Third Annual Graduate Student Poster Exhibition

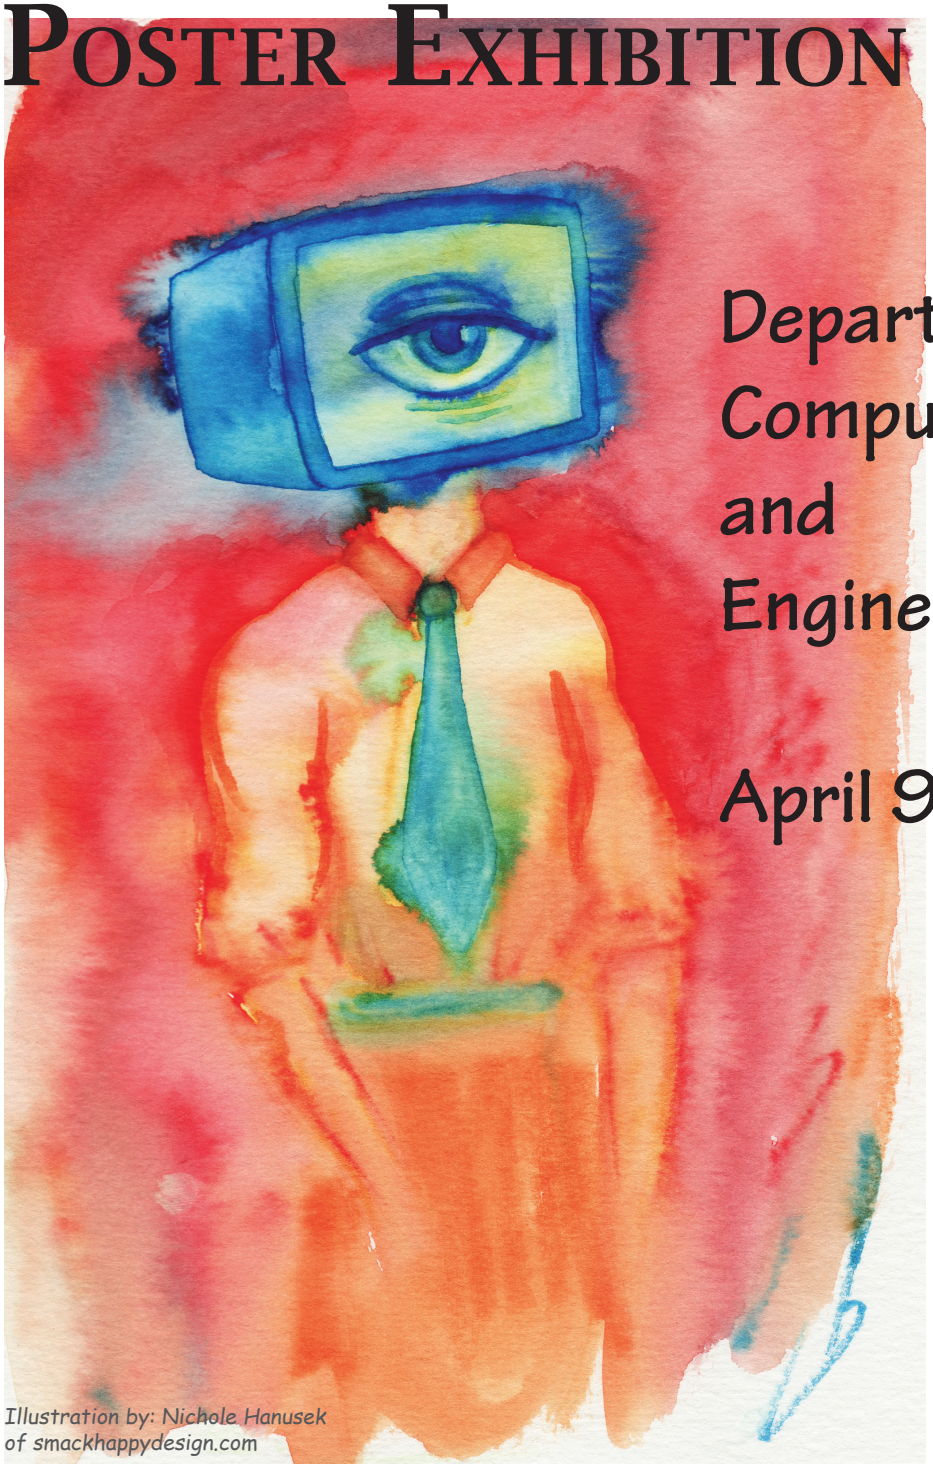

Illustration by: Nichole Hanusek
of smackhappydesign.com

Department of Computer Science and Engineering

April 9, 2009

WWW.CSE.OHIO-STATE.EDU

Welcome!

I would like to warmly welcome you to our Third Annual Graduate Student Research Poster Exhibition. When this program started, I believed it would be a great success and experience has proved this true.

Our students continue working hard and excelling in their areas of research. Please enjoy this opportunity to learn of their investigations. Talk with them, query, and challenge them; this is a learning and practice opportunity for them as well.

Thank you for attending and we look forward to seeing you next year.

Xiaodong Zhang
Robert M. Critchfield Professor
CSE Department Chairman

2009 Cojmputer Science & Engineering Graduate Student Research Poster Exhibitioners

\# Position in the Gallery - Name of Presenter
Year in PhD Program
Presenter's Advisor(s)
Poster Title
Abstract of Research Presented

1. - **Xiaoning Ding**
3rd Year
Xiaodong Zhang
BP-Wrapper: A System Framework Making Any Replacement Algorithms (Almost) Lock Contention Free
In high-end database systems, the execution concurrency level rises continuously. A new challenge for buffer management is to retain its scalability in responding to the highly concurrent data processing demands.
The page replacement algorithm, a major component in the buffer management, can seriously degrade the system's performance if the algorithm is not implemented in a scalable way. A lock-protected data structure is used in most replacement algorithms, where high contention is caused by concurrent accesses. This problem may not exist or can be tolerated in an environment of low concurrency, thus has not been given enough attention for a long time. We propose a system framework, called BP-Wrapper, that (almost) eliminates lock contention for any replacement algorithm without requiring any changes to the algorithm. In BP-Wrapper, we use batching and prefetching techniques to reduce lock contention and to retain high hit ratio. The implementation of BP-Wrapper in PostgreSQL version 8.2 adds only about 300 lines of C code. It can increase the throughput up to two folds compared with the replacement algorithms with lock contention when running TPC-C-like and TPC-W-like workloads.

