Steps Toward a Science of Service Systems

Jim Spohrer, Paul P. Maglio, John Bailey, and Daniel Gruhl
IBM Almaden Research Center

The service sector accounts for most of the world's economic activity, but it's the least-studied part of the economy. A service system comprises people and technologies that adaptively compute and adjust to a system's changing value of knowledge. A science of service systems could provide theory and practice around service innovation.

Over the past three decades, services have become the largest part of most industrialized nations' economies. Yet there's still no widely accepted definition of service, and service productivity, quality, compliance, and innovation all remain hard to measure. Few researchers have studied service, and institutions have paid little attention to educating students in this area.

The *service economy* refers to the service sector, one of three main economic categories, in addition to service activities performed in the extractive and manufacturing sectors. The growth of the service sector has resulted in part from the specialization and outsourcing of service activities performed inside manufacturing firms (for example, design, maintenance, human resources, customer contact specialists). According to a recent National Academy of Engineering report,[1] the service sector accounts for more than 80 percent of the US gross domestic product, employs a large and growing share of the science and engineering workforce, and is the primary user of IT. The report suggests that academic researchers ought to begin to focus on service businesses' needs by:

- adapting and applying systems and industrial engineering concepts, methodologies, and quality-control processes to service functions and businesses;
- integrating technological research and social science, management, and policy research; and
- educating and training engineering and science graduates prepared to deal with management, policy, and social issues.

One approach is to develop a general theory of service with well-defined questions, tools, methods, and practical implications for society. Some see economics, operations research, industrial engineering, management of information systems, multiagent systems, or the science of complex systems as the appropriate starting point for such a general theory. Others contend that the pervasiveness of services, such as government, education, healthcare, banking, insurance, IT and business services, creates a need for many specific engineering, management, or applied science disciplines.

We believe the solution lies in between those two approaches. Toward this end, we're cultivating an interdisciplinary effort called Service Science, Management, and Engineering—the application of scientific, management, and engineering disciplines to tasks that one organization (service provider) beneficially performs for and with another (service client). SSME aims to understand how an organization can invest effectively to create service innovations and to realize more predictable outcomes.[2-5] With information and business services the service economy's fastest-growing segments—and with the rise of Web services, service-oriented architectures (SOA), and self-service systems—we see a strong rela-

0018-9162/07/$25.00 © 2007 IEEE

tionship between the study of service systems and the more established study of computational systems.

SERVICE AND SERVICE SYSTEMS

Service can be defined as the application of competences for the benefit of another,[6] meaning that service is a kind of action, performance, or promise that's exchanged for value between provider and client. Service is performed in close contact with a client; the more knowledge-intensive and customized the service, the more the service process depends critically on client participation and input, whether by providing labor, property, or information.[7]

Individuals, families, firms, nations, and economies all represent instances of service systems.

Service systems comprise service providers and service clients working together to coproduce value in complex value chains or networks.[8] Providers and clients might be individuals, firms, government agencies, or any organization of people and technologies. The key is that providers and clients work together to create value—the client owns or controls some state that the provider is responsible for transforming according to some agreement between provider and client.[9]

Educational service systems

Consider universities as service providers that aim to transform student knowledge through agreements, relationships, and other exchanges among students and university faculty, including courses offered and taken, tuition paid, and work-study arrangements. Typically, students don't bear the complete cost of educational transformations. Rather, individuals, corporations, nonprofit organizations, and government sponsors help support universities. This financial support lets universities invest in infrastructure and other resources, offsetting the difference in the actual cost and the tuition that the market can bear.

Although potentially beneficial to everyone involved, this economic arrangement results in a service equation that's more complex than that of a single, unambiguous service client. Rather than managing a single coproduction relationship, universities manage coproduction relationships among multiple clients, each of whom might or might not know or care about the others or about their relative needs and expectations.

The student, who experiences the service firsthand, is likely to use qualitative measures to judge service quality, whereas a corporate or government supporter might rely more on collective quantitative measures, such as standardized test scores and number of graduates. Over time, universities have developed sophisticated processes and organizations to manage their complex service relationships. A university that excels in all these service relationships—producing expected or better-than-expected outcomes across the range of stakeholders—develops a reputation for excellence, thus generating even more interest from high-potential students and employees. The best get better.

