

cac bulletin

The Newsletter of the Center for Autonomic Computing

Vol. 2, Holiday 2010



In this issue:



Message from the Director

2

CAC at Mississippi State University

Feature article about the new CAC site at Mississippi State University, established in July 2010 and directed by Ioana Banicescu.

3

Spotlight on People at CAC

We profile select CAC faculty members: Professor Ali Akoglu (UA), Research Scientist Abhishek Dubey (Vanderbilt University), Professor Shantenu Jha (RU) and Professor Andy Li (UF).

4

Research Projects: Data Centers and Cloud Computing

Featured projects in data centers and cloud computing highlight CAC's research efforts in large-scale and distributed computing.

5

Research Projects: Networking and Intercloud Computing

Featured projects use networking tools to facilitate computation across multiple cloud computing resources.

6

Research Projects: Applications and Services

Featured projects focus on security management and resource allocation for Web services platforms.

7

Research Projects: Cybersecurity and Reliability

Featured projects include securing energy-critical infrastructure and a making scientific workflows using distributed computing resources more fault-tolerant.

8

CAC News

Noteworthy events, facility improvements and awards

9

Face to Face

Industry-university interactions and technical talks by CAC researchers

11

CAC Publications

12

Save the date! The next CAC review meeting will be held in April 2011. Stay tuned for details about the upcoming meeting and Center happenings at <http://www.nsfcac.org>. Interested in becoming an industrial partner? See page 10 of this newsletter for details!

Message from the Director



Dear Friends,

As 2010 comes to a close, I am excited for what we have accomplished this year and for what 2011 has in store for the CAC center. This issue of the CAC newsletter brings you some news, projects and plans that, I hope, will lead you to share in my excitement.

Let me start by announcing the addition to CAC of a new academic site at Mississippi State University (MSU). After one year of planning activities, MSU faculty and students led by Ioana Banicescu and Sherif Abdelwahed received official approval and funding by the NSF to join the Center in July of this year. You can read more about the CAC MSU site in this newsletter. I also wish to announce and welcome two new industrial members—The US Army Corps of Engineers' Engineer Research and Development Center (ERDC) and MobiLaps—with whom everyone at CAC looks forward to new cooperative initiatives.

We held our second 2010 biannual review meeting at the University of Florida on October 5-6. This was a successful meeting with many firsts. We had the largest number of presentations (18) and the largest number of demos (16) ever presented in a review meeting. In addition to copies of presentation slides, demos are now also made available for replay at the CAC website (www.nsfcac.org). This is a new feature that we plan to preserve for future meetings for the benefit of our members—all presentations and posters can be reviewed on the website, and all posters have a companion demo which can also be replayed on the web.

Looking forward, CAC faculty and members of the Industrial Advisory Board identified strategic technical areas of emphasis and growth for CAC activities. They are cybersecurity, cloud computing and intercloud computing. As illustrated in the figure below, these strategic areas are a natural outgrowth of the topics of strength and concentration of CAC activities in the last few years—datacenter management and high-performance computing, security and reliability, and virtual networking and services.

Already reflecting our strategic plan, we organized a workshop on Intercloud Computing at the University of Florida on September 30 and October 1, attended by participants from industry and academia. One outcome of this workshop is a plan for the deployment of an intercloud-computing testbed across the CAC university sites. One of the objectives of this testbed is to conduct research whose results will inform issues related to cloud interoperability standards—expect more on this in future newsletters and meetings.

In closing, I would like to invite industrial members to contribute to future issues of this newsletter by sharing their technical views, news and other information items they wish to provide because of their relevance to autonomic computing and CAC people. As always, do not hesitate to contact me at fortes@ufl.edu with your ideas for newsletter articles, questions or comments.

Happy reading and best wishes for the new year.

Sincerely,

José A. B. Fortes, CAC Director

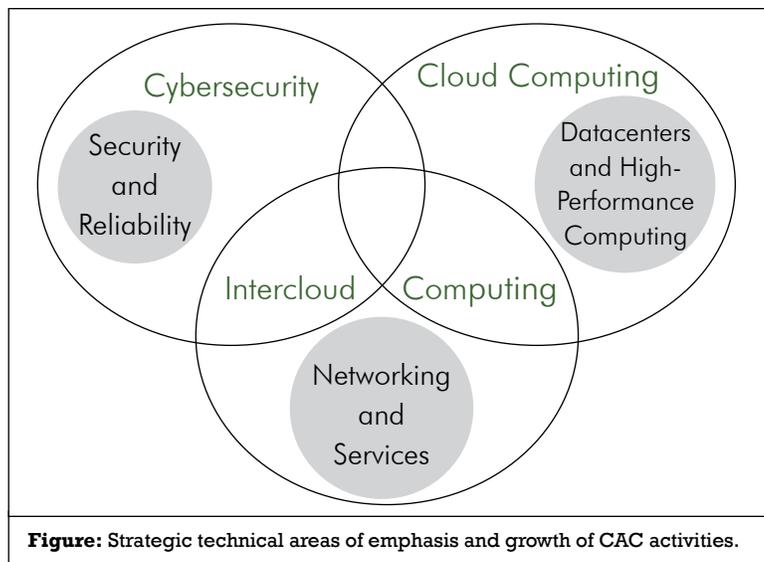


Figure: Strategic technical areas of emphasis and growth of CAC activities.

Feature: CAC at Mississippi State University

A fourth university site at Mississippi State University joined the Center late this summer. The new site, under the leadership of Ioana Banicescu and Sherif Abdelwahed, brings several new industry memberships, including new CAC member company, the US Army Corps of Engineers' Engineer Research and Development Center (ERDC). Microsoft, a long-standing member of several CAC sites, also supports CAC at MSU.

The director of the MSU site is Ioana Banicescu. Ioana received her undergraduate degree from Polytechnic University in Bucharest, Romania, and her M.S. and Ph.D. degrees from Polytechnic University in New York, and is currently a professor in MSU's Department of Computer Science and Engineering. Ioana became interested in CAC through her collaboration with other CAC researchers; she has collaborated in the past with CAC Director José Fortes and co-Directors Manish Parashar and Salim Hariri. She submitted a proposal for an NSF planning grant in 2008 to look into the possibility of establishing a CAC site at MSU with co-PI Sherif Abdelwahed. Sherif is a professor in the Department of Electrical and Computer Engineering and directs a research program in autonomic computing and model-based design and analysis of distributed systems.



Photo: Dr. Ioana Banicescu (left) and Dr. Sherif Abdelwahed (right), lead faculty members at the new CAC site at Mississippi State University.

Dr. Banicescu anticipates that her involvement in CAC will have benefits beyond advancing individual research projects. "At MSU," she said, "we considered that being a new NSF CAC site would attract collaborators for research (from other universities) and development (from industry); consequently, that will increase the visibility of the State of Mississippi in the IT industry arena. We hope that this will generate interest among IT companies to invest or even have a presence in the state of Mississippi."

The MSU group received a planning grant from NSF in 2009 to seek out industry partners and gauge levels of support for a new CAC site at MSU. An exploratory meeting for the MSU site was held in October 2009 on the Mississippi State campus in Starkville. This meeting was held in conjunction with CAC's fall 2009 semiannual review meeting, allowing prospective members of the MSU site an opportunity to get acquainted with CAC research and giving CAC personnel and members an overview of autonomic computing at MSU. After a successful planning meeting, the NSF grant establishing the CAC site at MSU was awarded in July 2010. "We're excited that the MSU site has become a

reality," José Fortes said of the addition of the new site, "and we are certain that the addition of the MSU site will increase the value of membership for all of CAC's industry partners."

