

PIRE: A Global Living Laboratory for Cyberinfrastructure Application Enablement

S. M. Sadjadi, S. Chen, S. Graham, S. Luis, Y. Deng
School of Computing and Information Sciences
Florida International University
11200 SW 8th St., Miami, Florida
(305) 348-2744

{sadjadi,chens,grahams,luiss,deng}@cis.fiu.edu

B. Furht
Dept. of Computer Science and Engineering
Florida Atlantic University
Boca Raton, Florida 33431
(561) 297-3486

borko@cse.fau.edu

P. Martinez
The Quantum Group, Inc.
Wellington, Florida

PMartinez@QuantumMd.com

N. Bowen
IBM T.J. Watson Research Center
Hawthorne, New York

bowenn@us.ibm.com

J. Caraballo
IBM Research
Boca Raton, Florida

jfc@us.ibm.com

ABSTRACT

This Partnership for International Research and Education (PIRE) is a 5-year long project funded by the National Science Foundation that aims to provide 196 international research and training experiences to its participants by leveraging the established programs, resources, and community of the Latin American Grid (LA Grid, an international academic and industry partnership designed to promote research, education and workforce development at major institutions in the USA, Mexico, Argentina, Spain, and other locations around the world). In return, PIRE will take LA Grid to the next level of research and education excellence. Top students, particularly underrepresented minorities, are engaged and each participant will receive multiple perspectives in each of three different aspects of collaboration as they work with (1) local and international researchers, in (2) academic and industrial research labs, and on (3) basic and applied research projects. PIRE participants will engage not only in computer science research topics focused on transparent cyberinfrastructure enablement, but will also be exposed to challenging scientific areas of national importance such as meteorology, bioinformatics, and healthcare. During the first year of this project, 18 students out of a pool of 68 applicants were selected; they participated in complementary PIRE research projects, visited 7 international institutions (spanning 5 countries and 4 continents), and published 9 papers.

Categories and Subject Descriptors

D.2.9 [Management]: Life cycle, Productivity, and Software process models; D.2.6 [Programming Environments]: Graphical environments, Integrated environments, and Interactive environments; D.2.2 [Design Tools and Techniques]: Modules and interfaces, Top-down programming, and User interfaces;

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee
Tapia'09, April 1-4, 2009, Portland, Oregon, USA.
Copyright 2009 ACM 978-1-60558-217-7/09/04...\$5.00.

General Terms

Management, Design, Economics, and Experimentation.

Keywords

Cyberinfrastructure, university/industry partnership, IT workforce.

1. INTRODUCTION

This Partnership for International Research and Education (PIRE) project engages top computer science students, particularly underrepresented minorities, in a unique and large-scale international partnership through an innovative research model. Our PIRE model of collaboration is based on the successful experiences from our Latin American Grid initiative (LA Grid, pronounced “lah grid,” <http://www.latinamericangrid.org>), an international academic and industry partnership designed to promote research, education and workforce development through collaboration with major institutions in Mexico, Argentina, Spain, and other locations around the world.

LA Grid has developed a *global living laboratory* where researchers are empowered to build new research partnerships and explore the synergies of their strengths in three thematic areas: cyberinfrastructure applications, cyberinfrastructure integration, and cyberinfrastructure enablement. The PIRE project aims to provide international research and training opportunities to its participants by leveraging LA Grid’s established programs, resources, and community; its activities are expected to take LA Grid’s living laboratory to the next level of research and education capacity by providing its participants with the support needed to prepare for and facilitate in-person collaboration.

LA Grid synergistically combines research, education and workforce development activities by identifying and promoting excellence via programs that increase student exposure to academic and industry research and development. The innovative LA Grid model is resulting in a strong and successful partnership. IBM’s LA Grid Scholars program provides its participating students with research, mentoring, and industry internship experiences. Our PIRE-enabled teams will draw upon the large pool of Hispanic students participating as LA Grid Scholars at

Florida International University (FIU) and Florida Atlantic University (FAU).

