Statement of Research/Creative Work

Topics of Interest: Distributed Systems, Software Engineering, Middleware Systems, Autonomic Computing, High-Performance Computing, Grid Computing, Cloud Computing, Pervasive Systems, Mobile Computing.

Summary: The main focus of my multi-disciplinary, and internationally spanned research lies in *automating the information technology* and providing *the right level of abstraction* for scientists and engineers of different disciplines (e.g., meteorologists, biologists, physicists, chemists, bio-medical engineers, mechanical engineers, etc.) to be able to perform their research in their own areas, where they are more comfortable, without the need to be distracted by and in some cases lost in the details of information technology. Currently, successful scientists and engineers who are working on complex problems are those who are also up-to-speed with the latest advances in information technology in general and specifically are experts in high performance computing; an area characterized as moving target.

I vision a day when scientists and engineers can perform their experiments in a timely manner without the need to constantly spend (or waste) their time on learning the latest high-performance computing technologies available out there. To make this vision a reality, as a PI or Co-PI, I have been involved in several small and large sponsored research projects in the past five years. To promote this vision through training new generations of engineers and scientists, I have included an educational and workforce training component in almost all my research projects. This new workforce can benefit from the results of my research and become more efficient in their areas of research by getting less and less involved with the unnecessary details of information technology.

Specifically, to provide *the right level of abstraction* to the meteorologists who conduct research in hurricane forecasting and to the weather officers who run weather simulations in a real-time fashion four times a day, we have developed a Hurricane Mitigation Portal that allows the weather researchers and officers to spend their time on the science part of weather research and forecasting, rather than on tedious and error-prone tasks of compiling, installing, and configuring the forecasting software as is currently the case. In addition, the portal enables the system administrators to easily add or remove computational resources (e.g., clusters, supercomputers, and virtual machines) that cab be allocated dynamically to the various simulation tasks based on preplanned high-level policies, or pressing emergent needs.

This portal is supported by our self-managing and Cluster-, Grid-, and Cloud-enabled workflow managers, meta-schedulers, resource managers, and application profilers. Unlike other approaches, our research in workflow management investigates distributed algorithms and peer-to-peer protocols that dynamically partition, map, and execute workflows and their corresponding subflows while providing cross-layer fault-tolerant and quality-of-service. We have a prototype based on BPEL and JSDL. Our research in meta-scheduling investigates interoperable protocols that allow existing meta-schedulers to interact. We have a prototype that enables interoperation of IBM's TDWB, BSC's eNanos, and Globus' GridWay. We have also developed an application-agnostic

methodology to mathematically model and predict the execution time of long-running applications. We have a primitive but highly usable model for one of the most popular Weather and Research Forecasting (WRF) simulators that quickly estimate the execution time of WRF on currently available clusters.

In summary, my research has been supported by governmental, industrial, academic organizations (e.g., NSF, IBM, Kaseya, TeraGrid, and FIU) for a total of $4.7 million ($1.7M as PI and $3M as Co-PI). I have published 61 papers (3 book chapters, 9 journal papers, 49 conference/workshop proceeding papers) mostly in premier journals and conferences in my research area. I have presented my research to the community through 34 conference and invited lecture presentations and also through numerous poster presentations. Currently, my students and I are actively collaborating and publishing papers with researchers from eight different countries (Argentina, Brazil, China, France, India, Japan, Mexico, and Spain).

Below is a selected list of my research projects:

- NSF PIRE: A Global Living Laboratory for Cyberinfrastructure Application Enablement (PI; $2.3M; 9/15/07–8/31/11; Co-PI for the first 2 years)

