synergy.cs.vt.edu

HPCwire January 10, 2012 Virginia Tech Unveils Energy-Efficient "HokieSpeed" Supercomputer

www.cs.vt.edu

People behind HokieSpeed

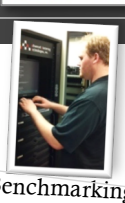

Benchmarking GPUs in a compute cluster

System 0.10 "Supercomputer"

Demonstrating resilient network connectivity

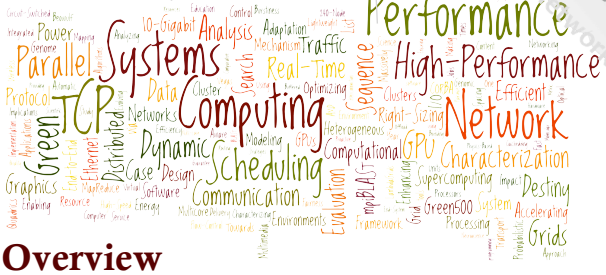

Power metering

Helping kids explore computer science

Making sense of 24TB of gene sequence data

Exploring AMD Fusion

Synergy & friends @ War Spur hiking trail

Overview

At the Synergy Lab, we conduct basic and applied research that explores a breadth of complementary intellectual activities that synergistically span all aspects of parallel computing – from systems software & middleware to tools & applications – in order to accelerate discovery and innovation.

Highlights

Best Paper Finalist at ACM/SPEC ICPE

Ph.D. student, Balaji Subramaniam, has his paper on energy-proportional computing selected as a *Best Paper Finalist* at the *4th ACM/SPEC Int'l Conf. on Performance Engineering*, 2012.

Virginia Tech unveils HokieSpeed at #11 on the Green500 and #96 on the TOP500 in November 2011

HokieSpeed debuted as the *highest-ranked commodity supercomputer in the U.S.* on the Green500, topping at 928.96 MFLOPS/watt.

Paper Selected Among Best in 20 Years of IEEE/ACM HPDC (1992 – 2012)

Lin, Ma, Archuleta, Feng, Gardner, Zhang, "MOON: MapReduce On Opportunistic eNvironments," *19th ACM Int'l Symp. on High-Performance Distributed Computing*, 2010.

Green500 Pioneer to Accelerate Mini-Drone Project

The Synergy Lab, along with collaborators at VT and NCSU, will use GPUs to accelerate the CFD simulation

NVIDIA Graduate Fellowship for GPU Computing

Ph.D. student, Ashwin Aji, received one of 12 NVIDIA Graduate Fellowships awarded worldwide for his work on parallelization methodologies and performance models for large-scale heterogeneous systems.

Research Areas

Systems

Systems software and tools; adaptive run-time systems, parallel computing, monitoring and measurement, emerging platforms in all forms and scales.

Networking

Architecture, performance measurement and evaluation, modeling, protocols, routing & fundamentals and network security.

Renaissance

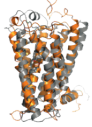

Accelerating BIG DATA apps, rational drug design, large-scale genomic search, facilitating CS education, human computer interaction and for cancer research.

NVIDIA CUDA RESEARCH CENTER

COMPUTE THE CURE

Green Computing

Power modeling and prediction, hardware and software controlled power management and benchmarking on platforms ranging from mobile processors and DSPs to supercomputers.

SUPERCOMPUTING in SMALL SPACES

THE GREEN 500

Research Publications

Visit the Pubs of the Synergy Lab
http://synergy.cs.vt.edu/publications.php

Those who fund our research

NSF
AIR FORCE OFFICE OF SCIENTIFIC RESEARCH
DEPARTMENT OF ENERGY
AIR FORCE RESEARCH LABORATORY
DEPARTMENT OF DEFENSE
HARRIS
AMD
CISCO
NATIONAL INSTITUTES OF HEALTH
CHREC NSF Center for High-Performance Reconfigurable Computing
IBM
NVIDIA
SUPERMICRO
Microsoft
VirginiaTech
JUNIPER NETWORKS
XILINX

Supercomputing in Small Spaces (SSS)

In early 2002, the SSS project unveiled its 240-node "Green Destiny" cluster, which consumed 3.2 kilowatts of power (i.e., two hairdryers) and delivered over 100 Gflops on the LINPACK benchmark, which would have placed it at #393 on the TOP500 List. Due to the audacity of Green Destiny, it ended up being featured in The New York Times, CNN, and PC World to name a few and currently resides in the Computer History Museum. Current research seeks to deliver green supercomputing via advances in architecture, algorithms, modeling, and adaptive software.

