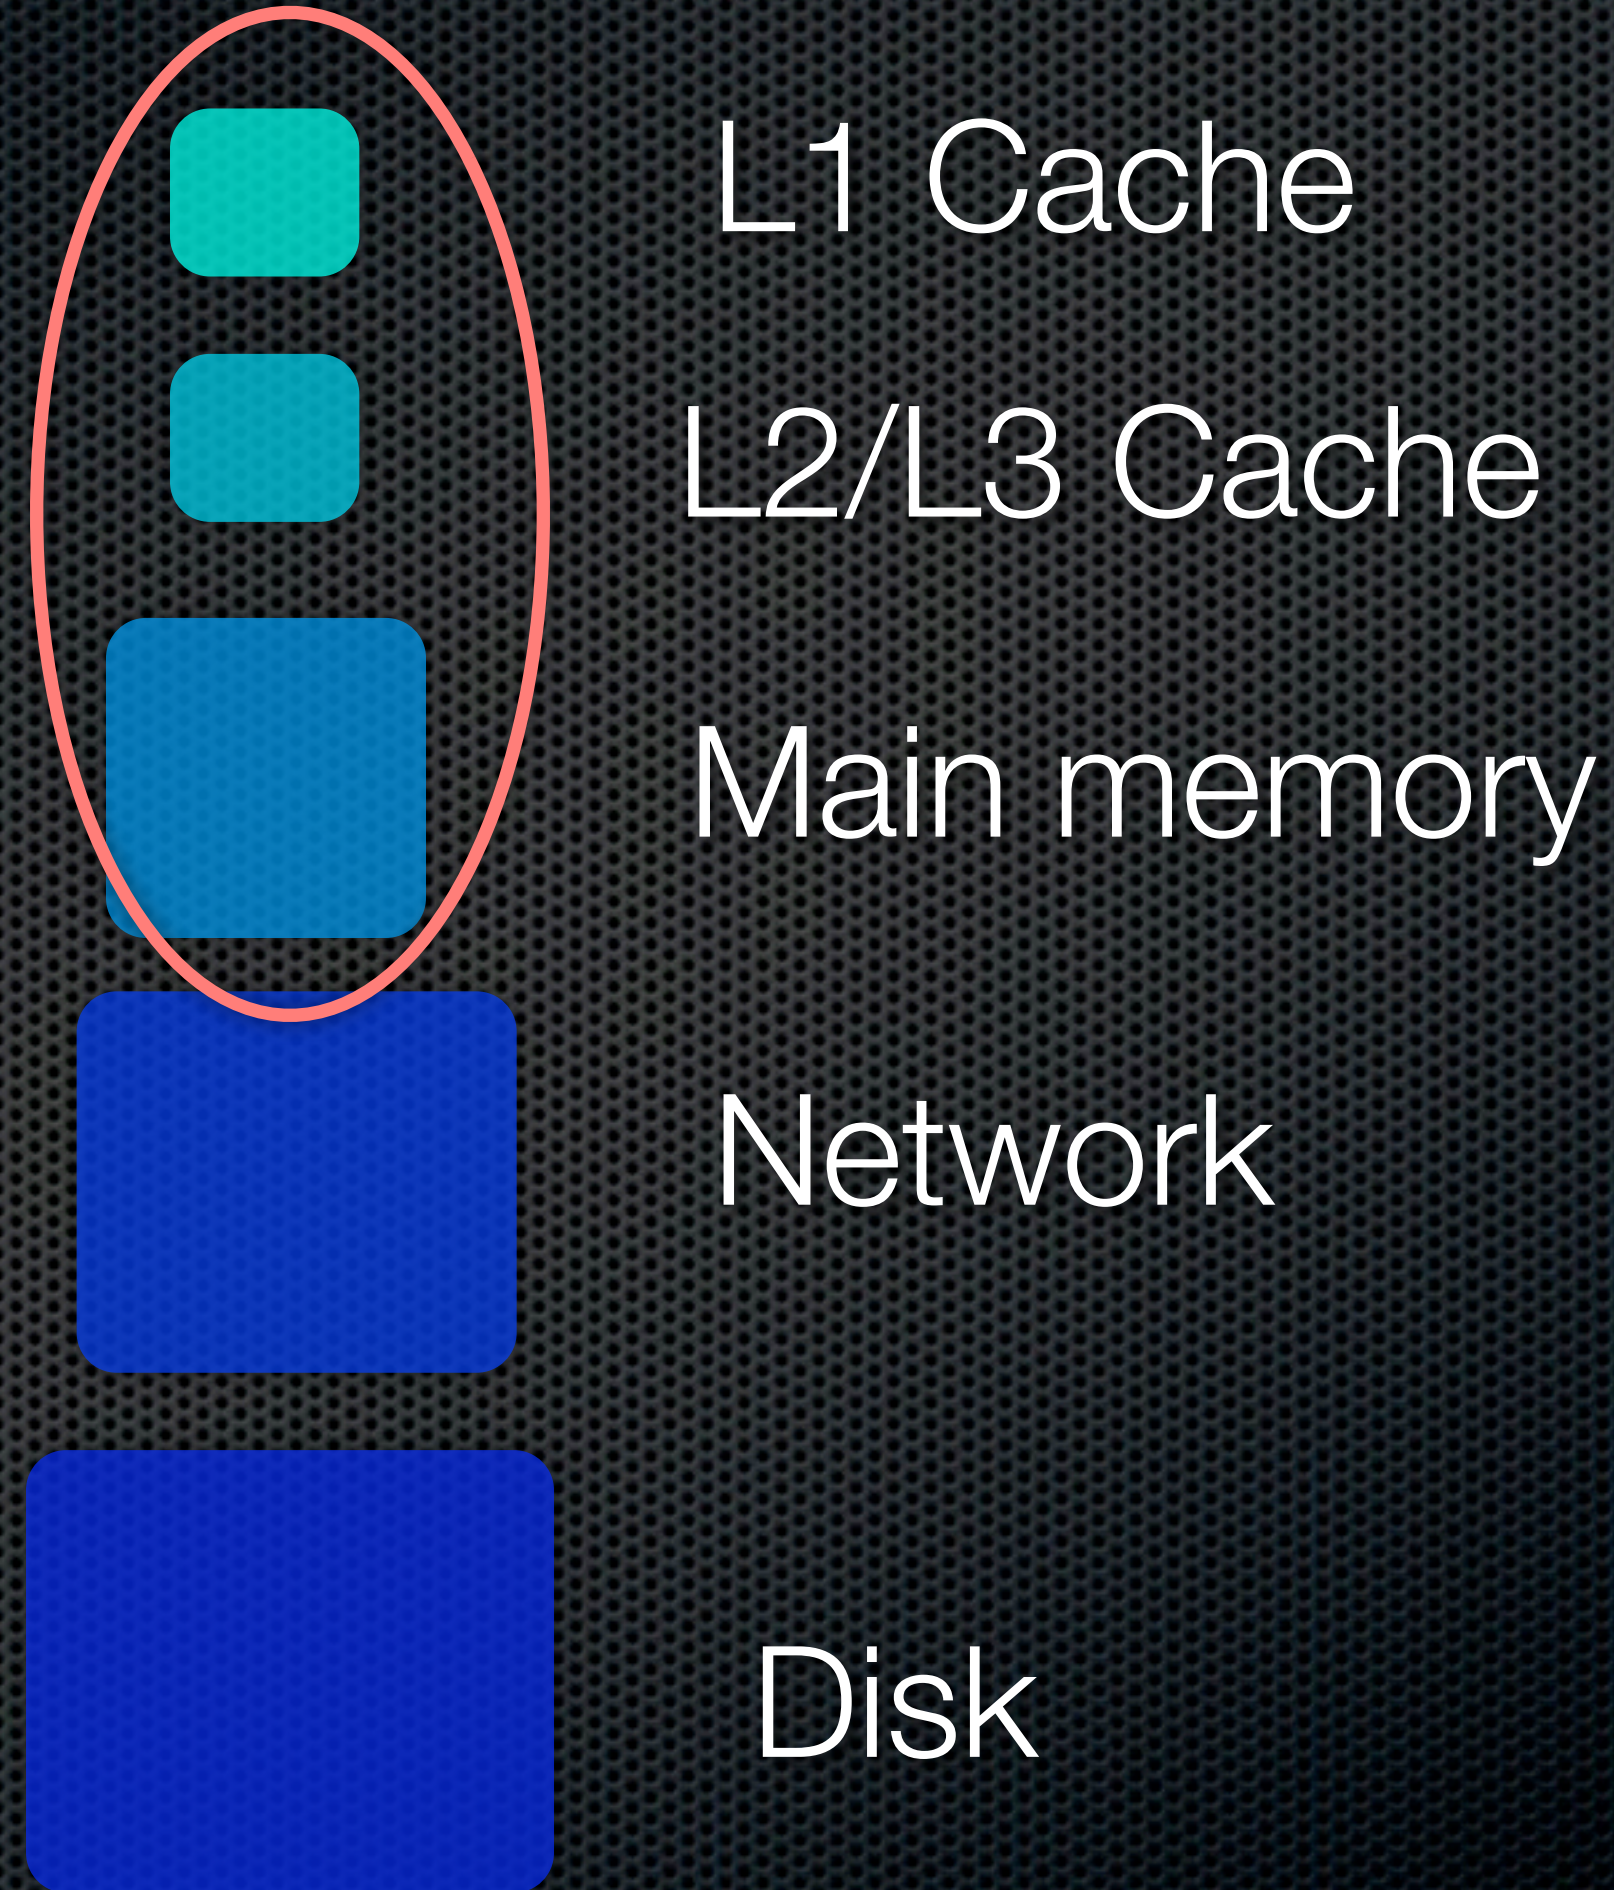
Motivation

Unique hardware

Algorithmic changes

Programming challenges

Demo



L1 Cache

L2/L3 Cache

Main memory

Network

Disk



GPU

Outline


Motivation

Unique hardware

Algorithmic changes

Programming challenges

Demo



- Many parallel processing units
- Large memory bandwidth
- Hide latency with multiple cores