  Project Summary: The goal of this PIRE project (http://pire.fiu.edu/) is to conduct research and provide international research and training experiences to its participants in the area of *Transparent Cyberinfrastructure Enablement* by leveraging the Latin American Grid's (LA Grid) established international collaborative programs, resources, and community. CyberInfrastructure (CI) aims to radically simplify the manner by which scientific and business domain experts develop, use, and maintain software applications over distributed computing resources. Our research aims to develop methodologies, platforms, and tools for better enabling CI applications in a way that eases the application development process and make resulting applications more adaptive to future changes of CI. Our approach is characterized as application-driven by basing and focusing our investigation on (1) supporting CI-enablement for a few carefully chosen critical application domains, *e.g.* weather modeling, life sciences, and healthcare, and (2) developing common methodologies, services and tools for developing CI-enabled applications in these domains. We hypothesize an enabling application development paradigm called Transparent Cyberinfrastructure Enablement (TCE), whose goal is to allow domain experts to effectively express the logic and software artifacts of domain applications while hiding the details of the CI architecture, software, and hardware stack.

  Results: This project is just finishing its second year. During the first year, 18 students traveled to 5 countries (Argentina, China, India, Mexico, and Spain), conducted research in 16 TCE-related projects, and published 19 papers (two journals, one book chapter, and 16 conference/workshop papers). They also, presented 18 posters during the FIU-FAU PIRE 2008 workshop. During the second year, 19 students traveled (or still travelling) to 5 countries (Brazil, China, France, India, and Spain).

- NSF CI-TEAM: Global CyberBridges (Co-PI; $765K; 10/2/06-6/31/10)

  Project Summary: The goal of the GCB project (http://cyberbridges.net/) is to produce a new generation of minority scientists and engineers capable of fully integrating

cyberinfrastructure (CI) into the research, education, professional, and creative processes of scientific disciplines. It establishes a base of sustainability through geographically expanding participation. It is designed to create a network of scientists and researchers that spans six research institutions spread over four regions and three countries: USA, China, Hong Kong, and Brazil.

Results: Results from first-year External Assessment Committee Report: Concluded that GCB program participation had been advantageous for the research of each team, and the program structure and classes worked well. The external assessment report says, "clear prospect for more rapid and computationally advanced research production." The assumptions about research problems and needs that justified the GCB project were validated by the experience of the teams. GCB has supported over sixty students as either fellows or team members. GCB has produced twenty-six peer-reviewed publications co-authored by GCB students at USA, Brazil, and China. Finally, the "Grid Enablement of Scientific Applications" course was developed and institutionalized by Dr. Sadjadi during the course of this GCB project.

- Latin American Grid (PI for three IBM-sponsored research projects; $80K in total; 9/1/06-present)

  Project Summary: *LA Grid*, pronounced "lah grid," is an international multi-disciplinary research community and virtual computing grid enabling institutions and industry to extend beyond their individual reach to facilitate collaborative IT research, education and workforce development. *LA Grid* is the first-ever comprehensive computing grid to link faculty, students, and researchers from institutions across the United States, Latin America and Spain to collaborate on complex industry applications for business and societal needs in the context of healthcare, life sciences and disaster mitigation.

  Under this umbrella project, I got three grants from IBM for the following projects:

  - Transparent Grid Enablement of the Weather Research and Forecast code (PI; $20K; 9/1/06-8/31/07)

    Project Summary: To mitigate the impacts of hurricanes, we need to provide users (e.g., residents, homeowners, business owners, and local, state, and federal governments) with sufficiently accurate information to allow them to plan accordingly. The WRF model Version 2.1.2 software distribution comprises about 360000 lines of source code.

  - Design and development of the LA Grid MetaScheduler at FIU (PI; $20K; 9/1/06-8/31/07)

    Project Summary: There is a need for matching and managing jobs on available resources in heterogeneous Grid computing environments.

    Our Solution: We are investigating a hierarchical and peer2peer approach to meta-scheduling and meta-brokering. Also, we have developed an adaptive framework for the execution of BPEL workflow processes for job flow management.

  - Grid Enablement of Hurricane Mitigation Applications (PI; $40K; 6/1/06-5/31/07)

Results: We have been able to successfully develop a prototype for the following components of our overall solution to Grid Computing:

- A Prototype for a Web-Based Portal for WRF Ensemble Forecast.
- A Prototype for a Job-Flow Management in Grid Computing Environments.
- A Prototype for a Meta-Scheduling in Grid Computing Environment.
- A Prototype for a Modeling the Execution Time of WRF.
- A Prototype for a Transparent Grid Enabler for Java Program.