The PIRE project is a university/industry collaboration between the FIU's School of Computing and Information Sciences, the FAU's Department of Computer Science and Engineering, the Barcelona Supercomputing Center (BSC, Spain), IBM Research Worldwide (China, France, India, Japan, USA), the Instituto Tecnológico y de Estudios Superiores de Monterrey (ITESM, Mexico), Tsinghua University (China), the University Nacional de La Plata (UNLP, Argentina), and the Universitat Politècnica de Catalunya (UPC, Spain). Each excels in cyberinfrastructure (CI) research over a range from basic to applied, and each provides resources that are leveraged as an agile shared cyberinfrastructure capable of supporting CI enablement and integration research, as well as CI-enabled applications such as hurricane mitigation, bioinformatics, and healthcare communication applications.

By coordinating closely with our partners, we conduct research experiments that integrate resources from supercomputers to desktop PCs. The PIRE research teams each consist of a group of world-class experts specializing in areas ranging from distributed computing theory to meteorology; each group builds on the research strengths of other groups in a synergistic fashion.

In the rest of this paper, we first introduce the PIRE collaboration model in Section 2. Next, we introduce the PIRE core research in Section 3. Major research and education activities are discussed in Section 4 followed by the report on our research contributions in Section 5. Finally, we conclude this paper in Section 6.

2. COLLABORATION MODEL

High-performing students, who have aspirations to earn advanced degrees, will be recruited into the PIRE Triangles from the LA Grid Scholars program (described above); most will have had prior research experiences via that program. Our PIRE Triangle Model, depicted in Figure 1, ensures that each participant will receive multiple perspectives in each of three different aspects of collaboration as they work with (1) local and international researchers, in (2) academic and industrial research labs, and on (3) basic and applied research projects. These Triangles leverage the diversity of our partners' research interests, their spread around the globe, and their nature. All PIRE students and researchers will participate in evaluations of the PIRE Triangle program to ensure the program is effectively enhancing student research experiences and contacts.




Figure 1. PIRE Triangle Model

The CI research envisioned by the LA Grid community, enabling scientific and business applications to use an agile shared CI, requires the participation of both academic and industry research participants. This research diversity allows us to capture the needs of these communities and to provide solutions that address these needs. By providing researchers with access to both international academic and industrial research experiences, we maximize the

professional exposure participants receive and increase their contact portfolio, thus enhancing a resource they will need in order to create additional international research collaborations.

LA Grid was co-founded through the efforts of FIU and high-ranking IBM corporate executives who support the initiative by engaging preeminent IBM researchers, distinguished engineers, and IT consultants to participate in LA Grid's activities. This corporate support creates visibility for the initiative throughout the company, which creates further opportunities for collaboration. The LA Grid partnership is already showing successful outcomes. Since January 2006, LA Grid has held three international research summits and one workshop, and has involved 35 LA Grid Scholars via internships, executive mentoring, and research projects. Over 40 faculty members at LA Grid institutions are participating in several joint research projects using our global living laboratory shared CI, which includes over 1,000 processors as well as computing capacity on the MareNostrum Supercomputer at the Barcelona Supercomputing Center with over 10,000 processors.

The PIRE-enabled international research program is targeted to provide at least 196 international research and education travel opportunities to its participants. Students will be able to experience research abroad, including trips to countries that are emerging rivals in the global Information Technology (IT) marketplace, such as India and China. The diversity of collaboration available to participants captures the spirit of our LA Grid partnership and our commitment to research excellence while enabling participants to create the lasting contacts they will use to create new international collaborations throughout their careers, thus effectively serving the PIRE Program's goals.

During Summer 2008, 18 students out of a pool of 68 applicants were selected. The student participants received cultural and language training prior to their trips and were given the option of receiving academic credit for their international research projects. Each student worked under the supervision of at least one faculty member at FIU/FAU and one faculty member at one of our international host institutions.

3. CORE RESEARCH

Cyberinfrastructure is an important enabler for expanding research discovery and industry applications. To realize CI's full potential, domain scientists need to be able to easily run their existing applications on the CI available to them; the scientists also need to be able to design their future applications in a way that allows them to take advantage of an ever-changing and growing CI. Large-scale software systems like the Weather Research and Forecasting (WRF) model have been developed to perform tasks such as emulating meteorological behavior. The scientists developing these models need to be able to concentrate on developing better models, not on re-developing their software to take advantage of CI.