The students' preparedness is crucial in determining the result of the service relationship. The better prepared that students are to learn, the more likely their transformations will meet expectations. Excellent universities are very selective in the students they accept, which functions as a kind of standardization of client inputs into the service-production system.

Universities have adapted to support complex relationships between service providers and clients, and they're now adapting to IT changes in how they package, deliver, manage, and measure education. Alternatives to traditional university education services now include remote teaching, self-paced learning, and online learning through role-playing games. In the end, we can't simply consider the university a service provider, but more like a complex adaptive system of people and technologies working together to create value (learning).

More precisely, we define a service system as a value-coproduction configuration of people, technology, other internal and external service systems, and shared information (such as language, processes, metrics, prices, policies, and laws). This recursive service system definition highlights the fact that service systems have internal structure (*intraentity services*) and external structure (*interentity services*) in which participants coproduce value directly or indirectly with other service systems. Individuals, families, firms, nations, and economies all represent instances of service systems. Individuals (who exchange service with external service systems) and the global economy (which contains many internal service systems that exchange service) represent special cases because most service systems have both internal and external service structures.

IT outsourcing service system

Consider IT outsourcing, which is a complex business-to-business service system. An IT outsourcing service provider offers to take over operation and maintenance of clients' IT investments and to do it better and cheaper than the clients can do it themselves. Thus, the provider aims to improve the efficiency of client IT operations, reducing cost over time by applying unique skills, experience, and capabilities more effectively than the client can.

Outsourcing service arrangements range from multi-billion-dollar deals in which the service provider takes over all of a Fortune 100 company's IT assets, to smaller deals in which the provider takes over a single functional area, such as help-desk operations or Web-server operations. There are many ways to structure outsourcing agreements, including transformational or nontransformational, single source or multiple source, and locating resources at the client's or provider's location. Negotiations determine a deal's structure, which is formally specified in a contract the provider and client jointly produce.

Professionals from the finance, legal, business operations, IT operations, and human resources departments contribute information in large outsourcing deals. A service-level agreement (SLA) specifies the metrics the client can use to monitor and verify the contract. The metrics match client business objectives to valid, quantifiable service provider performance indicators. IT outsourcing SLAs often include provider commitments to perform some activity within an agreed-to amount of time (for example, resolve high-severity IT problems in less than 60 minutes), or to maintain a minimal service-availability level (for example, no more than 120 minutes downtime per unit-month). Though SLAs are conventional and useful, achieving them is just one measure of client satisfaction.

Service-system characteristics

What are the simplest types of service? Reducing the application of competence to a list of instructions that one service system can communicate to another is a "tell me" type of service. The client can request and then use the instructions to gain the benefit of the provider's competence (say, through self-service). Thus, a conversation is a "building block" type of service in which two systems exchange self-service executable competence of benefit to both. Self-help books are an example of providers trying to reduce service to a set of instructions. More sophisticated service categories include "show me," "help me," and "do it for me." IT outsourcing is an example of a "do it for me" type of service.

Most of the time, real-world competences of great value aren't simple. It isn't possible to reduce complex competences to a list of easily executed instructions (consider riding a bike, flying a plane, or transforming a business supply chain). Some service systems might not have all the requisite skills to execute the instructions or it might just be physically impossible for a system to perform the service independently at the current technology level. Some services lose their significance when privileged entities don't perform them (for example, a vendor conducting an elevator safety inspection rather than an authorized city agency).

Some competences might have side effects and associated risks to other service systems if not executed properly; thus, they might require certification as well as proof of responsibility in dealing with unintended consequences (for example, drivers' licenses and car liability insurance). A general theory of service must clarify the characteristics of service systems and service competences that we see in everyday life—and must also clarify the value of different kinds of knowledge in diverse contexts as judged by diverse stakeholders.

