

**Final Report for Period:** 10/2007 - 09/2008**Submitted on:** 12/14/2008**Principal Investigator:** Schwan, Karsten .**Award ID:** 0326396**Organization:** GA Tech Res Corp - GIT**Submitted By:**

Eisenhauer, Greg - Co-Principal Investigator

**Title:**

ITR: Collaborative Research: Morphable Software Services: Self-Modifying Programs for Distributed Embedded Systems

**Project Participants****Senior Personnel****Name:** Schwan, Karsten**Worked for more than 160 Hours:** Yes**Contribution to Project:****Name:** Pu, Calton**Worked for more than 160 Hours:** Yes**Contribution to Project:****Name:** Pande, Santosh**Worked for more than 160 Hours:** Yes**Contribution to Project:****Name:** Eisenhauer, Greg**Worked for more than 160 Hours:** Yes**Contribution to Project:****Name:** Balch, Tucker**Worked for more than 160 Hours:** Yes**Contribution to Project:****Post-doc****Graduate Student****Name:** Agarwala, Sandip**Worked for more than 160 Hours:** Yes**Contribution to Project:****Name:** Chen, Yuan**Worked for more than 160 Hours:** Yes**Contribution to Project:****Name:** Fryman, Josh**Worked for more than 160 Hours:** Yes**Contribution to Project:****Name:** Ganev, Ivan**Worked for more than 160 Hours:** Yes**Contribution to Project:**

**Name:** Ghosh, Mrinmoy

**Worked for more than 160 Hours:** Yes

**Contribution to Project:**

**Name:** Kong, Jiantao

**Worked for more than 160 Hours:** Yes

**Contribution to Project:**

**Name:** Kumar, Sanjay

**Worked for more than 160 Hours:** Yes

**Contribution to Project:**

**Name:** Nathuji, Ripal

**Worked for more than 160 Hours:** Yes

**Contribution to Project:**

**Name:** O'Hara, Keith

**Worked for more than 160 Hours:** Yes

**Contribution to Project:**

**Name:** Poellabauer, Christian

**Worked for more than 160 Hours:** Yes

**Contribution to Project:**

**Name:** Sundaragopalan, srika

**Worked for more than 160 Hours:** Yes

**Contribution to Project:**

**Name:** Zhang, Kun

**Worked for more than 160 Hours:** Yes

**Contribution to Project:**

**Name:** Zhuang, Xiaotong

**Worked for more than 160 Hours:** Yes

**Contribution to Project:**

#### **Undergraduate Student**

**Name:** Reiss, Charles

**Worked for more than 160 Hours:** Yes

**Contribution to Project:**

REU

#### **Technician, Programmer**

#### **Other Participant**

#### **Research Experience for Undergraduates**

**Organizational Partners****University of Arizona**

Linked project.

**Other Collaborators or Contacts****Activities and Findings****Research and Education Activities:** (See PDF version submitted by PI at the end of the report)**Findings:** (See PDF version submitted by PI at the end of the report)**Training and Development:**

REU support associated with this grant has funded an undergraduate student who has since gained admission to a prominent PhD program.

**Outreach Activities:**

Our robotics faculty and students have been involved in numerous demonstrations of the GNATs and in order to educate the public about robotics research have given demonstrations to a range of audiences (e.g. middle-school and high-school students, high-tech executives).

**Journal Publications****Books or Other One-time Publications**

Kun Zhang, Tao Zhang, and Santosh Pande, "Binary translation to improve energy efficiency through post-pass register re-allocation", (2004). conference publication, Published Collection: Proceedings of the Fourth ACM International Conference on Embedded Software (EMSOFT) Bibliography: pages 74-85. Sept.