CI aims to radically simplify the manner by which scientific and business domain experts develop, use, and maintain software applications over distributed computing resources. However, the current tools used to create CI applications follow a bottom-up approach and present two problems: (1) they are either *too generic* and do not provide the right level of abstraction to allow experts in diverse domains to easily "code" their application logic

(e.g., CORBA, Web Services, and CASE Tools); or (2) they are *too specific*, in most cases following a stove-pipe development process, resulting in rigid and expensive solutions that do not promote the reuse of commonalities across domains.

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 (hence “top-down”) 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. In our approach, we factor out common services that can be reused across domains. This will ensure that our tools have broad significance and utility to a range of applications, thus avoiding the tendency for tools to be too generic to be effective. Specifically, we address the following key challenges:

1. *High-level Visual Interactive Development Environment (IDE)*: What is the appropriate IDE for domain experts to easily specify the logic of their applications? Do IDEs targeted for different domains have many common properties? Is it possible to develop a common IDE that supports multiple domains? Are the workbenches being developed and used by domain experts today the right solution?

2. *Automated Code Generation and Software/Hardware Reuse*: How can we automate the generation of executable code from a high-level specification provided by a domain expert? How can we reuse the existing software and hardware components and map abstract specifications to concrete resources in order to execute the application?

3. *Hiding the Heterogeneity of CI Architectures*: How do we hide the details of heterogeneous CI architectures and resources and provide a virtualized interface for application development while addressing efficient resource utilization?

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. This TCE paradigm will serve as the foundation for the study of application development methodologies, platforms, and tools that will significantly ease CI-enabled application development (hence broadening CI utilization) and make applications more portable and adaptable to future changes of CI. We believe that CI utilizing the TCE paradigm will be agile, flexible, and capable of serving a broad set of scientific and business communities.

Under the LA Grid project, FIU, FAU, and their partners have established a globally-integrated research and education program to respond to the above challenges and to realize TCE. We have taken a divide-and-conquer approach that allows us to tackle the above challenges in parallel. For the first challenge, we have identified three nationally-important domains of *CI Applications*. For the second challenge, we have established a number of *CI Integration* projects that enable aggregation and discovery of data, visualization of data, and weaving of high-level CI enablement services into the logic of domain-specific applications. For the

third challenge, we have established several *CI Enablement* projects that provide a layer of abstraction on top of heterogeneous CI architectures by offering high-level services.

PIRE supports the largely underrepresented students who are engaging in these state-of-the-art and forward-looking research projects to extend their research and training experience by sending them abroad to collaborate with our international partners. PIRE provides its participants with a unique opportunity to engage not only in computer science research topics focused on *cyberinfrastructure enablement and integration*, but also exposes them to challenging *cyberinfrastructure applications*. A brief introduction to these projects in their respective categories follows.

Cyberinfrastructure Applications. We have selected three challenging scientific and nationally important CI-application areas, meteorology, bioinformatics, and healthcare, as the initial target areas of our research.

1. *Hurricane Mitigation Applications*: We aim to mitigate the impact of hurricane landfalls by providing accurate and timely information to enable effective planning [1-4]. We seek to develop a high-level modeling platform that allows meteorologists to more easily use the Weather Research and Forecasting (WRF) model and to develop new meteorology models adaptable to different CI configurations. PIRE researchers will address challenging questions such as the expressive power of workflow languages [5], and the efficiency of the generated WRF code.

2. *Bioinformatics Applications*: We model and analyze biological systems, especially systems involving genetic material, to understand the root cause of several disorders and diseases in human body. In the past, to accomplish their tasks, biologists have been forced to learn a wide range of CI technologies. PIRE seeks to design and develop specialized visual tools for biologists so that they can easily develop and use bioinformatics applications without the need to become CI experts.

3. *Healthcare Communication Applications*. There are many complex needs for communication and data sharing in healthcare and telemedicine. These communications not only involve different media, e.g., voice, video, text, and images, but also need to follow certain workflow processes, ensure patient privacy, and allow exchange of data from distributed sources and medical information systems. PIRE aims to develop a systematic modeling approach and visual modeling language/tools that allow us to easily describe various healthcare communication services in the form of *communication schemas*. With the support of Communication Virtual Machine (CVM) technology [6], such a schema can be automatically synthesized and executed to realize complex healthcare communication services on demand, and transparent of, the underlying CI infrastructure, without the need for costly custom system development.

