CAC bulletin

The Newsletter of the Center for Autonomic Computing

Vol. 3, Summer 2011

In this issue:

Save the date! The next CAC review meeting will be held October 19-21 at Biosphere 2 in Tucson, Arizona. Stay tuned for details about the upcoming meeting and Center happenings at http://www.nsfcac.org.

Message from the Director

Dear Friends,

This is the Summer 2011 newsletter of the Center for Autonomic Computing (CAC) in which you will find the latest news about people, projects and activities at the four sites of the CAC—University of Florida, University of Arizona, Rutgers University and Mississippi State University. The Center is now in its fourth year of operation and its members are as busy as ever—new projects are being launched, unique infrastructure is being deployed, and new students are joining just as the 2011-2012 academic year is set to start.

The first 2011 biannual review meeting took place in Dallas on April 4-6. It was hosted by CAC industrial member Xerox—on behalf of the entire CAC family, my thanks to Xerox for hosting the meeting. As in past meetings, we continued to have live demos of the research results and software produced by CAC researchers, in addition to reports on ongoing projects and discussions of new proposed projects, some of which you can read about in this newsletter. Last but not least, we had presentations by Xerox researchers Nathan Gnanasambandam and Andres Quiroz about Xerox activities related to cloud computing. CAC members are able to review the proceedings of this and previous review meetings at the CAC web site.

The second 2011 biannual review meeting will be hosted by the University of Arizona at the Biosphere on October 19-21. This is the second time we will hold the meeting at the Biosphere—judging from the success of the first meeting and how everyone enjoyed the facilities, we are looking forward to going back to that unique and inspiring research facility. We are already planning for the 2012 review meetings; if you are interested in hosting such a meeting or have questions on how it could take place at your site, please do not hesitate to contact me.

This newsletter includes for the first time a communication from one of the CAC industrial members. My thanks to Don Cox from Raytheon for his very interesting article on his research experience with the Center. This is a feature we will include in every newsletter and I reiterate my invitation to other CAC members to share their views and favorite initiatives with the readership of this newsletter.

CAC continues to be successful in attracting supplemental funding for its activities. During this summer, the National Science Foundation has funded a CAC proposal to the competetive Fundamental Research Program for research on Unified Cloud Computing and Management; a collaborative award with the AIST laboratory in Japan on IT Virtualization for Disaster Mitigation and Recovery; and an Accelerating Innovation Research award in collaboration with the Center for Child Injury Prevention Studies (CChIPS) at The Children's Hospital of Philadelphia (CHoP) on Evaluation of Online Health and Wellness Promotion Applications. In addition to expanding the technical portfolio of the CAC, these awards also increase the economic leverage that make industry membership in CAC such an attractive value.

In closing, I reiterate my invitation for contributions from our CAC collaborators to the content of this newsletter. The newsletter now reaches upwards of 200 selected readers in industry and academia working in the broad areas of cloud computing, cybersecurity and intercloud computing from an autonomic computing perspective. Under this scope, contributions can be in the form of short articles, new initiatives, news about people, etc. As always, feel free to contact me at fortes@ufl.edu with any questions or comments.

The next CAC newsletter will be published in December 2011—until then, enjoy the rest of the summer and have a productive second half of 2011.

Sincerely,
José A. B. Fortes, CAC Director

Government, Industry & Academia: a Successful Partnership at the CAC

This feature article is contributed by Dr. Don Cox, Principal Engineer at Raytheon Missile Systems, Inc. in Tucson, AZ. Don is the chair of CAC's Industry Advisory Board (IAB) and graduated this year with his Ph.D. from the CAC University of Arizona site. In the following, Don discusses the outcomes of his doctoral thesis research and the benefits for industry members of collaboration with the Center.

Abstract

For almost four years, Raytheon Company and the National Science Foundation-sponsored Center for Autonomic Computing (CAC) at the University of Arizona have collaborated on the development of autonomic computing technology, the application of which is focused on critical infrastructure protection (CIP). Critical infrastructures are defined as the basic facilities, services and utilities needed to support the functioning of society. This relationship has matured to the point that Raytheon is co-sponsoring five additional projects covering a broad range of basic autonomic research topics: Application Security Management, Trust Agents, Network Defense, Autonomia, MANET Anomaly Detection and Game Theory Based Risk and Impact Analysis.

This researcher's work centered on the application of autonomic computing to protect industrial control systems commonly found in the critical infrastructure. We began with the thesis that infrastructure control systems can employ autonomic computing technology to detect anomalies and mitigate process disruption. In this article, we explore the progress of CIP research at the CAC, the successful relationship forged by the CAC with industry sponsors, and the beneficial impact on individual researchers.

Introduction

Critical infrastructures are defined as the basic facilities, services and utilities needed to support the functioning of society. A partial list includes: electrical power generation and distribution systems; telecommunications; transportation; water and wastewater treatment; and government services. For over six thousand years, civil engineers have built mechanical, and more recently, electrical systems to bring the needed services to make mankind more comfortable, secure and productive. During the 20th century, Western society's economic prosperity and technological advances have resulted in the construction of the most advanced societal infrastructures in the world.

Before 1980, few infrastructure process control systems were computerized. Process control was exerted via direct-wired analog or other electro-mechanical controllers. When microcomputers became available for infrastructure use, control engineers quickly incorporated these labor- and cost-saving control engines (Programmable Logic Controllers (PLC)) to automate infrastructure processes. Supervisory Control and Data Acquisition systems (SCADA) interconnect the various individual process controllers into an integrated control system where operators maintain (through human/machine interfaces (HMI)) situational awareness in addition to automating operational data collection and reporting generation.

Transported from the factory floor to electrical substations and other infrastructure facilities, industrial control systems were designed with a focus on performance, with little concern regarding cyber security. The fact that, after 30-plus years, the majority of the control systems designed and installed in the 1980s remain in operation with only minor maintenance and modification, is a testament to the engineering, quality and reliability of infrastructure control system technology.

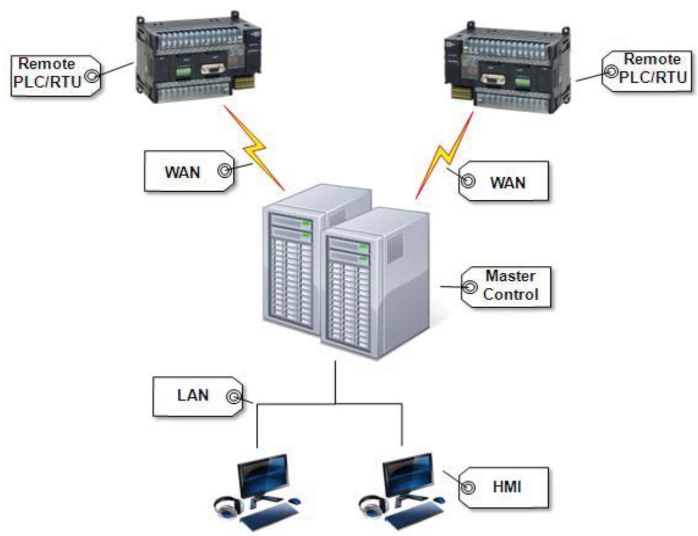

Figure: A typical SCADA system

However, as more sophisticated computers and communication technology becomes available to the industry, the longevity of early computerized control systems is creating serious, unforeseen vulnerabilities in the areas of cyber and physical security [1].

Cyber Threats and Defense

Few perpetrators of crime or mischievousness would consider causing damage to the infrastructures they depend on for water, power, etc., as disruption of one or more of these infrastructures would result in economic loss and potentially loss of life. Before the advent of the Internet, direct physical damage was infrequently experienced. By the closing of the century, new threats had appeared. Perpetrated by agents intent on destroying our society, control systems of our critical infrastructures are now experiencing malicious cyber attacks [2].

Information technology engineers have been moderately successful in defending business data networks against the plethora of cyber attacks that daily assault industry, government and academia [3]. These defensive and recovery efforts are good, and are needed in business and data processing networks. When a malicious agent is detected, technicians can simply reload the affected computer from clean backup copies of operating systems, business applications and critical data. Typically, this restoration takes many hours and the careful attention of IT professionals. This scenario is played out daily in many companies [4]. However, IT solutions to cyber insult are not effective in infrastructure control systems.

Many critical processes within our infrastructures are continuous (e.g., electric power, wastewater treatment, etc.) and cannot be interrupted at will without consequence to industry and the public [5]. As evidenced by the STUXNET worm, harm to infrastructures' physical components may occur before system operators can determine the source [6]. Researchers and developers have been lax over the years in addressing the cyber vulnerabilities of infrastructure control systems common to our critical infrastructures. Failure to protect the critical infrastructures from cyber assaults will result in physical, economic and social impacts, extending from the local level to the national level [7].

Autonomic computing, a systemic view of computing modeled after self-regulating biological systems, is an emerging sub-discipline of computer engineering. Conceived in 2001 to address the increasing complexity of computerized systems, the goal of autonomic computing is to develop computers that are capable of managing themselves. The basic functionality required of self-managed computerized systems is: self-configuration, self-protection, self-healing and self-optimization [8]. To enable these functions, critical operational attributes are used in a closed-loop feedback configuration that acts in response to changing behaviors within the controlled systems. Anomaly-based detection techniques, decision fusion, autonomic risk and impact analysis, and autonomic control algorithms are key to the successful application of this technology to infrastructure control systems.

