35

Coordinating User and Device Behavior in Wireless Grids

Lee W. McKnight[1], William Lehr[2], and James Howison[1]

[1] School of Information Studies, Syracuse University
[lmcknigh|jhowison]@syr.edu
[2] Massachusetts Institute of Technology wlehr@mit.edu

Summary. The evolution of computing is characterized by decentralization and decreasing institutional control over resources. Wireless Grids, that is, fixed and mobile ad-hoc resource sharing networks, are challenging environments in which users strategic behaviors are crucial to system performance. We review the mechanisms employed to regulate strategic behavior online–technical, social, legal and economic–and discuss trends in their operation and application in distributed wireless grid computing.

35.1 Introduction

Computing and communication networks have evolved from centralized, hierarchical systems under the management of a single entity toward decentralized, distributed systems under the collective management of many entities. Intelligence has shifted to edge-nodes, which increasingly are capable of acting as autonomous agents making complex decisions to create, deliver, or receive services [28, 46, 51]. Grid computing historically focused on the large-scale sharing of computing resources such as software, hardware, databases and data sources [14, 15]. Wireless grids organized as ad hoc networks of hardware, software, and content resources represent the epitome of this evolution from centralized systems toward ad hoc cognitive and cooperative networks–that is, what we call wireless grids–at the edge [18, 31, 32].

This article discusses the implications of this change for system and service design for distributed network applications including wireless grid applications. We identify some of the academic literatures that are likely to be increasingly relevant for adapting to these new challenges. In Section 35.2 of this article, we provide a stylized overview of the evolution of computing networks to wireless grids, to explain why the need to design for strategic behavior is becoming increasingly critical. We then briefly summarize the critical characteristics of wireless grids, as identified by our preliminary research on this issue [18, 32, 37]. Section 35.3 reviews the four principal mechanisms–technical, social, legal and economic–that are relevant for coordinating behavior in wireless grids and other distributed computing networks. Section 35.4 argues that these mechanisms evolve through the life-cycle of a technology and describes current trends in this evolution. The chapter concludes by sketching our

F.H.P. Fitzek and M.D. Katz (eds.), Cognitive Wireless Networks, 679–697.

future work on considering the implications of our analysis for the design of wireless grids.

35.2 From Systems Management to Grid Coordination

From a systems management perspective, a change is underway which is akin to the transition in the Soviet Union in the 1990s from a centrally-planned socialist system to a decentralized capitalist economy in Russia. Centralized command and control as modes of coordination are giving way to new mechanisms for allocating resources and moderating behavior [29]. Distributed ownership and decentralized control are raising new challenges for assuring system security and reliability. New network management mechanisms need to draw increasingly from the social, political/legal, and economic models of coordination used elsewhere in society. As with any significant change, there are both risks and new opportunities that must be better understood.

Traditional communication networks were designed on the basis of centralized, hierarchical control. In the 1960s, users connected to mainframe computers using dumb terminals. In such an environment, controlling and coordinating the behavior of edge-nodes was relatively simple and security protection could be handled largely by admission control. In the early days, computing resources were firmly under the control of a select cadre of IT professionals.

With the emergence of distributed processing and smart terminals in the 1970s, the problem of allocating resources and controlling the behavior of edge nodes became more complicated. However, most computing networks were still under the control of centralized network management supported by the power of management over employees.

In the 1980s, with the emergence of personal computing, Local Area Networks (LANs) and Wide Area Networks (WANs), computing and communications became increasingly integrated and distributed. A greater share of network intelligence was located in a continuously growing set of edge nodes. The heterogeneity of behavior that needed to be managed became even greater. Additionally, IT resources were increasingly under the direct control of end-users with much more diverse IT expertise. Corporate data managers now had to contend with non-IT specialists moving PCs among offices and loading or modifying application software in ways that were hard to monitor and manage. The resource allocation and coordination problem continued to grow more complex.

In the 1990s, the commercial emergence of the Internet expanded data communications and computing to a mass market, and increasingly provided a platform for interconnecting networks around the globe. The Internet's end-to-end architecture which facilitated peer communications among nodes stood in marked contrast to the traditional telecommunications networks which were based on hierarchical, centralized network management [28, 46, 51]. In the Internet, control is distributed to edge-nodes. However, the potential chaos that such a transition risked was moderated because key resources (e.g., DNS and routing infrastructure) were largely under the control of corporate data managers and carriers descended from the traditional telecommunications networks.

This technical architecture was mirrored by changes in industry structure and the policy environment. Traditional telecommunication and computing networks tended to be owned and managed end-to-end by a single entity (e.g., a carrier network or a corporate enterprise network). When these networks interconnected, these occurred at well-defined locations under bilateral (or multilateral) peering points. In the case of telecommunication carriers, the operation of these networks was also subject to substantial government regulation which constrained both the pricing and technical terms under which services were offered and interconnected.

Network management and ownership in the Internet, by contrast, is distributed among a global collection of heterogeneous end nodes, some of which are single computers or devices, while others are large networks in their own right. The diversity in ownership and computing/communication technology reflected in these edge-nodes raises the coordination problem to a new level of complexity. The Internet is also much more open than traditional network environments. The open, distributed nature of the Internet has facilitated the proliferation of computing in business and society, and contributed to dramatic growth in the ICT sectors and the global economy as a whole, but it has also raised problems for system designers.

Computer and network designers can no longer assume that systems will be owned and managed by a single entity with a single, coordinated set of goals. Increasingly, nodes are capable of self-interested behavior that can impact overall system performance in unpredictable and potentially adverse ways. The diversity of ownership in networking resources gives rise to diversity in strategic interests.

Coordinating behavior among nodes in a distributed network where all participants share common strategic interests is a difficult but well-defined problem for decision science. However, in an Internet-style environment, network management requires coordination among agents that are likely to have divergent capabilities and strategic interests. Resource allocation and control becomes a 'microeconomic' coordination problem. That is, whereas decision science provides a toolset for determining the optimal solution to single agent (common objective) problems, microeconomics provides a language/framework for studying the interactions of self-interested, strategically-independent agents. Its tools include the study of market behavior and game theory. Of course, many other academic disciplines also offer insights that are helpful in understanding how to design for strategic behavior, including computer science (parallel processing, ad hoc networks, and artificial intelligence), sociology and psychology, political science (including understanding interest group behavior and motivation), legal theory, and biology (especially evolutionary systems).