Cyberinfrastructure Integration. Every CI application involves complex data analysis and model development from huge amounts of data. It becomes very costly if all the applications need to develop their own software packages. We aim to develop a reusable, scalable, and reconfigurable software framework to address these issues.

1. *Data Mining Software Tools*: Our goal is to design algorithms for analyzing large amounts of spatial and/or temporal data and to

develop software toolkits that can apply these algorithms for mining domain-specific data.

2. *Visualization Software Tools*: It is well-recognized that information-rich visualization can help users understand and interpret patterns in data [7]. Recent visualization research projects lack the necessary support required for generating efficient on-demand visualizations in a resource-efficient and timely manner. We plan to design tools that enable (1) the creation of a transparent network between the users and the distributed locations where the data is situated, and (2) the visualization and rendering of distributed large data volumes, since only the reduced data sets and/or the discovered data/features need to be transferred across the network.

3. *Transparent Grid Enablement*: Current standards for cluster and grid programming have made significant progress in hiding the heterogeneity of the underlying devices, network protocols, and middleware layers from application developers. However, developers are still expected to develop complex parallel algorithms and programs. PIRE will investigate a new programming paradigm called Transparent Grid Enablement (TGE), which enables a *separation of concerns* in the development and maintenance of the non-functional concerns (i.e., the parallel code) and the functional concerns (i.e., the application logic) of scientific applications.

Cyberinfrastructure

Enablement. CI enablement includes several projects aimed to hide the heterogeneity of underlying CI hardware and software, and to provide CI integration/application projects with virtualized and interoperable CI services.

1. Autonomic Resource Management:

Our goal is to support CI applications with robust and self-regulating mechanisms for resource management. These mechanisms provide high-level quality-of-service (QoS) support for application execution on top of CI that has heterogeneous resources with different computation, storage, and network performance parameters.

2. *Meta Scheduling and Job Flow Management*: We aim to support CI applications with resources located and managed in different logical and/or physical domains using meta-scheduling and job flow management techniques. We address the issues of interoperability in different approaches to meta-scheduling and job flow management to realize an infrastructure that provides a uniform access to resources across different institutions.

3. *Communication Virtual Machine (CVM)*. We aim to develop a CI-transparent communication technology that will support the automatic generation of communication solutions. CVM represents a paradigm shift on how communication applications are conceived and delivered. Its design draws from model-driven

engineering [8, 9], communication middleware [10], and software virtualization and service-oriented architectures. CVM supports a modeling approach, in which complex, and possibly domain specific, communication services can be modeled as communication schemas. These schemas can be instantiated, negotiated, synthesized, and executed on-demand to support various forms of communication and data exchange, thus avoiding the need to develop and deploy costly custom communication systems. PIRE will further pursue in-depth study of various CVM design issues, including its interfacing with heterogeneous architectures and the composition of underlying CI-infrastructure.

Synergy among the research projects

It has been our strategy to select the projects for our LA Grid international research teams in such a way that the projects are mutually complementary and collectively support the goal of transparent CI application enablement, following the fore-mentioned application-driven approach. As our research progresses, the technologies and their associated tools developed in these projects will form CI application enabling platforms at different levels of abstraction. Although the projects described above are self-contained and are valuable in their own rights, the synergy between the different projects discussed above is

illustrated in Figure 2. First, the projects in the CI Applications Layer will identify the requirements and provide methodologies, frameworks, and modeling tools that enable domain experts to model, design, and code their applications with minimal attachment to the underlying CI. In designing and developing these “application interfaces,” common high-level, application-oriented functions, components, and tools are extracted and packaged into the CI Integration Layer, which provides services that can be reused by other applications in the targeted domains. The