The MSU site has begun research on several new projects in collaboration with CAC @ MSU members Microsoft and ERDC. According to Sherif Abdelwahed, "our collaboration with Microsoft is focused on applying model-based feedback techniques to improve the performance and cost effectiveness of large scale data centers, particularly when using Microsoft server HPC edition." Two projects with ERDC are also underway. These projects, which focus on automated security management in distributed computing systems hosting web-services and developing a model-integrated computing framework to support the development of autonomic computing systems, are described on page 7.

Facilities at MSU

MSU is home to the High Performance Computing Collaboratory (HPC²), a coalition of research groups and centers using high-performance computing to advance the state of the art in computational science and engineering. Computing facilities at HPC² include the Talon supercomputer, an IBM iDataPlex cluster that was ranked the ninth most energy-efficient supercomputer in the world on the June 2010 Green500 list and the Raptor supercomputer, ranked the 18th-fastest computer at any U.S. academic site in June 2010.



Photo: Facilities at MSU. The Talon supercomputer (left) was ranked the 331st most powerful in the world in June 2010 (photo credit: MSU University Relations). Below, administrators inspect the Raptor supercomputer at MSU's High Performance Computing Collaboratory (HPC²) (photo credit: MSU/Roger Smith).



Spotlight on CAC People



Ali Akoglu
University of Arizona

Ali Akoglu has been a faculty member in the Department of Electrical and Computer Engineering at the University of Arizona since August 2005. He leads the Reconfigurable Computing Laboratory, guiding a research program focusing on reconfigurable computing and specializing in application-specific reconfigurable architecture design space.

Ali earned his Ph.D. in Computer Science from Arizona State University and his B.S. degree in Computer Science from Purdue University. He is a collaborator on the cross-cloud computing project, which links resources at UA, RU and UF to create a testbed to enable research on autonomic computing in virtualized, distributed data centers.



Abhishek Dubey
Vanderbilt University

Abhishek Dubey is a Research Scientist in the Institute for Software Integrated Systems at Vanderbilt University. He has a close relationship with CAC's Mississippi State University site, currently pursuing research on autonomic computing through collaboration on several CAC projects at MSU on resource management and multi-tier enterprise systems. He also conducts research in fault-adaptive control technology for software in hard real-time systems.

Abhishek received his M.S. and Ph.D. degrees in Electrical Engineering from Vanderbilt University and his B.S. degree from the Institute of Technology at Banaras Hindu University. He has published over 20 research papers and is a member of IEEE.



Shantenu Jha
Rutgers University

Shantenu Jha is a new faculty member joining the CAC site at Rutgers University. He comes to Rutgers from Louisiana State University, where he is a Research Professor in the Department of Computer Science and the Director for Cyberinfrastructure Development at LSU's Center for Computation and Technology (CCT). Shantenu is also a visiting researcher at University College of London's Centre for Computational Science and a theme leader at the e-Science Institute in Edinburgh.

He is currently writing a book on "Abstractions for Distributed Applications and Systems: A Computational Science Perspective" to be published by Wiley in the next year.



Xiaolin (Andy) Li
University of Florida

Xiaolin (Andy) Li is a new faculty member in UF's Department of Electrical and Computer Engineering. He is director of the Scalable Software Systems Laboratory (S³Lab), with a mission to explore innovative ideas and algorithms, and design next-generation software components, middleware and services towards scalable, self-managing, high-performance, and trustworthy parallel and distributed systems.

He received his Ph.D. degree in Computer Engineering from Rutgers University. He is the recipient of a 2010 CAREER Award, one of the National Science Foundation's most prestigious awards, recognizing junior faculty for outstanding research and excellence in education.

CAC Research Projects: Data Centers and Cloud Computing

Abnormal Resource Usage Detection via Thermal Profiling of Datacenters

PI: Dario Pompili [pompili@cac.rutgers.edu]

Lead researcher: Eun Kyung Lee [eunkyung_lee@cac.rutgers.edu]

Datacenters are a growing component of society's IT infrastructure, including services related to health, banking, commerce, defense, and education. This has led to an increase in the size of datacenters and the density of servers. Annual energy and administration costs of today's datacenters amount to billions of dollars; the high energy consumption also translates into excessive heat dissipation, which, in turn, increases cooling costs and servers' failure rates. Moreover, the increasing demand on computational infrastructure increases the possibility of anomalies (i.e., misconfigurations, hardware failures and possible external attacks).

Each datacenter has a unique "thermal footprint" that reflects how temperature changes in time and space under different workload types and distributions. This footprint depends on datacenter-specific characteristics such as the ventilation system, the layout of racks, the density of the machines, and the heterogeneity of resources across racks. Knowledge of the thermal footprint, which will enable both efficient cooling to reduce energy cost and detection of anomalies, can be extracted by observing the usage of hardware resources (i.e., CPU, memory, disk, and network switches) by measuring the temperature inside the servers, and by monitoring performance of the ventilation system. However, the cost for gathering and processing this information is huge in large datacenters. Also, if servers are compromised during attacks, such data may not be reliable in terms of making remedial decisions.

The aim of this project is i) to exploit a heterogeneous sensing infrastructure to thermally profile a datacenter and capture its thermal footprint, and ii) to develop learning algorithms that use this footprint in order to detect abnormal resource usage due to illegitimate workload caused by policy infringements, attacks and failures. In synergy with CAC researchers, Rutgers University is developing efficient computing and communication solutions that will enable distributed networks of heterogeneous low-power sensors (e.g., temperature, humidity, airflow, infrared cameras) to function not only as passive measurement devices monitoring thermal

phenomena in a datacenter but also as intelligent data processing instruments capable of data quality assurance, statistical synthesis, and hypotheses testing as they stream data from the physical environment to the computational world.

Resource Usage Estimation for Cloud Computing Applications

PI: José Fortes [fortes@ufl.edu]

Lead researcher: Andréa Matsunaga [ammatsun@acis.ufl.edu]

The adoption and availability of Infrastructure-as-a-Service (IaaS) clouds brings new challenges to users and datacenter administrators. From the user perspective, clouds with the best cost-benefit need to be selected, taking into consideration a variety of provisioning models (e.g., immediate, bidding, best-effort and advanced reservation) and the SLAs offered by IaaS providers. From the perspective of datacenter administrators, optimal performance in scheduling user requests can be achieved if resource usage information is made available before provisioning the infrastructure.

In the context of predicting spatiotemporal utilization of resources, this project investigates the following: the process of collecting meaningful historical data from past requests, the suitability of several machine learning techniques for prediction of resource usage, potential improvements of existing machine learning algorithms, the impact of features on prediction accuracy, and the impact of accurate estimations on resource selection and scheduling.

One result of our work is a novel prediction technique, called Predicting Query Runtime Regression (PQR2), based on the adaptive method called Predicting Query Runtime (PQR). PQR is based on a binary tree in which every node has a 2-class classifier selected from a pool of classifiers. The tree is built repeating a split algorithm recursively until certain thresholds are reached, namely the following: minimum prediction accuracy of a node, minimum number of examples in a node after skipping a percentage of examples from the borders of the range, and minimum range. The PQR2 algorithm has been applied to scenarios where resources are used to run bioinformatics applications (BLAST and RAXML); results demonstrated that PQR2 can provide the best accuracy in scenarios requiring prediction of processing, memory and disk usage; our results also made the case for increasing the number of attributes typically used for resource prediction.