GPU

Outline

Motivation

Unique hardware

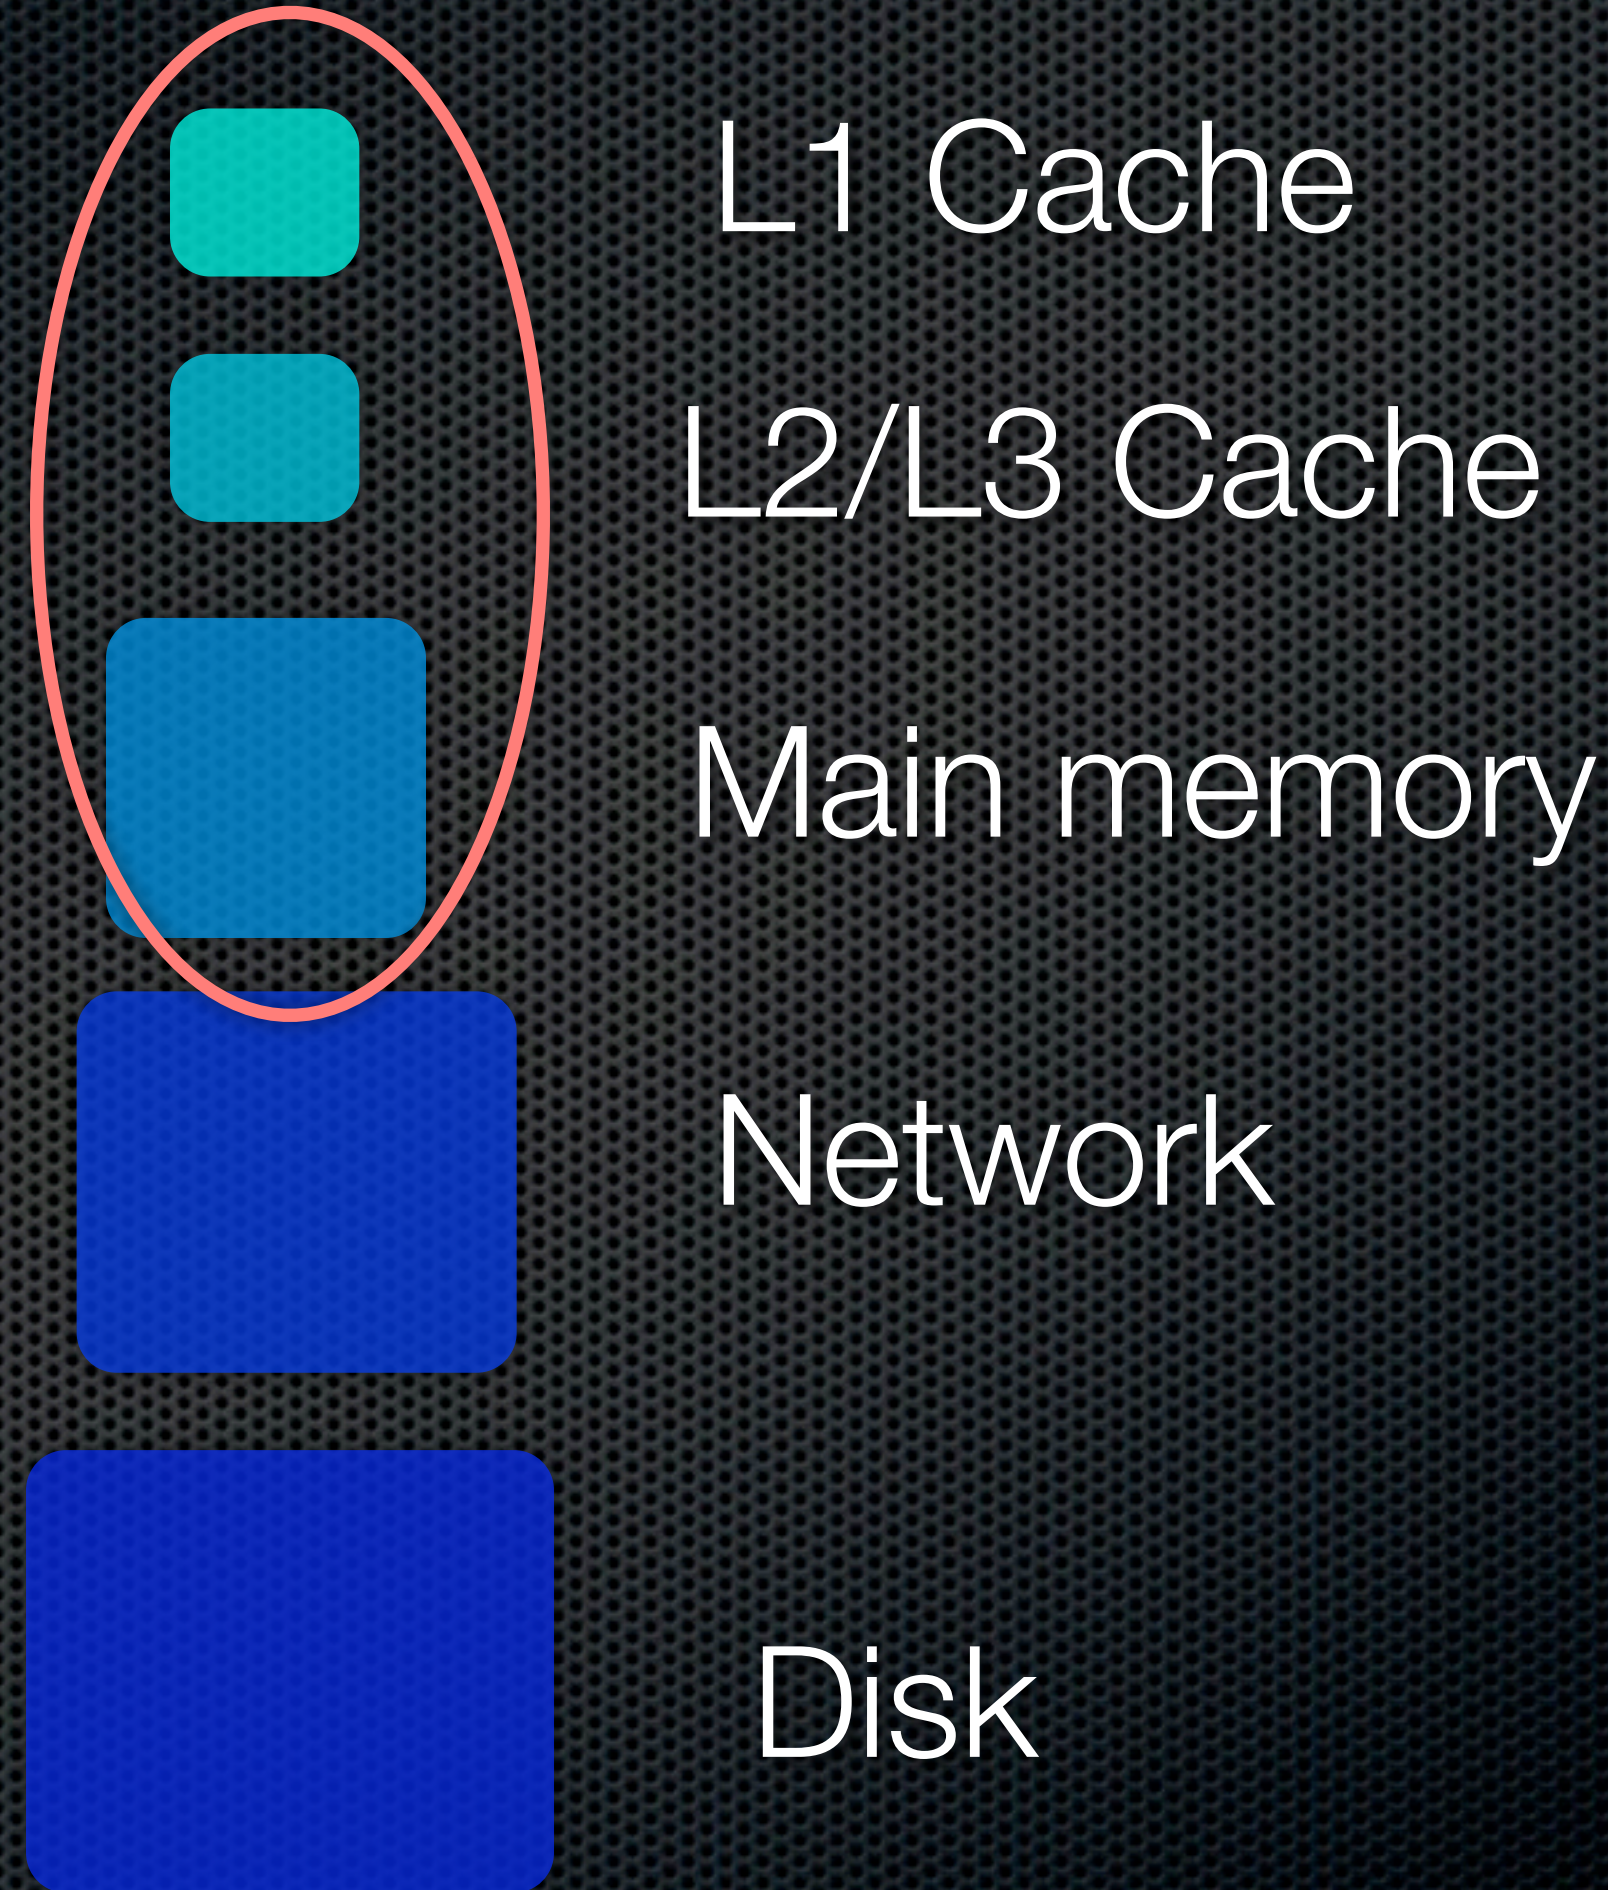
Algorithmic changes

Programming challenges

Demo



- Acoustic modeling
- RTM
- Sparse matrix ops
- Spectral factorization



- Many parallel processing units
- Large memory bandwidth
- Hide latency with multiple cores