2. - **Lifeng Sang**
5th Year
Anish Arora
Dialog Codes for Secure Wireless Communications
We investigate the feasibility of achieving perfect secrecy in wireless network communications without shared secrets. We introduce a secure coding problem in which not only the sender but also the receiver participates in the coding. In essence, the receiver's role is to selectively jam the sender's transmission at the level of bits, bytes, or packets. We then design a class of secure codes, which we call dialog codes, for diverse channel models and receiver models. Our codes are simple and efficient, with only O(1) complexity in both the encoding and the decoding process, and achieve optimal coding rate in some channel models.

3. - **Youding Zhu**
6th Year
Rick Parent
Online Transfer of Human Motion to Humanoids
Motion retargeting problem has been well studied and several off-line solutions exist based on optimization approaches that rely on pre-recorded human motion data collected from a marker-based motion capture system. From the perspective of human robot interaction, there is a growing interest in online motion transfer, particularly without using markers. Such requirements have placed stringent demands on retargeting algorithms and limited the potential use of off-line and pre-recorded methods. To address these limitations, we present an online task space control theoretic retargeting formulation to generate robot joint motions that adhere to the robot's joint limit constraints and self-collision constraints. Motion retargeting problem has been well studied and several off-line solutions exist based on optimization approaches that rely on pre-recorded human motion data collected from a marker-based motion capture system. From the perspective of human robot interaction, there is a growing interest in online motion transfer, particularly without using markers. Such requirements have placed stringent demands on retargeting algorithms and limited the potential use of off-line and pre-recorded methods. To address these limitations, we present an online task space control theoretic retargeting formulation to generate robot joint motions that adhere to the robot's joint limit constraints and self-collision constraints.