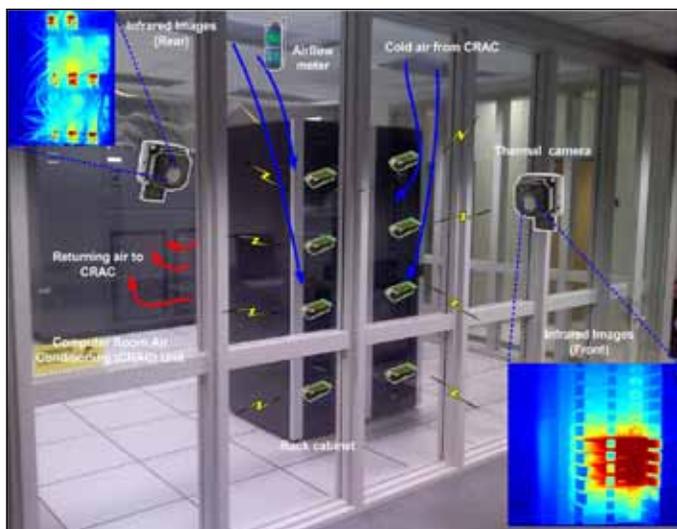


Figure: At far left, a schematic of the temperature-profiling equipment deployed at the Rutgers University datacenter. At near left, a floor plan of the machine room depicting airflow patterns.

CAC Research Projects: Networking and Intercloud Computing

Communication Management of Cross-cloud Computing

PI: Maurício Tsugawa [tsugawa@acis.ufl.edu]

The growing availability of commercial and scientific clouds is exposing the general public to unprecedented distributed computing power. Users of these resources need to determine which clouds best suit their needs, and potentially combine services offered by different clouds. A key component required for cross-cloud computing is that of a (virtual) network to connect nodes provisioned by different clouds. Although solutions to extend client networks into cloud networks exist, these approaches are designed and implemented considering a particular cloud or type of cloud. As a result, enabling communication among resources in different clouds becomes a very challenging task.

This project builds on ViNe, network virtualization middleware developed at the University of Florida, which enables high-performance IP network overlays. ViNe uses connectivity-recovery techniques to overcome connectivity limitations imposed by cloud providers to connect resources (VMs) across WANs and clouds. The goal of this project is to extend ViNe and incorporate autonomic computing techniques to make it self-managed and transparent to both cloud users and providers. Research and development are being conducted for: (a) a self-configured virtual network system where IP overlays are established with minimal intervention by users and providers; (b) network support for cross-cloud VM migration; (c) monitoring of overlay network performance (latency and throughput); (d) network services differentiation (e.g., prioritizing high-throughput applications on high-bandwidth overlay links and latency-sensitive applications on low-latency overlay links); and (e) dynamic overlay performance tuning (e.g., use of parallel links to take advantage of high-bandwidth backbones, use of compression techniques, and use of alternate overlay routes).

We have demonstrated that ViNe can connect VMs across multiple clouds to efficiently run massively parallel bioinformatics applications. The experiments used two testbeds: the NSF-funded FutureGrid in the U.S. and Grid5000 in France. A 750-VM (1500-core) Hadoop cluster was deployed across three FutureGrid sites and four Grid5000 sites. Insights on communication management needs of cross-cloud computing gained from this work are being used in the development of new ViNe components.

Autonomic Workflow Management on Hybrid Clouds Using CometCloud: A Demonstration at the ACS ITO Client Symposium, October 2010

PI: Manish Parashar [parashar@cac.rutgers.edu]

Lead researcher: Hyunjoo Kim [hyunjoo@cac.rutgers.edu]

As infrastructures that federate private and public clouds are being increasingly used by enterprises to meet their IT requirements, autonomic workflow management frameworks become essential to effectively execute business process workflows on these hybrid and dynamic platforms while satisfying complex and highly heterogeneous requirements as well as constraints for performance, cost, privacy, etc. Such autonomic frameworks can dynamically select an optimal mix of resource classes based on QoS and

resource requirements, policies and constraints; can monitor the execution of applications and services; and can adapt

both resources and services to ensure that these requirements and constraints are continually satisfied. Adaptations may involve scaling up, down or out using autonomic Cloudbursts and can handle problems including unanticipated workload bursts or resource failures.

Rutgers and Xerox collaborated to develop and deploy such an autonomic workflow management framework on top of the CometCloud autonomic cloud federation engine (www.cometcloud.org), and demonstrated it at the ACS ITO Client Symposium in Orlando, Florida, held October 11-13, 2010. This effort demonstrated the following capabilities:

1. Policy-driven synthesis of customized hybrid clouds using the dynamic, on-demand cloud federation services provided by CometCloud. Specifically, the hybrid infrastructure in the demonstration dynamically integrated private clouds at Rutgers, ACS and the Amazon EC2 public cloud with approximately 600 cores.
2. Programming abstractions and autonomic management mechanisms to support policy driven scale-up/down/out based on application requirements as well as system state.
3. Deployment of real-world enterprise application workflows executing on a real federated hybrid cloud infrastructure.

This demonstrated the benefits of hybrid cloud infrastructure to support heterogeneous and dynamics workloads, the feasibility of on-demand cloud bridging, and the ability of the CometCloud autonomic cloud federation framework to robustly provide these capabilities in a production environment.

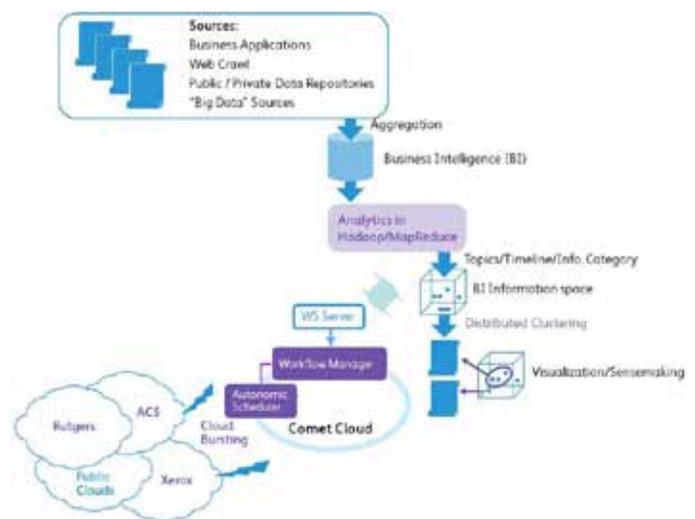


Figure: An illustration of the scenario used in the demonstration.

More Networking and Intercloud Computing

Go to nscac.org to learn about other CAC projects related to networking, including:

- End-to-end Services for P2P Virtual Networks in Collaborative and Cloud Computing Environments
- Self-organizing IP-over-P2P Overlays for Virtual Networking
- Autonomic Workflow Management in Dynamically Federated, Hybrid Cloud Infrastructures Using CometCloud.