What's the Problem?

In August 2003, the commercial electric power grid across the Northeast and Midwest United States and Ontario, Canada failed. A cascading overload, resulting from a utility in Ohio failing to keep trees trimmed out of power lines, ultimately resulted in the shutdown of 265 power plants and the blackout of over 50 million electric power customers. Because the grid safety cut-outs operated as designed, there was no significant equipment damage. Most customers regained power within days; however, some infrastructure plants required two weeks to regain full capacity.

That said, critical national security systems failed for lack of electric power: U.S. border check systems were not fully operational, water and sewage plants shut down, gas stations stopped working, and rail service was curtailed. Most cellular phone providers, radio stations and television stations found their backup power systems were insufficient and failed. The blackout is estimated to have caused economic losses of $7 to $10 billion [9]. Additionally, other sectors are experiencing significant economic losses due to down-time [10].

In a letter to the 111th Congress, U.S. House of Representatives, Committee on Homeland Security, the co-sponsors of The Critical Electric Infrastructure Protection Act, dated May 6, 2009, stated the following (in part):

> "Intentional and unintentional control system failures on the electric grid can have a significant and potentially devastating impact on the economy, public health and national security of the United States. For a society that runs on power, the discontinuity of electricity to chemical plants, banks, refineries, hospitals, and water systems present a terrifying scenario. Economist recently suggested that the loss of power to a third of the country for three months would result in losses over $700 billion." [11]

Since 2000, there has been an increase in cyber attacks on the operating systems of major critical infrastructure facilities, specifically the power grid. It was reported in the Wall Street Journal that senior Central Intelligence Agency official Tom Donohue told a meeting of utility company representatives in New Orleans that cyber attacks had taken out power equipment in multiple regions outside the U.S. The outages were followed with extortion demands [12].

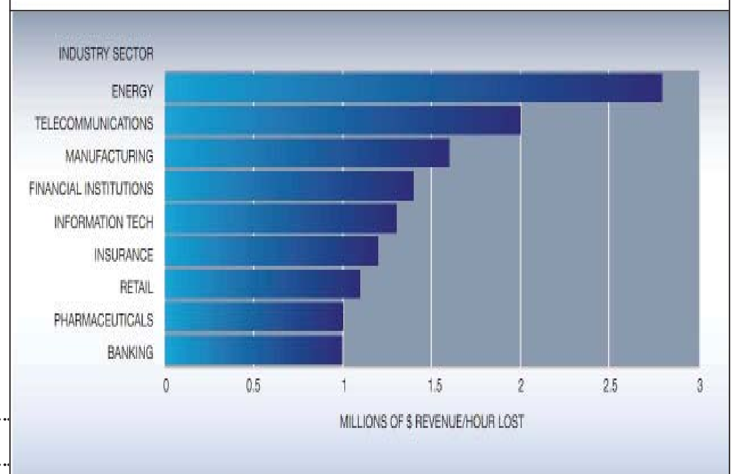

Data from *IT Performance Engineering and Measurement Strategies: Quantifying Performance Loss*, Meta Group, Stamford, CT (October 2000).

Figure: Infrastructure down-time revenue losses (2000)

Succinctly stated, the control systems of modern critical infrastructures are vulnerable to disruption from natural disaster, accidental or negligent operation and maintenance, and intentional assaults from malicious agents. Existing critical infrastructure vulnerabilities are placing at risk the continued existence and safety of Western society as we enjoy it today. It is therefore imperative that new technology be developed and rapidly implemented to protect the control systems of our critical infrastructures from failure, as well as mitigate disruption in the event of natural or man-made assault.

CAC/Industry Cooperation

The CAC research activities are aimed to develop technology ultimately of use to the industry members. Guided by an Industry Advisory Board (IAB) staffed by representatives from each industry member, the CAC university research centers select research topics relevant to the interest of industry.

In 2008, CAC founding IAB member Raytheon (Missile Systems, Tucson, AZ), a leading defense contractor, recognized the vulnerabilities of the critical infrastructures, and sponsored research to develop new and innovative technology to protect and defend the industrial control systems commonly used to control these systems. Another member of the IAB, Avirtek, Inc. (Tucson, AZ), shared this interest and created a partnership[1] to sponsor this research. Graduate students were recruited to participate in this activity. Under the direction of this author, UA Site Director Salim Hariri, industry partners and graduate students coalesced into a research team capable of successfully completing the research.

A Bold Thesis

To answer the need to protect critical infrastructure control systems from cyber assaults we focused our work proving the thesis:

> "Critical infrastructure control systems can employ autonomic computing technology to detect and mitigate operational disruption caused by natural disaster, accident or malicious cyber assaults."

Specifically, it is hypothesized that autonomic computing algorithms can be applied to infrastructure control and digital communications networks to: 1) detect and respond to anomalies caused by cyber assaults; and 2) detect and block infrastructures controller commands that if executed would result in process disruption.

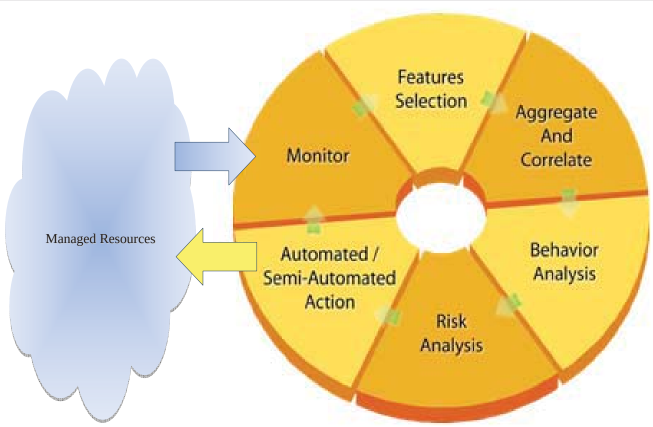

Figure: Expanded autonomic computing model

Acknowledging a functional overlap, we condense and bound autonomic computing functionality as belonging to four basic categories: self-configuring, self-healing, self-optimizing and self-protecting. We now lay a foundation based on two assumptions: 1) any accidental or intentional cyber intrusion will cause an anomaly, and 2) infrastructure control systems can execute process control outputs that will disrupt the process under control.

Accepting these assumptions, it can be hypothesized that: 1) computerized control systems can be developed having the capability to use operational process metrics to create a representative process model that accurately reflects the instantaneous state of the process under control; 2) that if a process model is created using data representing the desired "normal" operational state, that the "normal" process model can be used as a standard to evaluate all process states; and 3) at any instant in process operation, by comparing the current process state with the normal process model, process anomalies will be detected and can be quantified.

Given this anomalous process state data, corrective control output may be chosen from a library of possible responses or policies, and implemented by the control system. If, for the purpose of experimentation, we accept this hypothesis as valid, it then becomes the foundation of an autonomic computing paradigm from which detection, analysis, policy-based mitigation action selection and implementation algorithms can be derived. Addressing all of the possible autonomic control functionality as part of this research is impractical and outside the resources available. We therefore constrained our research to elements of autonomic self-protection functionality.

Research and Experimentation

The goal of the work was to develop autonomic computing technology to protect critical infrastructure control systems from cyber assaults. The work associated with this research thesis contributed the following to the body of computer and autonomic systems engineering knowledge:

Initial data collected by a critical examination of the critical infrastructures of Western society is used to identify the a) taxonomy; b) interrelationships and interdependencies; and c) control systems architecture. This effort clarified the societal impact of infrastructure failure; identified the "most" critical infrastructures; as well as categorized common control system architectures.

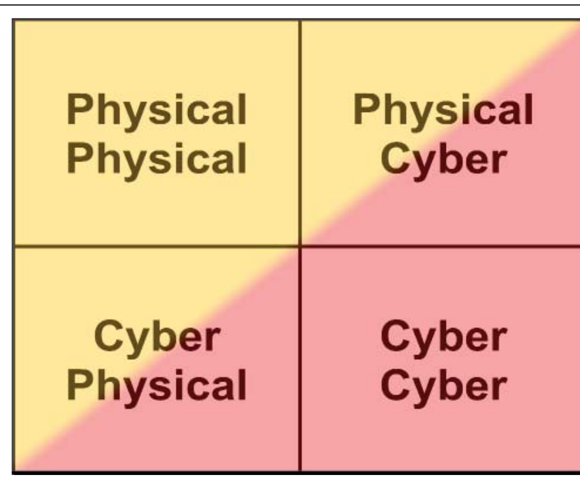

Figure: Attack methods/effects ontology

Using the data from the above analysis, we developed an ontology that maps infrastructure assault methodology with perpetrators' motivation and goals. Reducing the data, we derive four categories: physical assault resulting in physical damage; physical assault that resulted in damage to the cyber infrastructure; cyber assault resulting in physical

[1] The term "partnership" is used in the context of mutually beneficial cooperation and should not be interpreted to imply a contractual relationship in the legal sense.

damage; and lastly cyber assault resulting in damage to the cyber systems or data.