Ivan Ganey, Greg Eisenhauer, Karsten Schwan, "Kernel Plugins: When A VM Is Too Much", (2003). conference publication, Published Collection: In Proceedings of the 3rd Virtual Machine Research and Technology Symposium Bibliography: (VM'04)

Mrinmoy Ghosh, Weidong Shi, and Hsien-Hsin S. Lee, "CoolPression - A Hybrid Significance Compression Technique for Reducing Energy in Caches", (2004). conference publication, Published Collection: In Proceedings of the IEEE International System-On-Chip Conference (SOCC-04) Bibliography: pp. 399-402, Santa Clara, California, September

Weidong Shi, Hsien-Hsin S. Lee, Mrinmoy Ghosh, and Chenghuai Lu, "Architectural Support for High Speed Protection of Memory Integrity and Confidentiality in

Multiprocessor Systems", (2004). conference publication, Published  
Collection: In Proceedings of the ACM/IEEE International Conference on Parallel Architecture and Compilation Techniques (PACT-04)  
Bibliography: pp.123-134, Antibes Juan-les-Pins, France, September

Weidong Shi, Hsien-Hsin S. Lee, Chenghuai Lu, and Mrinmoy Ghosh, "Towards the Issues in Architectural Support for Protection of Software Execution", (2004). conference publication, Published  
Collection: In the Workshop on Architectural Support for Security and Anti-Virus (WASSA-04)  
Bibliography: in conjunction with the 11th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS XI), pp.1-10, Boston, October

K.J. O'Hara and T.R. Balch, "Distributed Path Planning for Robots in Dynamic Environments Using a Pervasive Embedded Network", (2004). conference publication, Published  
Collection: Proceedings of Third International Conference on Autonomous Agents and Multi-Agent Systems  
Bibliography: July

K.J. O'Hara and T.R. Balch, "Pervasive Sensor-less Networks for Cooperative Multi-Robot tasks", (2004). conference publication, Published  
Collection: Proceedings of 7th International Symposium on Distributed Autonomous Robotic Systems  
Bibliography: June

K.J. O'Hara and T.R. Balch, "Pervasive Embedded Networks for Supporting Multi-Robot Activities", (2004). conference publication, Published  
Collection: Proceedings of the AAAI-04 Workshop on Sensor Networks  
Bibliography: July

K. O'Hara and T. Balch and E. Dodson and V. Bigio, "The GNATs programmer's guide", (2004). manual, Published  
Bibliography: [http://borg.cc.gatech.edu/gnats/doc/gnat\\_manual.pdf](http://borg.cc.gatech.edu/gnats/doc/gnat_manual.pdf)

K.J. O'Hara and V.L. Bigio and E.R. Dodson and A. Irani and D.B. Walker and T.R. Balch, "Physical Path Planning Using the GNATs", (2005). conference publication, Published  
Collection: Proceedings of ICRA 2005  
Bibliography: April

T.R. Balch and V.L. Bigio and E.R. Dodson and A. Irani and K.J. O'Hara and D.B. Walker, "The GNATs -- Building Pervasive Embedded Networks to Support Mobile Robots", (2005). conference publication, Submitted  
Collection: Proceedings of ICRA 2005  
Bibliography: <http://borg.cc.gatech.edu/gnats/doc/icra-2005-hardware.pdf>

K.J. O'Hara and T.R. Balch, "The GNATs -- Low-Cost Embedded Networks for Supporting Mobile Robots", (2005). conference publication, Published  
Collection: Proceedings of the Third International Multi-Robot Systems Workshop  
Bibliography: <http://borg.cc.gatech.edu/gnats/doc/ohara-mrsw-2005.pdf>

Kun Zhang  
Santosh Pande, "Efficient Application Migration under Compilers Guidance", (2005). Conference publication, Published  
Collection: Proceedings of ACM SIGPLAN/SIGBED 2005 Conference on Languages, Compilers, and Tools for Embedded Systems (LCTES'05)  
Bibliography: June

Ripal Nathuji  
Karsten Schwan, "Reducing System Level Power Consumption for Mobile and Embedded Platforms", (2005). conference publication,

Published

Collection: Proceedings of the International Conference on Architecture of Computing Systems (ARCS 2005)

Bibliography: March

Ripal Nathuji

Bala Seshasayee

Karsten Schwan, "Combining Compiler and Operating System Support for Energy Efficient I/O on Embedded Platforms", (2005). conference publication, Published

Collection: Proceedings of the International

Workshop on Software and Compilers for Embedded Systems (SCOPES 2005)