GPU

Outline

Motivation

Unique hardware

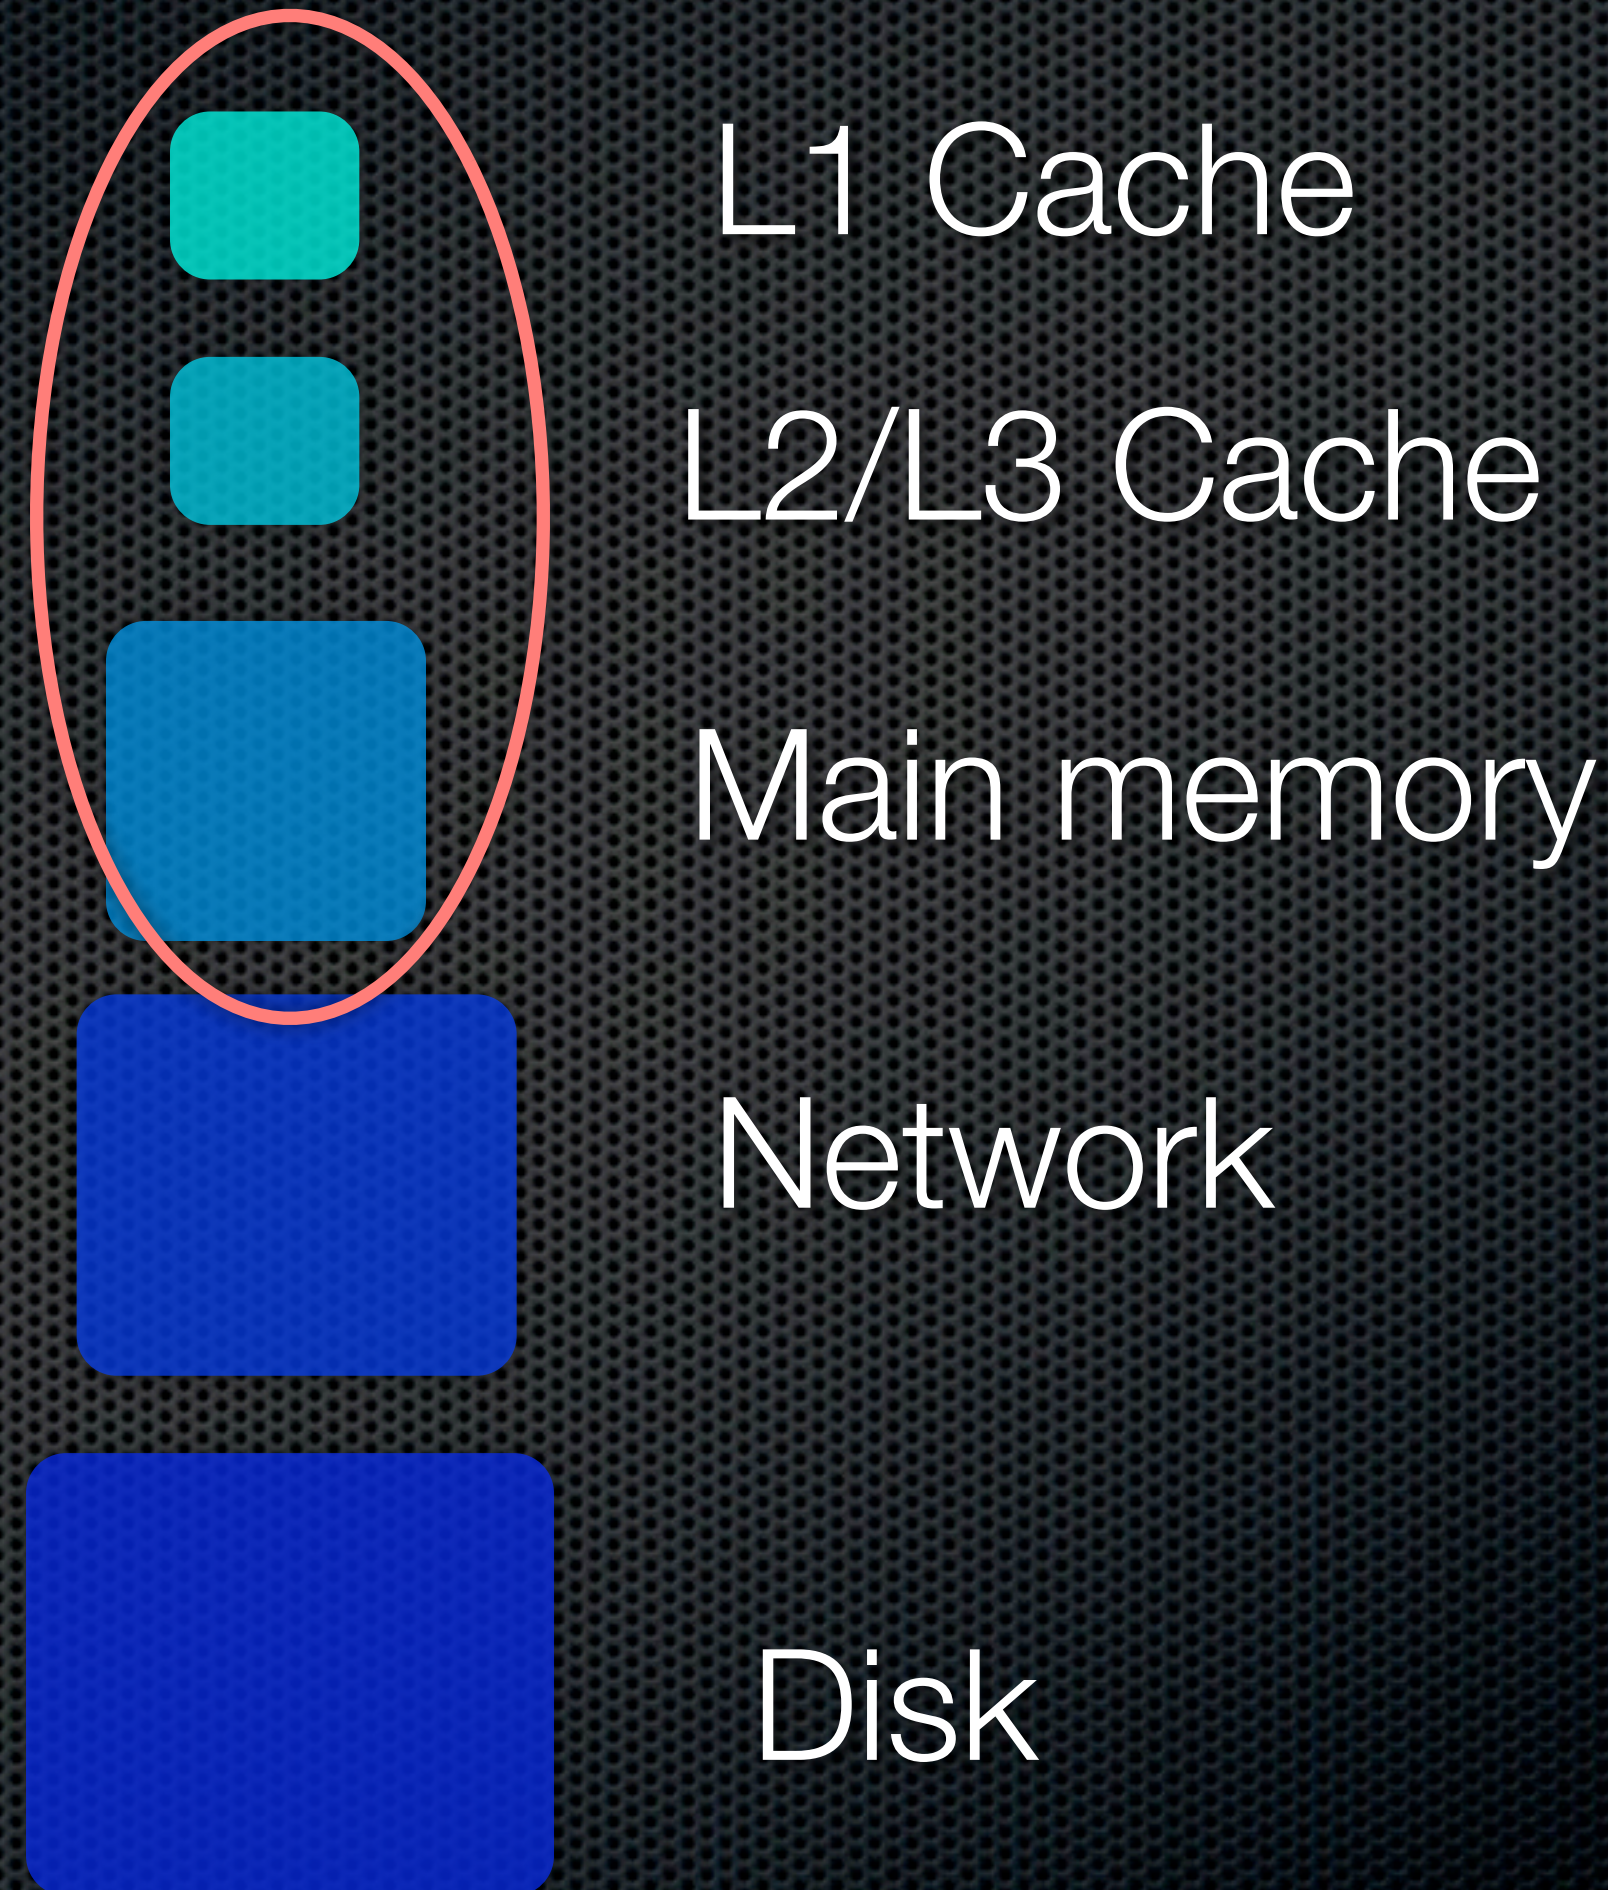
Algorithmic changes

Programming challenges

Demo



- ✦ Acoustic modeling
- ✦ RTM
- ✦ Sparse matrix ops
- ✦ Spectral factorization



- ✦ Exceed GPU memory at your own peril
- ✦ Compiler can be challenging
- ✦ 80GB/s sounds like a lot....