4. - **Ai Chen**
5th Year
Ten-Hwang (Steve) Lai
Measuring and Guaranteeing Quality of Barrier-Coverage in Wireless Sensor Networks
Sensors may fail due to various reasons. As more and more sensors fail, certain desired properties such as barrier coverage will diminish and eventually fall below a desired level. In such a case, the network will have to be repaired. It is therefore desirable to have mechanisms to monitor network properties. In this paper, we are interested in measuring the quality of barrier coverage. If the measured quality is short of a desired value, we further identify all local regions that need to be repaired. We also discuss how to actually repair a region.

5. - **Lei Ding**
5th Year
Mikhail (Misha) Belkin
Object Segmentation and Recognition from Images
This poster addresses image segmentation and recognition at the object level, and outlines a system architecture that enables synergies between semi-automatic segmentation and recognition. In particular, object segmentation is

achieved by transductive learning on hypergraphic image models with user supplied strokes. The object recognition component utilizes object-class models learned from labeled image databases, and can be enhanced by exploring the contextual constraints among objects. We show preliminary results demonstrating the usefulness of the proposed system.

6. - **Jason Kirschenbaum**
5th Year
Bruce Weide
Investigations in Tuning Proof Assistants for Program Verification
Tool advances in the theorem proving community have help create interactive proof assistants that can perform many reasoning steps automatically with guidance from users for help with subtle parts of a proof. We focus on the problem of proving verification conditions (VCs) generated from practical programs, fully automatically, using a nominally interactive proof assistant. More specifically, we show how an interactive proof assistant can be used to prove practical VCs by verifying extensions of abstract data types including Stacks, Sets, and Queues.

7. - **Bruce M. Adcock**
4th Year
Bruce Weide
The End of Debugging as We Know It
Using automatically generated verification conditions (VCs) from Resolve programs, we pass these on to one or more automated theorem provers. An overarching goal is to fully verify that the code meets the specifications given. This changes the development process used by the programmer; testing has become superfluous, and likewise standard debugging is unnecessary. This work focuses on how to use the results of trying to verify code to better aid the programmer. We examine how both the VCs and the theorem provers themselves allow the programmer to understand where their code is incorrect.

8. - **Na Li**
5th Year
David Lee
Malicious Node Conviction in Untrustworthy Distributed Network
Based on the information of network node accusing and testifying operations with a confidence measure of the assessment, we design algorithms to convict malicious nodes in an untrustworthy distributed network. We formalize the conviction problem into general mathematical models and present iterative procedures for network node assessment.

9. - **Matt Boggus**
3rd Year
Roger Crawfis
Procedural Creation of Three Dimensional Caves
Creating a three dimensional model of a cave is expensive. Spatial data obtained from surveys is extremely coarse and scanning an entire cave is a time consuming process that requires specialized equipment. We propose alternative methods to create three dimensional models of caves (for use in visualization and entertainment applications) by taking advantage of domain knowledge from Speleology, which is the study of caves. During this research we developed a spatial data structure suitable for editing a cave interactively along with operations to procedurally create models according to individual passage type and generalized cave patterns.

10. - **Kelly Yackovich**
3rd Year
Jay Ramanathan
Dynamic CitiScapes Architecture for Service Composition within Complex Systems
Based on our case study with a large US city, we have identified challenges with handling non-routine service requests that prevent the City from moving rapidly toward its eGovernment strategic goals. We address how the CitiScapes architecture can support the City's most crucial, challenging and strategic needs: streamlining the response and integration of different aspects of service requests in combination emergency and non-emergency scenarios. This support consists of three major aspects: meeting the need for human-supported dynamic service assembly in non-routine scenarios, enabling new methods of collaboration such as mobile device and GIS coordination, and enabling performance tracking, traceability, and transparency into business processes.