Bibliography: September

Christian Poellabauer

Tao Zhang

Santosh Pande

Karsten Schwan, "An Efficient Frequency Scaling Approach for Energy-Aware Embedded Real-Time Systems", (2005). conference publication, Published

Collection: Proceedings of the International Conference on Architecture of Computing Systems (ARCS'05)

Bibliography: Innsbruck, Austria, March

Christian Poellabauer

Leo Singleton

Karsten Schwan, "Feedback-Based Dynamic Frequency Scaling for Memory-Bound Real-Time Applications", (2005). conference publication, Published

Collection: Proceedings of the 11th Real-Time and Embedded Technology and Applications Symposium (RTAS)

Bibliography: March

Leo Singleton

Christian Poellabauer

Karsten Schwan, "Monitoring of Cache Miss Rates for Accurate Dynamic Voltage and Frequency Scaling", (2005). conference publication, Published

Collection: Proceedings of the 12th Annual Multimedia Computing and Networking Conference

Bibliography: San Jose, CA, January

Chinnakrishnan S. Ballapuram

Hsien-Hsin S. Lee

Milos Prvulovic, "Synonymous Address Compaction for Energy Reduction in Data TLB", (2005). conference publication, Published

Collection: Proceedings of the ACM/IEEE International Symposium on Low Power Electronics and Design (ISLPED-05)

Bibliography: San Diego, California, August

Joshua B. Fryman, "Softcache Architecture", (2005). Thesis, Published

Collection: Ph.D. Thesis

Bibliography: College of Computing, Georgia Institute of Technology

Ripal Nathuji

Keith J O'Hara

Karsten Schwan

Tucker Balch, "CompatPM: Enabling Energy Efficient Multimedia Workloads for Distributed Mobile Platforms", (2007). conference publication, Accepted

Collection: Proceedings of the ACM Multimedia Computing and Networking Conference

Bibliography: (MMCN)

Himanshu Raj  
 Balasubramanian Seshasayee  
 Keith J. O'Hara  
 Ripal Nathuji  
 Karsten Schwan

Tucker Balch, "Spirits: Using Virtualization and Pervasiveness to Manage Mobile Robot Software Systems", (2006). conference publication, Published

Collection: The Second IEEE International Workshop on Self-Managed

Networks, Systems and Services

Bibliography: (Self-Man)

Balasubramanian Seshasayee

Karsten Schwan, "Mobile Service Overlays: Reconfigurable Middleware for MANETs", (2006). conference publication, Published

Collection: Proceedings of the First International Workshop on Decentralized Resource

Sharing in Mobile Computing and Networking

Bibliography: Bala

Leo Singleton

Ripal Nathuji

Karsten Schwan, "Flash on Disk for Low-power Multimedia Computing", (2007). conference publication, Accepted

Collection: ACM Multimedia Computing and Networking Conference}

Bibliography: (MMCN)

Keith J. O'Hara

Victor Bigio

Shaun Whitt

Daniel Walker

Tucker Balch, "Evaluation of a Large Scale Pervasive Embedded Network for Robot Path Planning", (2006). conference publication, Published

Collection: IEEE International Conference on Robotics and Automation

Bibliography: (ICRA)

Keith J. O'Hara

Ripal Nathuji

Himanshu Raj

Karsten Schwan

Tucker Balch, "AutoPower: Toward Energy-Aware Software Systems for Distributed Mobile Robots", (2006). conference publication, Published

Collection: IEEE International Conference on Robotics and Automation

Bibliography: (ICRA)

Keith J. O'Hara

Tucker Balch, "Mobility and Pervasiveness in

Physical Computing Systems", (2006). conference publication, Published

Collection: Workshop on Mobility and Scalability in Wireless Sensor Networks (MSWSN) as part of the International Conference on Distributed Computing in Sensor Systems

Bibliography: (DCOSS)

Kun Zhang

Santosh Pande, "Minimizing Downtime in Seamless Migrations of Mobile Applications", (2006). conference publication, Published

Collection: In Proceedings of the ACM SIGPLAN conference on languages, compiler, and tools for embedded systems