Regardless of how competence leads to action and value, coordination and governance require shared information. Three key types of shared information are language, laws, and measures. Without some form of language, signaling, or standard encoding of information, systems would find coordination difficult, leading to missed opportunities for innovation or efficiency gains.[10]

Some competences might have side effects and associated risks to other service systems if not executed properly.

Provisioning sophisticated service and maintaining complex service systems requires laws and contracts. Typically, every service system has a governing authority that seeks to ensure that all those in the service system can communicate in shared languages and abide by shared laws. In firms, it's the CEO and board of directors; and in nations, it's government leaders and agencies, as well as shared legal documents and enforcement agencies.

Prices are one type of measure of the value of services exchanged within or between service systems. Often, standardizing the sets of measures used within and between service systems improves the productive capacity of the system by eliminating unneeded transaction costs and improving coordination. Language, laws, measures, and other types of shared information evolve over time as service systems invest to improve productivity, quality, compliance, and innovation.

BACKGROUND FOR A THEORY

The components of a service system are people, technology, internal and external service systems connected by value propositions, and shared information (such as language, laws, and measures). So a theory of service systems should explain what service systems are and aren't, how they arise and evolve, the relation between internal and external service systems, and the role of people, technology, value propositions, and shared information in the system. But what motivates our choice of service system components?

Shelby Hunt referred to seven types of a firm's resources[11] that map well to the service-system components. Richard R. Nelson and Sidney G. Winter[12] distinguished between physical and social technology, with physical technology mapping to the traditional notion of technology, and social technology mapping to people, other service systems, and shared information.

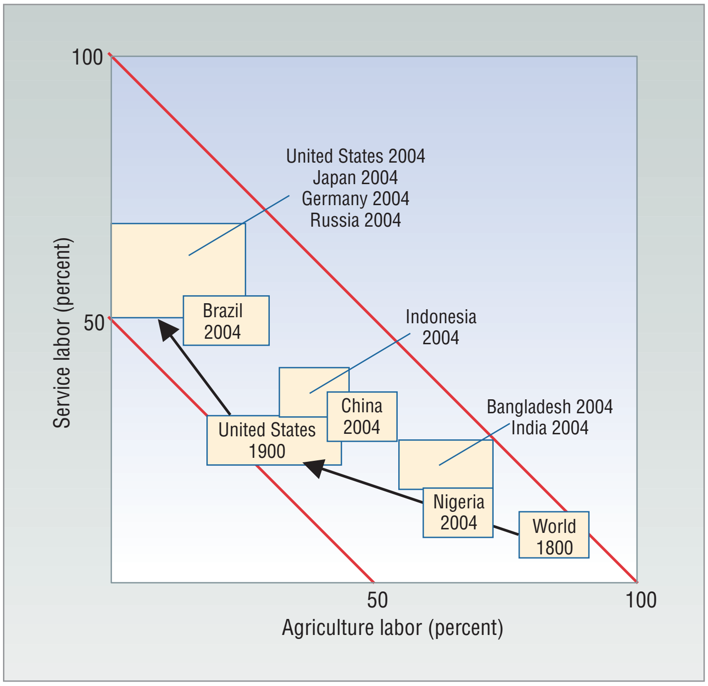

Figure 1. Global shift. Labor is shifting from agriculture to services in all nations, including the 10 with the largest populations.

The capabilities required to provision a service between service systems are distributed among people, technology, other service systems, and shared information. Douglas Engelbart[13,14] made similar distinctions when he referred to basic human capabilities coevolving with a human system and tool system. The result of the coevolution is a capability infrastructure that can augment knowledge workers and improve organizations' collective intelligence.

Carliss Y. Baldwin and Kim B. Clark[15] analyzed the coevolution of tools and human systems for the computer industry. They identified six modular operators—splitting, substituting, augmenting, excluding, inverting, and porting—that cover the full range of operations that service-system designers can do to service systems or any other type of human-designed artifact or system.

Baldwin and Clark examined the short-term and long-term economic impact of a modularity decision. Hunt[11] described the role of the entrepreneur and innovation in the context of a general theory of competition and the disequilibrium-provoking impact that innovation produces. What emerges is a picture of service systems with a complex internal service-system structure embedded in ecosystems with a complex external service-system structure.