CAC Research Projects: Applications and Services

Model-based Framework for Autonomic Resource Management of Web Services Platforms

PI: Sherif Abdelwahed [sherif@ece.msstate.edu]
 Lead researcher: Rajat Mehrotra [rm651@msstate.edu]
 Large-scale DoD systems such as the Global Information Grid (GIG) are hosted on heterogeneous and networked servers that execute in a dynamic and uncertain operating environment, caused by factors such as time-varying user workload and incomplete knowledge of the system operating state. These enterprise applications and systems are vulnerable to hardware and software component failures, operator errors (e.g., misconfigured software or hardware components), as well as malicious attacks. Therefore, achieving the stringent quality-of-service (QoS) goals expected of such systems is a challenging task, requiring a comprehensive approach to performance control, failure detection and diagnosis, and finally, failure recovery.

This project will develop and validate a fault-adaptive control framework (FACT) to manage distributed computing systems. The central idea here is to integrate control and diagnosis processes within a common model-based framework that will continually optimize system behavior in response to both changes within the system (e.g., hardware/software failures and configuration changes) as well as to external changes in the operating environment (e.g., time-varying workload intensity). We will implement this framework and, as a case study, use it to solve performance management problems in a virtualized execution environment hosted on a large blade-server cluster.

Automated Model-based Security Management of Web Services

PI: Sherif Abdelwahed [sherif@ece.msstate.edu]
 Lead researcher: Qian Chen [qc34@msstate.edu]
 Distributed computing environments are complex systems that aggregate vast quantities of information in order to manage their

computational loads in real time. The variety of network protocols and network interfaces applied on these systems introduce the potential for illicit cyber-penetration of a distributed computing system. Research is required to identify efficient management algorithms that detect faults through device failure and due to cyber penetration, identify the portion of the distributed computing system affected by a fault or network penetration, and react to faults and/or network penetrations by isolating the affected portion of the distributed computing system to prevent cascading failure events.

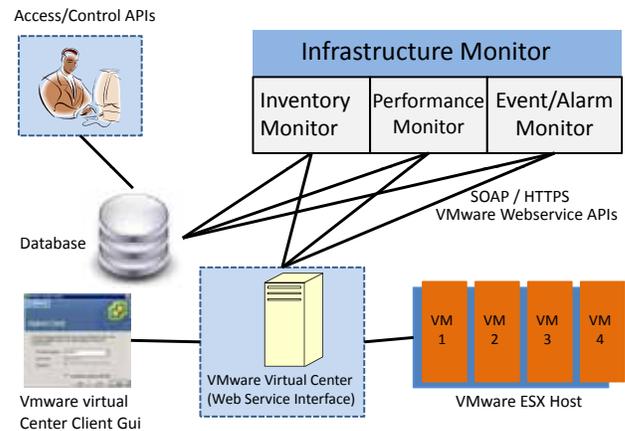


Figure: Virtual machine infrastructure security management

In this project, we propose the development of an autonomic performance and security management technology that integrates system control, optimization, and security analysis tasks into a common model-based framework that enables distributed computing systems under normal conditions to adapt efficiently to variations in load requirements and identify and mitigate potential security intrusions and maintain functionality. We will develop a computationally efficient model for distributed computing systems which provides effective estimates of system behavior through aggregation of load data and stochastic probability of network intrusion for each critical interface in the modeled distributed computing system. We will develop an integrated security management and adaptive performance control framework that unifies the two tasks of managing systems performance throughout normal and under security threat operation settings. We will validate the performance and functionality of the integrated autonomic performance and security management system.

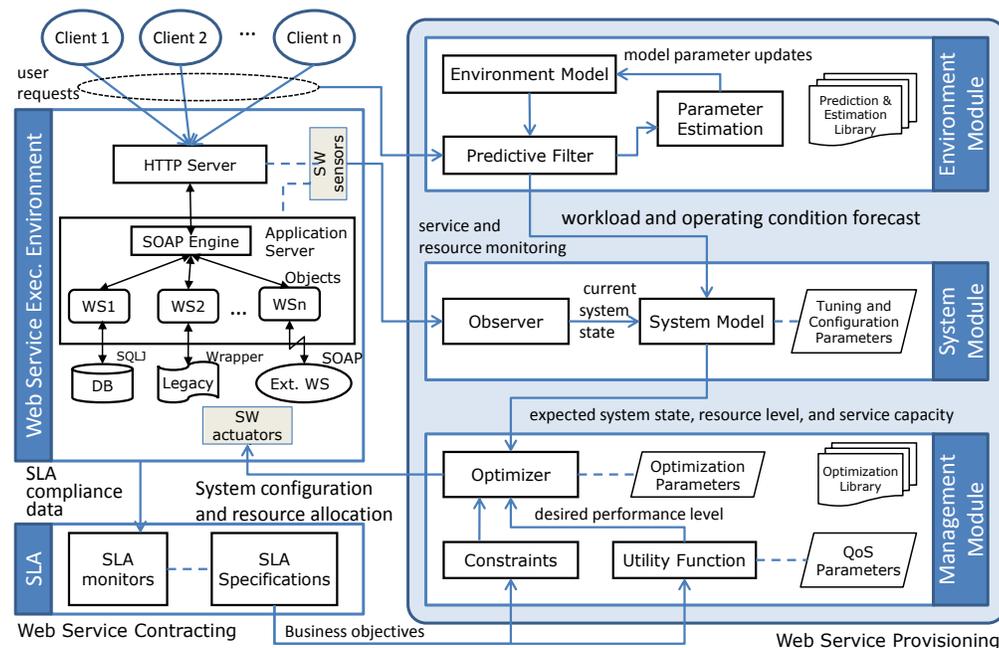


Figure: High-level architecture of the model-based autonomic Web services framework

tems which provides effective estimates of system behavior through aggregation of load data and stochastic probability of network intrusion for each critical interface in the modeled distributed computing system. We will develop an integrated security management and adaptive performance control framework that unifies the two tasks of managing systems performance throughout normal and under security threat operation settings. We will validate the performance and functionality of the integrated autonomic performance and security management system.

CAC Research Projects: Cybersecurity and Reliability

The successful operation of cyberinfrastructure requires the ability to detect previously unseen attacks in real time, to avoid system failures, to provide accurate risk and impact analyses, and to take appropriate action. The following CAC projects are geared towards providing secure and reliable computing services without the need for costly and time-intensive human intervention.

Self-healing Scientific Workflows with Autonomia

PI: Ali Akoglu [akoglu@ece.arizona.edu]

Lead researcher: Arjun Hary [arjunh@ece.arizona.edu]

Scientific workflows (SWFs) are widely used for managing scientific data (observational data, streaming sensor data, etc.). Kepler is a popular open-source SWF which simplifies the construction of complex data flow models and provides seamless access to distributed computing resources. As the complexity of the workflow applications running on heterogeneous distributed systems increases, fault management becomes a critical design issue for large-scale scientific and engineering applications. Due to the long execution times of these applications, it is important that they are fault tolerant—i.e., the workflow application can recover from faults without the need to restart the application from the beginning.

This project aims to extend Kepler's capabilities with Autonomia for supporting fault-tolerant scientific workflows (FT-SWF) using a checkpoint mechanism that takes corrective measures in an autonomic manner without the user's involvement whenever a fault is detected. Autonomia is a control and management environment that provides software wrappers to enable any hardware or software component to self-manage its operations. We have employed Autonomia as a monitoring layer over Kepler SWF to detect faults during workflow execution and resolve them according to fault-management policies at run time. Kepler uses the "actors" programming abstraction to specify the processes and "directors" to specify when these processes occur. By using Autonomia tools, we are able to extend Kepler to achieve a dynamic workflow where the properties of actors and directors are sensed and modified dynamically at run time without having to interrupt normal execution. To the best of our knowledge, this is the first approach to adding autonomic operations to Kepler.