Bibliography: (LCTES)

Jiantao Kong

Karsten Schwan

Patrick Widener, "CameraCast: Flexible Access to Remote Video Sensors", (2007). conference publication, Accepted

Collection: Multimedia Computing and Networking  
Bibliography: (MMCN'07)

Dong Hyuk Woo  
Mrinmoy Ghosh  
Emre Ozer  
Stuart Biles

Hsien-Hsin S. Lee, "Reducing Energy of Virtual Cache Synonym Lookup using Bloom Filters", (2006). conference paper, Accepted  
Collection: ACM/IEEE International Conference on Compilers Architecture and Synthesis for Embedded Systems  
Bibliography: (CASES-06)

Mrinmoy Ghosh  
Emre Ozer  
Stuart Biles

Hsien-Hsin S. Lee, "Efficient System-on-Chip energy Management with a Segmented Bloom Filter", (2006). conference paper, Published  
Collection: In Proceedings of the 19th International Conference on Architecture of Computing Systems (ARCS-06)  
Bibliography: pp 283-297, Frankfurt/Main, Germany. March,

Ripal Nathuji and Canturk Isci and Eugene Gorbatoov, "Exploiting Platform Heterogeneity for Power Efficient Data Centers", (2007).  
conference paper, Published  
Collection: ICAC '07: Proceedings of the Fourth International Conference on Autonomic Computing  
Bibliography: IEEE Computer Society

Mrinmoy Ghosh and Hsien-Hsin S. Lee, "DRAMdecay: Using decay counters to reduce energy consumption in DRAMs", (2006). conference  
paper, Published  
Collection: In Proceedings of the 3rd Watson Conference on Interaction between Architecture, Circuits, and Compilers (P=AC2)  
Bibliography: october

Mrinmoy Ghosh and Hsien-Hsin S. Lee, "Virtual Exclusion: An Architectural Approach to Reducing Leakage Energy in Caches for  
Multiprocessor Systems", (2007). conference paper, Published  
Collection: To appear in the 13th IEEE International Conference on Parallel and Distributed Systems  
Bibliography: December

Mrinmoy Ghosh and Hsien-Hsin S. Lee, "Smart Refresh: An Enhanced Memory Controller Design for Reducing Energy in  
Conventional and 3D Die-Stacked DRAMs", (2007). conference paper, Published  
Collection: To appear in the 40th ACM/IEEE International Symposium on Microarchitecture  
Bibliography: December

### **Web/Internet Site**

#### **URL(s):**

<http://cc.gatech.edu/~kjohara/MORPH/>  
<http://borg.cc.gatech.edu/gnats>

#### **Description:**

The Morph project web page.

### **Other Specific Products**

### **Contributions**

#### **Contributions within Discipline:**

The compiler-based code mobility framework makes it tractable to use mobility of code to our advantage removing barriers in its deployment. Code mobility is a very important asset

in data driven, wireless, mobile systems due to many reasons.

Our energy reduction and analysis results are significant because they can be incorporated into platforms in order enable them for consumers by providing the types of battery lifetimes required for the respective applications. Therefore, these types of technologies can provide benefits in a variety of environments via platforms such as medical devices, robots, and digital assistants.

#### **Contributions to Other Disciplines:**

Easing code mobility will allow development of mobile applications and infrastructure rapidly for mobile-commerce etc.

#### **Contributions to Human Resource Development:**

REU support for this grant has aided in gaining admission for one of our undergraduates to a good graduate program and hopefully to promising research career.

#### **Contributions to Resources for Research and Education:**

#### **Contributions Beyond Science and Engineering:**

#### **Categories for which nothing is reported:**

Any Journal

Any Product

Contributions: To Any Resources for Research and Education

Contributions: To Any Beyond Science and Engineering

The faculty members and students involved in the Morph project performed performing research on a variety of fronts, including those detailed below. Much of the research funded under this grant was completed prior to the original expiration date in 2007. An extension was granted in order to expend remaining REU funding. One of the funded students, Charles Reiss, has now moved on to graduate school at UC Berkeley and will hopefully develop into a solid contributor to our field. His work at Georgia Tech leveraged the mobile code platform developed in this grant and applied it in a middleware domain.

