



Wireless Grids

Distributed Resource Sharing by Mobile, Nomadic, and Fixed Devices

Wireless grids, a new type of resource-sharing network, connect sensors, mobile phones, and other edge devices with each other and with wired grids. Ad hoc distributed resource sharing allows these devices to offer new resources and locations of use for grid computing. This article places wireless grids in context, explains their basic requirements, and provides an example implementation that uses a wireless grid for distributed audio recording. Finally, it introduces articles in this special issue on wireless grid architectures and applications.

**Lee W. McKnight
and James Howison**
Syracuse University

Scott Bradner
Harvard University

Mobile, nomadic, and fixed wireless devices form new types of resource-sharing networks called wireless grids.¹ This issue of *IEEE Internet Computing*, beginning with our extended editorial introduction, presents an overview of current research in the area. We hope this issue of *IC* will serve to map the wireless grid landscape, to help others explore it in greater detail in the future.

Grid computing lets devices connected to the Internet, overlay peer-to-peer networks, and the nascent wired computational grid dynamically share network-connected resources. The wireless grid

extends this sharing potential to mobile, nomadic, or fixed-location devices temporarily connected via ad hoc wireless networks.² As Figure 1 shows, users and devices can come and go in a dynamic wireless grid, interacting with a changing landscape of information resources.

The dynamic case of unknown devices creates special challenges, which explains why the articles in this special issue focus on the topic. Following Metcalfe's law, grid-based resources become more valuable as the number of devices and users increases. The wireless grid makes it easier to extend grid computing to large

numbers of devices that would otherwise be unable to participate and share resources. While grid computing attracts much research, resource sharing across small, ad hoc, mobile, and nomadic grids draws much less. In fact, some readers will consider the topics in this special issue to fall outside their definition of grids. We – and the authors in this issue – believe otherwise, and will seek to convince those readers.

Wireless Grids

In some ways, wireless grids resemble networks already found in connection with agricultural, military, transportation, air-quality, environmental, health, emergency, and security systems. A range of institutions, from the largest governments to very small enterprises, will own and at least partially control wireless grids. To make things still more complex for researchers and business strategists, users and producers could sometimes be one and the same. Devices on the wireless grid will be not only mobile but nomadic – shifting across institutional boundaries. Just as real-world nomads cross institutional boundaries and frequently move from one location to another, so do wireless devices.^{1,3}

The following classification offers one way to classify wireless grid applications.² Our intent is to illustrate current conceptual frameworks for wireless grids research, and to encourage readers of this special issue to think of their own computational models, architectures, and applications for wireless grids.

- Class 1: Applications aggregating information from the range of input/output interfaces found in nomadic devices.
- Class 2: Applications leveraging the locations and contexts in which the devices exist.
- Class 3: Applications leveraging the mesh network capabilities of groups of nomadic devices.

The three classes of wireless grid applications conceptualized here are not mutually exclusive. Understanding more about the shareable resources, the places of use, and ownership and control patterns within which wireless grids will operate might assist us in visualizing these future patterns of wireless grid use.

What's New About Wireless Grids

Wireless grids offer a wide variety of possible applications. They can reach both geographic locations and social settings that computers have not