11. - **Qi Gao**
5th Year
Feng Qin
First-Aid: Surviving and Preventing Memory Management Bugs during Production Runs
Memory bugs in C/C++ programs severely affect system availability and security. This poster describes First-Aid, a lightweight runtime system that survives software failures caused by common memory management bugs and prevents future failures by the same bugs during production runs. We implemented it on Linux and our evaluation with real-world memory bugs shows that First-Aid is effective to common memory bugs and only incurs very low overhead.

12. - **Xiaole Bai**
5th Year
Dong Xuan
Authentication in Malicious Environments
Authentication is a key issue in the field of security. This poster presents the works on daily human computer authentication, computer human interactive proof, and peer authentication issues at Wireless Sensor Networks. These issues are studied in particular when possible attacks are considered.

13. - **Matthew Koop**
5th Year
D.K. Panda
Network-Aware Messaging for Ultrascale Supercomputers
Supercomputers of today are designed with many separate computing elements. The numbers of these elements can range from tens of thousands or even hundreds of thousands. In order to harness the power of all of these disparate elements and solve a single problem, they need to be connected through high-performance networks and able to send messages to each other. To achieve the best performance we design an adaptive hybrid messaging approach using different network hardware features. Evaluation has been performed on 4096 cores and shown to have up to 40% better performance over traditional designs.

14. - **Gopalakrishnan Santhanaraman**
7th Year
D.K. Panda
Designing High Performance and Scalable One-Sided Communication Middleware on Modern Interconnects
High-end computing (HEC) systems are enabling scientists and engineers to tackle grand challenge problems in their respective domains and make significant contributions to their fields. As these systems continue to grow in scale, the performance that applications can achieve depends heavily on their ability to avoid explicitly synchronized communication with other processes in the system as well as the capability to overlap computation and communication. One-sided programming models are gaining importance and popularity as they have provide many programming constructs that enable implicit communication using one-sided communication operations. At the same time modern interconnects like InfiniBand provides a rich set of network primitives like Remote Direct Memory Operations (RDMA), Remote Atomics, Scatter/Gather capabilities etc. My work explores designs a High Performance and Scalable One sided Communication subsystem in MPI by effectively harnessing the capabilities of modern Interconnects.

15. - **Zizhan Zheng**
4th Year
Prasun Sinha
Alpha Coverage: Bounding the Interconnection Gap for Vehicular Internet Access
Internet access via roadside WLAN access points (AP) has been demonstrated to be a feasible solution to provide opportunistic data service to moving vehicles. Using an in situ deployment, however, such a solution does not provide worst-case performance guarantees due to unpredictable intermittent connectivity. On the other hand, a full coverage solution is not very practical due to the prohibitive deployment and operational cost. We introduce a new notion of intermittent coverage for mobile users, called alpha-coverage, which provides worst-case guarantees on the interconnection gap while using significantly fewer APs than full coverage.

16. - **Ren-Shiou Liu**
4th Year
Prasun Sinha
Distributed Routing and Rate Control in Perpetual Sensor Networks
Various forms of energy, such as solar, wind, and thermal,can be harvested from the environment to extend the lifetime of sensor networks. A network based on energy harvesting can operate perpetually, by using its energy judiciously. Several applications based on data collection require high and fair data rate from all nodes in the network. However, dissimilar recharging rates of nodes and their time-varying nature pose new challenges. This poster presents a distributed and adaptive algorithm that computes the rate of data collection for each node and routes for the packets, with the goal of optimizing a network utility function.

