High Performance Computing Modernization Program (HPCMP) Summer 2011 Puerto Rico Workshop on Intermediate Parallel Programming & Cluster Computing

in conjunction with the National Computational Science Institute (NCSI)/ SC11 Conference

SHODOR

EARLHAM COLLEGE

Jointly hosted at Polytechnic U of Puerto Rico and U Oklahoma

and available live via videoconferencing (streaming video recordings coming soon)

HAMPTON UNIVERSITY

Widener University

Sponsored by DOD HPCMP, SC11/ACM, NCSI and OK EPSCoR

Intermediate Parallel Programming & Cluster Computing

Storage Hierarchy Review

SHODOR

EARLHAM COLLEGE

Sponsored by
DOD HPCMP,
SC11/ACM,
NCSI and
OK EPSCoR

Josh Alexander, University of Oklahoma
Ivan Babic, Earlham College
Ken Gamradt, South Dakota State University
Andrew Fitz Gibbon, Amazon.com
Mobeen Ludin, Earlham College
Tom Murphy, Contra Costa College
Henry Neeman, University of Oklahoma
Charlie Peck, Earlham College
Stephen Providence, Hampton University
Jeff Rufinus, Widener University
Luis Vicente, Polytechnic University of Puerto Rico
Aaron Weeden, Earlham College

Sunday July 31 – Saturday August 6 2011

Widener University

Since it is a review, lets review

When it comes to Intermediate Parallel Programming,

- What is the Storage Hierarchy?
- What is important relative to parallel programming?

What is the Storage Hierarchy?

- Progression from fast, small, expensive storage (registers) to slow, big, cheap storage (disk/network)
  - register (~350GB/sec)
  - -> Level 1 cache (~250GB/sec) (~ 70% of register)
  - -> Level n cache (balances effect of Level 1 cache and RAM)
  - -> RAM (~10GB/sec) (~ 3% of register)
  - -> Disk/network (~ 100MB/sec) (~ .03% of register)
- Parallel architectures -> complicated storage hierarchies
  - Shared memory and thread local
  - Memory local to each computational node
  - GPU Register, Local, Shared, Global, Constant, Texture

What is important relative to parallel programming?

- Performance comes from overlapping CPU and IO
- Storage is advancing at slower rates than CPUs
  effective data movement is more and more important
- Problem and algorithm help inform what to do
- Worth doing if can solve problems otherwise unreachable

Getting a drink
An Analogy

- I live in East Bay about 35 miles from San Francisco
- RAM is walking 50 feet to the refrigerator
- Lev 1 cache is walking 100 feet to the garage
- RAM is driving 1/3 mi to the grocery store
- Disk/Network is driving to San Francisco

Thanks for your attention!

Questions?