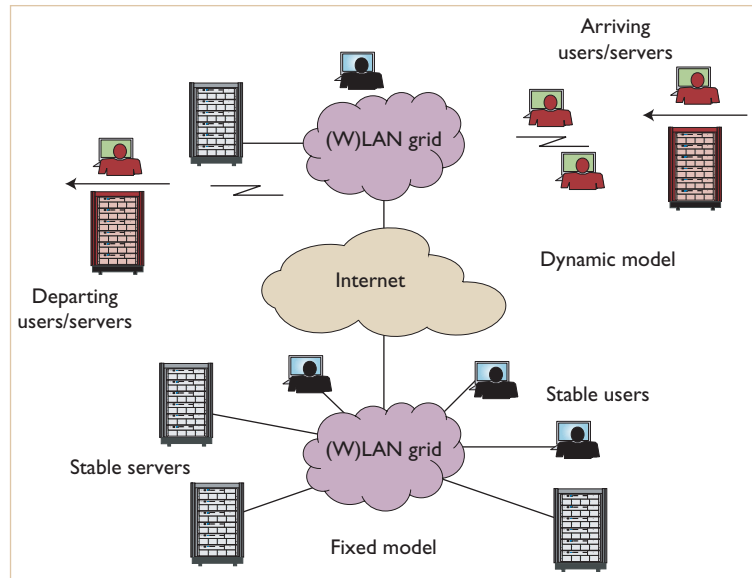


Figure 1. Dynamic and fixed wireless grids. Here we see two types of wireless grids: those composed of unknown mobile users and devices engaged in ad hoc resource sharing and service creation in a particular location, and those composed of components with known identities managed within a stable institutional structure.

traditionally penetrated. Wireless grids present three novel elements:

- new resources,
- new places of use, and
- new institutional ownership and control patterns.

Wireless devices bring new resources to distributed computing. In addition to typical computational resources such as processor power, disk space, and applications, wireless devices increasingly employ cameras, microphones, GPS receivers, and accelerometers, as well as an assortment of network interfaces (cell, radio, Wi-Fi, and Bluetooth). One important class of devices is *sensors*, which can supply information on temperature, health, or pollution levels, to name just a few.

People increasingly take wireless devices with them to new places, in both their personal and professional lives. The numbers of those devices that include sensors are growing. In fact, the pervasive mobile phone is developing into a super-sensor. From shopping malls to medical disaster areas, sporting events, and warehouse floors, wireless devices – and the sensors in them – are on the verge of becoming ubiquitous. Wireless grids present an opportunity to leverage available resources by enabling sharing between wireless and non-wireless resources.

The evolution from mainframes to the PC and to handhelds mirrors the history of institutional change from centralization to decentralization.⁴ Before the 1980s, only large institutions owned computing resources, and they used them for specific purposes. The PC revolution radically altered this situation, yet networks (LANs) still clearly reflected the administrative control of organizations. The Internet's emergence again broadened the ownership and control of devices on networks, introducing consumers and individual owners to the mix. Emerging peer-to-peer overlay networks, which create useful applications from heterogeneous devices, each with its own owner, demonstrate the power of resource-sharing. P2P also shows the difficulties and risks of designing network applications for a cross-institutional, heterogeneous environment. Indeed, P2P networks are overlay networks over which the institutional owners of the physical infrastructure often have only limited control: user-owned devices utilizing unlicensed spectrum can remove even this token of central control. Wireless grids represent the epitome of these transitions.

New resources, locations, and institutional structures present the opportunity, and the challenge, of wireless grids. This introduction to wireless grids focuses on technical approaches to wireless resource sharing. Our prototype wireless grid application demonstrates the practicality of this theoretical approach. We provide a brief technical description of our modest initial application, a distributed audio-recording prototype immodestly named DARC* (Distributed Ad Hoc Resource Coordination, pronounced "dark star"). We also briefly address the more general economic and policy issues with wireless grids to illustrate the potential costs and benefits.

Wireless Grids in Context

Wireless grids emerged from a combination of the proliferation of new spectrum market business models, innovative technologies deployed in diverse wireless networks, and three related computing paradigms: grid computing, P2P computing, and Web services.

Wired grids are typically aggregations of fixed resources between known institutions, be they academic or corporate, in high-trust and relatively static environments. Fixed wireless grids, like the one in Figure 1, borrow from the wired grid model. To participate in the Grid using the current Globus software, for example, a machine must create and

preregister an X.509 certificate. These static, trusted environments stand in stark contrast to the situation facing the wireless grid. Foster and Imanichi describe the wired-grid focus as "integration of substantial resources to deliver non-trivial qualities of service within an environment of at least limited trust."⁵ Nonetheless, the aims of grid computing are the same as those that underlie our efforts: "flexible, secure, and coordinated resource sharing among dynamic collections of individuals, institutions, and resources."⁶

P2P networks, such as Napster, Gnutella, and Kazaa, have characteristics in common with wireless grids. They must arrange coordinated sharing among heterogeneous devices, across unreliable network connections, with little or no prearrangement and little or no warning of site failure. While grid computing had to create a "persistent, standards-based service infrastructure,"⁵ P2P computing must address resource sharing in the face of unreliable networks and end devices.