17. - **Benjamin Schroeder**
4th Year
Rick Parent with Marc Ainger of the OSU School of Music
Real-time Sounding Plates for Rigid Body Simulations
Sound acts together with visual information to produce a whole impression of a scene. We present a technique that combines a finite difference-based model of sounding plates and a standard rigid body simulation to produce coupled video and audio. Our technique uses a dynamic simulation of contact, and reproduces attack sound variation due to material properties as well as damping due to resting contact between bodies. In addition, vibration information from the audio simulation is used to produce secondary motion of rigid bodies. The simulation runs at interactive rates and may be reconfigured on the fly.

18. - **Kaushik Sinha**
6th Year
Mikhail Belkin
Semi-Supervised Learning with Sparse Eigenfunction Bases
We present a new framework for semi-supervised learning involving sparse eigenfunction bases of data-dependent convolution operators. We show that 1) when the \ emph{cluster assumption} holds, i.e., high density regions are sufficiently separated by a low density valley, each high density area corresponds to a unique representative eigenfunction. Linear combination of such eigenfunctions form good classification functions. 2) By minimizing a L_{1} penalized convex loss function a sparse classifier involving these representative eigenfunctions can be consistently recovered using only a few labeled examples. Experimental results on a number of real-world data-sets, shows that our method is competitive with the state of the art.

19. - **Teng-Yok Lee**

4th year

Han-Wei Shen

Visualizing Time-Varying Features with TAC-based Distance Field

To analyze time-varying data sets, tracking features over time is often necessary to understand the underlying physical process. Here we propose a new framework to visualize time-varying features and their motion. We model a time-varying feature as a time series or Time Activity Curve (TAC), and use the Dynamic Time Warping (DTW) distance metric to compare the time series on each data point to the given feature and to estimate when the feature occurs.

Based on the derived distance field and the estimate of time, several visualizations can be derived to highlight the spatiotemporal properties of the feature.

20. - **Feng Chen**

3rd Year

Xiaodong Zhang

Hystor: A Hybrid Storage System Delivering SSD Performance at HDD Cost

With fast technical improvement of flash memory based Solid State Drives (SSDs), the storage performance is expected to be significantly improved. However, due to relatively high prices, low capacity, and limited lifecycle, SSDs will not entirely replace hard disk drives (HDDs) but will play a critical supporting role in large storage systems. A major challenge is how to efficiently integrate high-speed SSDs into existing HDD-based storage systems to organize a cost-effective high performance hybrid storage system. In this poster, we present the design and implementation of such a system, called Hystor, which achieves a set of optimization objectives from both system deployment and algorithm design perspectives.

21. - **Jonathan Eisenmann**

3rd Year

Rick Parent with Matt Lewis of ACCAD.

Interactive Evolutionary Design for Motion Variety

We present an intuitive method for novice users to interactively design custom populations of coherent, heterogeneous motion for crowd animation, from one input motion clip, thus amplifying an existing database of possible motion. User-created deformers are used by a genetic algorithm to manipulate animation channels, creating new variants on the input motion. Our design environment allows the user to traverse the space of possible motions, presents the user with populations of motion, and gradually converges to a desirable set of solutions. Each generated motion sequence undergoes a motion filtering process subject to user-specified, high-level metrics to produce a convincing result.

22. - **Issam Safa & Chunajiang Luo**

4th year & 2nd year

Yusu Wang

Gradiant Estimation in Eigen Space

In this work we propose to estimate the gradient of a real scalar function f in eigen space. More concretely given a compact 2-manifold M in R^3, and a real scalar function f, we use the Laplace-Beltrami operator to embed the manifold in an eigen space defined by the operator's spectral decomposition. We then proceed to show that gradient computation in eigen space preserve the critical points of f, as well as establishing a relation between the eigen gradient and the real gradient. The new results are then applied to the computation and smoothing of jacobi sets of M.

23. - **Mark Keck**

6th Year

James Davis

On Finding Static Occluding Structures with Few Views

We present a novel method for extracting occluding structures from video when the number of distinct views (cameras) is small.