FPGAs

Outline


Motivation

Unique hardware


Algorithmic changes

Programming challenges

Demo



L1 Cache



L2/L3 Cache



Main memory



Network



Disk



Bob Clapp

SEP Meeting 2009

FPGAs

Outline

Motivation

Unique hardware

Algorithmic changes

Programming challenges

Demo



L1 Cache

L2/L3 Cache

Main memory

Network

Disk

- ✦ Stream everything
- ✦ Very large (MBs) L1 cache like storage (block rams)
- ✦ Complete pipelining
- ✦ High op count no problem



Bob Clapp

SEP Meeting 2009

FPGAs

Outline

Motivation

Unique hardware


Algorithmic changes

Programming challenges

Demo



- ✦ Acoustic modeling
- ✦ Sparse matrix ops



L1 Cache

L2/L3 Cache

Main memory

Network

Disk

- ✦ Stream everything
- ✦ Very large (MBs) L1 cache like storage (block rams)
- ✦ Complete pipelining
- ✦ High op count no problem



FPGAs

Outline

Motivation

Unique hardware


Algorithmic changes

Programming challenges

Demo



- ✦ Acoustic modeling
- ✦ Sparse matrix ops



- ✦ Cost
- ✦ Programmability
- ✦ Some problems are difficult to express with streaming