Web services provide a third computing framework relevant to our approach to the wireless grid. Despite being so overhyped, at heart Web services simply facilitate remote access to resources. The idea that access to a pool of services occurs between a provider and a customer, or between two business partners, constrains Web services. Yet the technologies at the center of Web services – SOAP, lightweight XML message-passing, and the flexible Web Service Definition Language (WSDL) – apply to any ad hoc sharing situation.

The wireless grid thus draws on at least three of the computing paradigms currently undergoing rapid development. Figure 2 places wireless grids in context, illustrating how they span the technical approaches and issues of Web services, grid computing, P2P systems, mobile commerce, ad hoc networking, and spectrum management. How sensor and mesh networks will ultimately interact with software radio and other technologies to solve wireless grid problems requires a great deal of further research, but Figure 2 at least captures many of the main facets of a wireless grid.

Wireless Grid Infrastructure

We can understand the challenges of the wireless grid infrastructure by breaking them down into three categories:

- physical-layer technologies and policy,
- requirements for network infrastructure, and
- enabling middleware.

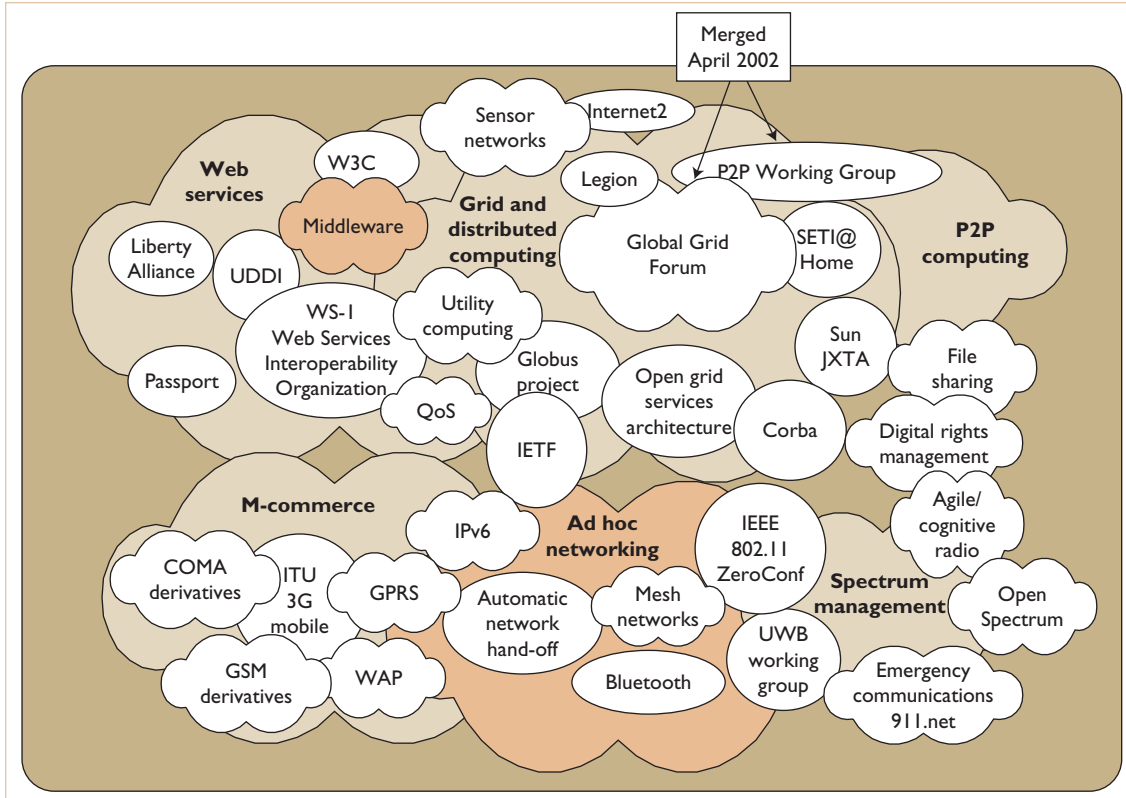


Figure 2. A wireless grid issues and standards map. The variety of issues and technologies illustrates how complex the needs of wireless grids are.