Moreover, computing/communication networks are becoming ever more important parts of our social (entertainment, cyber communities), economic (eCommerce), and political lives (eGovernment). In this environment, network design and management cannot be separated from the legal, political, social, and economic institutions governing human interactions in other spheres. Unsurprisingly, as computer networks become more central to our lives, the modes in which we regulate our lives in other spheres will become more relevant for how computer networks operate.

The openness of Internet-type networks allows businesses, their suppliers, and consumers to communicate and interact freely. The distributed and flexible architecture allows resources to be combined and used in novel ways, encouraging innovation and enhancing capabilities [9, 34]. However, this also increases the problem of protecting systems from myriad challenges ranging from viruses, denial of service attacks, intellectual property infringement (including protecting copyright in an era of resource sharing systems), and the abuse of privacy. Fraud and theft are also more common as criminals follow the money onto the information grid. Assuring system reliability and managing quality of service for diverse applications (delay-sensitive voice and file transfers on a shared network) is also more complicated when the identity, capabilities, and goals/incentives of end-nodes are not pre-configured and controllable.

The emergence of Peer-to-Peer (P2P) networks, such as Napster, Gnutella, Freenet and BitTorrent to name a few as well as computational networks such as Seti@Home and distributed.net reflects a reassertion of the end-to-end architectural model of the Internet and illustrates the importance of user behavior to system performance. In P2P networks the resources making up the network, storage space, routing and computational cycles are voluntarily provided by individual end-users with little or no institutional connections or trust. Shneidman and Parkes argue that, "perhaps the key defining characteristic of a peer to peer network is that one cannot distinguish between strategic nodes and the network infrastructure" [54]. Yet this risks overstatement as P2P networks are properly called overlay networks to emphasize that they run over the existing institutionally owned and managed infrastructure. This overlay nature gives leverage to attempts to centrally control peer to peer activities which we describe below as significant for the emergence of wireless grids.

The growth of wireless accelerates these trends because it increases opportunities for computing to become ubiquitous (always available, always connected), expands the heterogeneity of networking resources that need to be managed (mobility management and wireless/wireline interconnection), and the shear number of sensor network end nodes that need to be managed (connected computers in everything from our bodies to clothes, appliances, cars, and walls).

Wireless grids represent the epitome of this transition. In a wireless grid, even more than in overlay P2P networks, the edge nodes are the network. Designers and network managers of an ad hoc wireless grid will need to anticipate the strategic behavior of the end-nodes that will comprise the network. The challenge will be two fold: first, end-nodes will have to be induced to contribute resources to the network; and, second, to behave while part of the network in a way that helps maximize the total net benefits realized by the network. For example, in a wireless grid network, edge devices will likely need to be induced to contribute computing/communication resources to process traffic from other edge nodes, while at the same time refraining from behavior that deteriorates the service offered to other users (e.g., excessive use of shared resources) [6, 55].

Wireless grids are emerging from the coalescence of a number of independent research efforts and industry trends (see Figure 35.1). There are important developments associated in each of these areas that are critical to the evolution of wireless grids, but a concrete overall view is yet to emerge. Wireless grids will not be a computing network separate from the social/economic framework in which they operate. Continued multidisciplinary research is needed to properly design wireless grid networks[1].

We are engaged in a collaborative project to design infrastructure for wireless grids and to understand the virtual markets whose emergence we anticipate. This

[1] Several conference papers were our first efforts to explicate and define the critical features of wireless grids. These may be found on the website www.wirelessgrids.net. Note especially [3, 8, 18, 20, 32, 44, 57].

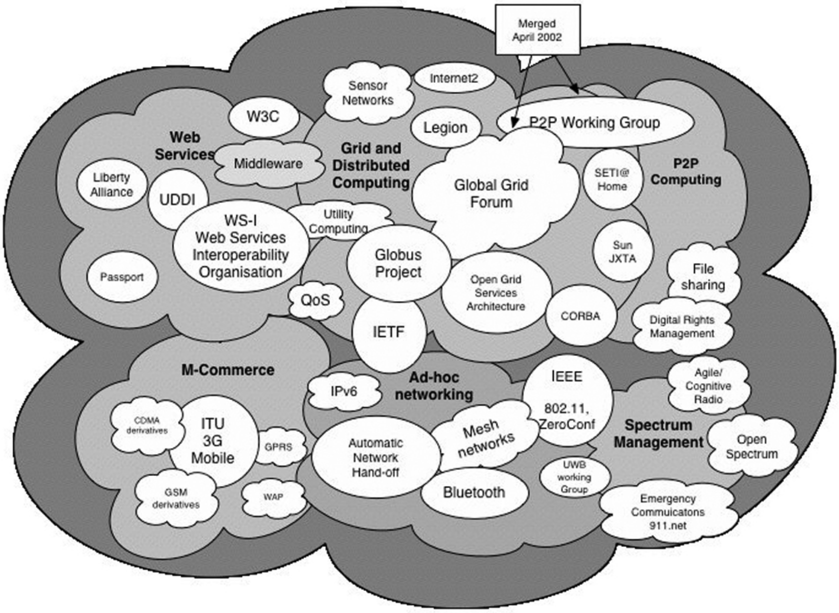

Figure 35.1. Research Issues and Industry Trends leading to wireless grids.

work will require the development of appropriate coordination mechanisms and will need to draw on and be compatible with coordination structures/institutions used elsewhere in networks and in society. In the next section we present a taxonomy of coordination frameworks, discuss their current realization on the Internet and their relevance for organizing behavior in distributed networks including wireless grids.

35.3 Coordinating Strategic Behavior in Distributed Networks

There are four prototypical ways in which to coordinate and allocate resources in distributed networks: (1) Technical; (2) Social (3) Legal and (4) Economic. Each of these is discussed further below, along with examples of their use in network system design and operation.

35.3.1 Technical

The traditional and still most common approach to network management is to use technical means to regulate behavior. Appropriate behavior can be 'hard-wired' into the network through hardware and software design. In biological systems, genetic coding may hard-wire in behavior and evolution can encourage and re-enforce behaviors that enhance a species prospects for survival.