Convey

Outline

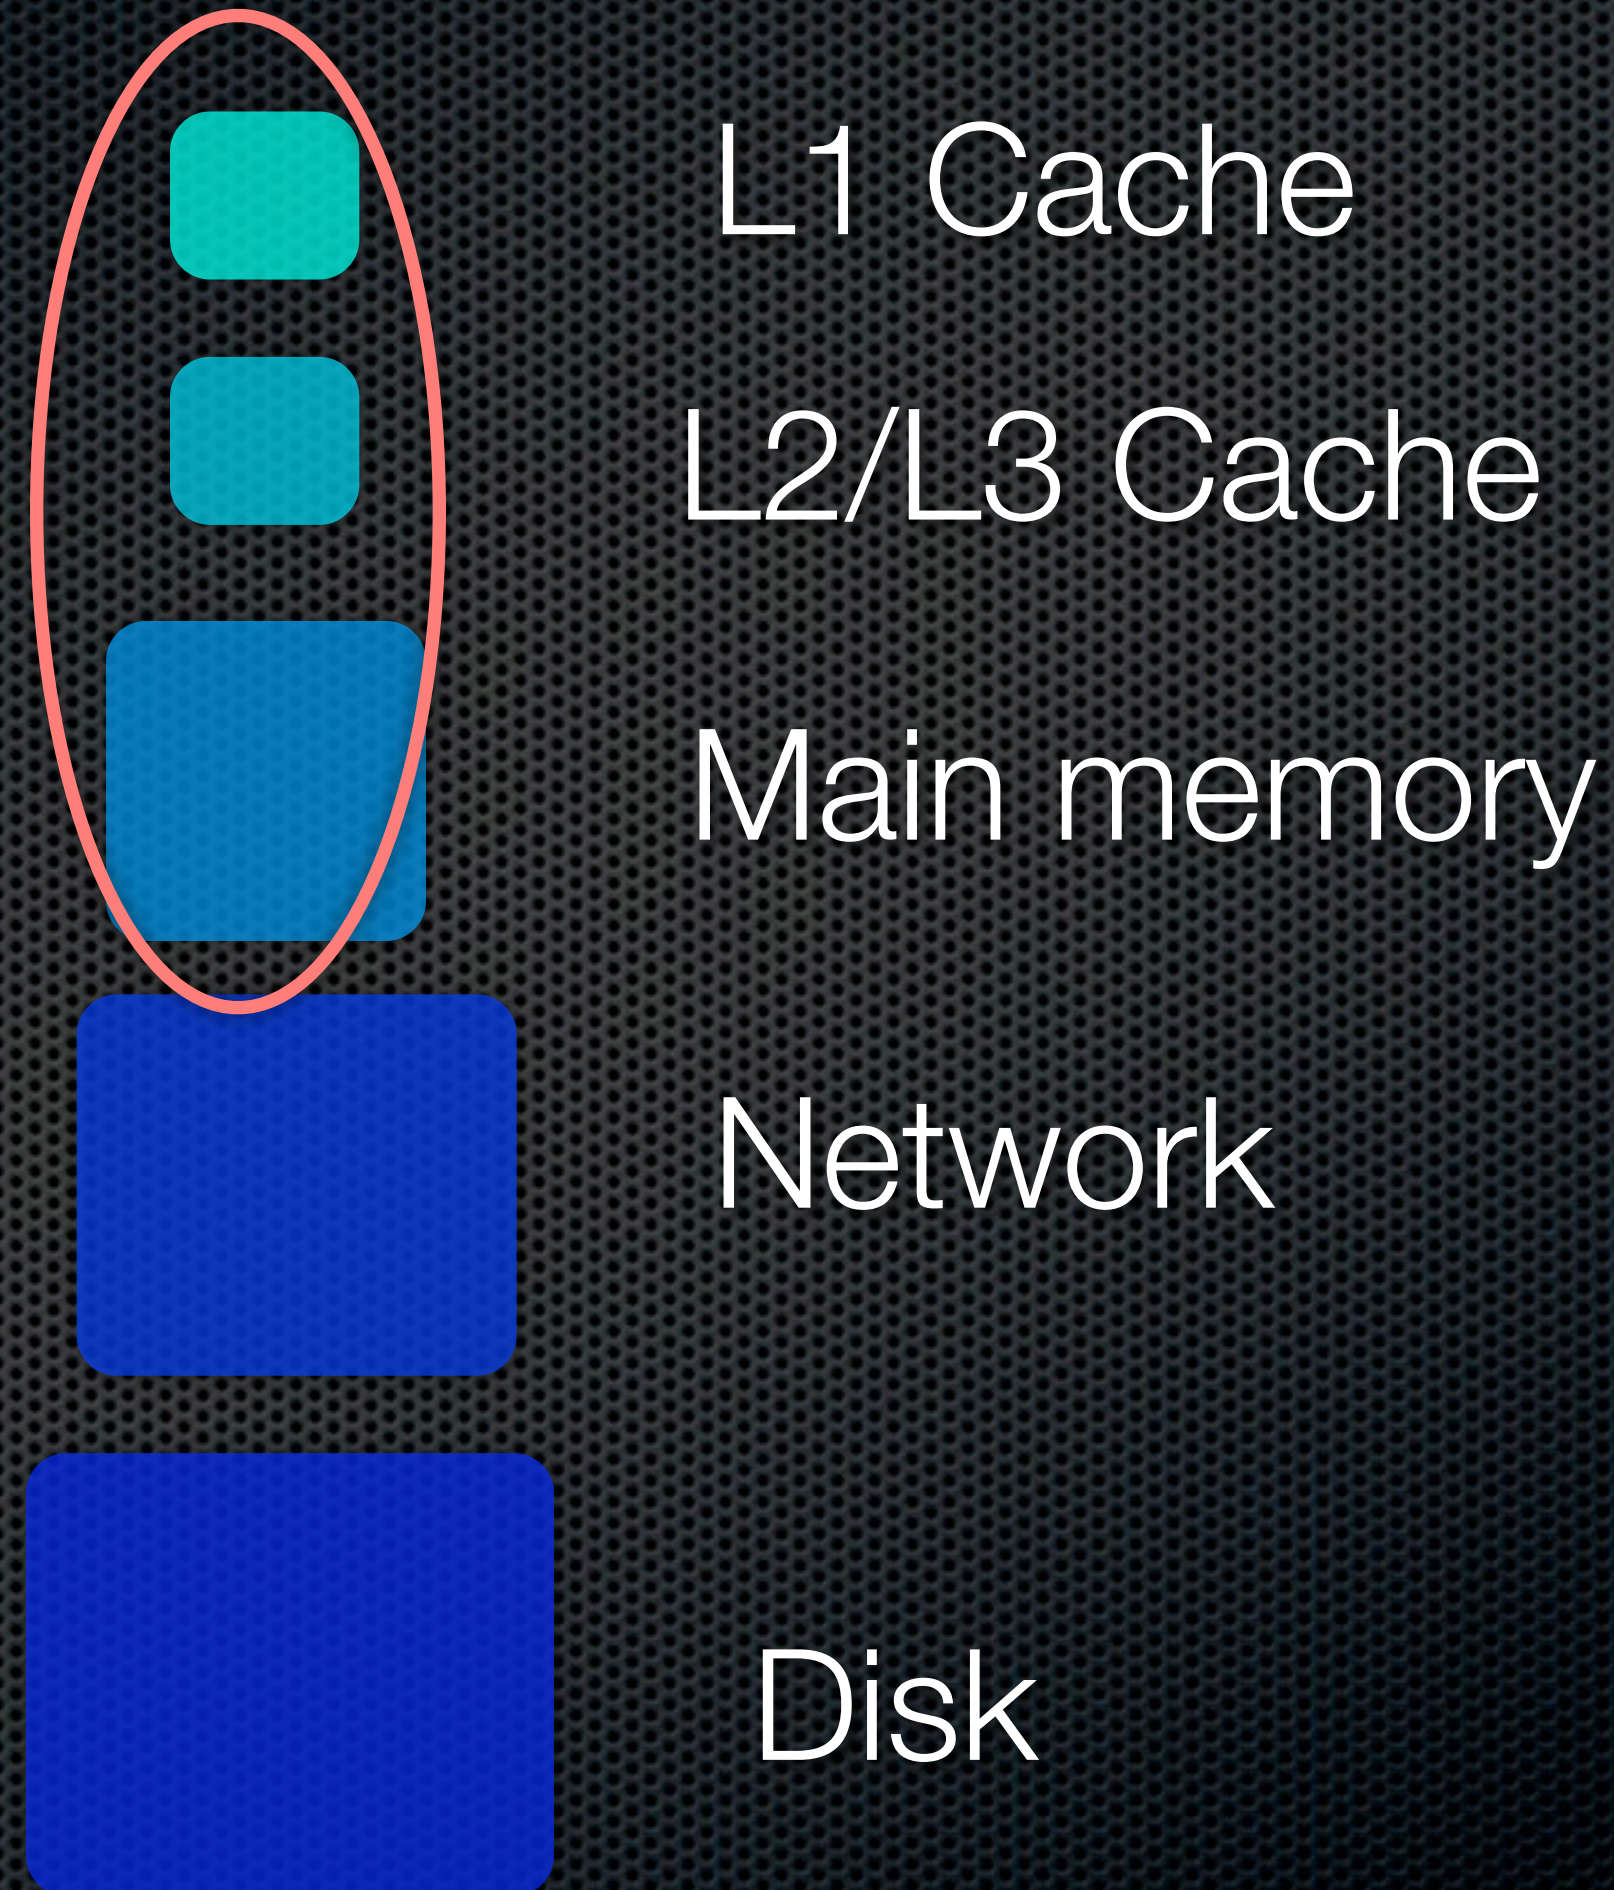
Motivation

Unique hardware

Algorithmic changes

Programming challenges

Demo



Bob Clapp

SEP Meeting 2009

Convey

Outline

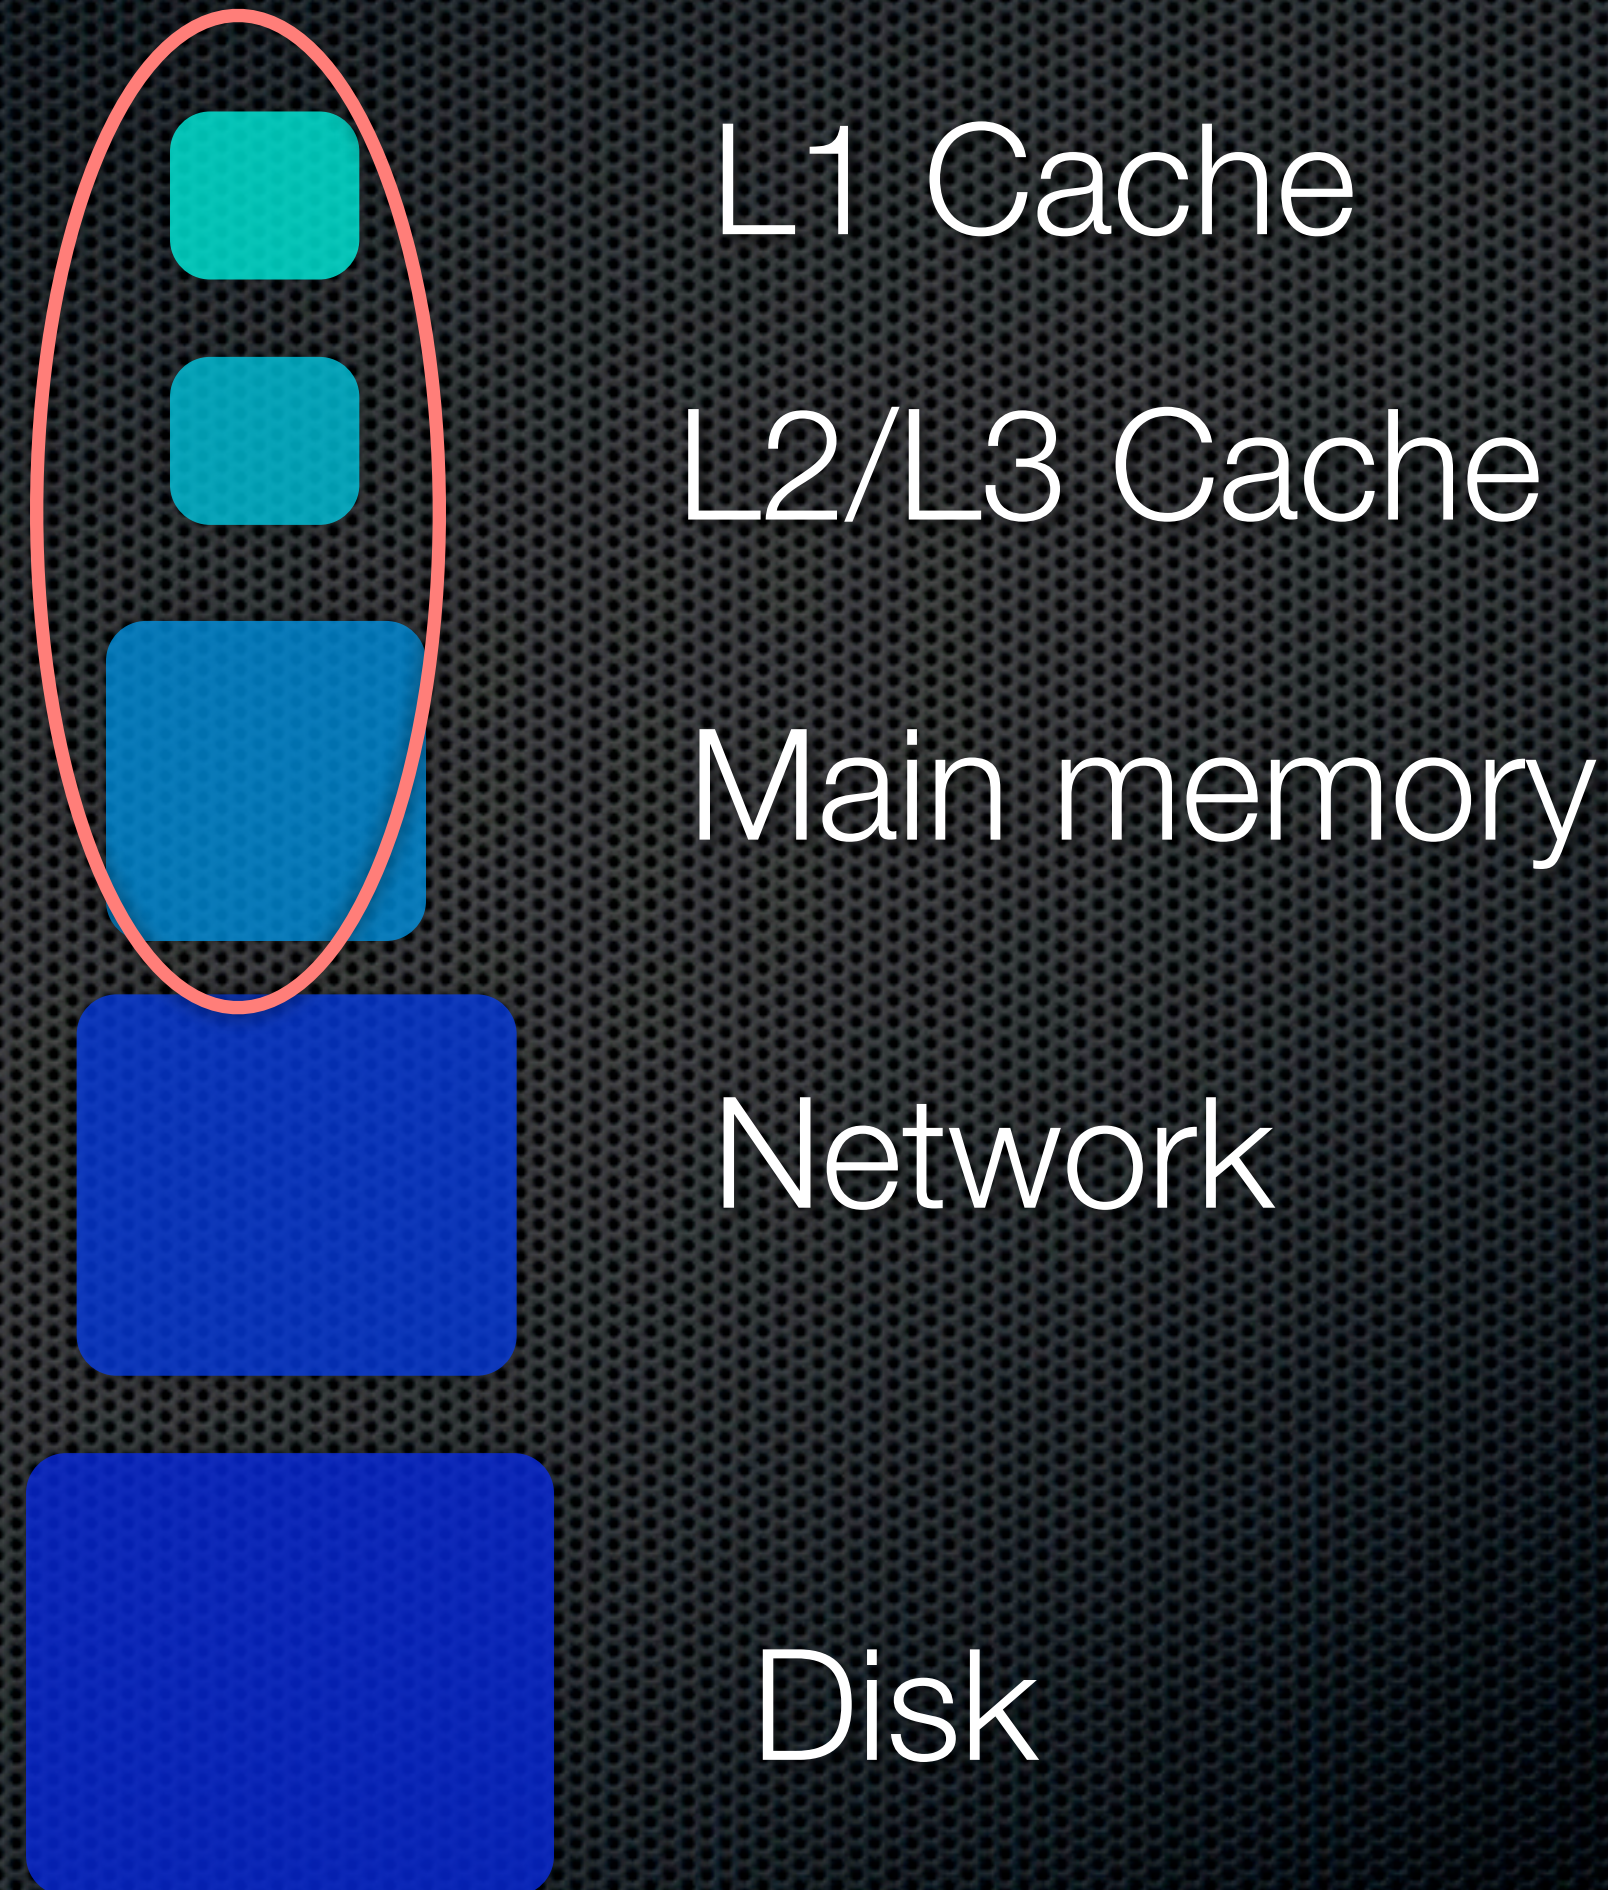
Motivation

Unique hardware

Algorithmic changes

Programming challenges

Demo



- ✦ FPGAs offer great potential but are difficult to program
- ✦ Create personalities to generalize FPGA computing
- ✦ Vector personality isn't a bad first choice



Bob Clapp

SEP Meeting 2009

What is the effect on algorithmic development?

Outline

Motivation

Unique hardware

Algorithmic changes

Programming challenges

Demo

- ✦ Offer up to two orders of magnitude improvement if
- ✦ Definition of “costly” isn’t as simple and can vary dramatically
- ✦ Some require significantly different programming styles



RTM boundary conditions

Outline

Motivation

Unique hardware

Algorithmic changes

Programming challenges

Demo

- ✦ Conventional CPU with high end disk roughly balances IO/compute
- ✦ GPU: Transfer of saved wavefield quickly becomes bottleneck
- ✦ Sicortex: Saving the wavefield is the bottleneck



Disk transpose

Outline

Motivation

Unique hardware

Algorithmic changes

Programming challenges

Demo

- Loop over desired output :iout{
 - Compute desired input sample: iin{
 - $out[iout]=in[iin]$



Disk transpose

Outline

Motivation

Unique hardware

Algorithmic changes

Programming challenges

Demo

- Loop over desired output :iout{
 - Compute desired input sample: iin{
 - $out[iout]=in[iin]$