We have demonstrated the capabilities of the self-healing Kepler implementation on a Windows 2008 HPC cluster using the Gross Primary Production (GPP) model, which is used to understand ecosystem-atmosphere carbon exchanges in dry land

environments. We have shown how to add the checkpoint/restart capability to the Kepler directors so that the workflow can resume its execution from the last checkpoint in case of a fault. Finally, we have evaluated the performance of the FT-Kepler on a distributed system with hardware- and software-based fault scenarios in terms of execution time, recovery time, and the checkpoint mechanism overhead. Our experimental results indicate that the checkpoint mechanism adds negligible overhead to the total execution time of the workflow and that, as the fault rate increases, the number of checkpoints should be increased. We plan to deploy FT-Kepler on the newly acquired IBM Blade System by CAC through NSF CRI grant that houses 168 processing cores.

Industrial Control Systems Cybersecurity Testbed

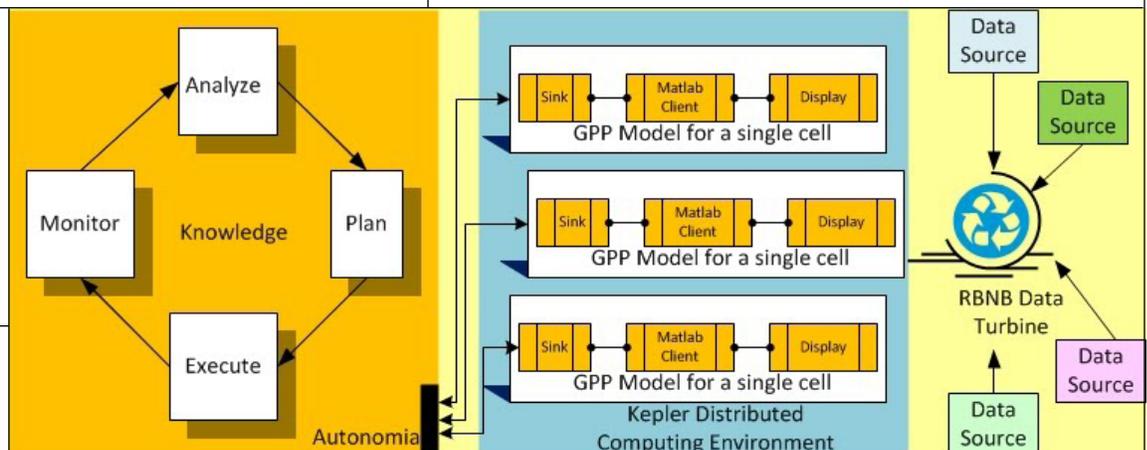
PI: Salim Hariri [hariri@ece.arizona.edu]

Lead researcher: Don Cox [dcox@raytheon.com]

Industries managing energy-critical infrastructures (e.g., power, water, gas and oil) are modernizing their Industrial Control Systems (ICS) using what is referred to as "Smart Grid" technology, which uses advanced computing and communications techniques to provide information that allows these services to operate far more efficiently. These infrastructures use Supervisory Control and Data Acquisition (SCADA) systems to control and manage their resources. Because SCADA systems were never designed with security in mind, they are extremely vulnerable to both internal and external attacks, and securing them is a challenging research problem. Consequently, SCADA networks become a prime target for cyber attacks due to the profound and catastrophic impacts they can have on our economy and all aspects of our life.

Due to ICS vulnerabilities and security problems, it is critically important to develop new innovative techniques that can accurately and in real-time detect cyber-attacks and penetrations targeting SCADA systems. To address this challenge, we are building a state-of-the-art ICS cybersecurity testbed at the University of Arizona CAC site. With this testbed, our goal is to integrate the autonomic management tools developed at CAC and member companies AVIRTEK and Raytheon so we can experiment with and evaluate innovative cybersecurity technologies to secure and protect ICS services. The testbed will also be a viable tool for teaching and training the next generation of the ICS workforce who will be expert in how to secure and protect our critical infrastructures, their resources and their services.

Figure: A data turbine feeds field data to Kepler SWF; composite Kepler actors run models or tasks in a distributed computing environment. Autonomia enables self-management of Kepler actors with continuous monitoring. The use of checkpoints with Autonomia enables self-healing and fault recovery.



CAC News

Membership

CAC welcomes two new industry members this year: *MobiLaps* and the *US Army Corps of Engineers' Engineer Research and Development Center (ERDC)*. *Hisham Kassab*, head of *MobiLaps' Network Appliance Software Division* is representing *MobiLaps* on the *Industry Advisory Board* and is collaborating with *University of Arizona* researchers on the *Network Monitoring System* project. *Cary Butler* is representing *ERDC* on the *IAB* and is collaborating with *Mississippi State University*.

Service

UF Site Director Renato Figueiredo has been selected to serve as an associate editor for the monthly scholarly journal *IEEE Transactions on Computers (TC)*. *TC* reaches a wide audience including researchers, developers and educators in the computer field.

CAC Director José Fortes has been appointed to the advisory committee for *NSF's Directorate for Computer and Information Science and Engineering (CISE)*. The *CISE* advisory committee provides the *CISE* with recommendations on support policies and facilitates *NSF's* responses to changes in the field.

Renato Figueiredo was selected to co-chair a subcommittee of the *NSF Task Force on Cyberlearning and Workforce Development*. The task force brings together a panel of experts to identify potential ways in which advanced computing and communications technologies might be leveraged to support learning, highlighting opportunities for further research.

Facilities

The computing machine room at *Rutgers University* has been entirely renovated. The renovated machine room at *Rutgers* includes a new 80-ton computer room air conditioner (CRAC) and a new state-of-the-art Dell cluster consisting of two Dell M1000E Modular Blade Enclosures, necessary interconnect/management infrastructure, and a supervisory node. Each enclosure is maximally configured with sixteen blades, each blade having two Intel Xeon E5504 Nehalem family quad-core processors at 2.0 GHz, forming an eight-core node. Each node (blade) has 6 GB RAM and 73

GB of local disk storage (10,000 RPM). The network infrastructure is comprised of an integrated 16-port Mellanox InfiniBand switch within each blade chassis, each switch having eight uplink ports and linked via eight InfiniBand lanes to the uplink ports on the switch in the other chassis. All blades have Mellanox Quad-Data-Rate (QDR) InfiniBand interface cards. There is also an integrated (redundant) 1-Gigabit Ethernet within each chassis, with two pairs of 10-Gigabit uplink capability in each chassis. In the aggregate, the new state-of-the-art cluster consists of 32 nodes, 256 cores, 192 GB memory and ~2.5 TB disk capacity, with a 20-Gigabit InfiniBand network and two 1-Gigabit Ethernet networks.

Did you know?

CAC facilities include dedicated space for visitors from member companies. If you are interested in visiting CAC, please contact a CAC faculty member.

Renovation on the new CAC suite at the *University of Florida* was completed in September. The office and meeting space is now home to five graduate students and one staff member. The reconfigurable suite also hosted the CAC semiannual review meeting and the Workshop on Intercloud Computing (see page 11 for more on these meetings). A weekly seminar is held in the new space to allow CAC students, faculty and visitors to speak about their research.