In computer and communication networks, standards and communication protocols limit the range of allowed designs and behaviors that may be encountered, thereby rendering system performance more predictable and controllable[2]. Of course, designing suitable protocols that do not unduly constrain flexibility is quite difficult, especially in a distributed environment. The key is to define open interfaces that provide sufficient assurance as to the functionality that will be supported to allow interoperability without dictating detailed implementation rules that might limit innovation.

In many cases, the determination of which behavior is consistent with optimal network performance will depend on local and system-wide conditions that may be changing dynamically. Allowing nodes autonomy to moderate behavior in response to local needs and conditions can enhance overall performance, but this local autonomy than creates the potential for strategic manipulation [26].

Over time, network design has moved to a layered architecture with well-defined interfaces supporting communication across layers. The trend towards technical standardization based on open interfaces has resulted in a number of important developments for the industry. For example, open interfaces can allow end-users to 'mix and match' components (e.g., like when consumers mix-and-match stereo components or software applications on personal computers) to create customized systems. The open interfaces can also enhance industry competition by supporting both system-level and component-level competition.

Industry standardization can also give rise to positive network externalities that expand demand and scale/scope economies that can lower industry costs. Because the choice of where to define interfaces and what technologies to accommodate has such important implications for industry economics, industry standardization is inherently strategic. Getting the industry to agree on what standard to adopt is often quite difficult. The process can be contentious, expensive, and slow. Indeed, the process may be slow precisely because the standard development organizations have adopted bureaucratic rules in order to protect standardization from strategic manipulation.

Even after a standard has been defined, assuring compliance can be quite difficult. The standard which allows a lot of implementation flexibility may not assure adequate interoperability. When the networks are owned and managed by a relatively few number of players "as was more often the case in traditional telephone networks" enforcing interoperability was relatively easier. Adoption and implementation of the standard can be managed centrally. On the Internet the adoption and implementation of standards is focused on the IETF/IESG Request For Comment process[3] which, as discussed below, is implemented through a voluntary process supported by an informal reputation system. In a wireless grid network, the proliferation of edge-nodes under autonomous control makes technical coordination much more difficult.

One approach that has been used to manage interoperability in the distributed control environment that characterizes unlicensed spectrum is to require equipment certification. This ex ante testing is used to certify that equipment will comply with the communication protocols that have been adopted. In the case of unlicensed spectrum use, the principle concern is that a transmitter will comply with limits on radiated power. Most other details concerning how the transmitter will behave

[2] Lessig makes a related argument in [27].
[3] See RFC 2555 for a summary of the development of the RFCs and their process.

are left unspecified and it is left up to users of the spectrum to adopt suitable communication protocols and strategies for contending with congestion.

The certification approach facilitates distributed and asynchronous deployment of network equipment, but it limits flexibility and becomes less tractable as radio transmitters and receivers become more adaptive and software-controlled. There are a number of reasons for this. First, certifying the behavior of software is inherently more difficult than for hardware. Second, power modulation represents an important option for managing spectrum use efficiently, and a priori power limits are overly restrictive. Third, the certification approach may tilt the industry playing field in favor of incumbents (e.g., established equipment makers), potentially harming innovative approaches.

In the Internet, the TCP/IP protocols manage congestion via statistical back-off: when nodes experience congestion, they slow down their transmissions randomly. This works quite well when networks are lightly loaded and its simplicity makes it easy (low cost) to implement in a distributed network. Nodes only need local information to self-regulate their behavior. The downside of this approach is that it does not support quality-differentiated services which are important once the network starts having to contend with traffic that has heterogeneous requirements (e.g., delay-sensitive telephony and delay-insensitive email) and intrinsic values (e.g., network control messages and music downloads).

While technical approaches to coordinating behavior based on standards and communication protocols or network etiquette will remain important, they are unlikely to be sufficient by themselves. For example, it is possible to tweak TCP/IP parameters to capture an excessive share of network resources. This was not a significant problem in the early days of the Internet when it was a government-subsidized network used mostly by academics. With the Internet's growing social and commercial relevance, the control of quality of service has moved beyond purely technical approaches.

35.3.2 Social

The second common mechanism employed to regulate strategic behavior in networks operates through the social networks in which actors are embedded. Professional and cultural ties provide leverage by which network managers, and participants, can punish undesirable behavior and reward behaviors supportive of the goals of the system. Social mechanisms often support and provide the leverage to enforce the behaviors encoded in the technical protocols and standards discussed above.

The social mechanisms of greatest interest are those that operate in two ways: through cognitive factors such as conscience (or morality) and social influences, especially reputation in the context of group membership [11]. While reputation concerns the opinions of others about an actor, conscience concerns the opinions of an actor about their own actions. Both mechanisms act to regulate behavior however from a network application designers perspective reputation is the most useful mechanism.

The basic proposition of behavior regulation through reputation is that because people care about their reputation they will not act in ways that damage it and will act in ways that enhance it[4]. Yet reputation, as a strongly socially contextual concept, has varying mechanisms, impacts on behavior and scaling properties[5].

One of the best known and most studied reputation system is the “feedback” mechanism employed by Ebay which allows buyers and sellers to exchange public information regarding their satisfaction with the transaction [48, 49]. This system aims to regulate the potentially selfish behaviors such as fraud or a bait and switch tactics. This system has been credited with the growth and rapid acceptance of the Ebay system and the ability for Ebay to avoid the need to provide costly dispute resolution systems or guarantees. The system also seems to provide desirable outcomes for sellers: Resnick et al. in [50] concluded that sellers with high reputation earned approximately seven percent more than low reputation sellers.