Improve performance by doing blocking



Disk transpose: In core

Outline

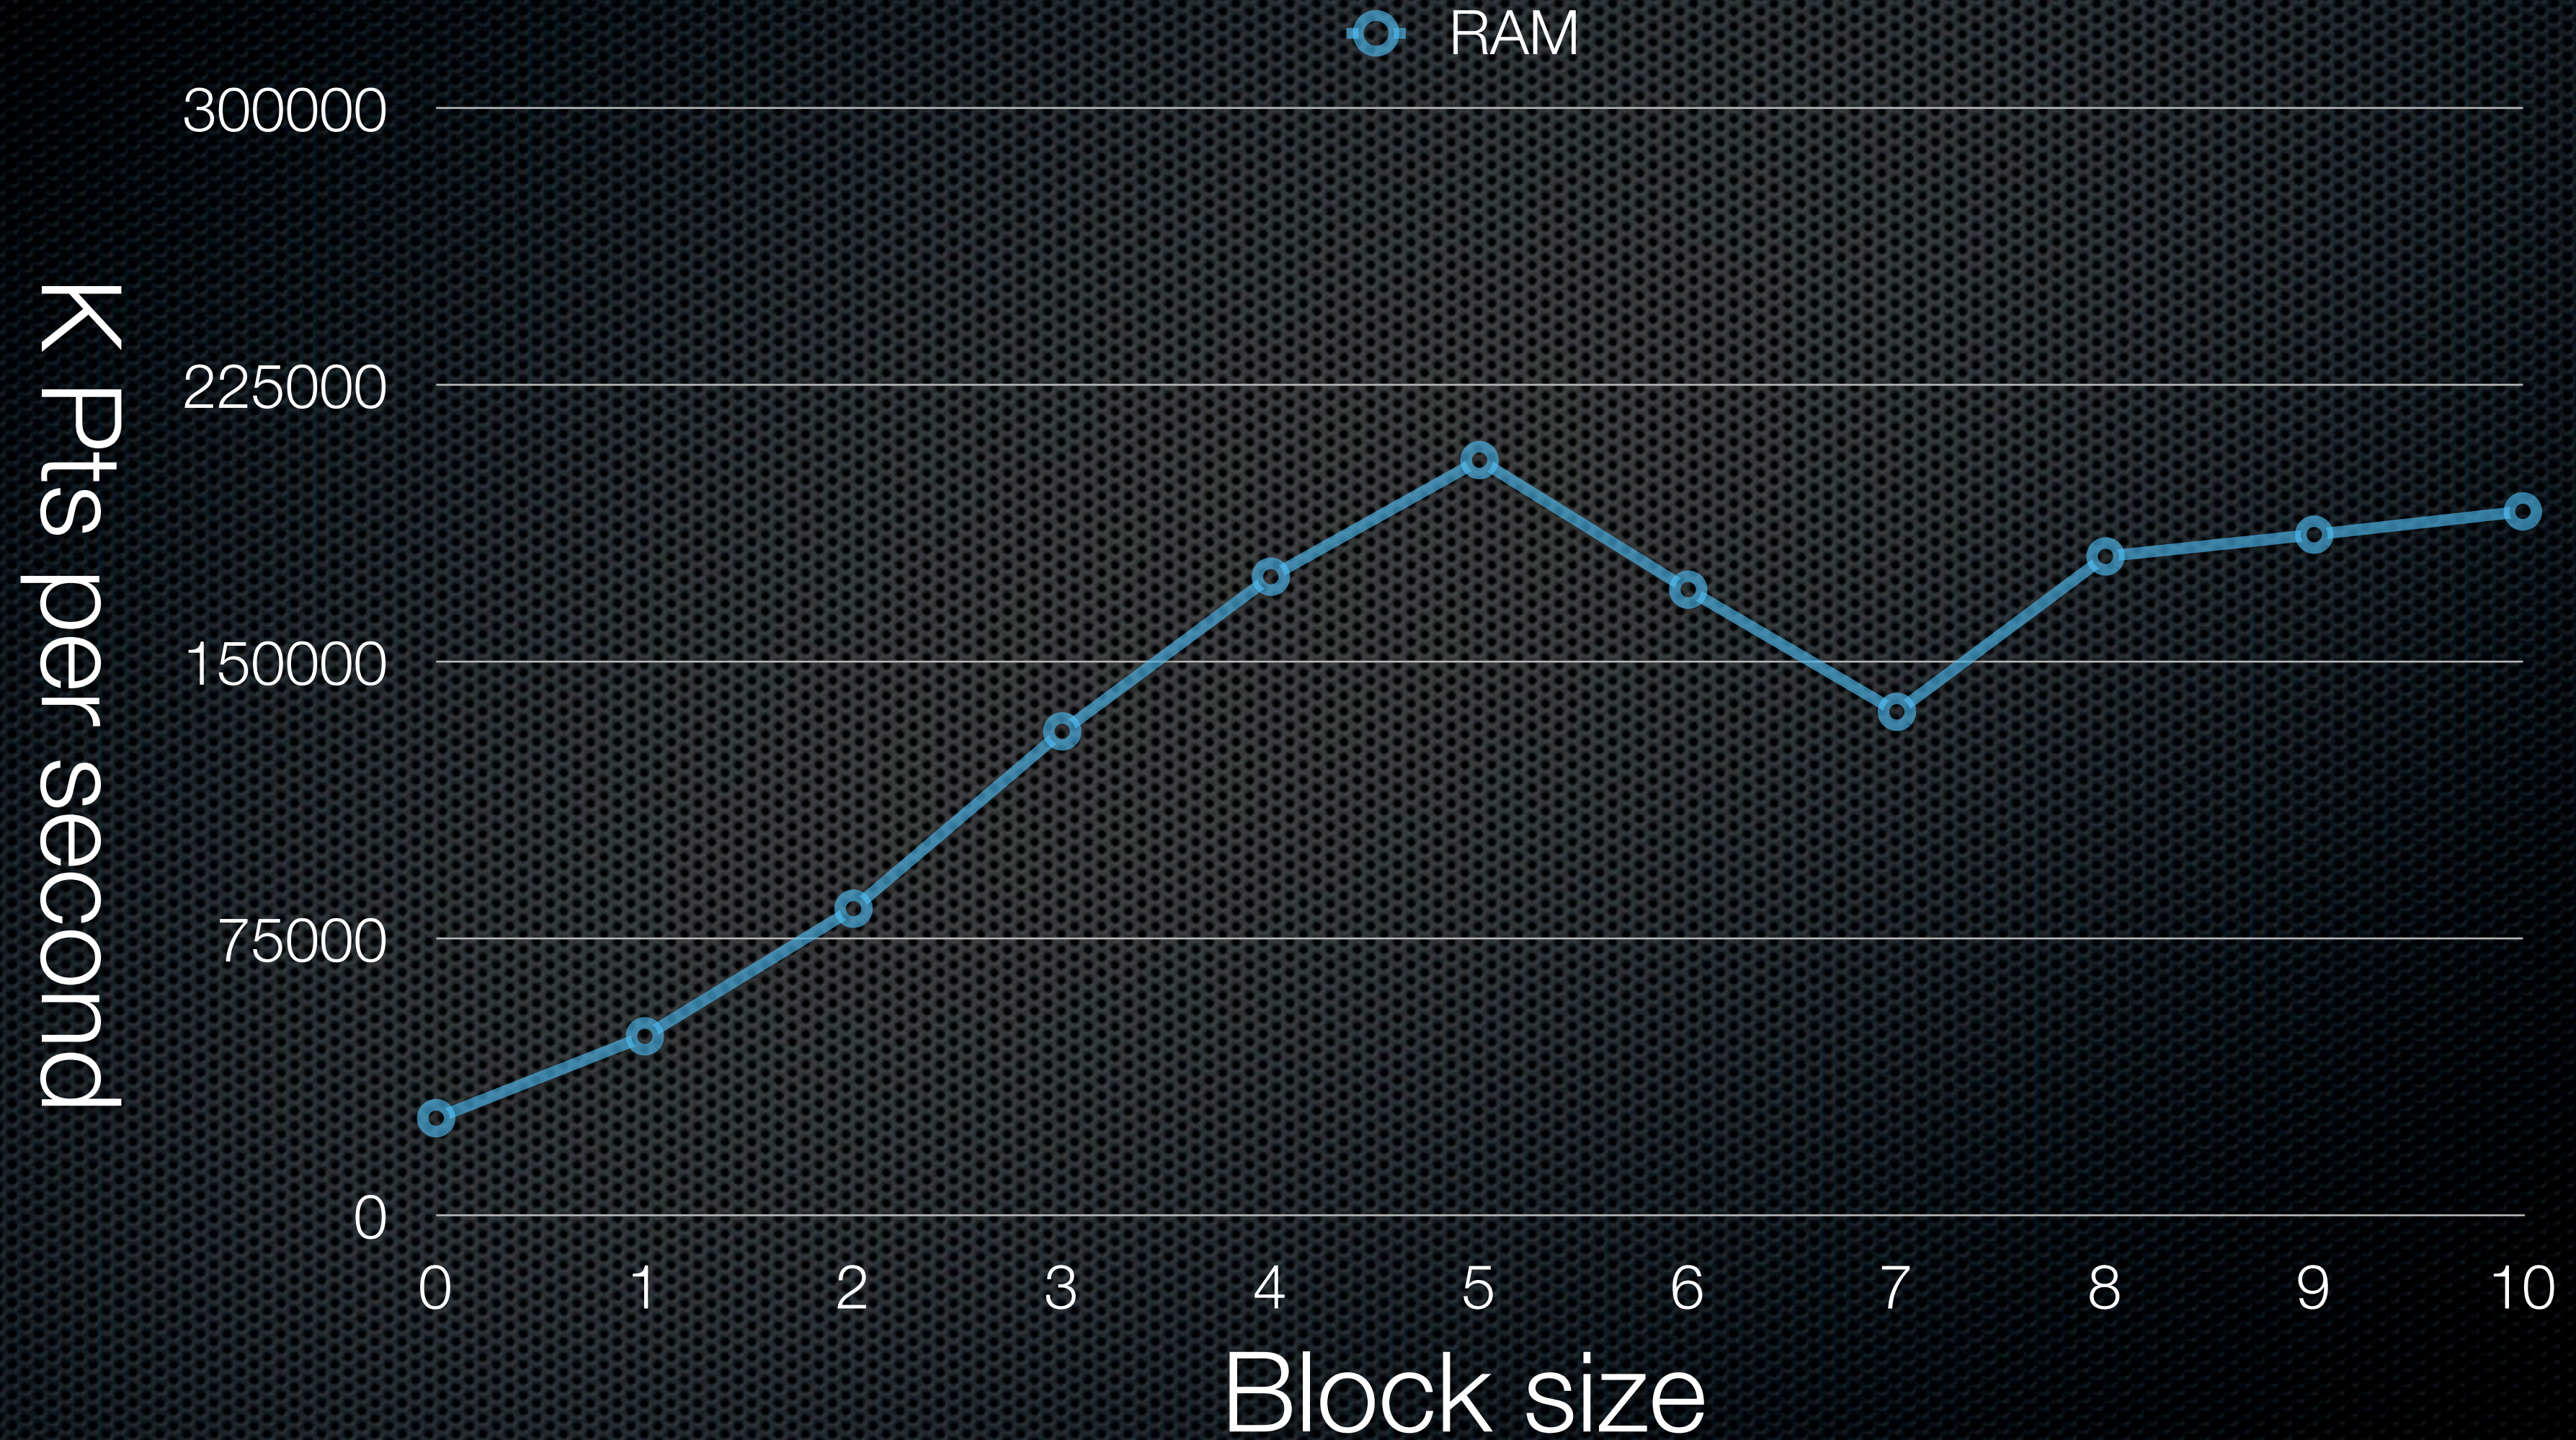
Motivation

Unique hardware

Algorithmic changes

Programming challenges

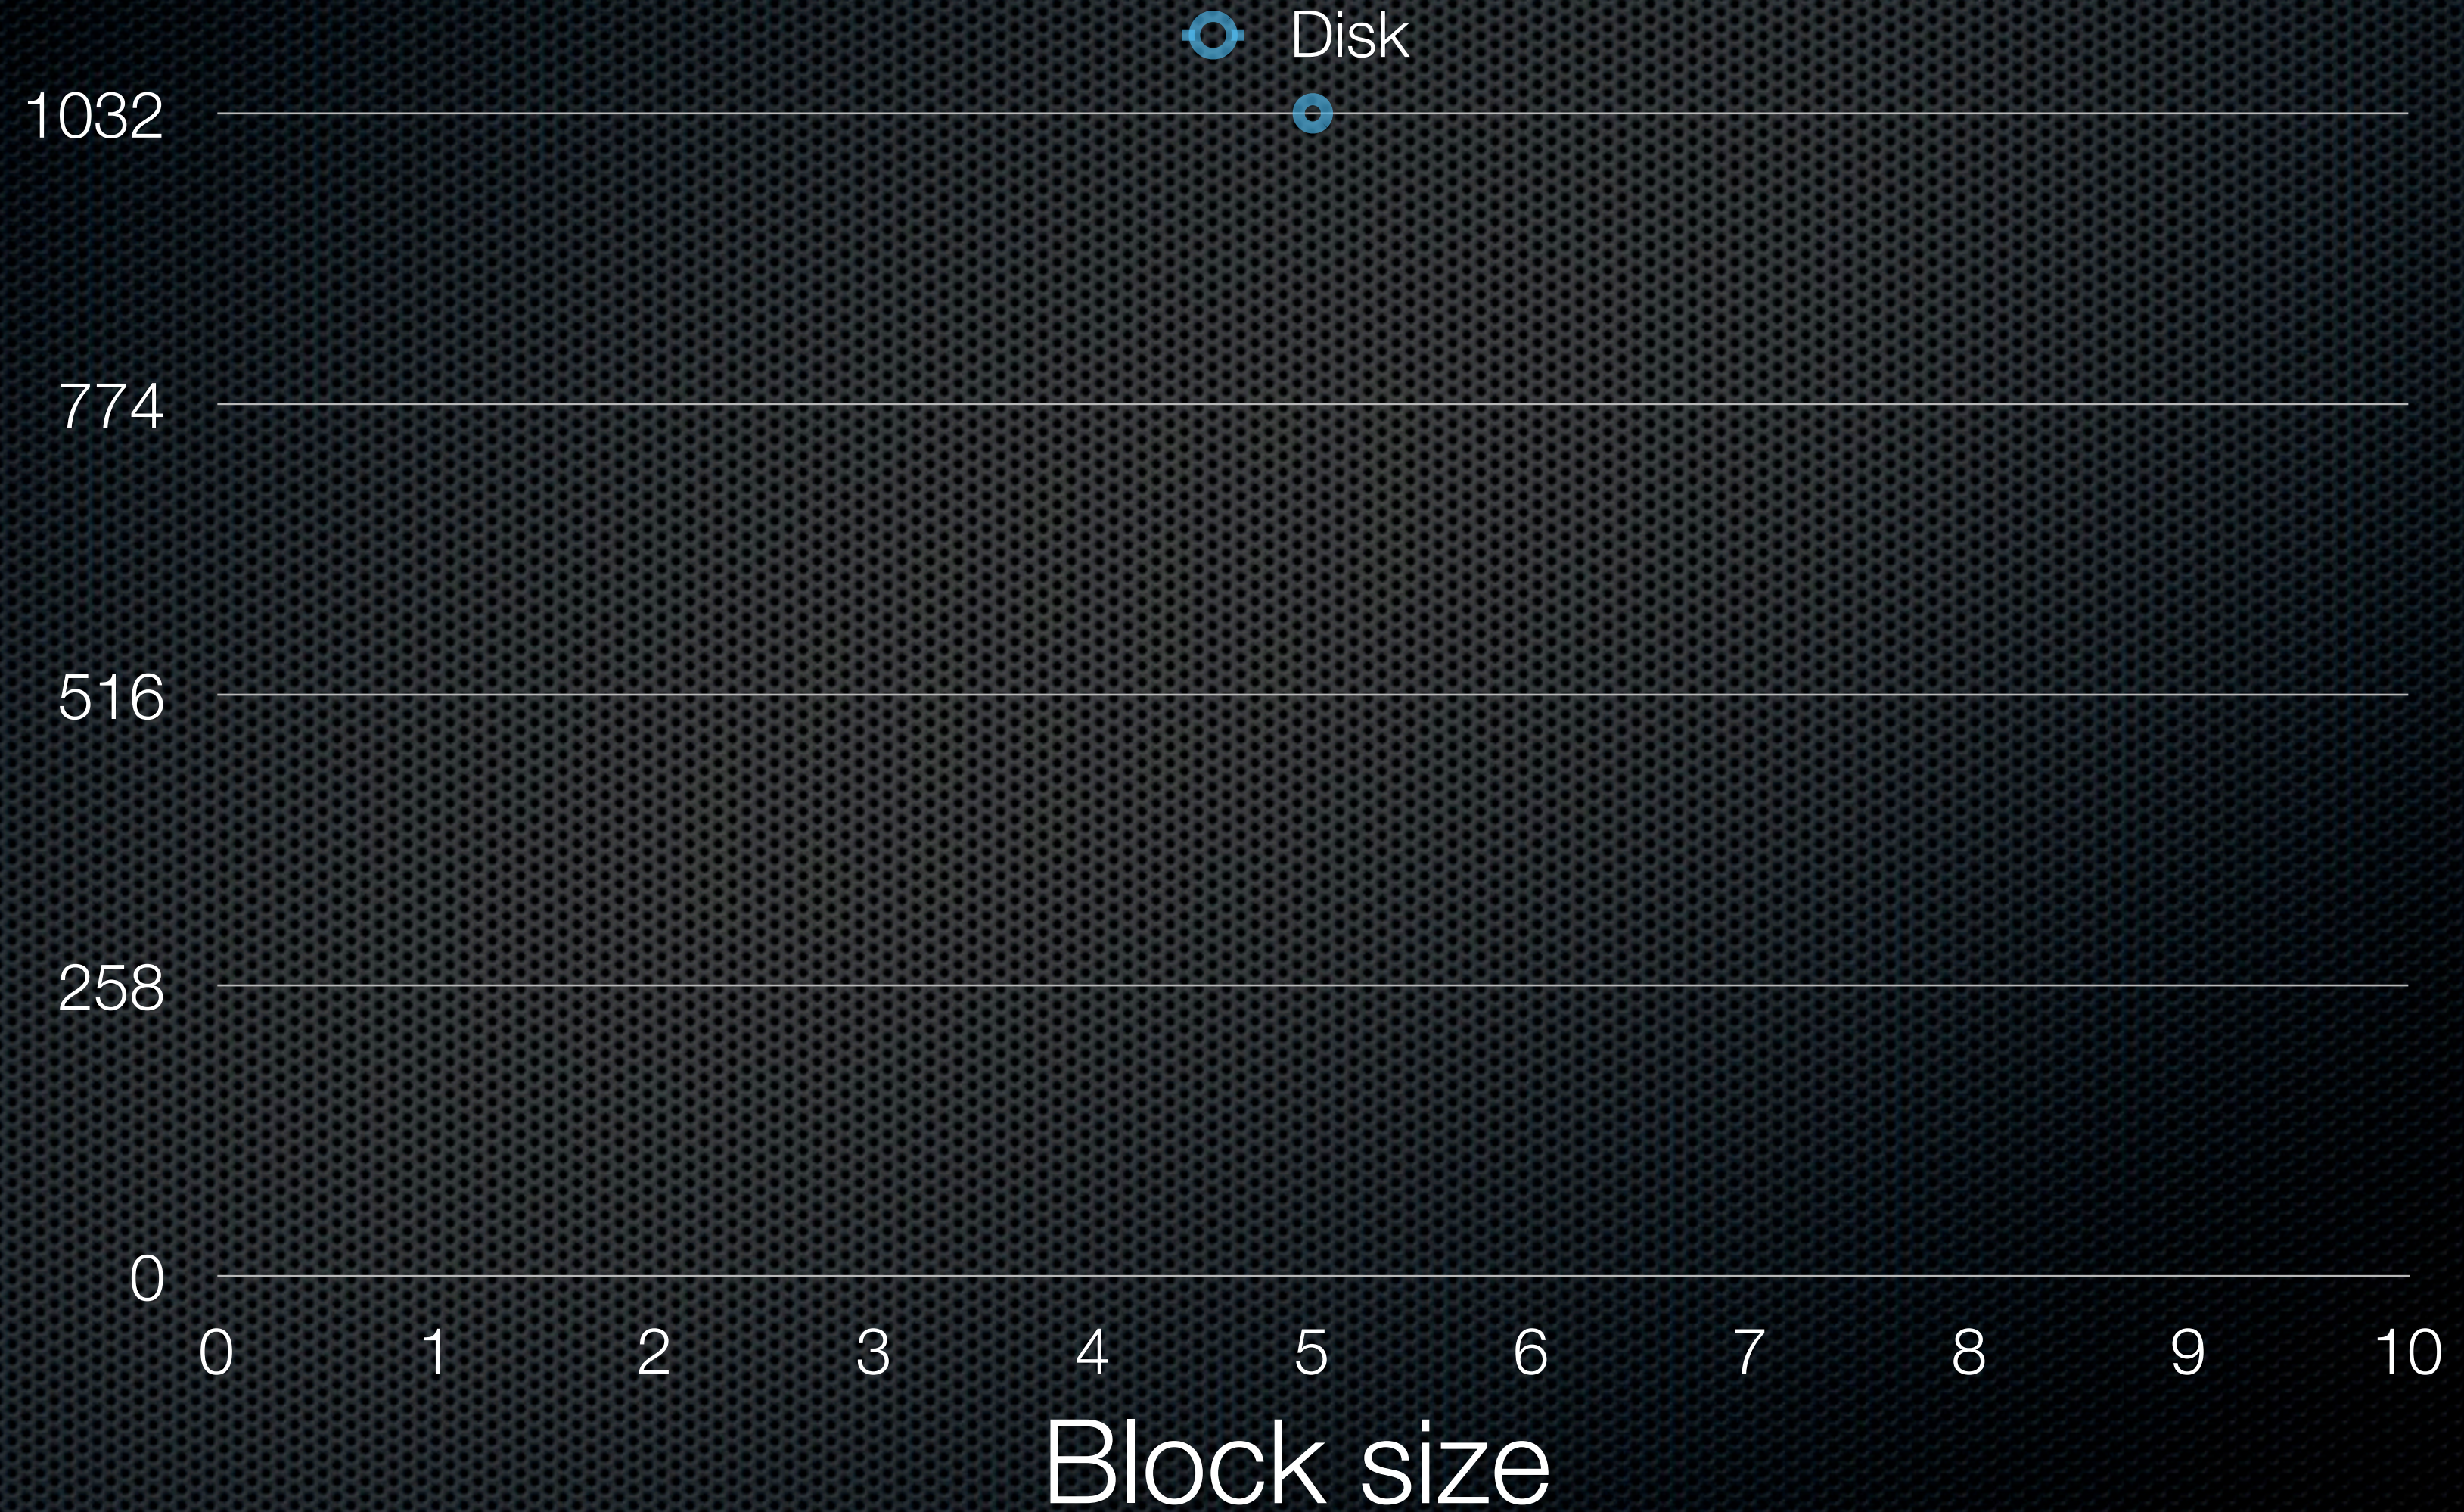
Demo



Bob Clapp

SEP Meeting 2009

Disk transpose: Out of core



Outline

Motivation

Unique hardware


Algorithmic changes

Programming challenges

Demo



Disk transpose: Out-of-core ECOram



Outline

Motivation

Unique hardware

Algorithmic changes

Programming challenges


Demo



Bob Clapp

SEP Meeting 2009

Disk transpose: Comparison



Outline

Motivation

Unique hardware

Algorithmic changes

Programming challenges

Demo



Bob Clapp

SEP Meeting 2009

Challenges

Outline

Motivation

Unique hardware

Algorithmic changes

Programming challenges

Demo

- ✦ Figuring out the optimal way to program varied hardware is asking too much to most researchers
- ✦ Algorithms tend to share a set of core routines
 - ✦ Explicit finite difference
 - ✦ FFT
 - ✦ Sparse matrix



PPL: Pervasive Parallel Lab

Outline

Motivation

Unique hardware

Algorithmic changes

Programming challenges

Demo

- ✦ Working on domain specific constructs
- ✦ Hide the underlying complexity from the end user by building wrapper routines that are able to optimally perform a given task on a given hardware
- ✦ Partnering with CEES for Geoscience applications



Bob Clapp

SEP Meeting 2009