In sum, several theories have identified the building blocks of service systems, but researchers have not yet developed a theory of service systems.

TOWARD A THEORY OF SERVICE SYSTEMS

A general theory of service systems should consist of three parts:

- *Science*–what service systems are and how to understand their evolution;
- *Management*–how to invest to improve service systems; and
- *Engineering*–how to invent new technologies that improve the scaling of service systems.

Science

The roles of people, technology, shared information, as well as the role of customer input in production processes and the application of competences to benefit others must be described and defined.

Globally, human labor is shifting to the service industries, as Figure 1 shows. Substituting technology for human labor in many agricultural and manufacturing processes is accelerating this trend. Human labor involves a range of physical, mental, and social actions. Machines can sometimes carry out routine physical and mental human actions more cost-effectively and precisely. People use ATMs, kiosks, e-commerce Web sites, and other self-service technologies rather than engaging in routine social interactions with other people.

Decomposing work into separable service activities that let labor move to low-cost geographies is also driving this trend. All employees in an organization render services to complete their tasks. When outsourcing components or substituting technology for components reconfigure work practices, it's not uncommon to refer to the reconfiguration as a decomposition of the work into separate services, as well as the creation of new services. With the rise of the Internet and Web services, enterprise architects are increasingly using SOA to flexibly integrate and dynamically reconfigure both human and computational services.

Knowledge intensity is increasingly a part of modern service value propositions. Nearly all service industries show growth in knowledge intensity, both through more skilled labor and more use of advanced technology. Even service industries, such as retail and hospitality franchises, with value propositions based on low-cost and generic skilled labor, invest heavily in advanced technological and organizational infrastructures that ensure productivity, quality, and compliance with standards.

Management

Understanding service system improvements and failures to improve is important to a theory of service systems, as it would enable effective management of service systems. We propose a triple-loop learning framework that is based on evaluating return on investments of transformation efforts aimed at improvement. The framework has three dimensions:

- *Efficiency* (plans). Things are done in the right way.
- *Effectiveness* (goals). The right things get done.
- *Sustainability* (relationships). The right relationships exist with other service systems.

Using shared information that ranges from news reports and polls to surveys and government and scientific studies, service systems can compare themselves to each other along efficiency, effectiveness, and sustainability dimensions. Using private internal information, service systems can compare their current state to previous states and identify historical trends in key performance indicators.

Well-known service systems with excellent reputations routinely receive value proposition proposals from other service systems to coproduce value. Reputation is often critically important to sustainability. The amount of shared information available to all service systems in an ecology of service systems enhances coordination and mutual sustainability.

Access to shared information that describes how to perform many different types of services can enhance the versatility of service systems and allow service systems to exploit self-service when there's a paucity of external value propositions choices. Efficiency concerns tend to push service systems toward overspecialization, while sustainability concerns tend to push service systems toward diversification and general competences. Effectiveness concerns tend to push service systems toward value propositions with the highest returns and longest expected time horizon for high returns.

Engineering

Under what conditions does a service system improve itself, and how can we design such systems that improve? Consider franchises and online e-commerce. In the past 50 years, franchises have transformed the business landscape. Local businesses have flocked to the franchise model, which enjoys economies of scale that local businesses don't. From the perspective of mobile service clients, franchises exploit shared information to provide a standard service experience and value proposition, independent of location.

E-commerce Web sites represent another recent advance related to scaling of service. Access to an online service, independent of geographic constraints, lets a service system more efficiently scale up internal and external service transactions. A system's ability to scale up depends on many factors. Most important is the nature of the resources that the service system is integrating to realize the competence being delivered as service.

Three types of key resources make up all service.

People. The more they're needed and the longer it takes to educate them or get them to competent performance, the more expensive human resources typically become. For example, each profession has only a limited number of people, and training more people with those professional skills takes time and educational investment.