## System Power Measurement, Analysis and Modeling

This research thrust has continued to explore embedded systems and devices, but in addition, it has begun to explore server systems and their power consumption. Thus, while continuing to construct a detailed understanding of power consumption and behavior patterns for a full embedded system such as a PDA device, it has also investigated server-class platforms, with specific focus on Intel platforms due to their market dominance. We developed power profiles for multiple generations of Intel platforms, starting with P4 platforms and moving to the most recent multi-core platforms being shipped. We have also developed new server-class benchmarks for these platforms.

## Cache-aware Execution

1. A key issue with server-class machines is the effective use (and the power implications thereof) of their cache infrastructures.

The main research thrust of this work is to reduce power consumption and increase performance by using appropriate methods for (1) tracking cache usage by threads and (2) then using such information to appropriately schedule threads. Ongoing work is developing new hardware mechanisms for tracking thread/cache affinity coupled with new scheduling techniques to exploit such information.

2. Reducing System-on-Chip power by reducing synonymous lookup using Bloom Filters.

Continuing the first thrust, we examine the L1 power issues due to address synonym in virtual cache designs, in which multiple virtual addresses can map to the same physical address for maintaining the correctness. The current solution is to look up all the synonymous cache sets serially to identify the potential hit at the expense of an elongated hit latency and larger power dissipation. We explored the design space of the Counting Bloom Filter to address this issue. An effective counting bloom filter will be able to positively identify the existence of address synonym to reduce the unnecessary accesses to the other potential synonymous sets in L1. Similarly, this scheme can reduce the access latency as well as the energy consumption in the L1.

The following two tasks were completed.

3. Reducing leakage power in caches via Virtual Exclusion

This research leverages two circuit techniques, gated Vdd and drowsy cache, and proposes a low cost, easily implementable architecture scheme. The Virtual Exclusion scheme saves leakage energy by keeping the data portion of repetitive cache lines off in the large higher level caches while still manages to maintain Multi-Level Inclusion, an essential property for an efficient implementation for conventional cache coherence protocols. By exploiting the existing state information in the snooping based cache coherence protocol with little hardware overhead, we anticipate to reduce leakage power for the repetitive cache lines present in the L2 cache.

4. Minimizing Periodic Refresh Energy in DRAMs using DRAM decaying control

This research exploit the access behavior of DRAM for possible power saving opportunities. Current DRAM controller refreshes a DRAM row periodically. During a refresh cycle, all the cells in a DRAM row will to be read and written back to (self-destructing). This refresh process has a big power and bandwidth overhead. We are developing an innovative method of using a decay counter for every row as an indicator to control the overhead of refresh. The basic concept is that a row that has been accessed due to a cache miss or eviction should not be refreshed again by the periodic refresh interval. We anticipate this new policy will reduce all the superfluous refresh activities.

## Managing Virtualized Resources

While previous work on energy preservation has focused on single resources or technologies, increasingly, platforms are being virtualized to better exploit platform resources. This means that former power management strategies with direct access to the resources being managed must now work via the hypervisor, accessing virtual vs. physical hardware. The VirtualPower approach developed in our research addresses this problem, with new mechanisms that can deal with (1) virtualized processor and other platform resources, (2) new mechanisms for changing the power behavior of virtualized resources that also exploit the management capabilities of underlying physical processors, (3) and flexibility in how power management is performed, with particular importance placed on exploiting the power management strategies already present in guest virtual machines.

An extension of VirtualPower now in progress concerns datacenter-wide power management strategies and mechanism support for implementing such strategies. Another extension concerns work on thermal issues and properties.

## Robotics

New techniques for power-aware multi-robot software systems have been developed, with specific emphasis on the mobility behavior and opportunities derived from such behavior. This has resulted in a distributed middleware for managing the power consumption properties of applications in domains like robotics, driven by and evaluated with representative robotics workloads. Power savings from autonomic middleware behaviors can be substantial, as demonstrated with said workloads running on portable platforms.