The ontology is used to determine the most effective insertion point for defensive technology within infrastructure control systems. The information is also important to the operators of infrastructures to understand the vulnerabilities inherent in operating and maintaining infrastructure systems.

Following the definition of the Attack/Effects Ontology, the team created a theoretical framework to support the creation of autonomic appliances for detection and mitigation of cyber assaults, and unauthorized or disruptive commands targeted at industrial control systems, typical of critical infrastructure facilities and processes. These appliances focused on detecting and defending the cyber/physical and cyber/cyber attacks.

The resultant technology consists, in part, of autonomic software developed at the CAC (University of Arizona) as part of this research. We leveraged Autonomic Network Defense system code from Avirtek (also a CAC industry member) to host our newly developed algorithms. Computerized models of several autonomic protection appliances were designed to demonstrate autonomic behavior in contemporary and legacy control network protocols.

Figure: Critical infrastructure protection testbed (as built)

To provide a controlled test environment, a unique Critical Infrastructure Protection Test Bed (CIPTB), was constructed. The CIPTB offers computer-in-the-loop integrated with hardware-in-the-loop capabilities for modeling and simulating critical infrastructures (infrastructures and control systems). A key functionality of the CIPTB is the capability to be isolated from inter- and intra-net connectivity to create a safe simulated environment whereby live malware could be released and controlled for scientific and operational analysis.

Contributions to Science

While personally satisfying for the research staff and offering a potentially high return on investment for industry members, this research also contributed significantly to the body of knowledge relating to autonomic computing and its application to protecting critical infrastructure control systems. A short list follows:

- Demonstrated that human control operator intervention is ineffectual to prevent or mitigate process disruption when digital command networks, are blocked by cyber assaults resulting in a denial of service conditions.
- Demonstrated that autonomic appliances can evaluate the temporal relevance of controller commands using a sliding time-gated window based on policy definitions.
- Integration of the multiple autonomic functionalities into a single control system protection appliance yields capabilities including detection of anomalous behaviors of TCP/IP digital communications networks and real-time analysis and discrimination of serial control network commands when used in SCADA and infrastructure control networks.
- Policy-directed response by autonomic protective appliance following detection of disruptive SCADA commands.
- Automated creation of cyber anomaly detection (parametric and visualization data) database supporting real-time and forensic analysis.
- Design and construction of an unique test bed for the development and demonstration of critical infrastructure protection technology based on autonomic computing paradigms [13].

Benefits of Collaboration

Aside from the technical results of this research, a foundational relationship was developed to the benefit of the industry partners and the researchers themselves. Industry partner Raytheon found the CAC (University of Arizona, Autonomic Computing Laboratory), under the direction of Co-director Dr. Salim Hariri, to be open and accessible.

Many progress review meetings were held at the university, opening channels of communications and opportunities for future research. Additionally, interested managers and engineers from Raytheon facilities located in Florida, Texas and Indiana traveled to Tucson, AZ to view demonstrations and understand the technology being developed.

Another benefit of conducting research within the CAC milieu is the educational advantages. Employed by Raytheon as a Principal Systems Engineer, this author was assigned to coordinate the CIP research. Coincidentally, this author was working on a Ph.D. in Electrical Computer Engineering. The timing of the Raytheon/CAC research permitted completion his Ph.D. research, which is detailed in this article. This author currently teaches an undergraduate engineering course at the University of Arizona as well as maintaining his employment at Raytheon.

Conclusion

This work successfully demonstrated through experimentation the research thesis: critical infrastructure control systems can employ autonomic computing technology to detect and mitigate the effects of operational disruption caused by natural disaster, accident or malicious cyber assaults. Additionally, new understanding of autonomic computing theory and application to infrastructure controls has been gained, as well as the development and construction of a unique and useful modeling and simulation test bed.

Corporate, government or academic membership in the NSF CAC is a unique research opportunity found in few institutions. All parties enjoy the affiliation and benefit from the research and educational opportunities. The return on investment for industry members is significant when considering the cost of standing-up research facilities the quality of that made available by the member universities.

All things considered, this author believes the most noteworthy benefit of association with the CAC is the education that graduate and undergraduate students receive from participating in CAC research projects.

The Author

Dr. Don Cox is a Principal Systems Engineer at Raytheon Missile Systems in Tucson, Arizona and sits as the Chair of the NSF CAC Industry Advisory Board. He has over 30 years experience in defense and commercial systems engineering and corporate management. He has coauthored four patents and earned a Ph.D. in Electrical & Computer Engineering at the University of Arizona. His papers on applications of autonomia for Global Information (GIG) security and management, and Modeling & Simulation in DOD Logistics have been published in the journals of the National Defense Industrial Association, Society for Modeling and Simulation International, and Simulation Interoperability Standards Organization. He is a frequent presenter at system and software workshops, conferences and symposia. Email contact: dcox@raytheon.com.

References

1. Brown, R.E. and Willis, H.L., "The Economics of Aging Infrastructure," IEEE Power and Energy Magazine, vol.4, no.3, pp.36-43, 2006
2. Weiss, Joseph, "Protecting Industrial Control Systems from Electronic Threats," New York, NY : Momentum Press, LLC, 2010.
3. Hines, Bill, "New technologies, New Temptations," [Online] March 22, 2006. [Cited: June 6, 2011.] http://www.ibm.com/developerworks/websphere/techjournal/0603_col_hines/0603_col_hines.html.
4. Manky, Derek "Top 10 Vulnerabilities inside the Network," Network World. [Online] November 8, 2010. [Cited: May 25, 2011.] http://www.networkworld.com/news/tech/2010/110810-network-vulnerabilities.html.
5. US DHS, "Cyber Security Assessments of Industrial Control Systems" Washington, DC : DHS Center for the Protection Of National Infrastructure, 2011.
6. McMillan, Robert, "Siemens: Stuxnet worm hit industrial systems," www.computerworld.com. [Online] September 16, 2010. [Cited: March 24, 2011.] http://www.computerworld.com.
7. Balavage, Catherine, "Ten Years After 9/11, Widespread Concerns About Terrorist Attacks on Airports, Mass Transit and Public Gatherings Persist," FrostMagazine.com. [Online] May 22, 2011. [Cited: May 25, 2011.] http://frostmagazine.com/2011/05/ten-years-after-911-wide-spread-concerns-about-terrorist-attacks-on-airports-mass-transit-and-public-gatherings-persist/.
8. IBM, "Autonomic Computing, Overview," [Online] May 25, 2011. [Cited: June 6, 2011.] http://www.research.ibm.com/autonomic/overview/.
9. U.S.-Canada Power System Outage Task Force, "Final Report on the August, 14 2003 Blackout in the United States and Canada: Causes and Recommendations," Washington, DC : U.S. Government, April 2004.
10. Ganek and Corbi, "The dawning of the autonomic computing era," IBM Systems Journal, vol. 42, no. 1, 2003.
11. Thompson, Benny G. and King, Peter T., 20090721142006-31951[1].pdf. ecopy, INC. [Online] July 21, 2009. [Cited: July 30, 2011.] http://chsdemocrats.house.gov/site documents/20090721142006-31951.pdf.
12. Gorman, Siobahan, "Electricity Grid in US Penetrated by Spies," Wall Street Journal, online. [Online] April 8, 2009. [Cited: June 30, 2011.] http://online.wsj.com/article/SB123914805204099085.html.
13. Mallouhi, M. and Al-Nashif, Y. and Cox, D., Chadaga, T. and Hariri, S., "A Testbed for Analyzing Security of SCADA Control Systems (TASSCS)," in Proceedings of 2011 IEEE PES Innovative Smart Grid Technologies, 2011. DoI:10.1109/ISGT.2011.5759169.

Moustafa Abdelbaky
Rutgers University

Moustafa is a Ph.D. student and a member of the CAC site at Rutgers University. His research interests are in parallel and distributed computing, cloud computing, and mobile development. He is currently working closely with researchers at IBM's Thomas J. Watson Research Center on developing a framework to provide HPC as-a-service (HPC Clouds) using supercomputing resources.

Moustafa has B.S. degrees in both Computer Science and Electrical and Computer Engineering from Rutgers University. He recently received the prestigious and highly competitive IBM Ph.D. Fellowship Award. He is pursuing his graduate degree under the direction of CAC faculty member Manish Parashar.

Zhitao Li
University of Arizona

Zhitao Li received his Bachelor of Science degree in Biomedical Engineering in 2010 from Huazhong University of Science & Technology in China.

He is currently working towards his Master of Science degree in ACL lab at the University of Arizona. In the lab, he is conducting research in the area of high performance computing.

Rajat Mehrotra
Mississippi State University

Rajat Mehrotra is a Ph.D. candidate in the Department of Electrical and Computer Engineering at Mississippi State University. His research interests include model-based performance management of distributed multi-tier computing systems hosted in physical as well as virtual environments through feedback-based control-theoretic approaches for optimal resource allocation. He also performs research in the area of model based autonomic security and power-aware high performance computing systems.

Rajat has published his research efforts through research publications and a book chapter. Before joining MSU, he worked as a senior engineer for three years in Alcatel-Lucent's Wireless R&D division for NGN networks.

Girish Venkatasubramanian
University of Florida

Girish Venkatasubramanian received his Ph.D. in Electrical and Computer Engineering from the University of Florida in 2011, under the supervision of UF Site Director Renato Figueiredo. His dissertation dealt with understanding TLB-related performance delays for virtualized workloads.