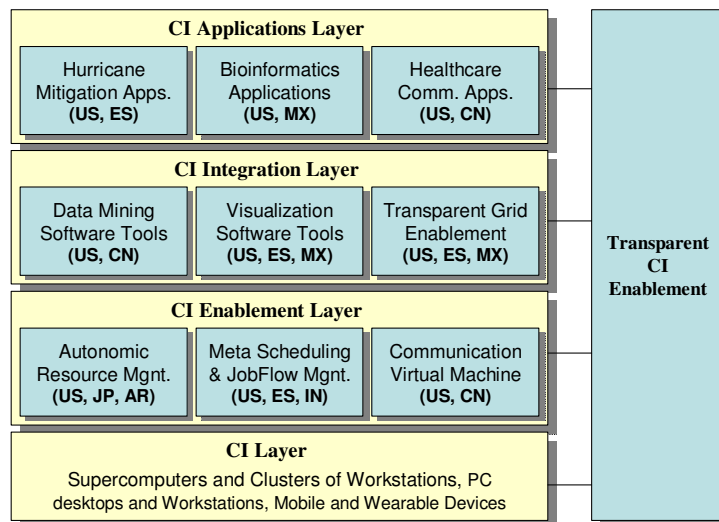


Figure 2: Relationship among the different projects. The International Country Codes are used to indicate the countries involved in each project: **US:** USA, **ES:** Spain, **MX:** Mexico, **CN:** China, **JP:** Japan, **IN:** India, and **AR:** Argentina.

projects in the CI Enablement Layer will provide system-level services and tools to support efficient and transparent management and utilization of heterogeneous CI resources through uniform virtualized interfaces. Such services will make applications even more adaptable to changes of the underlying infrastructure.

4. RESEARCH & EDUCATION

The PIRE faculty members completed five international travels to China, Argentina, Spain, Mexico, and India. As an outcome of these travels and extensive meetings and international teleconferences, nine international partners (namely, Barcelona Supercomputing Center, Universitat Politècnica de Catalunya, IBM China Research Lab, IBM India Research Lab, Instituto Tecnológico y de Estudios Superiores de Monterrey, Tsinghua

University, Universidad de Guadalajara, Universidad Nacional de La Plata, and Universitat Pompeu Fabra), which span five countries, pledged to receive our students.

We received 68 applications from FIU and FAU students. Of these, 27 students were selected as PIRE participants and 26 students accepted their roles as such. These students participated in a semester-long cultural and language training program that focused on the country where they will perform their research. The training sessions, which used a mixture of in-person and on-line sessions, focused on a small survival vocabulary including technical terms and useful phrases. Proper cultural behavior, business etiquette, and differences in manners were also addressed in each course.

Eighteen of the 26 PIRE participants finished their first PIRE travels during the summer of 2008. Out of the total of 18 students, 9 students were from FIU and 9 were from FAU; 12 students were graduate and 6 were undergraduate students. Two performed collaborative research on the *Pattern Based Fault-Tolerance at Workflow Management Systems* and *Compass: Cost of Migration-aware Placement in Storage Systems* projects at IBM's India Research Lab in New Delhi. Two performed collaborative research on the *Business Constraint Based Testing Generation* and *Multimedia Database and Information Management* projects at IBM's China Research Lab in Beijing. They were joined by four students who worked on *Data Mining for Video Transcoding and Encoding*, *Systematic Study of Emerging Solutions and Standards for Multimedia on the Web*, and *A Model Driven Engineering Framework for Migrating Web Applications to Mobile Devices* at Tsinghua University. Three students worked on *Bioinformatics Algorithms for High Performance Computing*, *Web Accessibility for the Hearing Impaired*, and *development and integration of several patterns for fault-tolerant and secure web services* projects at the UNLP in La Plata, Argentina. Four students worked on *Gene Selection for Cancer Classification*, *Data Mining and Machine Learning Cross Validation over Distributed Networks*, and *The Grid Enablement and the Performance Analysis of the Weather Research and Forecasting (WRF) code* projects at the Barcelona Supercomputing Center and with faculty of the UPC in Barcelona, Spain. The final three students traveled to UdG in Guadalajara, Mexico and worked on the *Web-Based Hurricane Weather Forecasting Portal* and *Improving the Efficiency in Genomic Analysis* projects. Each student's trip lasted 8 to 10 weeks. Students posted regular blog items at the pire.fiu.edu website; these posts focus on the research they performed, the relationships they built, and the travel experiences that they had. Four additional students are expected to travel in the Fall of 2008 and four more in the Spring of 2009. The next cohort of PIRE participants will be recruited in Fall 2008 and will receive cultural and language training in Spring 2009.