- IT Automation (Sole PI.; $290K in total supported by Kaseya International Shared Services, Sarl; 10/01/08-06/01/10)

  Project Summary: The increasing reliance of industrial, academic, and governmental organizations to their information technology (IT) supported services and the need for the high availability of these services in today's globalized market have led to an increasing demand for IT professionals who are capable of administrating and managing these services efficiently. The conventional and reactive approaches to IT management such as *break-fix* and *block-time*, which respond to IT-related problems only after they already interrupted services, are being replaced by preventive and proactive *IT automation* approaches, which monitor and prevent/correct potential problems before they occur. Currently our curriculum does not include such a course to instill IT automation knowledge to our undergraduate IT students. I propose to develop a course that trains our IT students how to proactively monitor, manage and maintain distributed IT infrastructure remotely, easily, and efficiently by using available integrated Web-based IT automation utilities. Example IT automation tools will be used to give students hands-on experience on auditing, assets and change management, network monitoring, OS imaging, patch management, help desk, remote control, user state management, end-point security, backup, and disaster recovery. In general, this IT Automation course prepares our IT students with a deep understanding of system administration tasks, comprehensive knowledge of different aspects of IT management and automation, and state-of-the-art solutions to remote and automated IT management.

  Results: A new IT Automation course was proposed by Dr. Sadjadi and was approved by FIU and is currently in the process of getting a permanent course number to be added to the FIU Course Catalog. This class was offered for the first time in Spring 2009 and was overwhelmingly well received by our IT students. I consider this class a success for the following reasons: out of the 30 students, seven of them could get a prestiguous and paid internship with Kaseya, five of them got approved to the PIRE program and spent 8 weeks of their PIRE internship in Brazil (all expenses paid), two student were hired by Kaseya, several are currently being interviewed, and two more have been taking an independent study with me during summer working on some interesting IT Automation research projects.

- NSF REU: Autonomic Computing Research at FIU (Sr. Inv.; $300K; 03/15/06-03/14/09)

  Project Summary: The overriding objective of the REU site (http://www.cis.fiu.edu/reu/) at Florida International University is to strengthen the pipeline of underrepresented students to graduate work in Computer Science. The Site provides REU students with a venue to directly participate in state-of-the-art research by exposing them to research opportunities, graduate school life, and work in large research collaboration settings. REU students participate in team-oriented projects

involving several areas of autonomic computing, including self-management in mobile & grid computing, testing of autonomic systems, autonomic communication systems, autonomic storage systems, data mining, and network monitoring and management. The Site's faculty continues their mentoring relationships with students after the completion of the summer program to ensure that each project is completed and that each participant has the support needed to make the transition to graduate school.

Results: Each year, this REU site provided 10 undergraduate students (4 from FIU; 6 from other universities) with a venue to directly participate in state-of-the-art research while exposing them to graduate school life and work in academia and industrial research labs for 12 summer weeks. The students participated in team-oriented projects involving various areas of autonomic computing and establish meaningful mentoring relationships with faculty, which continue well beyond the 12 weeks of summer research program. Over 50% of 2006 and 2007 participants have co authored at least one article in a variety of journals, conferences and workshops (10 publications) and 40% of students have either presented at a conference. Additional publications are under preparation as a part of follow through activities. Our REU site has achieved its goals in that 100% of our students have experienced the daily life in graduate school.

- Communication Virtual Machine

Project Summary: Today's communication solutions follow a stovepipe development of communication solutions results in rigid technology, limited utility, lengthy and costly development cycle, difficulty in integration and innovation.

Our Solution: Not to develop yet another communication solution. CVM provides a user-centric, model-driven approach for conceiving, synthesizing and delivering communication solutions across application domains.