At the physical level, where wireless devices share spectrum, the Open Spectrum⁷ campaign aims to replace the current FCC licensing model with rules-based public access to the airwaves. This effort and new methods of spectrum management are crucial to making wireless grids a reality.

At the networking layer, researchers are looking at power efficiency and coverage in wireless networks.⁸ Wireless devices are, by their mobile nature, often battery-powered; power consumption is therefore a crucial issue. Distributed applications can assist in saving power by offloading processing or data management to devices connected to the power grid, or to networks that require less transmitter power to reach.¹

In the following sections, we outline the development challenges the middleware research community has identified so far. The five abstract requirements for ad hoc resource sharing are resource description, resource discovery, coordination, trust establishment, and clearing.

Resource Description

Resource description is a basic requirement for resource sharing. Before any group of devices can

discuss needs and available resources, it must first agree on the manner in which to describe the resources. For example, if a group of devices wishes to share processor cycles, it must first be able to describe the processing requirements and capabilities. A variety of schemes provide for resource description; none of the schemes covers all resources, but taken together, they define most of the shared resources.

Different resource-sharing systems undertake the task of resource description in different ways. For example, P2P music-sharing networks utilize the filename and ID3 metadata tags.

Resource Discovery

Standards such as the IETF's ZeroConf, Microsoft's Universal Plug and Play, the Grid Resource Description Language (GRDL), standardized service-specific definitions using WSDL, and bandwidth descriptions from various QoS standards include resource-description protocols. Using these description vocabularies, devices can formulate their needs and publish their resources. This is the abstract process of resource discovery.

Different resource-sharing systems accomplish

this step in different ways. For example, the Open Grid Services Architecture (OGSA) system utilizes a Web service-style IndexService to gather resource descriptions published by each service instance. The Web services community has defined UDDI and its associated protocols for making a database of services available. Systems such as JXTA employ Zero-Conf, best known as Apple's Rendezvous, to discover resources using multicasts over the LAN.

Coordination Systems

Coordination systems allow one device to utilize another device's resources, or permit the pooling and scheduling of resources. Each device can use a range of familiar mechanisms to share different resources. For example, to share disk space, a device might use NFS, Samba, or WebDav as the coordination mechanism. On the other hand, to share processor cycles, a device might use the distributed.net client (available at www.distributed.net) as the coordination mechanism.

The new resources that nomadic devices make available raise new challenges in coordination of resource sharing. For example, people are still developing the coordination systems devices for sharing and aggregating sensor data. Facilitating the sharing of screens and audio-visual inputs will require new coordination mechanisms. Right now, even having laptop users share a projector as an output device during a meeting calls for much fiddling with cables. We have argued elsewhere¹ that a new sharing protocol could serve as a metaprotocol to support resource description, discovery, coordination, and trust establishment, but have not yet begun to test and evaluate how such a protocol might perform in practice.

Trust Establishment

Clearly, a resource-sharing transaction requires trust. Indeed, each element in a transaction might even require a different type of trust establishment. For example, before you provide a description of your available resources, you might want to be certain of the identity of the device or user that you are talking with, perhaps via an institutionalized identity system such as a public-key infrastructure or Kerberos. However, this trust establishment could differ significantly from what a cycle-aggregation system such as SETI@Home requires in order to have assurance in the results of code executed on other clients of the distributed system.