Occlusions are inferred from a set of per-voxel, per-view probabilistic models, encoding how likely it is that a particular 3D area (voxel) is occluded from each view. The major contributions of the work are (1) a procedure for learning static occlusions that incorporates temporal and shape-based constraints, and (2) a novel method for recovering 3D structure from learned occlusion areas.

24. - **Jeremy Morris**

5th Year

Eric Fosler-Lussier

Automatic Speech Recognition using Conditional Random Fields

The current state-of-the-art paradigm for automatic speech recognition (ASR) makes use of statistical models - specifically statistical sequence models - to create a probabilistic model for spoken language.

Typically the statistical model chosen for this task is a Hidden Markov Model (HMM) - a model that uses a top-down (generative) approach to recognizing speech. This work examines the idea of replacing the HMM with a bottom-up (discriminative) statistical sequence model known as a Conditional Random Field (CRF) - a model that has attributes that make it more attractive than HMMs for processing speech sequences.

25. - **Zhimin Yang**

4th Year

Dong Xuan

E-SmallTalker: A System for Mobile Social Networking in Physical Proximity

Social gap prevents people from fully utilizing face-to-face interaction in physical proximity. To shorten this gap, we propose E-SmallTalker, a mobile phone-based purely distributed communication system that encourages people to

communicate with each other. E-SmallTalker automatically generates short, user-controlled "small-talk" messages based on their common interests, friends, and experiences, and prompts results to its owner for further interaction. E-SmallTalker protects users' privacy by encoding information with Bloom filters. It also achieves user transparency by exchange messages without building a Bluetooth connection. Our experiments on real-world Bluetooth-enabled phones show E-SmallTalker is promising for facilitating face-to-face social interactions in physical proximity.

26. - **Venu Satuluri**
3rd Year
Srinivasan Parthasarathy
Scalable Graph Clustering Using Flows
We present a multi-level algorithm for clustering graphs that simulates probability flows on the graph to tease out its cluster structure. The graph is first successively coarsened to a manageable size, and a small number of iterations of flow simulation is performed on the coarse graph. The graph is then successively refined, with flows from the previous graph used as initializations for brief flow simulations on each of the intermediate graphs.
At the end the high-flow regions are clustered together, with regions without any flow forming the natural boundaries of the clusters. We show promising experimental results.

27. - Shansi Ren & Enhua Tan
3rd Years (Shansi Ren graduated WI 2009)
Xiaodong Zhang
Design, Implementation, and Evaluation of a Topology-Aware and Infrastructure-Independent BitTorrent Client
BitTorrent (BT) has carried out a significant and continuously increasing portion of Internet traffic. We present our design, implementation, and evaluation of a topology-aware BT system, called TopBT, to significantly improve the Internet resource utilization without degrading user downloading performance. The unique feature of TopBT client lies in that a TopBT client actively discovers network proximities (to connected peers), and uses both proximities and transmission rates to maintain fast downloading while reducing the transmitting distance of the BT traffic and thus the Internet traffic.
We have implemented TopBT based on widely used open-source BT client code base, and made the software publicly available. By deploying TopBT and other BitTorrent clients on hundreds of Internet hosts, we show that on average TopBT can save about 25% download traffic while achieving a 15% faster download speed compared to several prevalent BT clients. TopBT has attracted a large number of downloads for a wide usage across the world. It has also been integrated into a major BT system for testing and a potential adoption in mainstream BT systems.

28. - **Timothy Miller**
4th
Radu Teordorescu
Flexible Redundancy in Robust Processor Architecture
The scaling of transistors in modern microprocessors to minute sizes makes them less predictable and less reliable. To ensure continued growth in chip performance, new solutions must be developed to deal with these challenges in resource-efficient ways. We propose a new reliable processor architecture that dynamically adapts the amount of protection to the characteristics of each individual chip and its runtime behavior. This architecture uses fine-grain redundancy, voltage scaling and timing speculation to adapt to variation and tolerate errors due to heat and particle strikes. The goal is to provide reliability with a minimum amount of resources.

