
P2P and the Promise of Internet EQUALITY

✂ BY PHILIP E. AGRE

Technologies often come wrapped in stories about politics. These stories may not explain the motives of the technologists, but they do often explain the social energy that propels the technology into the larger world. In the case of P2P technologies, the official engineering story is that computational effort should be distributed to reflect the structure of the problem. But the engineering story does not explain the strong feelings P2P computing often evokes. The strong feelings derive from a political story, often heatedly disavowed by technologists but widespread in the culture: P2P delivers on the Internet's promise of decentralization. By minimizing the role of centralized computing elements, the story goes, P2P

systems will be immune to censorship, monopoly, regulation, and other exercises of centralized authority.

This juxtaposition of engineering and politics is common enough, and for an obvious reason: engineered artifacts such as the Internet are embedded in society in complicated ways. I propose to use the case of P2P computing to analyze the relationship between engineering and politics—or, as I want to say, between architectures and institutions. By “architecture” I mean the matrix of concepts designed into a technology. Examples include the concepts that underlie the von Neumann serial processor, the distinction

between clients and servers, and entity-relationship data models. By “institution” I mean the matrix of concepts that organizes language, rules, job titles, and other social categories in a given sector of society. For example, in the institutions of medicine, these concepts include “patient,” “case,” and “disease.” Architectures and institutions often are related, and systems analysts work largely by translating institutional concepts into system architectures.

It can be difficult to distinguish P2P computing from the Internet in general [7]. For example, is email an example of P2P? Fortunately, there is no need to pin down a definition.

Exploring the tension between the engineering story of rationally distributed computation and the political story of institutional change through decentralized architecture.

Instead, I want to explore the tension between the engineering story of rationally distributed computation and the political story of institutional change through decentralized architecture. If the world were simple, then these stories would coincide. In reality, the relationship between architectures and institutions is exceedingly complicated. I will briefly present four theories of this relationship, each associated with a 20th century theorist of institutions.

Veblen

Thorstein Veblen, a Norwegian-American economist and social critic, wrote during the Progressive Era, when tremendous numbers of professional societies were being founded, and he foresaw a society organized rationally by professionals rather than through the speculative chaos of the market. Veblen was impressed by a profession's ability to pool knowledge among its members, and he emphasized the collective learning process through which industry grows (for more on Veblen's theory, see [6]).

Some historical background will be useful. Although some professions grew slowly from medieval guilds, a great explosion of new professions around 1900 was facilitated by new communications and transportation infrastructures. These new infrastructures created economies of scale in the production and distribution of industrial goods, thus permitting the rise of large corporations and the division of intellectual labor that made professional specialization possible. They also supported professionals in organizing societies, editing and distributing publications, traveling to conferences, and so on. In fact, the first modern professions were organized by railroad workers, who pioneered the use of the new infrastructures [2]. Soon afterward, numerous other professional groups followed.

A profession combines elements of centralization and decentralization. For example, its members work for many different organizations; they have little formal authority over one another. The profession exists largely to support decentralized processes of networking and collective learning. But the profession's publications and conferences are administered centrally. Infrastructures also combine elements of centralization and decentralization. The telegraph and railroad industries moved (both for economic and political reasons) from fragmentation to oligopoly. Yet the functionality these industries provided—carrying things from point A to point B in a standardized way—supported decentralized social processes.

Because Veblen advocated neither capitalism nor socialism, he is usually regarded as an outlier in political history. Yet the Internet is making Veblen's theory more relevant than ever. Because 20th century infrastructures

were so cumbersome, only a small portion of the population could use them to organize professions, and the organizations that resulted were bureaucratic. However, numerous social groups now use the Internet to hold discussions, edit publications, organize meetings, and build social networks. As new service layers are added to the Internet, a complex array of architectures and institutions will become available to everyone. Some of these may resemble the professions of Veblen's day, but they might also support collective learning in other ways.

Hayek

Friedrich Hayek was an Austrian economist who provided intellectual ammunition for the fight against communism. His most famous argument is that no centralized authority could possibly synthesize all of the knowledge participants in a complex market use [5]. But Hayek was not an anarchist. He argued a market society requires an institutional substrate that upholds principles such as the rule of law. A productive tension is evident in Hayek's work; he is attracted to notions of self-organization, but he is also aware that self-organization presupposes institutions generally and government institutions in particular.

Hayek's work, like Veblen's, challenges us to understand the definition of "center." Observe, for example, that institutions, like architectures, are often organized in layers. Legislatures and courts are institutional layers that create other institutions, namely laws, that rest upon them. Law itself has layers; contract law is a layer, and so are individual contracts. The architecture of the Internet is also organized in layers. Do the more basic layers of institution and architecture count as centers? Yes, if they must be administered by a centralized authority. Yes, if global coordination is required to change them. No, if they arise in a locality and propagate throughout the population. At least sometimes, then, centralization on one layer is a precondition for decentralization on the layers above it. Complex market systems, for example, need their underlying infrastructures and institutions—the legal system, the stock exchange, eBay, containerized shipping—to be coordinated and standardized. Yet this kind of uniformity has often been imposed by governments and monopolies. Therefore, the conditions under which decentralized systems can emerge are complicated. The question is how the dangers of centralization can be minimized.