Any negotiation process must be able to ensure that the transaction partners' identities don't change

during the process or between steps. In other words, resource-sharing transactions must resist man-in-the-middle attacks. Therefore, we add as a requirement the use of a system that can assure the partners of at least a persistent anonymous identity. Mobile IP systems face a similar problem: they must be sure that the device requesting a change in the redirection of packets is the same as the device that initiated the communication. The Internet draft on Purpose-Built Keys proposes such a system.⁹

Clearing Mechanisms

A key component of a resource-sharing transaction is a clearing mechanism that establishes conditions under which a partner device or group of devices will extend access to the requesting device or group. This includes mechanisms that typically fall within an authorization process, but extend well beyond it. We use the term "clearing" in its economic sense, referring to the action (usually payment) required to "clear" or settle a market transaction. Most resource-sharing systems don't currently implement complex conditions for accessing resources. Typically, systems grant access to resources based on the ability of a device, service, or user to prove membership to an appropriate class. Common clearing protocols are Kerberos and X.509, which the OGSA toolkit uses.

Resource-sharing systems, such as P2P networks, that are outside the institutional environments in which schemes like Kerberos operate have innovated in this area and typically employ some form of quid-pro-quo exchange with their users. This sometimes means requiring users to run the full client — meaning that, in exchange for access to the network, they share the files they download. Other systems, such as BitTorrent, reserve superior service for network nodes that also provide uploads.¹⁰ We also anticipate resource-sharing situations in which bartering available resources for the period of their use will make economic sense, as in an effective power-trading network. The range of clearing mechanisms can clearly extend to complex market situations.

The Economics of Wireless Grids

Designers of resource-sharing applications that use cross-institutional nomadic devices face the problem of managing strategic behavior in their applications. This is no trivial matter, as failure to do so can lead to resource exploitation similar to that of spam, where the self-interest of the spammers can threaten the entire email infrastructure.

The market provides a potent mechanism for

coordinating behavior in distributed systems. Competitive markets, when they work, provide an efficient mechanism for allocating resources without presuming any common interest among resource producers or consumers. Unfortunately, the ideal of perfect competition seldom occurs in the real world.

Despite decades of study, substantial theoretical and practical challenges remain for economic systems and computer network models.¹¹ One of the key developments in this field is the recognition that the systems can't simply provide efficient resource allocation, but must also manage participants' incentives and strategies, as well as users' expectations.¹² So far, the market-controlled approaches to resource allocation in grids¹³ have failed to adequately address these strategic issues.¹⁴

The problem that clearing-mechanism design presents in the face of strategic activity is one of truthful revelation. Systems designers can build excellent resource-sharing systems as long as they can rely on the resource announcements and advertisements – if agents or their devices lie, all bets are off and system performance will be unreliable at best. The solution, or part of it, rests in incentive-compatible or *strategy-proof* computing, which utilizes game-theoretic mechanism design under the research banner of distributed algorithmic mechanism design.¹⁵ Work on trading agreements in quality of service also comes into play.³

Cost-modeling approaches, which examine the relative costs of local performance against those of distributed performance, can assist in understanding when it is cost-effective to engage in resource-sharing behavior.¹⁶

Applications

As a proof of concept and a demonstration of ad hoc distributed resource sharing, the Wireless Grids research team at Syracuse University is building DARC*. The system lets devices with no prior knowledge of each other collectively record and mix an audio signal such as a concert, speech, lecture, or emergency event. The project demonstrates the potential of wireless grids and distributed ad hoc resource sharing to harness the combined ability of mobile devices in social contexts outside the expected environments for computing.

Individual devices might face barriers to creating quality recordings – design limitations such as mono input, for example, or locational disadvantages such as being positioned at the far edge of the stage. DARC* enables the devices to collectively eliminate these disadvantages, turning the spa-