5. RESEARCH CONTRIBUTIONS

A brief summary of the research contributions made by this first group of students follows.

CI Integration: Out of the three main CI integration areas that we proposed, so far we have produced some contributions in the Transparent Grid Enablement.

For Transparent Grid Enablement, we realized that there is a pressing need to provide a range of users with accurate and timely

information that can enable effective planning for and response to potential hurricane landfall, so to mitigate the disastrous impact of hurricanes. There is a need for the Grid-enablement of the WRF code such that it can utilize resources available in partner organizations. A preliminary prototype of our Grid-Enabled WRF based on ensemble weather forecasting is now available. The transparent grid enablement methodology that is one of the major outcome of this research will be applicable to grid enablement of similar scientific applications.

CI Enablement: Below, we summarize our contributions in the three areas of our proposed CI enablement.

For Autonomic Resource Management, we identified that in a Grid computing environment, resources are shared among a large number of applications. Brokers and schedulers find matching resources and schedule the execution of the applications by monitoring dynamic resource availability and employing policies such as first-come-first-served and back-filling. To support applications with timeliness requirements in such an environment, brokering and scheduling algorithms must address an additional problem - they must be able to estimate the execution time of the application on the currently available resources. We propose a modeling approach to estimating the execution time of long-running scientific applications. The modeling approach we propose is generic; models can be constructed by merely observing the application execution "externally" without using intrusive techniques such as code inspection or instrumentation. The model is cross-platform; it enables prediction without the need for the application to be profiled first on the target hardware. To show the feasibility and effectiveness of this approach, we developed a resource usage model that estimates the execution time of a weather forecasting application in a multi-cluster Grid computing environment. We validated the model through extensive benchmarking and profiling experiments and observed prediction errors that were within 10% of the measured values.

For Job Flow Management, we realized that the execution of job flow applications is a reality today in academic and industrial domains. Current approaches to execution of job flows often follow proprietary solutions on expressing the job flows and do not leverage recurrent job-flow patterns to address faults in Grid computing environments. We provided a design solution to development of job-flow managers that uses standard technologies such as BPEL and JSDL to express job flows and employs a two-layer peer-to-peer architecture with interoperable protocols for cross-domain interactions among job-flow managers. In addition, we identified a number of recurring job-flow patterns and introduce their corresponding fault-tolerant patterns to address runtime faults and exceptions. Finally, to keep the business logic of job flows separate from their fault-tolerant behavior, we used a transparent proxy that intercepts job-flow execution at runtime to handle potential faults using a growing knowledge base that contains the most recently identified job-flow patterns and their corresponding fault-tolerant patterns.

For Meta Scheduling, we realized that the major hurdle on accessing shared computing resources from cooperating organizations is their heterogeneous meta-scheduling frameworks. To address this problem, we proposed a hybrid approach, combining hierarchical and peer-to-peer architectures for flexibility and extensibility of these systems. We introduced a set of protocols to allow different meta-scheduler instances to

communicate over Web Services. Interoperability between three heterogeneous and distributed organizations (namely, BSC, FIU, and IBM), each using different meta-scheduling technologies, was demonstrated under these protocols and resource models using the first version of our protocol.

For Communication Virtual Machine (CVM), a prototypical intuitive communication modeling language (CML) for modeling user communication requirements has been developed. We developed two equivalent variants of CML: the XML-based (X-CML) and the graphical (G-CML). The former is the version that CVM understands and processes, while the latter is the user-friendly graphical form. G-XML is analogous to the E-R diagram in the database domain. These two variants can be automatically converted to each other. A CVM prototype has been implemented using the following technologies. A Web-based user interface has been deployed with the Opera 8.5, a voice-enabled browser. This prototype enables creation, modification, and use of communication schema instance using voice commands. The technologies used at the browser side are HTML, Javascript for dynamic effects and the program logic, and XHTML+Voice for voice conversation. Part of the Javascript code uses AJAX technology (Asynchronous JavaScript and XML) to make web requests and responses in the background, without having to refresh the web pages. The rest of the CVM layers, including another user interface (currently being the most updated version of our UI), are implemented in Java, deployed on each node. JAIN SIP and Java Media Framework (JMF) are used for control and data communications, respectively. Finally, we used SER (SIP Express Router) server for registration and presence and Asterisk for connection to PSTN and audio mixing.