Programming paradigms: Streaming

Outline

Motivation

Unique hardware

Algorithmic changes

Programming challenges

Demo

- ✦ Useful paradigm for high op count per sample operations
- ✦ Working with Nvidia and Convey in attempt to see if this programming paradigm can be instituted on their hardware



Conclusions

- ✦ Many new hardware platforms available
- ✦ All are solving different portions of the data transfer problem
- ✦ Require new thought on types of algorithms and programming styles



Relative Cost of Imaging Algorithms

Outline

Motivation

Unique hardware

Algorithmic changes

Programming challenges

Demo



Contact Information



Spancion Sales Contacts:

- **Jeff Shiner**
 - 972-234-6687 (office)
 - 214-707-1944 (cell)
- **Tom Mailhot**
 - 408-616-5384 (office)
 - 408-221-0578 (cell)
- **Website:** www.spancion.com/ecoram

Appro Sales Contact:

- **Anthony Kenisky, 512-751-2713 or akensisky@appro.com**
 - Anthony Kenisky, VP of Sales
 - Office: 800-513-3338 X207
 - Cell: 512-751-2713
- **Website:** www.appro.com

Relative Cost of Imaging Algorithms

Outline

Motivation

Unique hardware

Algorithmic changes

Programming challenges

Demo

Relative Cost of Imaging Algorithms

Outline

Motivation

Unique hardware

Algorithmic changes

Programming challenges

Demo



RTM with different boundary conditions

Outline

Motivation

RTM

Time
Reversal

Saving
Boundaries

Random
Boundaries

Example