Girish obtained his B.E. degree in Electrical and Electronics Engineering from PSG College of Technology, India and his M.S. in Electrical and Computer Engineering from UF.

Girish was awarded the UF International Center's Certificate of Achievement for Outstanding Academic Performance in 2004, 2005, 2008 and 2011. He was selected as an Outstanding International Student by the UF International Center in 2006.

Girish is currently at Intel, where he works on virtualization software for binary translation.

FARE-SHARE: A Resource Provisioning Framework for Mobile Grids

PI: Dario Pompili [pompili@cac.rutgers.edu]
Lead researchers: H. Viswanathan [vhari@eden.rutgers.edu]
E. K. Lee [eunkyung_lee@cac.rutgers.edu]

The mobile computing revolution is characterized by two clearly visible trends as shown in the figure below. The first trend is the increase in mobile devices' computational capabilities aimed at improving user experience. The second trend is the growing popularity of the cloud computing paradigm, which pushes processing and storage (required for running applications) to remote servers on the Internet while retaining only a light front-end on the local device. The former is market- and demand-driven and is dictated primarily by users' purchasing behavior as well as by recent innovations in human-computer interaction, while the latter can be attributed to the gains that cloud computing provides in terms of lower infrastructure costs and reduced time-to-market for innovative applications (by offering infrastructure and platform as a service). These diverging trends have rendered the powerful mobile devices heavily under-utilized.

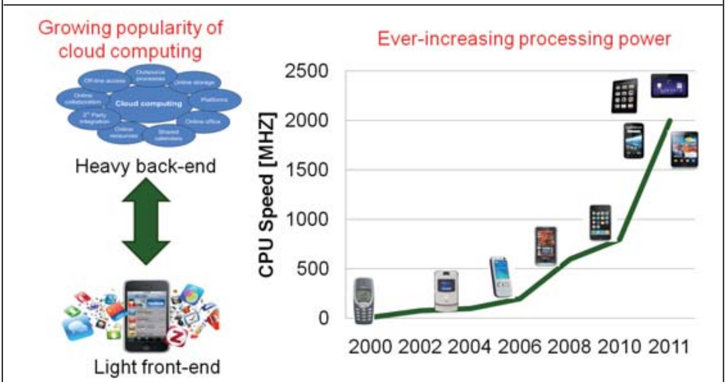

Figure: Two diverging trends characterizing the mobile computing revolution

FARE-SHARE is aimed at organizing these powerful smart hand-held devices into a mobile wireless grid [1] so that their communication and collective computational capabilities can be harnessed to enable innovative data- and compute-intensive mobile applications. The response time, quality, and relevance of data- and compute-intensive mobile applications can be drastically improved through mobile grid computing. Applications that will benefit from mobile grid computing include distributed rainfall and flood-risk estimation [2], distributed wireless channel estimation, distributed target detection and tracking, estimation of pollution level using real-time air-quality measurements (as shown at bottom left), and content-based distributed multimedia search and sharing [3], just to name a few. The communication cost involved in enabling these applications using the conventional approach [4]—aggregating large amounts of sensor data at a server for centralized computation—is prohibitive. Also, the resource capabilities of a single device may be insufficient to process all the data and produce meaningful results in realistic time bounds.

FARE-SHARE is an efficient autonomic resource provisioning framework that addresses the major research challenges associated with mobile grid computing: namely, discovery and provisioning of computing resources. FARE-SHARE strives to minimize computational load on individual mobile devices by exploiting parallelism while incurring the minimal communication cost and, hence, energy expenditure for supporting parallelism among multiple devices. It identifies two entities in a mobile grid, namely a broker and mobile devices, which may play the role of 1) consumers, which generate service requests, 2) data providers, which provide sensor data, or 3) resource providers, which provide computational, storage, and communication resources for processing the data. The broker is a remote entity whose role is to process the requests from the consumers, to determine the set of resource providers that will process the data (aided by a resource provisioning engine), and to distribute the workload among resource providers.

As the roles played by the mobile devices in mobile grid can change, FARE-SHARE relies on a credit-based system to ensure that a consumer does not gain an undue advantage by heavily consuming without providing any service. Also, such a system allows for credit transfers between multiple applications thus widening the applicability of the framework. As there will be concurrent application requests generated by a number of consumers, FARE-SHARE will multiplex these requests and "share the fare fairly" among the consumers.

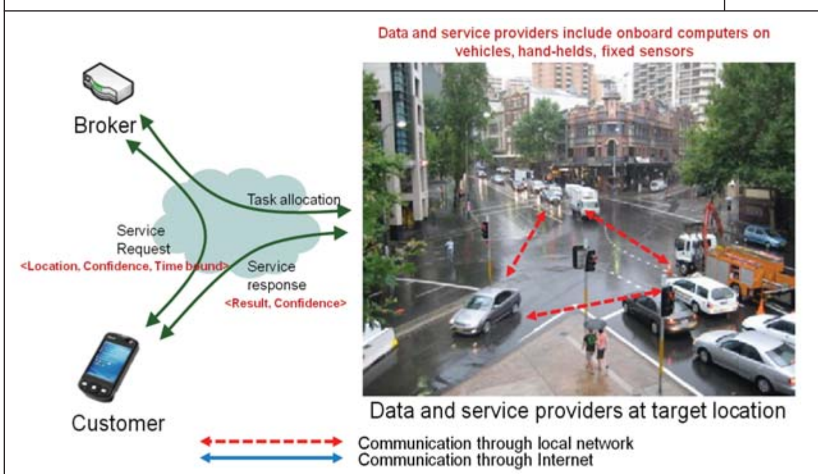

Figure: Example application—estimating pollution levels using real-time air-quality data

References:

1. G. Li, H. Sun, H. Gao, H. Yu, and Y. Cai, "A Survey on Wireless Grids and Clouds," in Proc. of IEEE International Conference on Grid and Cooperative Computing (GCC), August 2009.
2. "Urban Flooding Preparedness Research Powered by Mobile Cloud," Available online at: http://computeinmotion.com/2011/02/urban-floodingpreparedness-research-powered-by-mobile-cloud/#more-68
3. "Hyrax: Cloud Computing on Mobile Devices using MapReduce." Available: http://oai.dtic.mil/oai/oai?verb=getRecord&metadataPrefix=html&identifier=ADA512601
4. D. Quercia, N. Lathia, F. Calabrese, G. D. Lorenzo, and J. Crowcroft, "Recommending Social Events from Mobile Phone Location Data," December 2010.

Autonomic cyber-physical system for Biosphere 2 Landscape Evolution Observatory research

PI: Salim Hariri [hariri@ece.arizona.edu]
Lead researcher: Venkata Krishna [vkn@arizona.email.edu]

The interaction of the hydrological cycle with other global changes is affecting ecosystems on Earth. The Landscape Evolution Observatory (LEO) is an interdisciplinary project at Biosphere 2 (B2) aimed at quantifying the interactions between soils and biological, hydrological, geological and atmospheric processes. To develop a good understanding of these interactions of processes, there is a pressing need for experimental infrastructures which can couple complex physical systems with hydrological models. The ability to control the physical systems based on the information from monitoring systems equipped with a variety of field instrumentation will considerably accelerate the discovery cycle by iteratively testing models against experimental systems. To date, there is no complete solution for controlling and managing computational models and experimental test beds in a closed-loop form with autonomic sensing, data-model fusion, autonomic workflows and online analysis capabilities. Such a solution will accelerate research and discovery for grand earth science challenges, as well as many other cyber-physical systems, such as intelligent transportation systems, greenhouses, etc.

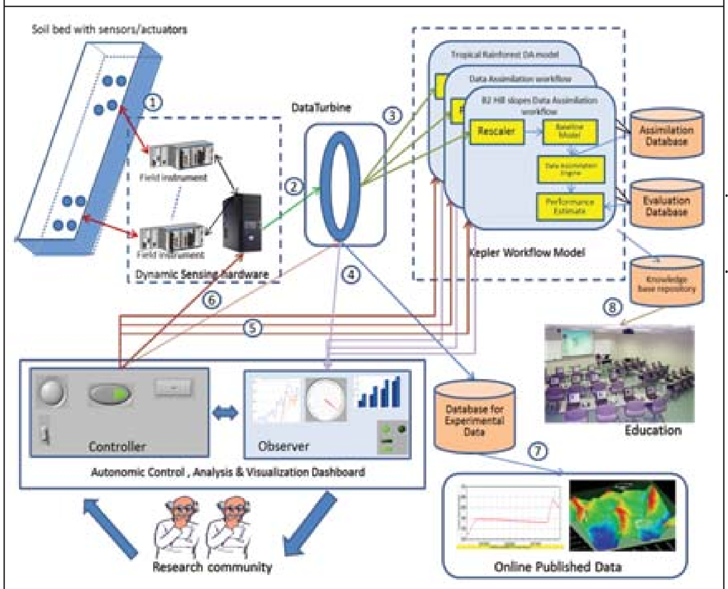

Figure: A schematic of the service-oriented cyberinfrastructure for earth science experiments at Biosphere 2.