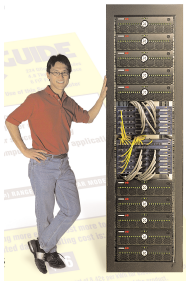

Green500

The Green500 provides rankings of the most energy-efficient supercomputers in the world. We raise awareness about power consumption, promote alternative total cost of ownership performance metrics, and ensure that supercomputers only simulate climate change and not create it.

Heterogeneous Parallel Computing, e.g., CPUs+GPUs in HokieSpeed

Recent trends have exposed the CPU as a "jack of all (computing) trades, master of none." To address this, heterogeneous computing systems with multiple types of brains have emerged to herald a new age in supercomputing. Building on our expertise, at Virginia Tech, we address a myriad of aspects in accelerator-based parallel computing from systems software to middleware and libraries to applications. (Much of this work grew out from the NSF Center for High-Performance Reconfigurable Computing or CHREC.)

mpiBLAST

mpiBLAST is a freely available, open-source, parallel implementation of NCBI BLAST. By efficiently utilizing distributed computational resources through database fragmentation, query segmentation, intelligent scheduling, and parallel I/O, mpiBLAST improves NCBI BLAST performance by several orders of magnitude while scaling to hundreds of processors. mpiBLAST is also portable across many different

MyVICE

This project seeks to leverage virtual environments to introduce the power of (parallel) computing to the classroom at an early age – first via LOGO for grades 1 – 2 and then transitioning to "programming with pictures" via Storytelling Alice and Scratch for grades 3 – 8.

CUDA Research Center

NVIDIA has named Virginia Tech (VT) as a CUDA Research Center, recognizing the university's research and development of general-purpose computing on graphics processing units (GPUs) across many disciplines. Such research centers are "at the forefront of some of the world's most innovative and important scientific research," according to the company. Prof. Feng heads the CUDA Research Center at VT.

Follow Us in the News …

New York Times: http://nyti.ms/X1DWDs
HPCwire: http://bit.ly/ZyX8MM
Wired: http://bit.ly/ZyXd 34
WVTF NPR: http://bit.ly/ZyXuTD

BIG DATA Analytics

Data is being generated at a rate that outstrips our ability to compute on that data. For example, DNA sequencers generate *10-fold* more data every 18 months while our compute capability only *doubles* every 18 months. This research aims to co-design across multiple layers (hardware, architecture, algorithms, and software) to bridge the above gap across disciplines.

Cloud Computing

With the simultaneous need for high-performance computing (HPC) and its associated expense, cloud computing offers an alternative means to access HPC, one without its exorbitant cost but at the expense of less control and arguably less performance. This work seeks to create a new generation of scalable analysis & management software for "client+cloud" environments to deliver HPC for BIG DATA and scientific computing.

Networking

Our aim is to develop models, architectures and protocols for next-generation networks. We have focused on performance evaluations of existing systems. In doing so we have implemented optimized systems and protocols as proof of concepts.

CHREC
NSF Center for High-Performance Reconfigurable Computing

The NSF Center for High-Performance Reconfigurable Computing consists of more than 30 institutions from academia, industry, and government with synergistic interests and goals in reconfigurable and adaptive computing for a broad range of missions – from satellites to supercomputers. The Synergy Lab is one of the two labs that lead CHREC activities at the Virginia Tech site.

Those who bring synergy to the lab

Faculty & Staff


W. Feng


M. Gardner


H. Lin

Students


Aniket


Ashwin


Balaji


Brian


Carlo


Jing

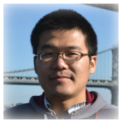
Kaixi


Konstantinos


Lokendra


Mariam


Marwa


Nabeel


Paul


Sriram


Tom


Umar


Vignesh