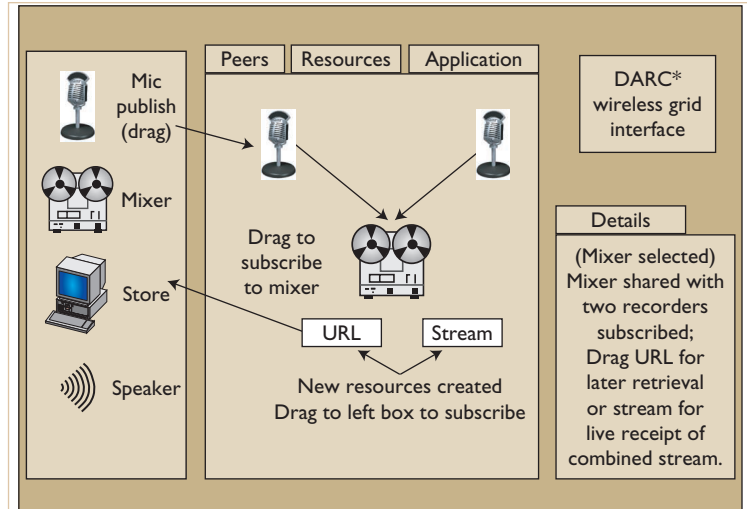


Figure 3. Example user interface for DARC*. Drag and drop allows users to easily publish their services and compose them into applications.

tial distribution into higher-quality stereo sound and, eventually, surround sound. The application would also help to create a total audio record of a large-scale event in which static microphones could record only nearby sounds. In addition to the resource-sharing elements, the development team has had to design the coordination mechanism for distributed music.

The virtual application consists of a *mixer* and two or more *recorders*. A user wishing to initiate a recording session elects to act as either or both a mixer and a recorder and waits for the involvement of a second *recording service*. Recording begins once two services have registered with the mixer. The recorders stream the recordings to the mixer, ensuring low resource requirements. The recorders then initiate a *listening service* to receive the mix back from the mixer. This initially takes the form of a password and URL for later retrieval.

Figure 3 illustrates the user interface for DARC*, which we intend to generalize for use in our middleware framework. One of the crucial challenges we face is creating an interface that will work on small-profile devices such as cell phones and PDAs while usefully scaling to traditional computers.

Once the recording is complete, the mixer's UI allows the user to ensure that the channels are in sync, apply desired digital effects, and then combine the mono streams into a stereo WAV file. The mixer reciprocates the recorders' service by streaming the WAV file back or providing a URL.

This simple demonstration incorporates each of the abstract elements of resource sharing we out-

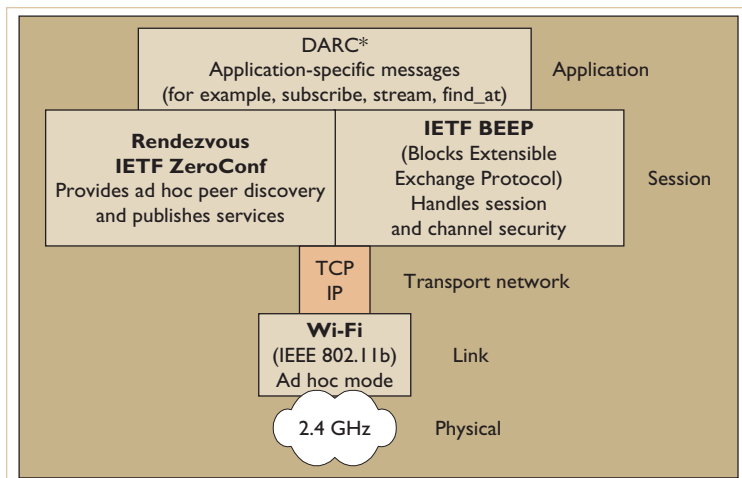


Figure 4. A full protocol stack for DARC*. The implementation architecture for this distributed audio-recording application illustrates the DARC* system.

lined earlier. Figure 4 shows the full protocol stack for DARC*. DARC* devices describe resources using a simple string, then discover them using multicast DNS (a joint effort between IETF's ZeroConf and DNS Extensions working groups). Coordination occurs through application-specific messages via the Blocks Extensible Exchange Protocol (BEEP), which manages session persistence and security, thus demonstrating trust management. The clearing mechanism we used is a simple exchange of an audio stream for a token to receive the resulting sound mix later. We implemented the entire application in Java, drawing on open-source implementations of Rendezvous and BEEP.

We intend to demonstrate this application on laptops, PDAs, and cell phones. We also intend to build a Web services-style implementation that utilizes WSDL and UDDI, which will facilitate integration with OGSA and wired grids. We are exploring the use of the Friend-of-a-Friend (FOAF) protocol for group formation and as a clearing mechanism.