## Compiler-based Code Migration Framework

We have already reported on this work last year, with no additional results attained during the past year. Instead, a new focus of this research has been to consider a new class of architectures relevant to this domain: heterogeneous multicore machines, a specific example being IBMs Cell processor.

## System Power Measurement, Analysis and Modeling

Extensive experimentation with multiple workloads and platforms has established the strong differences existing across different machine architectures in terms of their performance/power profiles. As reported in a recent paper published in [4], differences even across multiple generations of a single vendors platforms (i.e., Intel) can be substantial. Moreover, it is well-established that typical datacenters contain multiple generations of machines, so that these differences are present in many modern settings.

## Low-power Architecture Support

We refer the reader to multiple publications that have appeared during the last year, including [3, 1, 2, 5].

## Coordinated Resource Management

A new insight from our research is the importance of coordinating the resource and power management (i.e., the management) actions taken across different machines and even across different levels of abstraction of a single machine. In recent joint work with researchers at HP Labs, this has resulted in the design of a coordination infrastructure for virtualized machine platforms, spanning firmware-level, to hypervisor-level, to VM-level management actions and policies. In experiments conducted at Georgia Tech, we have found multiple beneficial effects of coordinated management actions, including, for power management, the ability to dynamically manage virtual machines according to the utility of the applications they run rather than according to lower level metrics like machine utilization. Experiments with representative multi-tier enterprise applications and in addition, with a prototype search engine, establish these facts.

## Virtual Platforms for Robotics

A new direction of our work is to apply virtualization concepts to describing the platforms on which robotics or embedded applications run. Toward this end, we have constructed a virtual camera device, i.e., a device that appears local to a guest operating system but is actually remote. Such a device can be accessed from anywhere without the need for complex application software. It can be shared by multiple parties, again without the need for application level software implementing such sharing. Finally, interesting new devices can be constructed from multiple physical devices that are thus virtualized, such as a multiple view camera actually composed of multiple physical cameras distributed across some space. We have constructed a basis for the runtime assembly and use of virtual platforms like that and have shown them viable with modern hardware and operating system platforms.

## Compiler-based Code Migration Framework

No further results are reported for this period.

## Power-Aware Scheduling

Results have now moved from power-aware to utility-aware scheduling, that is, to performing scheduling such that power goals are met so as to maximize the utility of applications. Outcomes include (1) that the quality of information produced by an application might be compromised to reduce power usage, (2) that lower vs. higher utility applications might give rise to different priorities concerning power reduction for virtual machines, and (3) that virtual machine migration methods can be used for both power savings and performance/cost advantages.

## References

- [1] Mrinmoy Ghosh and Hsien-Hsin S. Lee. Virtual exclusion: An architectural approach to reducing leakage energy in caches for multiprocessor systems. To appear in the 13th IEEE International Conference on Parallel and Distributed Systems.
- [2] Mrinmoy Ghosh and Hsien-Hsin S. Lee. Dramdecay: Using decay counters to reduce energy consumption in drams. In *In Proceedings of the 3rd Watson Conference on Interaction between Architecture, Circuits, and Compilers (P=AC<sup>2</sup>)*, October 2006.

- [3] Mrinmoy Ghosh and Hsien-Hsin S. Lee. Smart refresh: An enhanced memory controller design for reducing energy in conventional and 3d die-stacked drams. To appear in the 40th ACM/IEEE International Symposium on Microarchitecture, December 2007.
- [4] Ripal Nathuji, Canturk Isci, and Eugene Gorbatoov. Exploiting platform heterogeneity for power efficient data centers. In *ICAC '07: Proceedings of the Fourth International Conference on Autonomic Computing*, page 5, Washington, DC, USA, 2007. IEEE Computer Society.
- [5] Dong Hyuk Woo, Mrinmoy Ghosh, Emre Ozer, Stuart Biles, and Hsien-Hsin S. Lee. Reducing energy of virtual cache synonym lookup using bloom filters. In *Proceedings of the ACM/IEEE International Conference on Compilers Architecture and Synthesis for Embedded Systems*, pages 179–189, October 2006.