Consider the case of the Internet. Despite its reputation as a model of decentralization, its institutions and architecture nonetheless have many centralized aspects, including the DNS and Microsoft's control over desktop software. The Internet Engineering Task Force is less centralized than most standards organizations, but it is still a center. Let us consider one aspect of the Inter-

net: the end-to-end arguments [11], which move complexity out of the network and into the hosts that use it. This approach maximizes flexibility: each layer in the Internet protocol stack is defined in a general way, and end users can create new layers atop the old ones. But it also shifts complexity away from the centralized expertise of network engineers, placing it instead on the desktops of the people who are least able to manage it. Much of the Internet's history, consequently, has consisted of attempts to reshuffle this complexity, moving it away from end users and into service providers, Web servers, network administrators, authentication and filtering mechanisms, firewalls, and so on. The P2P character of the TCP/IP protocols has remained much the same; the reshuffling takes place mostly on other layers. Thus, a decentralized network can support centralized services, and vice versa. For example, the asymmetrical client/server architecture of the Web sits atop the symmetrical architecture of the Internet.

In moving toward a decentralized ideal, the P2P movement must confront the types of centralization that are inherent in certain applications (see [10]). For example, if users contend for the same physical resource, some kind of global lock is needed. Most markets have this property [12]. Some mechanisms do exist for sharing a resource without an architecturally centralized lock, such as the backoff algorithms of Ethernet and TCP. Complete centralization is not the only option. Still, it is a profound question how thoroughly the functionality of a market mechanism like Nasdaq, eBay, or SABRE can be distributed to buyer/seller peers.

North

For the economic historian Douglass North, an institution can be understood by analogy to the rules of a game [9]. An institution, in this sense, is conceptual structure that allows people to coordinate their activities. In particular, it creates incentives, such as the profit motive, that tend to channel participants' actions. North, like Hayek, takes markets as his major example, and he interprets history as a steady march toward free markets. He also (like Hayek) argues that the rules of the game change only slowly and incrementally, and he tries to explain how that evolution takes place.

For North, institutions do not imply the top-down imposition of inflexible rules. "Institutions" does not mean "organizations." Quite the contrary, institutions are distributed throughout the whole population. Banks, for instance, are organizations, but the institution of banking includes everybody who has a bank account. The institution is a web of relationships, including all of the customs, skills, and strategies that weave the web together. The motor of institutional evolution is not efficiency but self-interest. To understand

how institutions evolve, one must analyze the players and their strategies. But even when a single player is powerful enough to constitute a "center," it has its effects only by interacting with the interests and strategies of the other players.

Consider, for example, the institutional context in which the ARPANET arose [1]. In their attempt to create a decentralized computer network, the program managers at ARPA had an important advantage: they controlled the finances for a substantial research community. They consciously created incentives that would promote their goals. They required their contractors to use the ARPANET, and they drove the adoption of email by methods such as being accessible only through that medium. But ARPA did not succeed by imposing an alien way of life. To the contrary, it tried to amplify the cultural forms already present in the community. Nor did the architectures and institutions of ARPANET-based research evolve as ARPA had planned. The user community had little interest in ARPA's vision of resource-sharing. Instead, the network's growth was unexpectedly driven by its users' enthusiasm for email.

Thus, despite ARPANET's centralized institutional environment, North's vision of incremental evolution of institutional rules through the interaction of contending interests describes its history quite well. (Mueller [8] applies North's theory to the history of ICANN.) This history had subtle consequences for the decentralized architecture it produced. Because of ARPA's authority, everyone took for granted that ARPANET's user community was self-regulating. This institutional feature is reflected in the poor security of the Internet's email standards. When the Internet became a public network, the old assumptions no longer applied. Chronic security problems were the result. Likewise, institutional context will probably be crucial for the development of P2P architectures.

Commons

John Commons was a Progressive Era economist who eventually trained many of the leaders of the New Deal. Guided by his union background, Commons [3] viewed every institution as working rules defined by collective bargaining. After all, every institution defines social roles (doctor, patient, teacher, student, landlord, tenant, and so on), and each social role defines a community (for example, the community of doctors and the community of patients). Commons argues that each community develops its own culture and practices, which eventually become codified in law.