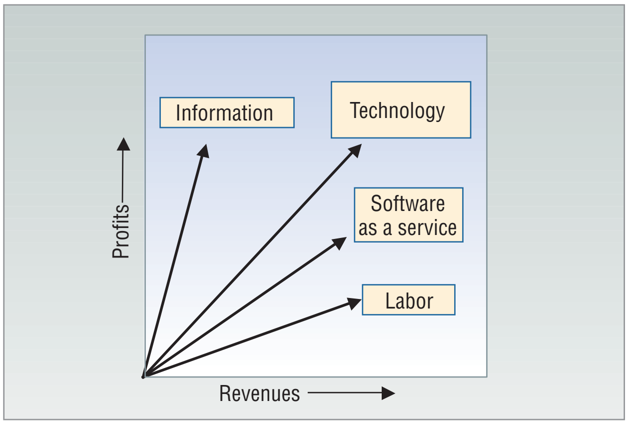

Figure 2. Revenue and profit scales. Revenue and profit scale differently in service systems that integrate different types of resources.

So scaling a service system that depends on human resources might require seeking out labor from another less expensive geography, repurposing and retraining people from another industry sector, or identifying demographic segments yet to join the labor force.

Technology. Technological resources are like most physical material supplies. Typically, the more one buys, the lower the price vendors demand. The incremental cost of the next unit of production is lower than the last. Thus, a service system can take better advantage of scale if it integrates technology or other types of physical material resources.

Shared information. Informational resources enjoy incredible scale efficiency because of the small incremental cost in duplicating them. Creating the next unit of an informational resource has virtually zero cost. Nevertheless, pirating and illegal copying can erode some of the advantages of scaling service systems based on informational resources.

Service systems integrate people, technology, and information resources in different proportions. As a result, each system is unique, resulting in situations in which revenue and profits scale differently. Figure 2 compares the revenue and profit-scaling properties of firms based on software, product, software as a service, and high- and low-skill labor-based services.

Service systems as computational systems

Because IT is such an important part of service systems today, we might ask how service systems are similar to and different from computational systems.

The main difference is people. The largest service system, the global economy, includes more than six billion people. Some large firms have hundreds of thousands of employees. People do a lot of the work—physical, mental, and social. Furthermore, unlike computational system components, we can't easily model and simulate the behavior of people doing work in service systems. For example, laws and policies only partially govern people. Even when citizens and employees know government laws and corporate policies, compliance isn't complete, which creates risk as well as opportunities.

So, perhaps if we model people as components with stochastic behavior, we could apply existing theories of computational systems to service systems. For example, the notion of trust is well developed in fields of computer science that deal with privacy protection and secure systems. But noncompliance creates opportunities as well as risks.

Many innovations break a rule or violate a policy. How can we tell the difference between cheating and innovation in a service system, where people informally and formally change rules and policies?

Service systems are complex adaptive systems made up of people, and people are complex and adaptive themselves. Service systems are dynamic and open, rather than simple and optimized. And there are many different kinds of value, including financial, relationship, and reputation.

Mechanism design theory, a new branch of theoretical computer science that integrates with game theory and economics, introduces the notion of a social utility function in the context of computational systems. A fundamental problem in economics and game theory hinges on the fact that sometimes individual and collective goals aren't aligned. The emerging field of incentive engineering, which human-capital management students in business schools are studying, addresses the problem of incentive alignment of individuals and larger groups.

Services science is an emerging field that seeks to tap into these and other relevant bodies of knowledge, integrate them, and advance three goals—aiming ultimately to understand service systems, how they improve, and how they scale.

A science of service can provide a foundation for creating lasting improvements to service systems. Nevertheless, we're only beginning this effort. Service science aims to understand and catalog service systems and to apply that understanding to advancing our ability to design, improve, and scale service systems for practical business and societal purposes. The study of service systems is an integrative, multidisciplinary undertaking, and many disciplines have knowledge and methods to contribute. Nothing is settled, and we still have much work to do.

Acknowledgment

We thank Michael Maximilien for helpful comments on an early draft.