29. - **Mukundan Sridharan**
6th Year
Anish Arora
Virtualization in Sensor Networks
The proliferation of sensing testbeds across the world has given rise to a number of interesting challenges, particularly in the domain of virtualizing, discovering, and federating sensor infrastructures. The problems could be classified into intra-network challenges such as network virtualization, isolation and network monitoring, and extra-network challenges such as architectures/frameworks for uniform resource specification, discovery, and experimentation. In this poster, we propose the KanseiGenie architecture that is compatible with the GENI framework for federating multiple sensing infrastructures. We also look at the specific problem of virtualizing the MAC layer in sensor networks and present some initial solutions.

30. - **Qian Zhu**
5th Year
Gagan Agrawal
Support Time-Critical Event Handling in Distributed Environments
There are many applications where a timely response to an important event is needed. Often such response can require significant computation and possibly communication, and it can be very challenging to complete it within the time-frame the response is needed.
At the same time, there could be application-specific flexibility in the computation that may be desired. The goal is to maximize a benefit function which captures what is most desirable to compute while satisfying the time constraints. In our research, we have studied autonomic parameter adaptation, resource allocation and fault tolerance to support such a goal in time-critical event handling.

31. - **Michael Andereck**
2nd Year
Rick Parent
Interactive Smoothed Particle Hydrodynamics
This project seeks to create an interactive tactile environment for fluid simulation. The original goal for the project involves bringing these simulations into the classroom. Rather than having wet (water), messy (mud), or possibly dangerous (mercury) fluids physically present, we can simulate the fluids using smoothed particle hydrodynamics, and the interactions using a light-sensor touch table.

32. - **Matthew Lang**
5th Year
Paul Sivilotti
Modular Verification of Maximality Properties
A maximal program is not only correct but can generate all behaviors permitted by its specification. We demonstrate that maximality is non- compositional and develop a technique for proving the maximality of composed systems. We illustrate our technique by proving the maximality of a solution to the strong fairness resource-allocation problem.

33. - **Brian Larkins**
5th Year
P. Sadayappan
Global Trees: A Framework for Linked Data Structures on Distributed Memory Parallel Systems
The Global Trees (GT) system provides a multi-layered interface to a global address space view of distributed tree data structures. The Global Trees system utilizes coarse-grained data movement to enhance locality and communication efficiency. The key benefits of using this system include efficient shared-memory style programming of distributed trees, tree-specific optimizations for data access and computation, and the ability to customize many aspects of GT to optimize application performance.

34. - **Muthu Manikandan Baskaran**
5th Year
P. Sadayappan
Automatic Parallelization and Optimizations for Modern Multi-core Architectures
Current trends in computer architecture exemplify the emergence of multiple processor cores on a chip. Significant challenges in such modern multi-core architectures are (1) effective exploitation of available parallelism, and (2) enhancement of memory performance to prevent it from being a bottleneck. It is also necessary to address performance-influencing factors that are specific to an architecture. In our work, we develop techniques for automatic parallelization for multi-core architectures and also develop optimizations that enable efficient execution of programs on multi-core systems. We provide effective automatic compiler support for managing on-chip and off-chip memory accesses. To effectively exploit parallelism in multi-core systems, we develop (1) an effective multi-level tiling approach for mapping computation and (2) an automatic parallelization approach for asynchronous, load-balanced parallel execution. We develop compile-time transformations to address GPGPU architecture-specific performance factors, and develop an end-to-end compiler framework for GPGPUs.