Commons, like Veblen, was an outlier in the history of political economy. He disagreed with Marx's vision of history as the inevitable victory of one social class, and

preferred a vision of collective bargaining among evenly matched groups. But he also differed from authors like Hayek and North, whose ideal society consists of little but the private dealings of individuals. In a subtle way Hayek's and North's worldviews are static: their utopia can only be achieved once institutional evolution stops, and they offer specific ideas about what this stopping-place should look like. Commons, by contrast, assumed that institutions evolve without end. The process of institutional change, therefore, lay at the center of his theory. His ideal was democracy, whether in government, industry, or any other institution, and he believed that collective bargaining in any context could evolve sophisticated working rules to govern individuals' dealings.

Commons' notion of collective bargaining sounds more centralized than he intended. His prototype was his experience of union-management bargaining, and he later applied that experience in developing policies such as workers' compensation by consulting with both unions and employers. But, as he understood, the picture of interest groups arguing across a table hardly captures the complexity of collective bargaining as a broad historical phenomenon. For one thing, interest groups participate in an "ecology of games" [4]—rule-making controversies in diverse venues, each interacting with the others in complex ways.

For example, consider the current war over music distribution. Napster caused a revolution in the institutions of music, and its subsequent decline should provoke reflection about that revolution's nature. Napster had a fatal flaw: although it afforded P2P sharing of music files, its centralized directory made it susceptible to legal attack. But Napster also had another, more subtle flaw: it did not provide a viable institution for allowing musicians to make a living. Some musicians can survive on live performances and merchandise, but most still rely on record sales. Of course, the record industry has few defenders as a mechanism for connecting musicians and audiences. But its dysfunctions arguably result from the intrinsic problems of marketing information goods. P2P file-sharing is an architecture looking for an institution, and any new institution will have to address these intrinsic problems.

The collective bargaining Napster has set in motion has at least three parties: musicians, fans, and record companies. As in every negotiation, each party has its own political problems—comprehending the situation, getting organized, settling differences, choosing representatives, coordinating actions, and so on. The negotiation takes place in many venues, including legislatures, courts, treaty organizations, contract negotiations, marketplaces, and standards organizations. New institutions will somehow result.

The post-Napster institutions of music distribution will presumably depend on new technologies. At the moment, most technical development is aimed at two models: P2P models that resist legal assaults and rights-management models that preserve existing economic models or migrate toward subscription models. It is unclear whether a thoroughly P2P architecture can survive, particularly if monopolies such as Microsoft change their own architectures to suit the record companies' needs. It is also unclear whether fans, with their own strategies, have any interest in the record companies' models. The most important question, however, is whether the new architectures and institutions of music provide musicians with reasonable career strategies. Their approach to collective bargaining is only now taking form.

Conclusion

What has been learned? Decentralized institutions do not imply decentralized architectures, or vice versa. The drive toward decentralized architectures need not serve the political purpose of decentralizing society. Architectures and institutions inevitably coevolve, and to the extent they can be designed, they should be designed together. The P2P movement understands that architecture is politics, but it should not assume that architecture is a substitute for politics. Radically improved information and communication technologies do open new possibilities for institutional change. To explore those possibilities, though, technologists will need better ideas about institutions. ■

REFERENCES

1. Abbate, J. *Inventing the Internet*. MIT Press, Cambridge, MA, 1999.
2. Alfred, D. and Chandler, Jr. *The Visible Hand: The Managerial Revolution in American Business*. Harvard University Press, Cambridge, MA, 1977.
3. Commons, J.R. *Institutional Economics: Its Place in Political Economy*. University of Wisconsin Press, Madison, WI, 1934.
4. Dutton, W.H. The ecology of games shaping communications policy. *Commun. Theory* 2, 4 (1992), 303–328.
5. Hayek, F.A. *Individualism and Economic Order*. University of Chicago Press, Chicago, 1963.
6. Hodgson, G.M. *Economics and Utopia: Why the Learning Economy Is Not the End of History*. Routledge, London, 1999.
7. Minar, N. and Hedlund, M. A network of peers: Peer-to-peer models through the history of the Internet. In A. Oram, Ed., *Peer-to-Peer: Harnessing the Power of Disruptive Technologies*, O'Reilly, 2001.
8. Mueller, M.L. *Ruling the Root: Internet Governance and the Taming of Cyberspace*. MIT Press, Cambridge, MA, 2002.
9. North, D.C. *Institutions, Institutional Change, and Economic Performance*. Cambridge University Press, Cambridge, MA 1990.
10. O'Reilly, T. Remaking the peer-to-peer meme. In A. Oram, Ed., *Peer-to-Peer: Harnessing the Power of Disruptive Technologies*. O'Reilly, 2001.
11. Saltzer, Reed, J.W., Reed, D.P. and Clark, D.D. End-to-end arguments in system design. *ACM Transactions in Computer Systems* 2, 4, (1984), 277–288.
12. Shirky, C. Listening to Napster, in A. Oram, Ed., *Peer-to-Peer: Harnessing the Power of Disruptive Technologies*. O'Reilly, 2001.

PHIL AGRE (pagre@ucla.edu) is an associate professor in the Department of Information Studies at the University of California, Los Angeles.