Computing facilities at the *University of Florida* site have been significantly increased since the summer. Improvements include installation or expansion of five clusters totaling 72 new machines and 796 cores. CAC @ *UF's* first 10-Gigabit/sec connection has been installed, providing a two-hop connection on to the *National LambdaRail* and *Internet2* networks. Two 60-amp and three 30-amp circuits have been run to power two new racks of equipment. The *UF* cluster associated with the nationwide *Future Grid* project is now online and conducting intercloud research. A cluster tied to the *Archer* computer architecture community project is now running and available for *Simics* and other simulation jobs. A cluster



Figure: Photos of the new computing facilities at the Rutgers University site. At far left is a close-up of the new state-of-the-art Dell cluster.



Photo: A weekly seminar held in the new CAC suite at the University of Florida site.

dedicated to the NUMAcloud project has added a marine bioscience lab to its list of customers running virtualized HPC workloads. The hardware acquired as part of the Center-wide cross-layer autonomies testbed is fully functional and has been used to survey several cloud middleware applications; it is just now being turned in to its final production configuration.

Visitors

Pierre Riteau, a Ph.D. student in Computer Science at the University of Rennes 1 in France, visited the CAC UF site this summer, where he collaborated with CAC researchers on intercloud computing. He was instrumental in the deployment of intercloud testbeds that included three FutureGrid sites in the US and Grid 5000 in Europe. "This summer at the University of Florida was very productive, and we did some amazing work. I look forward to continuing my collaboration with CAC," Riteau said of his visit.

Awards

MSU Site Director Ioana Banicescu and MSU faculty member **Sherif Abdelwahed** each received the State-Pride Award and Hearin Faculty Excellence Award from MSU's Bagley College of Engineering. The StatePride campaign is a four-year fundraising effort and recognizes faculty for outstanding performance in teaching, research and service. The Robert M. Hearin Foundation of Jackson, Mississippi sponsors the Faculty Excellence Awards, which are awarded annually in recognition of professional excellence in teaching and research.

Professor Manish Parashar is among the recipients of the 2010 IBM Faculty Award. This is a prestigious award that recognizes outstanding faculty and promotes innovative and collaborative research in disciplines of mutual interest between IBM and university faculty. The IBM Faculty Awards Program is a competitive world-wide program. Dr. Parashar's specific award is to support research on "autonomic data management for data-intensive application workflows."

In its November 2010 meeting, the Board of Directors of IEEE conferred the grade of IEEE Fellow upon **Professor Manish Parashar** of the Rutgers University "for contributions to parallel and distributed computing." According to IEEE, "the grade of IEEE Fellow recognizes unusual distinction in the profession, and shall be conferred only by invitation of the Board of Directors upon a person of outstanding and extraordinary qualifications and experience in IEEE-designated fields, and who has made important individual contributions to one or more of these fields." The number of new

Fellows in any one year cannot exceed one-tenth percent of the total voting IEEE membership.

CAC @ UF student **Prapaporn Rattanatamrong** was awarded the Engineering Outstanding International Student Award by UF's International Center. This award is given in recognition of excellence in research, teaching, and/or service and was presented to Ms. Rattanatamrong at a ceremony on November 18. UF's College of Engineering awarded a scholarship to CAC @ UF student **Selvi Kadirvel** at the same ceremony for her excellent contributions and service to the College.



Photo: CAC @ UF students Selvi Kadirvel and Prapaporn Rattanatamrong (second and third from left, respectively) being recognized at the UF International Center awards ceremony. They are pictured with Dr. David J. Sammons, Dean of the International Center (left) and Ms. Debra Anderson, Coordinator of International Student Services (right).

Become a CAC industrial member

CAC members collaborate with and advise researchers to create a diverse, industry-relevant research program. Members are afforded access to leading-edge developments in autonomic computing and to knowledge accumulated by academic researchers and other industry partners. The annual membership fee of \$35K allows industry partners to reap the full benefits of CAC's full research program, with an operating budget of over \$1 million annually. To inquire about membership in CAC, please contact Center Director Jose Fortes at fortes@ufl.edu.

Face to Face

CAC researchers are committed to promoting the Center, addressing industry-relevant research problems, and furthering collaboration. We achieve this in part by attending conferences to increase the visibility of CAC projects and attract new members, and by visiting CAC member companies to boost industry participation. This section of the Bulletin highlights the meetings between CAC personnel and the world at large.

September 2010

CAC @ UF and the Advanced Computing and Information Systems (ACIS) Lab hosted the [Workshop on Intercloud Computing](#) in Gainesville, Florida on September 30 & October 1. The primary objective of the workshop was to introduce academics and industry representatives interested in intercloud computing to an effort which the ACIS and CAC are organizing in conjunction with the IEEE Cloud Computing Standards Study Group (CCSSG). One of the goals of this effort is to develop a worldwide testbed for cloud-to-cloud interoperability. The workshop included a presentation by CCSSG Vice-Chair David Bernstein, of Huawei Corporation.

October 2010

CAC @ UF hosted a semiannual meeting in Gainesville, Florida on October 5 & 6. Forty-two attendees from CAC universities, member companies and representatives of prospective industry members gathered to hear updates on CAC projects, discuss new directions for CAC research, and review Center operating procedures.



Photo: Attendees of the CAC semiannual review meeting in Gainesville. Pictured from left to right are Vida Scarpello, Salim Hariri, Gary Whitted, Patti Duett, José Fortes, Dario Pompili, and Xiaolin Li.

Researchers from the CAC site at Rutgers University attended the [Affiliated Computer Services Information Technology Outsourcing \(ACS ITO\) Client Symposium](#) with member company Xerox October 11-13 in Orlando, Florida. Xerox, ACS and Rutgers showed a live demo at the symposium on cloud bursting using CometCloud on a hybrid computing infrastructure. The demo is the product of a CAC research project and is considered in detail on page 6. CAC @ Rutgers also demonstrated Xerox printing technology using clouds and presented Autonomic Workflow Management in Dynamically Federated, Hybrid Cloud Infrastructures using CometCloud at a shared booth in the symposium's Tech Pavilion.

CAC @ UF students Pierre St. Juste, David Wolinsky and Prapaporn Rattanatamrong attended the [CollaborateCom](#) conference held October 9-12 in Chicago. Prapaporn gave an invited talk entitled "BMI Cyberworkstation: a Cyberinfrastructure for Collaborative Experimental Research on Brain-machine Interfaces."

Pierre and David gave presentations titled "SocialDNS: A Decentralized Naming Service for Collaborative P2P VPNs" and "On the Design of Autonomic, Decentralized VPNs," respectively.

CAC @ UF postdoctoral researcher Andréa Matsunaga attended a workshop entitled "[Cyberinfrastructure and the Dimensions in Biodiversity—Planning for Success](#)," held in Madison, Wisconsin October 13-15. The workshop convened an interdisciplinary group to recommend policies for the development of cyberinfrastructure supporting integrative research in biodiversity sciences.

November 2010

CAC faculty and students attended the [2010 Supercomputing conference \(SC10\)](#), held in New Orleans, LA November 13-19. CAC hosted a booth in the exhibit hall as part of the Center's efforts to attract top-quality graduate students and recruit new industry members. Center Director José Fortes, UF Site Director Renato Figueiredo and MSU Site Director Ioana Banicescu were all in attendance. CAC @ MSU student Srishti Srivastava presented a poster titled "Optimizing the performance and robustness of dynamic scheduling for scientific applications with large number of timesteps" at the conference.