LEO experiments are characterized by the large amounts of data streaming and processing required to effectively study hydrological cycles and their interactions with the environment. The Ring Buffer Network Bus (RBNB) Data Turbine is an open-source initiative to develop an infrastructure for data management, processing and synchronization. Kepler is another open-software tool for scientific workflow management. In our earlier work, we integrated the Data Turbine with Kepler to enable multi-channel data streaming from our experimental testbed into computational models. In parallel, we also created a physical experimental testbed inside Biosphere 2. The test bed consists of 30 sensors that monitor physical properties such as the weight, water flow, and conductivity with an operational infrastructure to simulate rain on the test bed. National Instrument's Compact RIO is programmed with LabView to collect and store data in real time. However, the experimental test bed and the workflow management components have not been integrated yet.

The most important goal of this project is to enable accelerated discovery cycles (ADCs) by coupling computational models with the unique experimental infrastructure of Biosphere 2. We propose the service-oriented cyberinfrastructure pictured at lower left for establishing a distributed system which integrates the physical infrastructure of B2 with data assimilation models executing on a virtualized cloud platform with hundreds of cores available at the Autonomic Computing Lab at the University of Arizona. Our objective is to develop strategies for robust data acquisition, autonomic decision making in sensing, data-model fusion, analysis and visualization to support the operations of the ADC.

This project involves the development of a fault-tolerant, autonomic and large-scale heterogeneous sensor network. In parallel to this activity, we will develop eco-hydrological models and the data model assimilation engine along with strategies to stream and process data using Data Turbine. We will finally develop a collaboration platform with portals and interfaces for data analysis and visualization. We will develop web tools for secure access to the experimental setup and models. Overall, the system will facilitate knowledge generation through educational outreach programs for the benefit of students and next-generation earth scientists.

Performance modeling of distributed multi-tier enterprise systems

PI: Sherif Abdelwahed [sherif@ece.msstate.edu]
Lead researcher: Rajat Mehrotra [rm651@msstate.edu]

Model-based techniques have been explored recently by researchers aiming to develop effective autonomic management techniques for multi-tier enterprise systems under uncertain and dynamic operating conditions. Additionally, the increasing size and complexity of these systems in typical dynamic and uncertain environments lead to more stringent requirements for application availability, reliability, QoS guarantee. The general aim is to minimize operational costs while maximizing a multi-dimensional QoS metric that typically includes service related factors such as response time, throughput, and reliability.

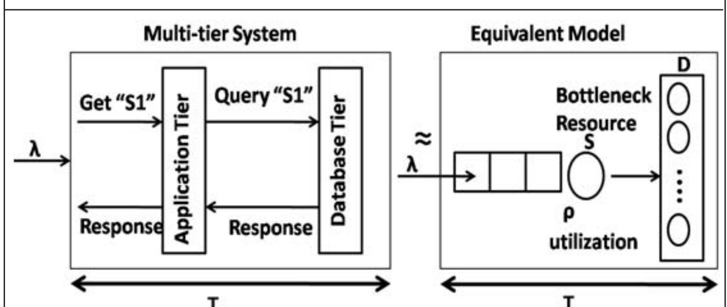

Figure: A queuing model for the two-tier system

In this project, we present a systematic approach to developing accurate models for distributed multi-tier enterprise systems. The proposed approach first identifies the system parameters through extensive experimentation and then defines the relationship among these parameters and identifies the underlying model structure of the system using regression methods and Bayesian estimation techniques. The developed model can then be used in various model-based autonomic management structures used for performance control, fault adaptation, and security management.

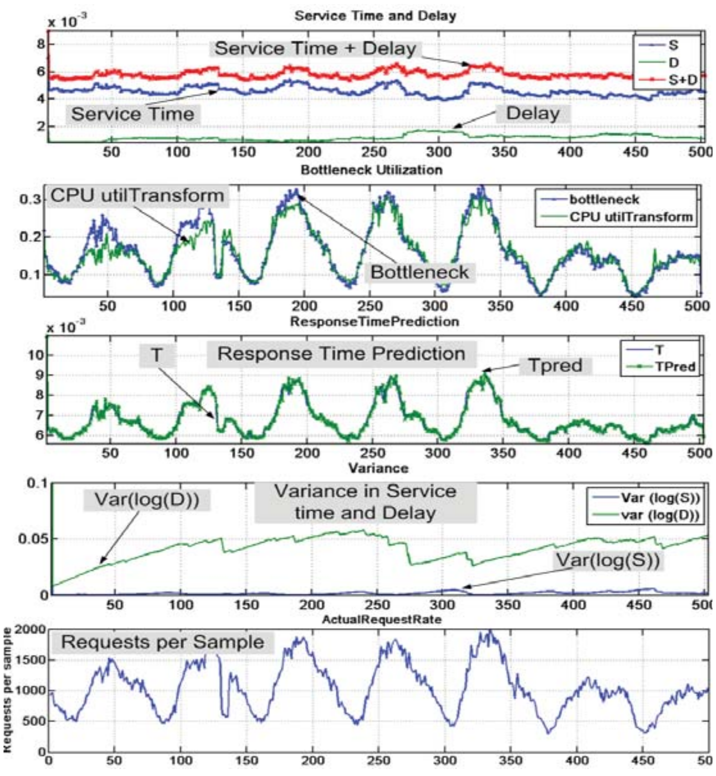

Figure: Offline Exponential Kalman output. Service time and delay are in millisecond range. Response time is specified in seconds.

The expected outcome of this project is a model-based online framework that can be utilized for any application and physical system by applying similar kinds of regression tests and building a predictive controller that can optimize the system operation with specified QoS objectives. Also, a generic mathematical model for a general class of distributed computing systems will be developed that will accurately reflect the run-time system behavior, performance, and reliability. Additionally, a simulation environment will also be developed that will utilize the above mathematical model for predicting system behavior. We have developed a systematic performance-monitoring approach for distributed systems by measuring system variables (physical/virtual CPU utilization and memory utilization), application variables (application queue size, queue waiting time, and service time), and performance variables (response time, throughput, and power consumption) accurately with minimum latency at a specified rate. Furthermore, we have shown that the proposed monitoring approach can be utilized to provide input to an application-monitoring utility to understand the underlying performance model of the system for successful on-line control of the distributed systems for achieving predefined QoS parameters. The proposed approach has low run-time overhead in terms of computation and memory resources and tracks system model parameters at run time with high accuracy. Moreover, the proposed approach is scalable for additional system parameters, and adapts quickly to the dynamic workload variations. Experiments show that the developed model captures the system behavior in varying environmental conditions with low variance.

Autonomic Cloud Management Services (ACMS)

PI: Salim Hariri [hariri@ece.arizona.edu]
Lead researcher: Farah Aldeen [falfay@email.arizona.edu]

As the internet and related applications continue to grow at a rapid rate, balancing the power consumption of memory and processor subsystems and the performance of server systems has become a primary concern for the HPC community. There is an urgent need for autonomic power management due to the expanding size of those systems. Power optimization for large-scale systems has not yet been solved due to a lack of formulated metrics and algorithms. Our goal is to extend our theoretical framework and general methodology for autonomic power and performance management for large-scale data centers in order to:

1. autonomically scale-up/down cloud resources (cores, storage, I/O devices) to meet multi-objective requirements (performance, energy, availability, etc.);
2. identify the cloud workload characteristics and their requirements at run-time;
3. identify the features to be monitored and analyzed at run-time for accurate characterization of the current operation region of cloud services; and
4. identify the structures (AppFlow, VmFlow, ServFlow, etc.) that can be used to steer the platform autonomic manager architecture. This approach will enable effective use of hardware resources under varying workload conditions so that idle hardware could be put to a low-power state and the power consumption of the system could be managed.

We have developed a framework to simulate the behavior of a computing system at the server, memory/processor subsystem and core levels. We designed a workload generator for CPU-, memory- and I/O-intensive operations. We have developed algorithms to autonomously scale CPUs and DRAMs allocated to the system for power management under various execution modes. We have developed mechanisms for monitoring and collecting data on features such as CPU load, memory usage, page faults, I/O reads and writes, etc. We have developed the AppFlow tool for tracking and predicting application and platform behavior. AppFlow successfully characterizes the dynamic behavior of applications with respect to key features being monitored. With AppFlow, we can track changes in the behaviors of applications and thereby proactively ensure the optimum performance/utility for the applications.

We will develop an autonomic management methodology to find an optimal VM configuration for a specific workload such that its service level agreements (SLAs) are satisfied.

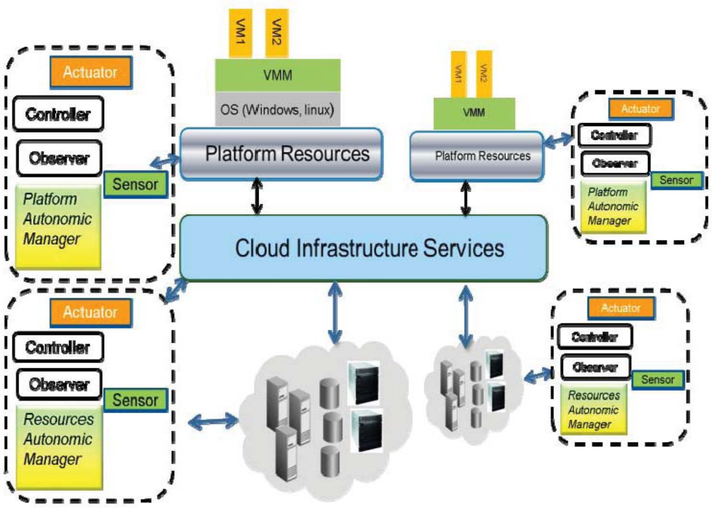

Figure: Schematic of the autonomic cloud management system.