References

1. US Nat'l Academy of Engineering, *"The Impact of Academic Research on Industrial Performance,"* Nat'l Academies Press, 2003.
2. H. Chesbrough and J. Spohrer, "A Research Manifesto for Services Science," *Comm. ACM*, July 2006, pp. 35-40.
3. P.M. Maglio et al., "Service Systems, Service Scientists, SSME, and Innovation," *Comm. ACM,* July 2006, pp. 81-85.
4. J. Spohrer et al., "Convergence and Coevolution: Towards a Services Science," M.C. Roco and W.S. Bainbridge, eds., *Nanotechnology: Societal Implications 1: Maximizing Benefits for Humanity*, Springer, 2006.
5. J. Spohrer and D. Riecken, "Special Issue: Services Science," *Comm. ACM,* July. 2006, pp. 30-32.
6. R.F. Lusch and S.L. Vargo, eds., *"The Service-Dominant Logic of Marketing: Dialog, Debate, and Directions,"* M.E. Sharpe, 2006.
7. S.E. Sampson and C.M. Froehle, "Foundations and Implications of a Proposed Unified Services Theory," *Production and Operations Management,* summer 2006, pp. 329-343.
8. J.M. Tien and D. Berg, "A Case for Service Systems Engineering,"*J. Systems Science and Systems Eng.,* Mar. 2003, pp. 113-128.
9. J. Gadrey, "The Misuse of Productivity Concepts in Services: Lessons from a Comparison between France and the United States," J. Gadrey and F. Gallouj, eds., *Productivity, Innovation and Knowledge in Services: New Economic and Socio-economic Approaches,* Edward Elgar Publishing, 2002.
10. B. Paton, "Two Pathways to Energy Efficiency: An Energy Star Case Study," *Human Ecology Rev.*, vol. 11, no. 3, 2004, pp. 247-259; www.humanecologyreview.org/pastissues/her113/paton.pdf.
11. S. Hunt, *A General Theory of Competition: Resources, Competences, Productivity, and Economic Growth,* Sage Publications, 2000.
12. R.R. Nelson and S.G. Winter, *An Evolutionary Theory of Economic Change,* Belknap/Harvard Univ. Press, 1982.
13. D.C. Engelbart, "Augmenting Human Intellect: A Conceptual Framework," summary report, Stanford Research Institute, Oct. 1962.
14. D.C. Engelbart, "Evolving the Organization of the Future: A Point of View," *Proc. Stanford Int'l Symp. Office Automation*, Ablex, 1980, pp. 287-296.
15. C.Y. Baldwin and K.B. Clark, *"Design Rules: The Power of Modularity,"* MIT Press, 2000.

***Jim Spohrer** is director of Almaden Services Research at IBM Almaden Research Center. His research interests are service research, artificial intelligence, and learning systems. He received a PhD in computer science from Yale University. He is a member of ACM. Contact him at spohrer@us.ibm.com.*

***Paul P. Maglio** is senior manager of service systems research at IBM Almaden Research Center. His research interests are service research, skill learning, and human interaction with IT systems. He received a PhD in cognitive science from University of California, San Diego. He is a member of the Cognitive Science Society and ACM. Contact him at pmaglio@us.ibm.com.*

***John Bailey** is a research scientist at IBM Almaden Research Center. His research interests are human factors in service systems, service science, and human-computer interaction. He received a PhD in human factors psychology from the University of Central Florida. He is a member of the Human Factors and Ergonomics Society and ACM. Contact him at baileyj@us.ibm.com*

***Daniel Gruhl** is a research scientist at IBM Almaden Research Center. His research interests are semantic computation, service systems, and distributed supercomputers He received a PhD in electrical engineering from the Massachusetts Institute of Technology. Contact him at dgruhl@us.ibm.com.*