Next Stages

The emergence of wireless grids parallels the historical trend that has seen computing shift from a hierarchical structure – in which computing was an organizationally controlled activity – to a situation in which the only guarantee is that individual users will follow their strategic interests. The initial developers of grid computing applications and architectures naturally focused on applications deployable within and across hierarchically controlled organizations such as super-

computing centers. To reach commercial markets and end users, designers of P2P networks and wireless grids must learn to leverage strategic trends toward mobility and nomadicity. Dynamic wireless grids cannot be controlled in the same manner as a traditional network design.

Application developers have an opportunity to draw on the new resources, interfaces, and locations that wireless devices provide. We have sketched the abstract requirements for ad hoc resource sharing and described a modest demonstration application. An abundance of research challenges remain in crafting engineering applications and appropriate radio technologies, as well as developing reliable clearing mechanisms and spectrum policy. As the articles in this special issue show, researchers are tackling applications of wireless grids in education, emergency communications, health, and other markets. Commercial opportunities might soon emerge, if early prototypes and experimental studies offer any indication.

Articles in This Issue

Unfortunately – or fortunately, depending on your point of view – a great deal of research and development remains before wireless grids achieve their full potential. Both of the articles in this special issue describe parts of the ongoing research and the associated challenges.

In “Integrating Wireless Sensor Networks with the Grid,” Mark Gaynor and his colleagues describe innovative medical and warehouse monitoring applications they’ve developed to address the challenges and architectures for integrating wireless and wired grids. Ad hoc networks would enable emergency response teams to quickly establish a monitoring network in mass-casualty events. In warehouses, smart-dust motes could monitor expensive purchases for damage throughout the supply chain. The team designed these initial applications without any standardized application-specific protocols. Instead, they propose using an hourglass model as a standardized API for integrating ad hoc sensing networks with existing wired grids.

Junseok Hwang and Praveen Aravamudham also pursue the goal of integrating wireless and wired grids. Their proposed proxy-based architecture could let even highly resource-constrained wireless devices access OGSA-compliant grid services through a policy-managed delegate. Based on this initial model, the authors suggest further work toward scalable wireless grid architectures spanning large and small grids, sensors, and supercomputers.

We commend these articles to you and look forward to many more addressing the challenges of wireless grid computing. We hope you enjoy the overview of wireless grids this introduction – and this special issue – present. ☐

Acknowledgments

This research was supported in part by US National Science Foundation (NSF) grant no. 0227879, the CASE Center at Syracuse University, BT, Cisco, Novell, Nokia, Frankston Innovating, and River's Edge. The authors gratefully acknowledge the helpful comments and contributions of academic, industry, and government colleagues in the NSF Partners for Innovation "Virtual Markets in Wireless Grids" (www.wirelessgrids.net) project, as well as the contributions of the DARC* (Distributed Ad Hoc Resource Coordination) development team: David Grandinetti, Ian Molloy, Garrett Wilkins, and Harshavardhana Kikkeri.