We will evaluate a manual trial and error method and control theory-based approach for matching SLA requirements with the platform resources. We will select an appropriate VM-template based on varying load during runtime with case based reasoning (CBR). After generating workload for real cloud computing applications and services, and implementing monitoring and data acquisition into AppFlow data structure (Spring 2011), we will implement the autonomic management algorithms developed as part of our simulation framework on the IBM Cloud Test-bed (Spring 2011). We will demonstrate the performance of the ACMS and investigate ways to enhance its capabilities with adaptive learning and dynamic reconfiguration at computing, memory and network levels (Fall 2011).

Using SAGA to develop Autonomic Scientific Applications

PI: Shantenu Jha [shantenu.jha@rutgers.edu]
Lead researcher: Yaakoub El Khamra [yelkhamra@gmail.com]

The development of simple and effective distributed applications that can utilize multiple distributed resources remains challenging. Additionally, many applications have irregular and highly variable resource requirements which are very difficult to predict in advance. As a consequence of irregular execution characteristics, dynamic resource requirements are difficult to predict a priori, thus rendering static resource mapping techniques such as workflows ineffective.

The dynamic resource utilization problem—predicting requirements and assigning optimal resources—can be addressed more efficiently using autonomic approaches. However, it is difficult to implement advanced application characteristics, such as autonomic behavior for high-performance scientific applications. Notwithstanding, there exists a large class of applications which could benefit immensely with support for autonomic properties and behavior.

Using SAGA, the Simple API for Grid Applications, it is easy to construct distributed autonomic workflow management systems. The workflow we are interested in developing is an ensemble Kalman filter workflow for automatic history matching/reservoir characterization. The individual ensemble members are instances of a parallel reservoir simulator that are run in container pilot jobs. The main features are:

- distributed: different ensemble members can run on different machines
- fault tolerant: it resurrects failed jobs
- autonomic: it can adapt to changes in workload (computational cost of ensemble members) and changes in queue wait time (using the Batch Queue Prediction tool)

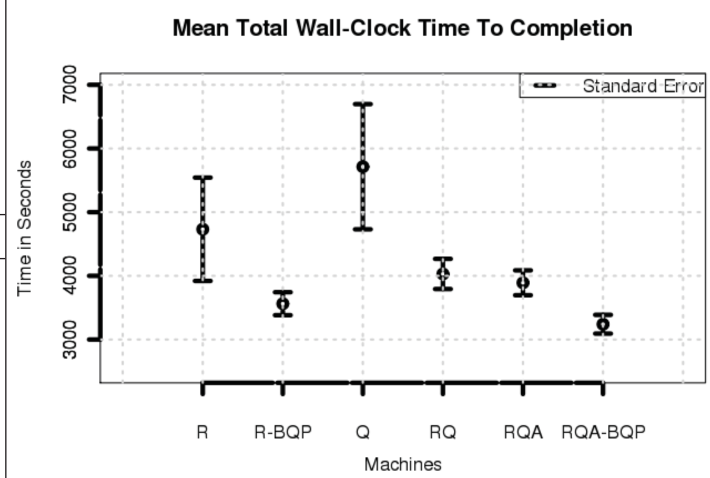

Figure: Time to completion for a task using different distributed autonomic workflow management systems.

The above figure shows the total time to completion in seconds for a typical reservoir characterization study. Using distributed computing reduces the total time to completion. Furthermore, using autonomic scheduling with BQP data, we are able to reduce the total time to completion by one quarter.

This work represents initial effort towards the development of SAGA-based middleware, services and tools to implement autonomic capabilities that can be used for scientific applications that will be deployed on national and international distributed cyberinfrastructure.

Personal device social networking with end-to-end trusted communication

PI: Renato Figueiredo [renato@acis.ufl.edu]
Lead researchers: Pierre St. Juste [ptony82@acis.ufl.edu]
Heungsik Eom [eomhes@gmail.com]

Mobile devices such as smart phones and tablets are increasingly becoming the platform of choice for many daily life tasks that computer users perform; for instance, the emerging application area of mobile augmented reality (MAR) capitalizes on the ubiquity of these devices to create rich user experiences. Future-generation devices are relevant to the platforms, software, and defense industries that form the core of the Center. This research addresses questions at the core of understanding the requirements of these platforms—such as to what extent computation should take place within or outside the platform and how the support for emerging modalities of communication among devices might impact platform design.

This project will research and develop prototypes of systems that address communication and information processing needs of future mobile augmented reality devices, including:

1. Trusted end-to-end connectivity among devices guided by social networks: a wide range of users of smart mobile devices use them on a regular basis in social networking platforms, such as Facebook, Twitter and Foursquare. With increases in wireless bandwidth, social platforms that enable device-to-device communication are poised to provide new modalities of augmented reality, where the information held in mobile devices of peers can provide additional context to be used in augmented reality. In order to tap into this potential, it becomes key to provide strong authentication and privacy in communication, in a way that is seamless and transparent to users. This task will leverage CAC research and prototypes on social virtual networks as a framework for device-device communication, and will analyze the implications in the design of the device platforms in order to support effective private virtual networking communication.
2. Cloud-provisioned data analytics/mining MAR tasks: in addition to real-time augmented reality tasks that execute on the mobile platform itself, there are several tasks that require computational and storage capacity beyond that of a mobile platform, and thus may be best served by more powerful computers. This task will research "cloud-bursting" scheduling approaches that account for the availability of resources offered by multiple resource providers, such as desktops of individual users and their social peers, low/mid-range servers of small/medium businesses, and private or public cloud infrastructures. This task will leverage CAC research and prototypes on self-organizing virtual networks and virtual appliances as a basis for seamless cloud provisioning.

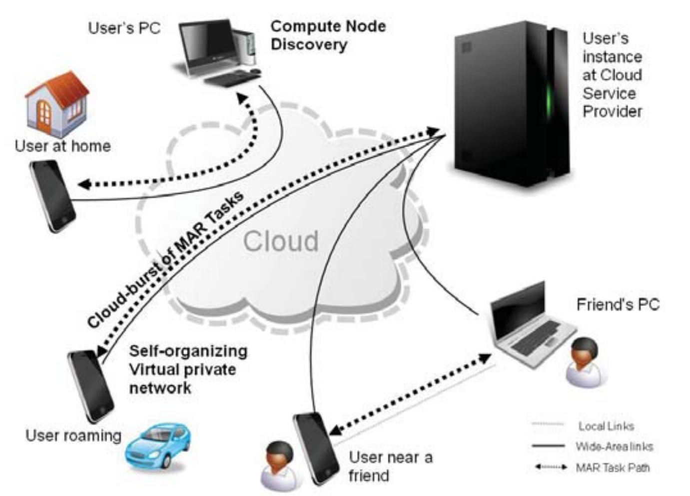

Figure: Our secure connection management framework transparently handles communication to various nodes (i.e., a cloud instance, a friend's PC, or a user's home PC). Once a connection is available, the mobile device utilizes it to offload computation for Mobile Augmented Reality tasks.

The outcomes of this project include prototype implementations of systems that allow seamless connectivity of MAR devices and off-loading analytics computation to cloud-provisioned resources.

Get involved!

CAC projects are developed in collaboration with industry participants to produce focused, industrially relevant research. CAC faculty and students gather with industry representatives twice a year to review progress on existing projects and to explore new research directions that best meet member companies' goals in the area of autonomic computing. Our next semiannual meeting will be held October 20 & 21 at Biosphere 2 in Tucson, AZ. We invite industry representatives to attend the meeting and tell us how CAC research can address your company's needs.

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.

January 2011

CAC researchers at MSU visited the Army Corps of Engineers' Engineer Research and Development Center (ERDC) in January for a project meeting. CAC MSU faculty member Sherif Abdelwahed gave a talk entitled "Automated Model-based Security Management of Web Services"; CAC MSU student Rajat Mehrotra presented "Model-based Framework for Autonomic Resource Management of Web Services Platforms".

April 2011

The CAC Spring 2011 Semiannual Meeting was held April 4-6 in Dallas, TX. The meeting was hosted by Xerox at its subsidiary, Affiliated Computer Services. Twenty-nine people from CAC universities and member companies gathered to hear updates on CAC projects, discuss new directions for CAC research, and review Center operating procedures.

Rutgers Site Director Dario Pompili gave a talk at the University of Connecticut's Computer Science and Engineering Department entitled "Inter-glider Underwater Communication and Coordination for Ocean Monitoring and Coastal Tactical Surveillance," on April 15.

CAC collaborator Abhishek Dubey presented a talk entitled "Large Scale Monitoring and Online Analysis in a Distributed Virtualized Environment" at the 8th IEEE Conference and Workshops on Engineering of Autonomic and Autonomous Systems (EASe 2011), held in Las Vegas, NV on April 27-29.