Reputation is also employed as a tool for combating email Spam through blacklists of mailservers known (or believed to) send spam. The best known of these is the Realtime Blackhole List (RBL)[6]. While far from perfect[7], these systems have helped reduce spam from operators of open mail relays. Open mail relays are tempting to self-interested systems administrators because they offer convenience in configuration and for their intended users who do not have to deal with authentication or changing outgoing mailservers when moving between networks and IP addresses. However open relays provide conduits for the senders of spam into the Internet mail infrastructure, an activity which causes significant inconvenience to end users and consumes significant amounts of network bandwidth. Listing a mail server in a blacklist is a statement that the server has a bad reputation and means that servers which subscribe to that list will not accept mail from the legitimate users of the server and will 'bounce' the messages with a statement that the users mail server is suspected of spamming. The operator of an open relay is therefore encouraged to adopt more system-friendly behavior through a combination of technical (blocking) and reputation (reports made to users of the server and other systems administrators that are embarrassing to admin of an open relay).

Reputation has been employed also to coordinate behavior in P2P files sharing systems. Here the system designers goal is to increase the quantity and quality of content available on the network. Accordingly, Gnutella and Kazaa both provide mechanisms to prioritize the downloads of clients that have established a good reputation for providing uploads. While these mechanisms are currently quite basic they are developing rapidly, for example, BitTorrent, which provides swarming downloads by re-using clients currently downloading from a server as parallel servers for other clients employs a version of the Tit for Tat strategy developed in formal analyses of the Prisoners Dilemma game [10].

---

[4] In [39], Moreton and Twigg discuss the similarities between reputation systems and markets in which actors are motivated by money. Economic mechanisms are considered in Section 35.3.4 below.

[5] A useful taxonomy of types of reputation and their characteristics is provided by [40]. In early 2003 an NSF funded workshop was organized to support and develop this field. Resnick and Dellarocas' summary of the workshop provides an excellent introduction [47].

[6] `http://www.declude.com/junkmail/support/ip4r.htm` lists over 90 known blacklist services.

[7] At the time of writing spam fighting blacklists where under sustained denial of service attacks believed to be launched by spammers.

Scholars are seeking to formalize reputation systems to support the development of P2P and distributed computing applications. For example, [1] and [24] propose to utilize a distributed data structure to store complaints about nodes in a P2P network and in an electronic market.

These reputation mechanisms share aspects of community-public goods. Because providing reputation reports is costly in time and resources, participants may have a natural preference for free-ride on the information provided by others. Thus a major challenge in building systems to formalize and extend social regulation through reputation is addressing the issues of incentive compatibility that arise. This challenge is addressed in [23] and [38].

Reputation has been utilized in promoting desired end-node behavior in distributed computing projects which are pre-cursors of Grid applications. The Seti@Home project leverages both reputation and conscience by compiling and making available statistics on the number of units that users have processed. This information is made prominent on the local client, leveraging conscience, and through league tables, periodically released on the Seti@Home site and lists which leverage reputation. Furthermore when interesting results are discovered the user or team who undertook the processing is highlighted despite the random distribution of work units. It is not clear what use this type of reputation is to the actors but the emergence of highly competitive teams (containing thousands of members) aiming to process the largest number of work units suggests that it is an effective motivator of desirable behaviors[8]. However, this motivator is far from unproblematic–cheating through altered software has been discovered within the Seti@Home system [54].

Social mechanisms rely on the strength of social ties or group identification to regulate behavior in networks. This mechanism is clearly limited by the growth and expansion of actors interconnecting through networks, which, by the sheer increase in numbers, reduces the effectiveness of both informal reputation systems and morality derived from group membership. In addition the rising financial rewards available through network misbehavior, such as Spamming, motivate actors to compare these rewards to the often less quantifiable reputation rewards. Nevertheless social regulation remains an important mechanism particularly in situations characterized by high levels of repeat interactions.

35.3.3 Legal

Legal and political systems are designed to regulate and enforce a wide variety of behavioral prescriptions and prohibitions in the interests of promoting the well-being of the broader community. Roman law (unitary law) and common law (Anglo-Saxon) legal traditions share many elements of commonality, while differing in their approaches to legal change and adaptation [27]. For the United States, the Communications Act of 1934 (as most recently amended in the U.S. Telecommunications Act of 1996) defines the legal framework for media and telecommunications systems and services. The Act includes detailed specification and regulatory guidelines for interconnection of networks.

Behaviors affecting the use of radio spectrum have been addressed primarily through legal means, including provisions for licensed and unlicensed (Part 15)

[8] A sample league table can be seen here `http://www.muskratgroup.com/kwsn/teams.html`

frequency use. While the recognition and treatment of property rights is fundamental to capitalist economies, including the unlimited right to exclude, this centralized legal framework for spectrum management has been undermined by the development of new spectrum sharing technologies. In other work we have critiqued the lack of foresight exhibited by the legislators who enacted that law [41]. Subsequent events, including the emergence of wireless grids, prove our point. Ubiquitous wireless grid environments will pose challenges to many areas of law and law enforcement as diverse heterogeneous market, policy, and user requirements must be simultaneously resolved in a shared resource environment.

The Digital Millennium Copyright Act has been employed to limit the behavior of network users. Section 1201 of this act was the basis for US vs ElcomSoft and Sklyarov in which a Russian programmer, Dimitry Sklyarov, was arrested on a visit to the US for providing a circumvention device able to remove the encryption from Adobe PDF files. ElcomSoft, his employer, also faced charges. The DCMA also forms the legal basis for actions designed to obtain evidence of copyright infringement from ISPs. In 2003 the RIAA obtained subpoenas against, amongst others, Verizons ISP, who was required to release the details of subscribers accused of sharing copyrighted music on P2P services. These subpoenas made possible the contributory copyright infringement suits made against over 260 individuals in 2003. The RIAAs stated strategy is to utilize the threat of such lawsuits to reduce the, from their perspective, undesirable behavior of users providing resources to P2P music file sharing networks[9].

Contract law has also been employed to regulate behavior online. The Terms and Conditions required of ISP customers usually contain acceptable use provisions which restrict activities considered to be undesirable by the network designers, such as running servers on home access accounts. These contracts also facilitate ISPs cooperating with law enforcement officials or legal subpoenas for evidence.