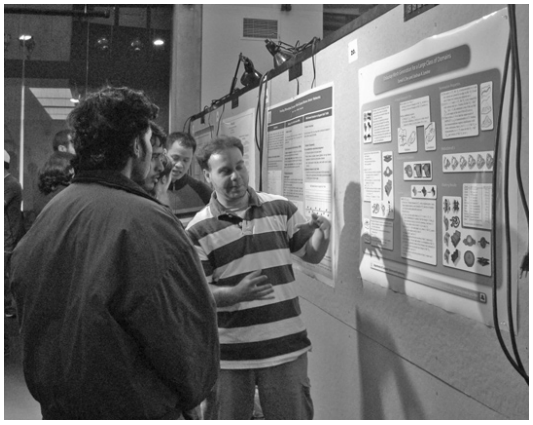

Presenters by Advisors

Gagan Agrawal
30. Qian Zhu

Anish Arora
2. Lifeng Sang
29. Mukundan Sridharan

Mikhail (Misha) Belkin
5. Lei Ding
18. Kaushik Sinha

Roger Crawfis
9. Matt Boggus

James Davis
23. Mark Keck

Eric Fosler-Lussier
24. Jeremy Morris

Ten-Hwang (Steve) Lai
4. Ai Chen

David Lee
8. Na Li

D.K. Panda
13. Matthew Koop
14. Gopalakrishnan Santhanaraman

Rick Parent
3. Youding Zhu
17. Benjamin Schroeder
21. Jonathan Eisenmann
31. Michael Andereck

Srinivasan Parthasarathy
26. Venu Satuluri

Feng Qin
11. Qi Gao

Jay Ramanathan
10. Kelly Yackovich

P. Sadayappan
33. Brian Larkins
34. Muthu Manikandan Baskaran

Han-Wei Shen
19. Teng-Yok Lee

Prasun Sinha
15. Zizhan Zheng
16. Ren-Shiou Liu

Paul Sivilotti
32. Matthew Lang

Radu Teordorescu
28. Timothy Miller

Yusu Wang
22. Issam Safa & Chunajiang Luo

Bruce Weide
6. Jason Kirschenbaum
7. Bruce M. Adcock

Dong Xuan
12. Xiaole Bai
25. Zhimin Yang

Xiaodong Zhang
1. Xiaoning Ding
20. Feng Chen
27. Shansi Ren & Enhua Tan

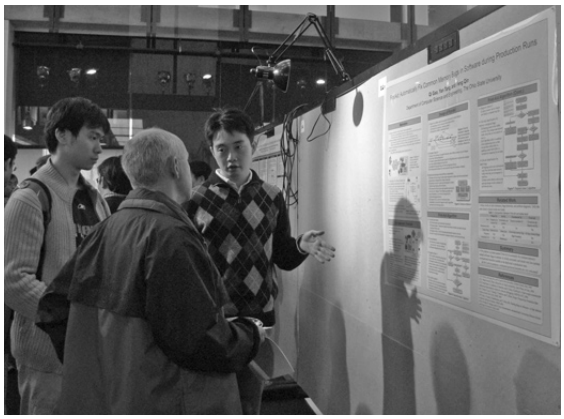

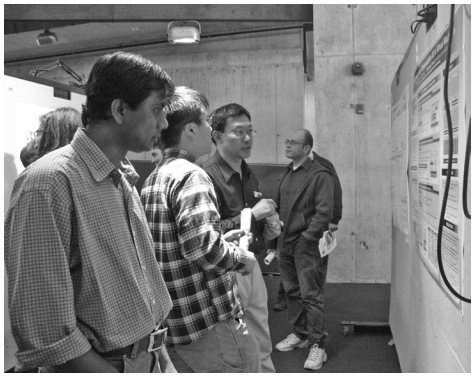

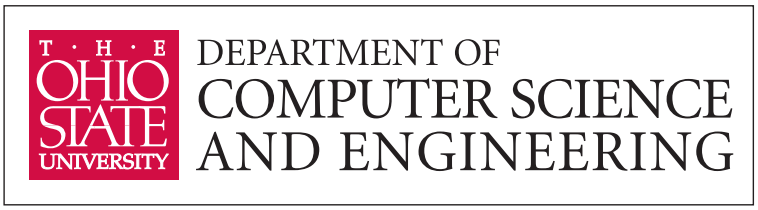
THE OHIO STATE UNIVERSITY
DEPARTMENT OF
COMPUTER SCIENCE
AND ENGINEERING