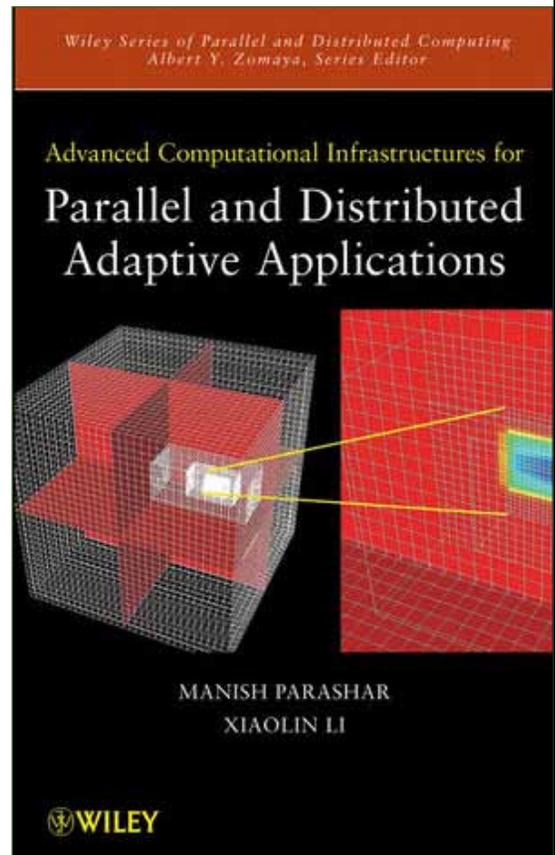
Photo: CAC Director Jose Fortes (left), system administrator Matt Collins (center) and UF Site Director Renato Figueiredo (right) at the CAC booth at the 2010 Supercomputing conference in New Orleans.

Jose Fortes and Shantenu Jha served as panelists at the panel on "Adopting Clouds" at the [2nd IEEE International Conference on Cloud Computing \(CloudCom 2010\)](#), held in Indianapolis November 30 to December 3. Their position statements can be heard and viewed at <http://www.computer.org/portal/web/csmediacenter/cloudcom2010>. UF student Selvi Kadirvel gave a presentation titled "Self-Caring IT Systems: A Proof-of-Concept Implementation in Virtualized Environments" at the conference on December 2.

December 2010

UF student David Wolinsky and Research Scientist Maurício Tsugawa attended the [The 2nd IEEE International Workshop on Management of Emerging Networks and Services \(IEEE MENS 2010\)](#), to be held in Miami December 6-10. David gave a presentation entitled "Grid Appliance—On the Design of Self-Organizing, Decentralized Grids." Maurício presented "User-level Virtual Networking Mechanisms to Support Virtual Machine Migration Over Multiple Clouds."

1. R. Mehrotra, A. Dubey S. Abdelwahed and A. Tantawi, "A Power-aware Modeling and Autonomic Management Framework for Distributed Computing Systems," to appear in the Handbook of Energy-Aware and Green Computing, 2011.
2. Venkata Krishna Nimmagadda, Ali Akoglu, Salim Hariri, and Talal Moukabary, "Cardiac Simulation on Multi-GPU Platform," to appear in the Springer Journal of Supercomputing.
3. M. Mallouhi, Y. Al-Nashif, D.Cox, T. Chadaga, and S. Hariri, "A Testbed for Analyzing Security of SCADA Control Systems (TASSCS)," to be presented at the second IEEE PES Innovative Smart Grid Technologies Conference (ISGT 2011), Anaheim, CA, January 2011.
4. S. Srivastava, F. M. Ciorba and I. Banicescu, "Employing a Study of the Deterministic Metrics to Assess the Robustness of Dynamic Loop Scheduling," in Proceedings of the 17th Annual IEEE High Performance Computing Symposium(HiPC), Goa, India, December 2010.
5. Jing Xu and José Fortes, "Multi-objective Virtual Machine Placement in Virtualized Data Center Environments," to be presented at the 2010 IEEE International Conference on Green Computing and Communications (GreenCom 2010), Hangzhou, China, December 2010.
6. Bithika Khargharia, Haoting Luo, Youssif Alnashif, Salim Hariri, "AppFlow: Autonomic Performance-per-Watt Management of Large-Scale Data Centers," in Proceedings of the 2010 IEEE/ACM International Conference on Green Computing and Communications (GreenCom2010), Hangzhou, China, December 2010.
7. David Wolinsky, Kyungyong Lee, P. Oscar Boykin, Renato J. Figueiredo, "On the Design of Autonomic, Decentralized VPNs," presented at CollaborateCom 2010, Chicago, IL, October 2010.
8. Pierre St. Juste, David Wolinsky, Kyungyong Lee, P. Oscar Boykin, Renato J. Figueiredo, "SocialDNS: A Decentralized Naming Service for Collaborative P2P VPNs," presented at CollaborateCom 2010, Chicago, IL, October 2010.
9. E. K. Lee, I. S. Kulkarni, D. Pompili, and M. Parashar, "Proactive Thermal Management in Green Datacenters," Springer Journal of Supercomputing, 2010.
10. H. Liu, E. K. Lee, D. Pompili, and X. Kong, "Application of Infrared Camera Networks for Datacenter Thermal Management," submitted for publication.
11. N. Mahadevan, S. Abdelwahed, A. Dubey, and G. Karsai, "Distributed Diagnosis of Complex Causal Systems Using Timed Failure Propagation Graph Models," presented at the IEEE Systems Readiness Technology Conference (AUTOTESTCON'10), Orlando, FL, September 2010
12. Francesc Guim, Ivan Rodero, Julita Corbalan, Manish Parashar, "Enabling GPU and Many-Core Systems in Heterogeneous HPC Environments using Memory Considerations," presented at the IEEE 12th International Conference on High Performance Computing and Communications, September 2010.
13. Pierre St. Juste, David Wolinsky, P. Oscar Boykin, Michael J. Covington, Renato J. Figueiredo, "SocialVPN: Enabling Wide-Area Collaboration with Integrated Social and Overlay Networks," Computer Networks, Vol. 54, No. 12. (26 August 2010), pp. 1926-1938.
14. Ivan Rodero, Juan Jaramillo, Andres Quiroz, Manish Parashar, Francesc Guim, S. Poole, "Energy-Efficient Application-Aware Online Provisioning for Virtualized Clouds and Data Centers," presented at the First Annual IEEE International Green Computing Conference, August 2010.
15. Ivan Rodero, Eun Kyung Lee, Dario Pompili, Manish Parashar, Marc Gamell, Renato J. Figueiredo, "Towards Energy-Efficient Reactive Thermal Management in Instrumented Datacenters," presented at the 5th Workshop on Virtualization in High-Performance Cloud Computing, August 2010.
16. C. Docan, M. Parashar and S. Klasky, "DataSpaces: An Interaction and Coordination Framework for Coupled Simulation Workflows," presented at the High Performance and Distributed Computing (HPDC'10), Chicago, Illinois, June 2010.
17. Arjun Hary, Ali Akoglu, Youssif AlNashif, Salim Hariri, and Darrel Jenerette, "Design and Evaluation of a Self-healing Kepler for Scientific Workflows," presented at the ACM International Symposium on High Performance Distributed Computing, Chicago, IL, June 2010.
18. A. Quiroz, Parashar, M., Gnanasambandam, N., and Sharma, N, "Autonomic Policy Adaptation using Decentralized Online Clustering," to appear in Proceedings of International Conference on Autonomic Computing (ICAC), June 2010.