35.3.4 Economic

The market's Invisible Hand provides another potent mechanism for coordinating behavior in distributed systems. Competitive markets, when they are operable, provide an efficient mechanism for allocating resources that do not presume any common interest among resource producers or consumers. Buyers and sellers, each seeking to maximize their individual welfare, will compete for scarce resources. Excess demand for resources drives market prices up, inducing consumers who value the resource the less than the current price to leave the market and inducing suppliers who can produce at lower cost to increase supply. Excess supply has the opposite effect. In the idealized competitive market, the atomistic buyers and sellers each act independently, ignoring their impact on the market price, yet collectively their distributed behavior drives the market to equilibrium. In the efficient equilibrium, supply and demand are balanced, resources are produced at the lowest possible cost, and allocated to the highest-value demand.

Unfortunately, the ideal of perfect competition is seldom realized in the real world; and even the ideal economic model is somewhat sketchy with respect to the dynamics of how a market approaches equilibrium [25]. Indeed, real world markets depend critically on the social, legal, and technical environment that shapes the way

[9] See `http://www.riaa.com/news/newsletter/090803.asp`

in which actors exchange information, negotiate for the exchange of goods, complete their transaction, and in the event of disputes, reconcile any difficulties.

The development of functional economic systems for computer networks has been studied for over thirty years. One of the key developments in this field is the recognition that the systems cannot simply provide efficient allocations of resources (as per [56, 59]) but must manage incentives and strategies of the participants (as per [12, 13, 53, 58]). Unfortunately this second step has proven to be difficult in both theory and implementation. This is clear from the market controlled approaches to resource allocation in Grids, summarized in [19], which fail to adequately address strategic issues.

Buyya et al. in [7] demonstrate both the usefulness of an economic approach to resource allocation within Grid computing environments and the difficulties faced. They developed and implemented a market-making scheme involving the interaction of consumer and producer agents undertaking a wide range of economic interaction models, including auctions and announced prices. This market-making scheme was able to demonstrate efficiency in the allocation of resources on the Grid.

However the model suffers from two key difficulties that will serve to illustrate the complexity of difficulties of implementing computational market system: bootstrapping and incentive management. The system suffers from a bootstrapping problem: Grid services provide the underlying services for a market designed to motivate the provision of Grid Services, The Grid computing environments provide necessary infrastructure including security, information, transparent access to remote resources, and information services that enable us to bring these two entities together [7, p. 2]. Without these basic requirements markets do not function effectively. Real world markets are embedded in social relations, not the least of which is the system of contract law and the enforcement mechanisms that support it.

Buyya et al.s approach, grounded as it is within traditional resource allocation literature, does not adequately address the strategic challenges of networked computing. Buyya et al. acknowledge this when they present their function for Resource Value,

> Resource Value = Function (Resource strength, Cost of physical resources, Service overhead, Demand, Value perceived by the user, Preferences)

And state,

> The last three are difficult to capture from consumers unless they see any benefit in disclosing their actual demand, preference, and/or resource value, which varies from one application to another. [7, p. 4]

If these parameters have to be truthfully disclosed to reach the desired resource allocation efficiencies then the system must be designed in such a way that it is to the agents benefit to reveal such private information. Otherwise the system is open to systematic under- or over-statement of private valuations and will not achieve the desired (and expected) efficiencies. This is the heart of Shneidman and Parkes recent criticism of the literature on economic analyses of Grid computing where they argue that recent papers on economic models for resource scheduling in scientific Grid computing have not explored issues of rationality [54, p. 6] (referencing [7]).

An important challenge for designers of wireless grid technology will be to design for virtual markets. The mechanisms for determining who participates in these markets, how information is exchanged, how participants negotiate for the exchange of resources, payment/compensation mechanisms, and monitoring/enforcement structures will all be critical elements that must be developed. These mechanisms must be incentive compatible. That is participants have to trust these mechanisms to behave as expected and in such a way that induces them to participate and elicits cooperative behavior that is also self-interested and selfish.

For prices to emerge and markets to function appropriately, it must be possible to define common resources using a collective and public language that can allow resources to be "commoditized". Participants have to know what they are negotiating for when they decide to purchase or supply a unit of computing or communication power. Figuring out what are the right ways to describe and quantify commodities and the terms and time limits for purchase/supply contracts will represent a difficult challenge. Eventually, we will need service level agreements for wireless grids [33].

There are a diverse range of market mechanisms in use. These range from free exchange (e.g., subsidized) to barter systems (exchange of goods without money) to the arm's length exchange (exchange for money with limited prior contact or on-going contractual relationship) to bilateral or multilateral exchange. All of these have been used in various contexts within modern communication networks.

For example, WiFi free nets and the enterprise networks provided to corporate employees or university students are often subsidized. Although they obviously cost the provider, the consumer does not directly pay for access to the resources. Network peering may be considered a form of barter exchange in which interconnecting carriers agree to exchange traffic at no charge. In the Internet, the lack of a more developed economic system has plagued multi-lateral 'free' peering with consistent congestion problems, leading most backbone carriers to move to bilateral peering.

Telephone service markets offer numerous well-developed versions of more advanced economic market systems. Traditionally, these were regulated as common carriers, which protected atomistic consumers from being discriminated against. Atomistic residential and small business consumers purchase service without term commitments according to regulated tariffs. The more competitive markets such as long distance services and cellular services are less heavily regulated. Consumers churn among alternative providers in response to more attractive price/quality offerings. The competition for consumers forces carriers to lower costs and enhance quality. Advertising and marketing help suppliers and consumers learn about available options.

Wireless grids are likely to make use of all of these market models as they develop. In anticipation, it would be useful to consider how to design for flexible market models that do not presume a particular industry structure or mode of exchange. A key element will be to design for market interfaces. These are most likely to occur via open interfaces that can be standardized so that the requisite information may be exchanged among parties that may be exchanging resources at arm's length. If the parties have an on-going relationship and shared common interests, than the market exchange interface may be quite simple (e.g., exchange within a single firm). Alternatively, if the relevant commodity can be provided in a market situation that approaches the competitive ideal, than again the market interface may be quite simple–the Invisible Hand of the market can supply coordination. More typically, the transaction will involve agents with potentially conflicting, self-interested goals and the designer will need to consider the game-theoretic aspects of exchange (e.g., asymmetric and incomplete information, reward/penalty structures, sequencing of actions, player strategy spaces, etc.).