IEEE

IRI-2007

Sponsored by:
The IEEE Systems, Man and Cybernetics Society

The 2007 IEEE International Conference on INFORMATION REUSE AND INTEGRATION

August 13-15, 2007

Hilton Hotel, Las Vegas, USA

http://www.sis.pitt.edu/~ iri07/

With rapidly increasing volumes of information in digital forms, we are constantly charged with the challenges of efficiency in information usage and knowledge extraction. *Information Reuse and Integration* (IRI) seeks to maximize the availability of information and creation of knowledge, and to reuse these information and knowledge in addressing new issues. IRI plays a pivotal role to capture, maintain, integrate, validate, extrapolate, and apply both information and knowledge to augment decision-making capacity in application domains. The IEEE IRI conference serves as a forum for researchers and practitioners from academia, industry, and government to present, discuss, and exchange ideas that address real-world problems with real-world solutions. The IEEE IRI will feature contributed as well as invited papers. Theoretical and applied papers are both included in this call. The conference program will include special sessions, open forum workshops and keynote speeches. Several funding agency program directors - including NSF, ONR, et al. - will present an open panel discussion entitled *Funding Opportunities in Information Reuse and Systems Engineering*. The conference includes, ***but is not limited to***, the areas listed below:

- Large Scale Data and System Integration
- Component-Based Design and Reuse
- Unifying Data Models (UML, XML, etc.) and Ontologies
- Database Integration
- Structured/Semi-structured Data
- Middleware & Web Services
- Reuse in Software Engineering
- Data Mining and Knowledge Discovery
- Sensory and Information Fusion
- Reuse in Modeling & Simulation
- Information Security & Privacy
- Automation, Integration and Reuse Across Applications
- Survivable Systems & Infrastructures AI & Decision Support Systems
- Heuristic Optimization and Search
- Knowledge Acquisition and Management
- Fuzzy and Neural Systems
- Soft/Evolutionary Computing
- Case-Based Reasoning
- Natural Language Understanding
- Knowledge Management and E-Government
- Command & Control Systems (C4ISR)
- Human-Machine Information Systems
- Biomedical & Healthcare Systems
- Homeland Security & Critical Infrastructure Protection
- Manufacturing Systems & Business Process Engineering
- Space and Robotic Systems
- Multimedia Systems
- Service-Oriented Architecture
- Autonomous Agents in Web-based Systems
- Information Integration in Grid, Mobile and Ubiquitous Computing Environment
- Systems of Systems
- Semantic Web and Emerging Applications
- Information Reuse, Integration and Sharing in Collaborative Environments

Instructions for Authors: Papers reporting original and unpublished research results pertaining to the above and related topics are solicited. Full paper manuscripts must be in English of length 4 to 6 pages (using the IEEE two-column template). Submissions should include the title, author(s), affiliation(s), e-mail address(es), tel/fax numbers, abstract, and postal address(es) on the first page. Papers should be submitted at the conference web site: **http://www.sis.pitt.edu/~iri07/**. If web submission is not possible, manuscripts should be sent as an attachment via email to either of the Program Chairs on or before the deadline date of March 25, 2007. The attachment must be in .pdf (preferred) or word.doc format. The subject of the email must be "IEEE IRI 2007 Submission." Papers will be selected based on their originality, timeliness, significance, relevance, and clarity of presentation. Authors should certify that their papers represent substantially new work and are previously unpublished. Paper submission implies the intent of at least one of the authors to register and present the paper, if accepted. Authors of selected papers that are also presented at the conference will be invited to submit expanded versions of their papers for review for publication in an approved special issue of the IEEE SMC Transactions, part C, on IRI to be published in 2008.

Important Dates:

February 11, 2007	Workshop/Special session proposal
March 25, 2007	Paper submission deadline
April 29, 2007	Notification of acceptance
May 20, 2007	Camera-ready paper due
May 20, 2007	Presenting author registration due
July 10, 2007	Advance (discount) registration
July 31, 2007	Hotel reservation (special discount rate) closing date

General Chairs

Stuart Rubin
SPAWAR Systems Center, USA
stuart.rubin@navy.mil

Shu-Ching Chen
Florida International University, USA
chens@cs.fiu.edu

Program Chairs

Weide Chang
California State University, USA
changw@ecs.csus.edu

James B. D. Joshi
University of Pittsburgh, USA
jjoshi@mail.sis.pitt.edu