CAC Director José Fortes gave a talk at the 5th International IEEE EMBS Conference on Neural Engineering titled "Towards Closed-Loop Brain-Machine Experiments across Wide-Area Networks", held April 28 to May 1 in Cancun, Mexico.

May 2011

CAC UF alumnus Ming Zhao gave a presentation entitled "Towards Simulation of Parallel File System Scheduling Algorithms with PFSsim" at the 7th IEEE International Workshop on Storage Network Architecture and Parallel I/O (SNAPI).

Rutgers Site Director Dario Pompili presented a talk entitled "On the Impact of Neighborhood Discovery on Geographical Routing in Wireless Sensor Networks" at the 34th IEEE Sarnoff Symposium, held in Princeton, NJ, May 3-4.

CAC @ RU student Ivan Rodero presented "Energy-Aware Application-Centric VM Allocation for HPC," at the High-Performance Grid and Cloud Computing Workshop (HPGC), held in conjunction with the 26th IEEE International Parallel & Distributed Processing Symposium in Anchorage, Alaska.

CAC UF student Kyungyong Lee gave a presentation titled "Parallel Processing Framework on a P2P System Using Map and Reduce Primitives" at the 8th International Workshop on Hot Topics in Peer-to-Peer Systems (HotP2P) on May 20, held in conjunction with IPDPS 2011.

June 2011

David Wolinsky, a recent graduate of the UF CAC site, presented "Experiences with Self-Organizing, Decentralized Grids Using the Grid Appliance" at the 20th International ACM Symposium on High-Performance Parallel and Distributed Computing (HPDC) in San Jose, CA, held July 8-11.

Dario Pompili presented a paper entitled "SILENCE: Distributed Adaptive Sampling for Sensor-based Autonomic Systems" at the 8th IEEE International Conference on Autonomic Computing (ICAC 2011), held June 14-17 in Karlsruhe, Germany.

José Fortes gave two talks at ICAC 2011, titled "Adaptive Virtual Resource Management with Fuzzy Model Predictive Control" and "A Multi-objective Approach to Virtual Machine Management in Datacenters," respectively.

Dario Pompili presented "QUO VADIS: QoS-aware Underwater Optimization Framework for Inter-vehicle Communication using Acoustic Directional Transducers" at the 8th IEEE Conference on Sensor, Mesh and Ad Hoc Communications and Networks (SECON), held in Salt Lake City, UT on June 27-30.

July 2011

Renato Figueiredo presented "On the Performance of Tagged Translation Lookaside Buffers: A Simulation-Driven Analysis" at the 19th International IEEE Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS).

José Fortes presented a paper entitled "Improved Real-Time Scheduling for Periodic Tasks on Multiprocessors" at the High-Performance Computing and Simulation Conference (HPCS 2011), held in Istanbul, Turkey, on July 4-6.

UF student Selvi Kadirvel presented a paper titled "Towards Self-Caring MapReduce: Proactively Reducing Fault-Induced Execution-Time Penalties" at HPCS 2011.

CAC MSU student Srishti Srivastava presented "Enhancing the Functionality of a GridSim-based Scheduler for Effective Use with Large-Scale Scientific Applications" at the IEEE International Sympoisum on Parallel and Distributed Computing (ISPDC2011) held July 6-9 in Cluj-Napoca, Romania.

August 2011

Prapaporn Rattanatamrong, a student at the CAC site at UF, presented a paper titled " Mode Transition for Inline Scheduling of Adaptive Real-time Systems on Multiprocessors" at the 7th IEEE International Conference on Embedded and Real-Time Computing Systems and Applications (RTCSA), held in Toyama, Japan on August 28-31.

September 2011

Heungsik Eom presented a paper at the 13th International Conference on High Performance and Communications (HPCC2011) entitled "SOLARE: Self-Orgarnizing Latency-Aware Resource Ensemble". HPCC 2011 was held September 2-4 in Banff, Canada.

Service

CAC Rutgers faculty member Dario Pompili served as the Technical Program Committee Co-Chair of the 2011 ACM Workshop on UnderWater Networks (WUWNet) (http://wu-wnet.org).

The ACM Publications Board has voted to confirm Prof. Manish Parashar as the next co-Editor-in-Chief of the ACM Transactions on Autonomous and Adaptive Systems (TAAS). His co-Editor-in-Chief is Dr. Franco Zambonelli of University of Modena and Reggio Emilia. This is a three-year appointment. ACM Transactions on Autonomous and Adaptive Systems is the premiere journal focused on foundational, engineering, and technological aspects of computing systems exhibiting emergent, autonomic and adaptive behavior.

Awards

Moustafa Abdelbaky was awarded a prestigious IBM Fellowship. The fellowship program honors Ph.D. students who have demonstrated interest in research in IBM's focus areas, and have been nominated by faculty members for excellence in research. IBM Fellowships are competitively awarded every year, and carry a one-year stipend. Congratulations, Moustafa!

Professor Dario Pompili, Assistant Professor of ECE at Rutgers and Director of the Cyber Physical Systems (CPS) Laboratory, has won the prestigious NSF Career Award for the project "Investigating Fundamental Problems for Underwater Multimedia Communication with Application to Ocean Exploration". This is a five-year grant in the amount of $599,825. More details on the project are given at: http://www.ece.rutgers.edu/node/321.

Rutgers University graduate student Moustafa AbdelBaky led a team that won first place at the IEEE SCALE Challenge. The objective of the SCALE Challenge, which has been held annually since 2008, is to highlight and showcase real-world problem solving using computing that scales. Mr. Abdelbaky's presentation demonstrated the effective use of ensemble applications on a geographically distributed federation of supercomputers. In particular, the demonstration showed an Enhanced Oil Recovery deployed on an HPC-cloud comprised of IBM Blue Gene/P supercomputers located in New York and Saudia Arabia. For more information, please see http://nsfcac.rutgers.edu/icode/scale. SCALE 2011 was held on May 26, 2011 in conjunction with the 11th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing.

Graduations

David Wolinsky, a student at the CAC UF site, has graduated with a Ph.D. and has been named a Postdoctoral Associate in Yale University's Department of Computer Science with Professor Bryan Ford. The project is known as Dissent and funded under the DARPA SAFER program. The purpose of the postdoc is to design and build next-generation disruption-resistant anonymous communication systems and then to perform rigorous experimentation using this prototype system.

CAC student Girish Venkatasubramanian graduated from the UF CAC site and went to work for member company Intel.

In the news

Rutgers Site Director Dario Pompili's research was featured in Rutgers' Daily Targum on April 18. The Targum article focused on Dr. Pompili's work on wireless sensor networks and their applications, specifically their use in monitoring smoker activity. The sensors, pictured below, are equipped with accelerometers and gyroscopes for detecting movement, and are programmed with an algorithm that can be taught to recognize movements that are particular to smoking.

According Dr. Pompili, the wireless sensors can be used to improve studies on smoking habits. These studies typically rely on self-reporting, which can be misleading. "There is a lot of bias in self-reporting," Dr. Pompili said. "Maybe the smoker smokes more because of the stress of the self-reporting, or he reports a lower amount."

The sensors developed for this project have a wide range of applications in addition to collecting data on smokers' habits. For instance, the sensors could be used to monitor the vital signs of athletes or combat troops. "You could non-invasively obtain the vital signs and predict fainting due to stress and many other problems," Dr. Pompili said.

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.