Understanding the need for incentive-compatible optimal design is much easier than explicating how this might be accomplished without complexity that will hinder the adoption of the mechanisms[10]. No single economic approach will be ideal for all circumstances. The appropriate economic design will depend critically on the other elements of the environment: the technical, social, and legal context in which participants will interact.

35.4 Interactions and Dynamics in Regulation

The process of change in the computing environment described in Section 35.3 gives rise to a particular dynamic among the four models described above.

Early in a technologies life cycle, technical and social coordination mechanisms are most useful and were clearly emphasized in the development of the Internet. They allow for the greatest level of innovation and utilize the familiarity and shared intentions of the development community as a trusted base to support this innovation. However, a natural byproduct of technology becoming more mainstream is that the range of parties that are interacting become less familiar to each other–there are less repeat interactions, less common expertise/knowledge/experience to induce conformation–so self-interested and potentially harmful behavior increases. The Internet is reaching this second phase. Simultaneously, the stakes of non-cooperation have risen sharply as businesses rely on Internet services to invest and risk real money.

This section briefly describes two currently developing responses to this situation. The first is a move to 'harden' technical regulation and to substitute law for the social regulation that had supported technical regulation. The second is an expansion in legal provisions relating to behavior online and the development of surveillance systems to support their operation.

35.4.1 Hardening Technical Regulation with Legal Enforcement

There are a number of current proposals that would strengthen technical regulation through both hardware and software initiatives and through legal means to mandate their use. This reflects a loss of confidence with current voluntary technical regulation.

"Trusted Computing" has been proposed as a solution to computer insecurity and viruses and the use of computers for copyright infringement–all areas of 'misbehavior' online. The Trusted Computing Group, an industry body lead by Intel and Microsoft, propose designing systems which are only able to run code which has been verified through a digital signature. The system would be incapable of running non-signed or altered code and network applications would be able to ascertain that their peers where running particular versions whose behavior could be relied upon [2, 52].

Trusted Computing would thereby create a technical mechanism would could be used to protect against the execution of virus code as well as to prevent the

[10] There is significant work underway in the area of Distributed Algorithmic Mechanism Design. See [16, 42, 45, 60]. However [43] reminds us of risks in decision marking complexity in online markets.

execution of infringing digital media behavior. This proposal sidesteps issues of providing incentives for desired behavior by recreating the institutional fence whose breakdown we examined in Section 35.2, by providing the technical hook for external control over the uses of computing devices.

In [52] Schecter et al. consider the ability of Trusted Computing to control end node behavior. They introduce this through an ironic demonstration that this capability could be used by P2P music sharing network designers to protect themselves from the attempts of content owners to disrupt the networks. Injecting corrupted content and flooding networks are tactics which have been adopted by the content industries and are, from the network designers point of view, undesirable and detrimental to system performance. Trusted Computing platforms would allow network clients to ascertain that a peer is running application code without these detrimental behaviors and to exclude misbehaving clients from the network. In [52] example clients are able to exclude clients designed to reduce network throughput by flooding bandwidth with extraneous traffic.

It is clear, then, that Trusted Computing would merely provide a technical hook for end node control but that market and legal provisions will determine how that hook is used.

A similar development can be observed in the TV broadcast industry where the digitization of content is viewed as creating opportunities for violations of copyright that would threaten the viability of the conversion to DTV. The Broadcast Protection Discussion Group (BPDG), an industry body charged with preventing this self-interested behavior has been proposed that there be a 'broadcast flag' attached to 'protected' content which would indicate that that content may not be used in certain ways, and that compliant devices be designed to respect this flag. This proposal has been incorporated into the ATSC standards as an optional part[11], however in August 2002 the FCC issued a Notice of Proposed Rulemaking[12] in which it stated that that it was inviting discussion on the question, "Should the FCC mandate that consumer electronics devices recognize and give effect to the broadcast flag?".

Legislative proposals such as The Consumer Broadband and Digital Television Promotion Act (CBDPTA), proposed in 2002 by Sen. Ernest Hollings, would mandate the use of copy protection scheme in any device that can "retrieve or access copyrighted works in digital form" and it has been suggested that this implies the legislation of Trusted Computing. These bills have so far not received broad support but reflect the trend of providing legal backing to the use of hardened technical standards for the regulation of online behavior. These are an attempt to return to the systems management paradigm discussed in Section 35.2.

35.4.2 New Legal Provisions and Their Surveillance Implications

The second response to declining trust, rising misbehavior and increasing stakes are efforts to utilize the civil and criminal justices systems to enforce desired behavior online. It is the nature of justice proceedings that they occur after infringing behavior in question and that admissible evidence of infringing behavior must be brought before a court. For this reason legislative proposals typically imply an increase in surveillance of online behavior.

[11] ATSC Standard A65/A.
[12] FCC Digital Broadcast Copy Protection MB Docket No. 02-230.

The recent high profile investigation, arrest and pending prosecution of the juvenile writers of the SoBig virus reflect the increasing use of law enforcement to combat online misbehavior that threatens the stability and performance of computer networks. These actions are quite distinct from the enforcement of laws relating to pornography or fraud because the social evil targeted is a decline in system performance which had hitherto been considered purely a technical matter. Many US States and European countries have, or are in the process of passing laws against unsolicited email, known as Spam which contain steep financial penalties. The DCMAs copyright provisions, as discussed above, are increasingly being used to target online behavior. Common to all these laws is the need to collect admissible evidence of infringing behavior online.

The Internet community has struggled with calls for lawful interception of internet traffic for the purpose of evidence collection [5]. In May 2000, after an internal debate, the IETF issued RFC 2804, IETF Policy on Wiretapping in which it writes, "The Internet Engineering Task Force (IETF) has been asked to take a position on the inclusion into IETF standards-track documents of functionality designed to facilitate wiretapping. This memo explains what the IETF thinks the question means, why its answer is "no", and what that answer means." [22, p. 1]. Yet despite this policy the issue has not subsided. Cisco Systems has made available, as an optional router software feature that must be specifically requested, the capacity to give access to data flowing through routers in a form specifically designed to be legally admissible [30]. Their initiative to publish this capability as an Internet Draft [4] indicates that this debate is far from closed and that the pressure to collect evidence to support Legal regulations of online behavior remains strong.