6. CONCLUSION

PIRE is an opportunity to bring together a global network of researchers, faculty members, and students currently distributed across global distances. PIRE will develop a body of theory, techniques, and tools that will facilitate a more effective and broader global CI. It will also help build collaborations between scientists from different nations with very different cultures, traditions, and infrastructures, thus demonstrating a model for making an international CI available for solving critical and nationally important complex problems. The preliminary results from the first year of this project is promising and we believe that our global living laboratory will have profound impact on the effectiveness of international collaborative research on cyberinfrastructure and on generations of researchers. The successful PIRE-enabled projects will provide a new paradigm for transparent cyberinfrastructure enablement supported by a concrete methodology and numerous tools and services useful to researchers and developers in academia and industry.

FIU is the leading Hispanic computer science degree producer among those ranked as Research Universities that have high or very high research activity by the Carnegie Foundation [11]. Presently, Hispanic participation in the US IT industry is the lowest among underrepresented groups; this represents a significant threat to US competitiveness [12]. Training Hispanics to participate in the US IT industry is imperative to the US's ability to compete in foreign markets.

PIRE is impacting FIU and FAU faculty doing research and teaching by allowing them to benefit from the established

international research collaborations as well as from the global living laboratory. The research findings as well as the tools and services developed in the research projects can be incorporated in computer science and IT curricula. This project will demonstrate a proof-of-concept of the effect of making such an international CI available for solving critical and nationally-important complex scientific problems. PIRE will have profound impact on the effectiveness of international collaborative research on CI and the creativity and effectiveness of diverse generations to come.

7. ACKNOWLEDGMENTS

This work was supported in part by the National Science Foundation (OISE-0730065, HRD-0833093, HRD-0317692, OCI-0636031, and IIS-0552555). Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not reflect those of the NSF and IBM.

8. REFERENCES

- [1] H. E. Willoughby. "Improvements in observations, models and forecasts." In *HURRICANE! Coping with Disaster*, edited by R.H. Simpson, AGU, 205-216. 2002.
- [2] P. Singh, N. Zhao, S. C. Chen, K. Zhang. "Tree Animation for a 3D Interactive Visualization System for Hurricane Impacts." *Intl. Conf. on Multimedia*, pp. 598-601, 2005.
- [3] S. C. Chen, S. Hamid, S. Gulati, N. Zhao, M. Chen, C. Zhang, and P. Gupta. "A Reliable Web-based System for Hurricane Analysis and Simulation." *Proceedings of the IEEE International Conference on Systems, Man and Cybernetics 2004*, pp. 5215-5220, October 10-13, 2004.
- [4] S. C. Chen, M.-L. Shyu, C. Zhang, W. Z. Tang, and K. Zhang. "Damage Pattern Mining in Hurricane Image Databases." *IEEE Intl. Conf. on Info. Reuse and Int.*, pp. 227-234. 2003.
- [5] O. Ezenwoye and S. M. Sadjadi. "Composing aggregate web services in BPEL." *Proceedings of the 44th ACM Southeast Conference (ACMSE 2006)*, March 2006.
- [6] Y. Deng, S. M. Sadjadi, P. J. Clarke, C. Zhang, V. Hristidis, R. Rangaswami, and N. Prabakar. "A communication virtual machine." *30th Annual International Computer Software and Applications Conference*, September 2006.
- [7] U. Fayyad, G. Grinstein, and A. Wierse (editors), *Information Visualization in Data Mining and Knowledge Discovery*, Morgan Kaufmann, 2001.
- [8] J. Bettin. "Model-driven software development: An emerging paradigm for industrialised software asset development." *Tech report, SoftMetaWare*, June 2004.
- [9] D. C. Schmidt. "Model-Driven Engineering." *IEEE Computer*, February 2006, 25-31.
- [10] D. C. Schmidt. "Applying Patterns and Frameworks to Develop Object-Oriented Communication Software." V. 1 of *Handbook of Prog. Lang.* MacMillan Computer Pub., 1997.
- [11] Carnegie Foundation for the Advancement of Teaching. "Doctoral/Research Universities-High Research Activity."
- [12] Computer Research Association. "Taulbee Survey Report 2006-2007."