CAC Publications

1. R. Mehrotra, A. Dubey S. Abdelwahed, and A. Tantawi "A Power-aware Modeling and Autonomic Management Framework for Distributed Computing Systems," in Handbook of Energy-Aware and Green Computing, S. Ranka and I. Ahmad Eds. Chapman and Hall/CRC Press, accepted, to be published in January 2012.
2. Ram Prasad Viswanathan, Youssif Al-Nashif, Salim Hariri, "Application Attack Detection System (AADS): An Anomaly Based Behavior Analysis Approach," in Proceedings of the 9th ACS/IEEE International Conference.on Computer Systems and Applications (AICCSA'11), December 2011 (accepted).
3. B. Chen and D. Pompili, "Prioritization of Patient Vital Signs Transmission Using Wireless Body Area Networks," E-Healthcare Systems and Wireless Communications: Current and Future Challenges, Mohamed K. Watfa, Ed., to be released October 2011.
4. P. Rattanatamrong and J. A. B. Fortes, "Mode Transition for Online Scheduling of Adaptive Real-Time Systems on Multiprocessors," The 17th IEEE International Conference on Embedded and Real-Time Computing Systems and Applications (RTCSA), Toyama, Japan, August 28-31, 2011.
5. S. Kadirvel and J.A.B. Fortes, "Towards Self-Caring MapReduce: Proactively Reducing Fault-Induced Execution-Time Penalties," The International Conference on High Performance Computing & Simulation (HPCS), Istanbul, Turkey, July 4-8, 2011.
6. P. Rattanatamrong and J. A. B. Fortes, "Improved Real-Time Scheduling for Periodic Tasks on Multiprocessors," The International Conference on High Performance Computing & Simulation (HPCS), Istanbul, Turkey, July 4-8, 2011.
7. H. Viswanathan, E. K. Lee, and D. Pompili, "Self-organizing Sensing Infrastructure for Autonomic Management of Green Datacenters," IEEE Network Magazine, July 2011.
8. Girish Venkatasubramanian, Renato Figueiredo and Ramesh Illikkal, "On the Performance of Tagged Translation Lookaside Buffers: A Simulation-Driven Analysis," in Proceedings of the 19th International IEEE Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS), July 2011.
9. Sristi Srivastava, Ioana Banicescu, Florina Ciorba and Wolfgang Nagel, "Enhancing the Functionality of a GridSim-based Scheduler for Effective Use with Large-Scale Scientific Applications", in Proceedings of the IEEE International Symposium on Parallel and Distributed Computing (ISPDC 2011), July 2011.
10. B. Chen and D. Pompili, "QUO VADIS: QoS-aware Underwater Optimization Framework for Inter-vehicle Communication using Acoustic Directional Transducers," Proc. of IEEE Conference on Sensor, Mesh and Ad Hoc Communications and Networks (SECON), Salt Lake City, UT, June 2011.
11. David Wolinsky and Renato Figueiredo, "Experiences with Self-Organizing, Decentralized Grids Using the Grid Appliance", in Proceedings of the 20th International ACM Symposium on High-Performance Parallel and Distributed Computing (HPDC), June 2011.
12. L. Wang, M. Zhao, J. Xu and J. A. B. Fortes, "Adaptive Virtual Resource Management with Fuzzy Model Predictive Control", The 6th International Workshop on Feedback Control Implementation and Design in Computing Systems and Networks (FeBID 2011) held in conjunction with The 8th International Conference on Autonomic Computing, Karlsruhe, Germany, June 14, 2011.
13. E. K. Lee, H. Viswanathan, and D. Pompili, "SILENCE: Distributed Adaptive Sampling for Sensor-based Autonomic Systems," in Proceedings of the IEEE International Conference on Autonomic Computing (ICAC), Karlsruhe, Germany, June 2011.
14. J. Xu, J. A. B. Fortes, "A Multi-objective Approach to Virtual Machine Management in Datacenters", The 8th International Conference on Autonomic Computing, Karlsruhe, Germany, June 14-18, 2011.
15. B. Chen and D. Pompili, "Advanced MAC Protocols for ZigBee Networks," ZigBee Network Protocols and Applications, C. Wang, T. Jiang, and Q. Zhang, Eds., Auerbach Publications, CRC press, May 2011.
16. H. Viswanathan, E. K. Lee, I. Rodero, D. Pompili, M. Parashar, and M. Gamell, "Energy-Aware Application-Centric VM Allocation for HPC," in Proceedings of High-Performance Grid and Cloud Computing Workshop (HPGC), Anchorage, AK, May 2011.
17. E. K. Lee, J. P. Varkey, and D. Pompili, "On the Impact of Neighborhood Discovery on Geographical Routing in Wireless Sensor Networks," Proc. of IEEE Sarnoff Symposium, Princeton, NJ, May 2011.
18. Yonggang Liu, Renato Figueiredo, Dulcardo Clavijo, Yiqi Xu, and Ming Zhao," Towards Simulation of Parallel File System Scheduling Algorithms with PFSsim," in Proceedings of the 7th IEEE International Workshop on Storage Network Architecture and Parallel I/O (SNAPI), May 2011.
19. Kyungyong Lee, Tae Woong Choi, Arijit Ganguly, David I Wolinsky,Oscar Boykin, Renato Figueiredo, "Parallel Processing Framework on a P2P System Using Map and Reduce Primitives," in Proceedings of the Eighth International Workshop on Hot Topics in Peer-to-Peer Systems (HotP2P), May 2011
20. R. Mehrotra, A. Dubey S. Abdelwahed, and A. Tantawi "A Power-aware Modeling and Autonomic Management Framework for Distributed Computing Systems," in Handbook of Energy-Aware and Green Computing, S. Ranka and I. Ahmad Eds. Chapman and Hall/CRC Press, to be published in 2011.
21. Girish Venkatasubramanian, Renato Figueiredo, Ramesh Illikkal and Donald Newell, "TMT—A TLB Tag Management Framework for Virtualized Platforms," accepted for publication in the International Journal on Parallel Programming (IJPP).
22. Venkata Krishna Nimmagadda, Ali Akoglu, Salim Hariri, and Talal Moukabary, "Cardiac Simulation on Multi-GPU Platform," to appear in the Springer Journal of Supercomputing.

23. M. Mallouhi, Y. Al-Nashif, D.Cox, T. Chadaga, and S. Hariri, “A Testbed for Analyzing Security of SCADA Control Systems (TASSCS),” the Second IEEE PES Innovative Smart Grid Technologyies Conference (ISGT 2011), Anaheim, CA, January 2011.
24. 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.
25. M. Tsugawa, P. Riteau, A. Matsunaga, J. Fortes, 'User-level Virtual Networking Mechanisms to Support Virtual Machine Migration Over Multiple Clouds', IEEE International Workshop on Management of Emerging Networks and Services,December 2010.
26. Jing Xu and José Fortes, ”Multi-objective Virtual Machine Placement in Virtualized Data Center Environments,“ in Proceedings of the 2010 IEEE International Conference on Green Computing and Communications (GreenCom 2010), Hangzhou, China, December 2010.
27. 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.
28. 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.
29. I. Rodero, E. K. Lee, D. Pompili, M. Parashar, M. Gamell, and R. J. Figueiredo, "Towards Energy-Efficient Reactive Thermal Management in Instrumented Datacenters," Proceedings of the IEEE/ACM International Conference on Energy Efficient Grids, Clouds and Clusters Workshop (E2GC2), Brussels, Belgium, October 2010.
30. 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.
31. D. Pompili, T. Melodia, and I. F. Akyildiz, "Distributed Routing Algorithms for Underwater Acoustic Sensor Networks," IEEE Transactions on Wireless Communications, vol. 9, no. 9, pp. 2934-2944, September 2010.
32. 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.
33. 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.
34. D. Pompili and I. F. Akyildiz, "A Multimedia Cross-layer Protocol for Underwater Acoustic Sensor Networks," IEEE Transactions on Wireless Communications, vol. 9, no. 9, pp. 2924-2933, September 2010.
35. Pierre St. Juste, David Wolinsky, P. Oscar Boykin, Michael J. Covington, Renato Figueiredo, “SocialVPN: Enabling Wide-Area Collaboration with Integrated Social and Overlay Networks,” Computer Networks, vol. 54, no. 12., pp. 1926-1938, August 2010.
36. 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.
37. 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.
38. B. Chen and D. Pompili, "Transmission of Patient Vital Signs using Wireless Body Area Networks," Mobile Networks and Applications (Springer), July 2010.
39. B. Chen, P. C. Hickey, D. Pompili, "A Trajectory-aware Communication Solution for Underwater Gliders using WHOI Micro-Modems," Proc. of IEEE Conference on Sensor, Mesh and Ad Hoc Communications and Networks (SECON), Boston, MA, June 2010.
40. I. S. Kulkarni and D. Pompili, "Task Allocation for Networked Underwater Vehicles in Critical Missions," IEEE Journal of Selected Areas in Communications, vol. 28, no. 5, pp. 716-727, June 2010.
41. E. K. Lee, I. S. Kulkarni, D. Pompili, and M. Parashar, “Proactive Thermal Management in Green Datacenters,” Springer Journal of Supercomputing, vol. 51, no.1, pp. 1-31, June 2010.
42. 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.
43. Arjun Hary, Ali Akoglu, Youssif AlNashif, Salim Hariri, and Darrel Jenerette, “Design and Evaluation of a Self-healing Kepler for Scientific Workflows,” the ACM Int'l Symposium on High Performance Distributed Computing, Chicago, IL, June 2010.
44. A. Quiroz, Parashar, M., Gnanasambandam, N., and Sharma, N, “Autonomic Policy Adaptation using Decentralized Online Clustering,” in Proceedings of the International Conference on Autonomic Computing (ICAC), June 2010.
45. 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).

46. 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.
47. 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.
48. 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.
49. 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.
50. 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.
51. 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.
52. 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.
53. 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.
54. 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.
55. Manish Parashar and Xiaolin Li, ”Advanced Computational Infrastructures for Parallel and Distributed Applications (Wiley Series on Parallel and Distributed Computing),“ Wiley, December 2009.
56. 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.
57. 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.
58. Katarzyna Keahey, Maurício Tsugawa, Andréa Matsunaga and José A. B. Fortes, "Sky Computing," IEEE Internet Computing, September/October 2009.
59. 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.
60. 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.
61. Jing Xu and José Fortes, “Multiobjective Virtual Machine Placement and Migration,” ACIS Technical Report TR-09-001, September 2009.
62. Jing Xu and José Fortes, “Optimization in Autonomic Data Center Resource Management,” ACIS Technical Report TR-09-002, September 2009.
63. 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.
64. Maurício Tsugawa, “On the Design, Performance, and Management of Virtual Networks for Grid Computing,” Ph.D. dissertation, [Gainesville, Fla.]: University of Florida, July 2009.
65. 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.
66. 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.
67. 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, vol. 1, p. 73, February 2009.

To request additional copies of this newsletter, please send an email to Julie at jooly@acis.ufl.edu