35.5 Conclusion and Implications for Wireless Grids

Despite the continuously increasing interests in wireless grid[13], research on wireless grids is scattered. This chapter has examined the evolution of computing from systems management within known institutional contexts to the decentralized and end-user centric model, i.e., from a user perspective [36]. We highlighted the increasing importance of self-interested strategic behavior and the need for network application designers to be able to promote desired behaviors while discouraging undesired behaviors. We argued that wireless grids are the epitome of these developments. We identified and examined four mechanisms for the regulation of strategic behavior–technical, social, legal and economic–and examined dynamics within them.

It is clear that designers of wireless grids and their applications will need to draw on all of these mechanisms. We expect a similar dynamic in the development of wireless grids to that which has occurred through the technology life cycle of existing network technologies. It is, therefore, sensible that the initial focus be on traditional technical and informal reputation systems which are better placed to promote innovation and experimentation. It is similarly sensible to design such that the evolution towards economically managed situations is easier and so that this evolution, and future innovation, is not foreclosed by premature hardened technical

[13] For instance, IEEE Internet Computing launched a special issue on wireless grids in 2004 guest edited by Lee McKnight, where researchers studied wireless grids via various perspectives [17,21,35].

and legal proscription. The best manner in which to accomplish this is a topic for our further research.

Acknowledgement. This article is based upon research supported in part by the National Science Foundation, Partners for Innovation program grant #0227879, involving faculty and students at MIT, Boston, Northeastern, Syracuse, and Tufts Universities, the Swiss Federal Institute of Technology (ETH Zurich), University of Colorado, Boulder, BTexact, Cisco Systems, Hitachi, Nokia, Novell, and other firms, in partnership with TeleCom City, TRITEC, the Museum of Science Boston, and the School of Information Studies and CASE Center at Syracuse University. See `www.wirelessgrids.net` for more information. The opinions expressed in this article are those of the authors. Any errors of fact or by omission are the responsibility of the authors and not the institutions with which they may be affiliated.

References

1. K. Aberer and Z. Despotovic. Managing trust in a peer-2-peer information system. In *the Tenth International Conference on Information and Knowledge Management*, 2001.
2. R. Anderson. Trusted computing FAQ. `http://www.cl.cam.ac.uk/users/rja14/tcpa-faq.html`, 2003.
3. D. Anius. Enhancing innovation within a regional wireless grid. In *The International Conference on Computer, Communication and Control Technologies*, Orlando, Florida, 2003.
4. F. Baker, B. Foster, and C. Sharp. Cisco architecture for lawful intercept in IP networks, Internet Draft (Informational). Technical report, IETF, 2003.
5. P. A. Branch. Lawful interception of the internet. Technical report, Centre for Advanced Internet Architectures, 2003.
6. L. Buttyán and J.-P. Hubaux. Stimulating cooperation in self-organizing mobile ad hoc networks. *Mobile Networks and Applications*, 83(5):589–592, 2003.
7. R. Buyya, H. Stockinger, J. Giddy, and D. Abramson. Economic models for management of resources in peer-to-peer and grid computing. In *the SPIE International Conference on Commercial Applications for High-Performance Computing*, Denver, USA, 2001.
8. B. Chen and C. H. Chang. Mobility impact on energy conservation of ad hoc routing protocols. In *International Conference on Advances in Infrastructure for Electronic Business, Education, Science, Medicine, and Mobile Technologies on the Internet*, L'Aquila, Italy, 2003.
9. D. Clark. The design philosophy of the darpa internet protocols. In *SIGCOMM '88: Symposium proceedings on Communications architectures and protocols*, pages 106–114, New York, NY, USA, 1988. ACM Press.
10. B. Cohen. Incentives build robustness in bittorrent. In *Workshop on Economics of Peer-to-Peer Systems*, 2003.
11. R. Conte and C. Castelfranchi. *Cognitive and Social Action*. UCL Press, 1995.
12. J. Crowcroft, R. Gibbens, F. Kelly, and S. Ostring. Modelling incentives for collaboration in mobile ad hoc networks. In *WiOpt '03*, 2003.

13. E. Drexler and M. Miller. *Ecology of Computation*, chapter Incentive engineering for computational resource management. Elsevier Science Publishers B. V., 1988.
14. Foster, C. Kesselman, J. Nick, and S. Tuecke. The physiology of the grid: An open grid services architecture for distributed systems integration. `citeseer.nj.nec.com/foster02physiology.html`, 2002.
15. Ian Foster, Carl Kesselman, and Steven Tuecke. The anatomy of the grid: Enabling scalable virtual organizations. *Int. J. High Perform. Comput. Appl.*, 15(3):200–222, 2001.
16. E. Friedman and D. Parkes. Pricing wifi at starbucks - issues in online mechanism design. In *Fourth ACM Conf. on Elec. Commerce*, 2003.
17. M. Gaynor, S.L. Moulton, M. Welsh, E. LaCombe, A. Rowan, and J Wynne. Integrating wireless sensor networks with the grid. *IEEE Internet Computing*, 8(4):32–39, 2004.
18. M. Gaynor, L. W.McKnight, J. Hwang, and J. Freedman. Wireless grid networks and virtual markets. In *the International Conference on Computer, Communication and Control Technologies*, 2003.
19. J. Gomoluch and M. Schroeder. Market-based resource allocation for grid computing: A model and simulation. In *1st International Workshop on Middleware for Grid Computing*, Rio de Janeiro, Brazil, 2003.
20. J. Hwang and P. Aravamudham. Proxy-based middleware services for peer-to-peer computing in virtually clustered wireless grid networks. In *the International Conference on Computer, Communication and Control Technologies*, Orlando, Florida, 2003.
21. Junseok Hwang and P. Aravamudham. Middleware services for P2P computing in wireless grid networks. *IEEE Internet Computing*, 8(4):40–46, 2004.
22. IESG IAB. IETF policy on wiretapping. Technical Report RFC 2804, IETF, 2000.
23. R. Jurca and B. Faltings. An incentive compatible reputation mechanism. In *the IEEE Conference on E-Commerce*, Newport Beach, CA, USA, 2003.
24. S. D. Kamvar, M. T. Schlosser, and H. Garcia-Molina. The eigentrust algorithm for reputation management in p2p networks. In *the Twelfth International World Wide Web Conference*, 2003.
25. T. F. Kompas. Traditional notions of equilibrium reconsidered. *Equilibrium*, 11:484–495, 2000.
26. W. Lehr. Standardization: Understanding the process. *Journal of the American Society for Information Science*, 43(8):550–555, 1992.
27. L. Lessig. *Code and Other Laws of Cyberspace*. Basic Books, 1999.
28. J. MacKie-Mason, S. Shenker, and H. Varian. Network architecture and content provision: An economic analysis. *Telecommunications Policy*, 20(3):203–217, 1996.
29. T. W. Malone, J. Yates, and R. I. Benjamin. Electronic markets and electronic hierarchies. *Communications of the ACM*, 30(6):484–497, 1987.
30. D. McCullagh. Inside Cisco's eavesdropping apparatus. `http://news.com.com/2010-1071-997528.html`, 2003.
31. L. W. McKnight, D. Anius, and O. Uzuner. Virtual markets in wireless grids: Peering policy obstacles. In *Proceedings of the International Conference on Computer, Communication and Control Technologies*, 2003.