19. Ivan Rodero, Juan Jaramillo, Andres Quiroz, Manish Parashar, Francesc Guim, "Towards Energy-Aware Autonomic Provisioning for Virtualized Environments," presented at The ACM International Symposium on High Performance Distributed Computing (HPDC) 2010, Chicago, IL, June 2010 (poster).
20. C. Docan, F. Zhang, M. Parashar, J. Cummings, N. Podhorszki and S. Klasky. "Experiments with Memory-to-Memory Coupling for End-to-End Fusion Simulation Workflows", presented at the 10th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid'10), Melbourne, Australia, May 2010.
21. Andréa Matsunaga and José Fortes, "On the Use of Machine Learning to Predict the Time and Resources Consumed by Applications," presented at the 10th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing, May 2010.
22. S. Srivastava, I. Banicescu, I. and F. M. Ciorba, "Investigating the Robustness of Adaptive Dynamic Loop Scheduling on Heterogeneous Computing Systems," in Proceedings of the IEEE International Parallel and Distributed Processing Symposium (IPDPS-PDSEC), April 2010.
23. R. Mehrotra, A. Dubey, S. Abdelwahed, and A. Tantawi, "Model Identification for Performance Management of Distributed Enterprise Systems," Technical Report ISIS-10-104, Institute for Software Integrated Systems, Vanderbilt University, April 2010.
24. C. Docan, M. Parashar and S. Klasky, "Enabling High Speed Asynchronous Data Extraction and Transfer Using DART," *Concurrency and Computation: Practice and Experience*, DOI:10.1002/cpe.1567, March 2010.
25. I. Banicescu, H. Lim, R. Carino and S. Kim, "A Parameter Study of a Hybrid Laplacian Mean-Curvature Flow Denoising Model," *Springer Journal of Supercomputing*, March 2010.
26. Selvi Kadirvel and José Fortes, "Towards IT Systems Capable of Managing Their Health," preliminary extended abstract, presented at the 16th Monterey Workshop, Redmond, WA, March 2010.
27. Shivangi Chaudhari, "Accelerating Hadoop Map-Reduce for Small/Intermediate Data Sizes using the Comet Coordination Framework," MS Thesis, Department of Electrical & Computer Engineering, Rutgers, The State University of New Jersey, 2009. Advisor: Manish Parashar.
28. Hyunjoo Kim, Yaakoub el-Khamra, Shantenu Jha, and Manish Parashar, "An Autonomic Approach to Integrated HPC Grid and Cloud Usage," presented at the 5th IEEE International Conference on e-Science, Oxford, UK, December 2009.
29. Manish Parashar and Xiaolin Li, "Advanced Computational Infrastructures for Parallel and Distributed Applications (Wiley Series on Parallel and Distributed Computing)," Wiley, December 2009.
30. Mahendra Kumar, Richard E. Newman, José Fortes, Dennis R. Durbin, Flaura K. Winston, "An IT appliance for remote collaborative review of mechanisms of injury to children in motor vehicle crashes," in Proceedings of the 5th International Conference on Collaborative Computing: Networking, Applications and Worksharing (CollaborateCom 2009), pp. 1-9, November 2009.
31. Jing Xu, Ming Zhao and José Fortes, "Cooperative Autonomic Management in Dynamic Distributed Systems," in Proceedings of the 11th International Symposium on Stabilization, Safety, and Security of Distributed Systems, pp. 756-770, November 2009.
32. Katarzyna Keahey, Maurício Tsugawa, Andréa Matsunaga and José A. B. Fortes, "Sky Computing," *IEEE Internet Computing*, September/October 2009.
33. Hyunjoo Kim, Manish Parashar, David J. Foran and Lin Yang, "Investigating the Use of Cloudbursts for High-throughput Medical Image Registration," presented at the 10th IEEE/ACM International Conference on Grid Computing (GRID2009), Banff, Canada, October 2009.
34. A. Quiroz, N. Gnanasambandam, M. Parashar, H. Kim, and N. Sharma, "Towards Autonomic Workload Provisioning for Enterprise Grids and Clouds," presented at the 10th IEEE/ACM International Conference on Grid Computing (GRID2009), Banff, Canada, October 2009.
35. Jing Xu and José Fortes, "Multiobjective Virtual Machine Placement and Migration," ACIS Technical Report TR-09-001, September 2009.
36. Jing Xu and José Fortes, "Optimization in Autonomic Data Center Resource Management," ACIS Technical Report TR-09-002, September 2009.
37. Bithika Khargharia, Salim Hariri and Mazin S. Yousif, "An Adaptive Interleaving Technique for Memory Performance-per-Watt Management," *IEEE Transactions on Parallel and Distributed Systems*, Vol. 20, no. 7, July 2009.
38. Maurício Tsugawa, "On the Design, Performance, and Management of Virtual Networks for Grid Computing," Ph.D. dissertation, [Gainesville, Fla.]: University of Florida, July 2009.
39. N. Gnanasambandam, M. Parashar, A. Quiroz, H. Kim, and N. Sharma, "Towards Autonomic Workload Provisioning for Enterprise Data Centers and Clouds," presented at Hot Topics in Autonomic Computing Workshop at ICAC 09, Barcelona, Spain, June 2009.
40. Yaser Jararweh, Arjun Hary, Youssif Al-Nashif, Salim Hariri, Ali Akoglu, Darrel Jenerette, "Accelerated Discovery through Integration of Kepler with Data Turbine for Ecosystem Research," presented at the 2009 IEEE/ACS International Conference on Computer Systems and Applications, May 2009.

41. A. Quiroz, N. Gnanasambandam, M. Parashar, and N. Sharma, "Robust Clustering for the Management of Self-monitoring Distributed Systems," Springer Journal on Cluster Computing, vol. 12, no. 1, p. 73, February 2009.
42. Pierre St. Juste, David Wolinsky, Jiangyan Xu, Michael Covington, Renato Figueiredo, "On the Use of Social Networking Groups for Automatic Configuration of Virtual Grid Environments," presented at Grid Computing Environments (with SC-2008), November 2008.
43. Jing Xu, Ming Zhao, José Fortes, Robert Carpenter, Mazin Yousif, "Autonomic Resource Management in Virtualized Data Centers Using Fuzzy Logic-based Approaches", Cluster Computing, vol. 11, no. 3, September 2008.
44. Renato Figueiredo, P. Oscar Boykin, Pierre St. Juste, David Wolinsky, "Social VPNs: Integrating Overlay and Social Networks for Seamless P2P Networking," presented at IEEE WETICE/COPS 2008. Best paper award, June 2008.
45. Salim Hariri, Yaser Jararweh, Yeliang Zhang, Talal Moukabary, "Physics Aware Programming Paradigm: Approach and Evaluation," in Proceedings of the 6th International Workshop on Challenges of Large Applications in Distributed Environments, June 2008.
46. Bithika Khargharia, Salim Hariri, Wael Kdouh, Manal Hourri, Hesham El-Rewini, Mazin Yousif, "Autonomic Power and Performance Management of High-performance Servers," presented at IEEE International Symposium on Parallel and Distributed Processing, Miami FL, April 2008.
47. Jing Xu, Ming Zhao, José Fortes, Robert Carpenter, Mazin Yousif, "On the Use of Fuzzy Modeling in Virtualized Data Center Management," in Proceedings of the Fourth international Conference on Autonomic Computing, June 2007.