References

1. L.W. McKnight and J. Howison, "Towards a Sharing Protocol for Wireless Grids," *Proc. Int'l Conf. Computer Comm. and Control Technologies*, Int'l Inst. of Informatics and Systems, Orlando, Fla., vol. 000648, 2003.
2. L.W. McKnight and M. Gaynor, "Wireless Grid Issues," *Proc. 8th Global Grid Forum (GGF8)*, Global Grid Forum, 2003; www.wirelessgrids.net/docs/draft-ggf-lwmcknight-wg-issues-0.pdf.
3. L.W. McKnight and W. Lehr, "Show Me the Money: Agents and Contracts in Service-Level Agreement Markets," *INFO*, vol. 40, no. 2, Feb./Mar. 2002, pp. 24-36.
4. L.W. McKnight, W. Lehr, and J. Howison, "Coordinating User and Device Behavior in Wireless Grids," *Information Systems Research*, submitted for publication, 2004.
5. I. Foster and A. Imanitchi, "On Death, Taxes, and the Convergence of Peer-to-Peer and Grid Computing," *Proc. Peer-To-Peer Systems II: 2nd Int'l Workshop (IPTPS '03)*, LNCS 2735, Springer-Verlag, 2003.
6. I. Foster, C. Kesselman, and S. Tuecke, "The Anatomy of the Grid: Enabling Scalable Virtual Organizations," *Proc. Euro-Par 2001 Parallel Processing*, LNCS 2150, Springer-Verlag, 2001, pp. 1-4.
7. D. Weinberger et al., "Open Spectrum FAQ," 2003; www.greaterdemocracy.org/OpenSpectrumFAQ.html.
8. B. Chen and C.H. Chang, "Mobility Impact on Energy Conservation of Ad Hoc Routing Protocols," *Proc. Int'l Conf. Advances in Infrastructure for Electronic Business, Education, Science, Medicine, and Mobile Technologies on the Internet (SSGRR '03)*, Scuola Superiore G. Reiss Romoli, Abruzzo, Italy, 2003.
9. S. Bradner, A. Mankin, and J. Schiller, "A Framework for Purpose-Built Keys," IETF Internet draft, June 2003; work in progress.
10. B. Cohen, "Incentives Build Robustness in BitTorrent,"

Proc. 1st Workshop on Economics of Peer-to-Peer Systems, SIMS Berkeley, 2003; www.bitconjurer.org/BitTorrent/bittorrentecon.pdf.

11. I.E. Sutherland, "A Futures Market in Computer Time," *Comm. ACM*, vol. 11, no. 6, June 1968, pp. 449-451.
12. E. Drexler and M. Miller, "Incentive Engineering for Computational Resource Management," *The Ecology of Computation*, B. Huberman, ed., Elsevier Science, 1988.
13. J. Gomoluch and M. Schroeder, "Market-Based Resource Allocation for Grid Computing: A Model and Simulation," *Proc. 1st Int'l Workshop on Middleware for Grid Computing (MGC '03)*, 2003; http://virtual.incc.br/mgc2003/cam_ready/MGC290_final.pdf.
14. J. Shneidman and D. Parkes, "Rationality and Self-Interest in Peer to Peer Networks," *Proc. Peer-To-Peer Systems II: 2nd Int'l Workshop (IPTPS '03)*, LNCS 2735, Springer-Verlag, 2003, pp. 138-148.
15. H.R. Varian, "Economic Mechanism Design for Computerized Agents," *Proc. 1st Usenix Workshop on Electronic Commerce*, Usenix Assoc., 1995; www.usenix.org/publications/library/proceedings/ec95/varian.html.
16. J. Grey, *Distributed Computing Economics*, tech. report MSR-TR-2003-24, Microsoft Research, 2003; [ftp://ftp.research.microsoft.com/pub/tr/tr-2003-24.pdf](http://ftp.research.microsoft.com/pub/tr/tr-2003-24.pdf).

Lee W. McKnight is an associate professor at Syracuse University's School of Information Studies and a research associate professor of computer science at Tufts University. His research interests include wireless grids, virtual markets, and mobile regions. He received a PhD in political science from the Massachusetts Institute of Technology. McKnight has coauthored and coedited five books from MIT Press, including the forthcoming *Internet Services*. Contact him at lmcknigh@syr.edu.

James Howison is a doctoral candidate in information science and technology at the School of Information Studies at Syracuse University's School of Information Studies. His research interests include distributed ad hoc computing and the social science of software engineering, focusing on free and open-source software development. He received his honors undergraduate degree in economics and politics from the University of Sydney. He was an invited presenter at the first O'Reilly peer-to-peer conference in 2000. Contact him at jhowison@syr.edu.

Scott Bradner is the university technology security officer in the Harvard Office of the Provost. His research interests include privacy and security issues. He is a board member of the American Registry for Internet Numbers (ARIN) and serves as the Secretary to the Internet Society (ISOC) board. He writes the weekly "Net Insider" column for *Network World*. Contact him at sob@harvard.edu.