32. L. W. McKnight and J. Howison. Towards a sharing protocol for wireless grids. In *Proceedings of the International Conference on Computer, Communication and Control Technologies*, 2003.
33. L. W. McKnight and W. Lehr. Show me the money: Agents and contracts in service level agreement markets. *INFO*, 40(2):24–36, 2002.
34. L. W. McKnight, P. M. Vaaler, and R. L. Katz. *Creative destruction: business survival strategies in the global Internet economy.* MIT Press, 2001.
35. Lee McKnight. Wireless Grids–Distributed Resource Sharing by Mobile, Nomadic, and Fixed Devices. *IEEE Internet Computing*, 8(4):2–10, 2004.
36. L.W. McKnight, R. Sharif, and L. van de Wijngaert. *Wireless Grids:Assessing a New Technology from a User Perspective*, pages 169–182. Springer, 2005.
37. W. McKnight, D. Anius, and O. Uzuner. Internet Services. Quality of Service in Grids, Networks and Markets. In press.
38. N. Miller, P. Resnick, and R. Zeckhauser. Eliciting honest feedback in electronic markets. In *Working Paper for the SITE02 workshop*, 2003.
39. T. Moreton and A. Twigg. Trading in trust, tokens and stamps. In *1st Workshop on the Economics of P2P systems*, Cambridge University, 2003.
40. L. Mui, M. Mohtashemi, and A. Halberstadt. Notions of reputation in multi-agents systems: a review. In *the first international joint conference on Autonomous agents and multiagent systems*, 2002.
41. W. R. Neuman, L. W. McKnight, and R. J. Solomon. *The Gordian Knot: Political Gridlock on the Information Highway.* MIT Press, 1997.
42. C. Ng, D. Parkes, and M. Seltzer. Strategyproof computing: Systems infrastructures for self-interested parties. In *1st Workshop on the Economics of P2P systems*, 2003.
43. Odlyzko. Internet pricing and the history of communications. *Computer Networks*, 36(5-6):493–517, 2001.
44. J. S. Park, A. Nanda, and J. Howison. Security challenges and countermeasures in wlans. In *the International Conference on Computer, Communication and Control Technologies*, Orlando, Florida, 2003.
45. D. C. Parkes. Iterative combinatorial auctions: Achieving economic and computational efficiency. Technical report, University of Pennsylvania, 2001.
46. D. Reed, J. Saltzer, and D. Clark. Active networking and end-to-end arguments. *IEEE Network*, 12(3):67–71, 1998.
47. P. Resnick and C. Dellarocas. Online Reputation Mechanisms: A Roadmap for Future Research. Technical report, First Interdisciplinary Symposium on Online Reputation Mechanisms, 2006.
48. P. Resnick and R. Zeckhauser. Trust among strangers in internet transactions: Empirical analysis of ebay's reputation system. *Advances in Applied Microeconomics, Elsevier Science*, 11, 2003.
49. P. Resnick, R. Zeckhauser, E. Friedman, and K. Kuwabara. Reputation systems: Facilitating trust in internet interactions. *Communications of the ACM*, 43(12):45–48, 2000.
50. P. Resnick, R. Zeckhauser, J. Swanson, and K. Lockwood. The value of reputation on ebay: A controlled experiment. Technical report, School of Information, University of Michigan, 2002.
51. J. H. Saltzer, D. P. Reed, and D. D. Clark. End-to-end arguments in system design. *ACM Transactions on Computer Systems*, 2(4):277–288, 1984.

52. S. E. Schecter, R. A. Greenstadt, and M. D. Smith. Trusted computing, peer-to-peer distribution, and the economics of pirated entertainment. In *The Second Workshop on Economics and Information Security*, 2003.
53. S. Shenker, D. Clark, D. Estrin, and S. Herzog. Pricing in computer networks: Reshaping the research agenda. *ACM Computer Communication Review*, 26(6):19–43, 1996.
54. J. Shneidman and D. Parkes. Rationality and self-interest in peer to peer networks. In *Proceedings of the 2nd International Workshop on Peer-to-Peer Systems*, 2003.
55. B. Strulo, J. Farr, and A. Smith. Securing mobile ad hoc networks – a motivational approach. *BT Technology Journal*, 21(3):81–89, 2003.
56. Sutherland. A futures market in computer time. *Communications of the ACM*, 11(6):449–451, 1968.
57. O. Uzuner and R. Davis. Digital fingerprinting for distribution volume tracking:intellectual property protection in wireless grids. In *the International Conference on Computer, Communication and Control Technologies*, Orlando, Florida, 2003.
58. H. R. Varian. Economic mechanism design for computerized agents. In *Workshop on Electronic Commerce*, 2000.
59. C. A. Waldspurger, T. Hogg, B. A. Huberman, J. O. Kephart, and W. S. Stornetta. Spawn: A distributed computational economy. *Software Engineering*, 18(2):103–117, 1992.
60. C. J. Woodard and D. C. Parkes. Strategyproof mechanisms for ad hoc network formation. In *1st Workshop on the Economics of P2P systems*, 2003.