WORKSHOP ON NETWORK DESIGN AND SOCIETAL VALUES
SEPTEMBER 24-25 2008
ARLINGTON, VIRGINIA

DRAFT REPORT

INTRODUCTION

Digital electronic networks have emerged as one of the most powerful and exciting technologies of the late 20th and early 21st centuries, embodying and promoting wide-ranging societal and individual aspirations to create, produce, communicate, buy, sell, organize, connect, associate, educate, learn, entertain, campaign, and collaborate on a local, community, national, and global scale.

One mark of a great technology is its capacity to transform and be transformed. This we have witnessed in the relatively short lifespan of digital electronic networks, as societies have reacted to them and, in turn, shaped and reshaped them in multiple iterative cycles of mutual transformation. For scientists and engineers, the challenges are legion. In this document, however, we report on some of the complex interactions between network science and technology and societal values, focusing on moral, political, and sometimes also cultural values.

BACKGROUND

The broad community of network scientists and engineers, in collaboration with the NSF Network Science and Engineering (NetSE) program, poses this challenge: to develop the fundamental principles and methodical knowledge that will help us understand large, complex networks, and help us better design such networks in the future. The scope of NetSE ranges from design and development of network technologies, to "network science" and to the relationships between and among both of these and with people and societies. As a step toward this ambition, scholars and researchers, both inside and beyond traditional science and engineering, have been invited by NSF and the NetSE Council (an external community organization helping to refine the NetSE objective) to participate in a series of workshops to think about key issues and approaches.

In this context, on September 24-25, 2008, the workshop on Network Design and Societal Values assembled a group of scholars and researchers in the humanities, social sciences, law and policy as well as scientists and engineers to identify promising (fertile) research in the humanities, social sciences, law and policy, past and potential, that connects the study of moral and political values with computer and information system design, development, and deployment.[1] Although the focus of the workshop was specifically on networks, the workshop sought to bring to bear the wealth of expertise and past work on the complex, mutual interplay between design of technology and political and social life. Participants also acknowledged the tension at the heart of this task: the attempt to design better networks in order to accommodate and support uses which network decisions may enable or precipitate.

But identifying promising research in relevant non-technical fields was not the only goal. At least as important was to identify past and potential research that could speak beyond

[1] See Appendix I for a list of participants.

the communities of its authors' academic origins to network scientists and engineers as well – results, questions, approaches, literatures, cases, and issues that might be meaningful to scientists and engineers, that might even influence the design of computer and information networks (e.g., hardware and software). What might these be? How might decisions in network design usefully and systematically take them into consideration? And, by the same token, what hard problems in network science and engineering might successfully migrate onto the agenda of the humanistic, social, and political study of technology? How might these problems stir and energize these areas?

THE REPORT

The report is inspired by ideas emerging from brief presentations by Workshop participants and from the discussion following these presentations, where several salient themes gelled.[2] Going into the workshop, participants were asked to prepare remarks not only about their own work but reflecting a line of research or scholarship in which they conceived their work to fit. This placement did not need to track traditional disciplinary boundaries (e.g., the names of their home departments) but could be associated with a set of questions, a particular method, an object, or objects of study, a set of issues, an annual conference, etc. Participants were asked to reflect on how this line of work contributed to a landscape of study of networks and societal values, for scientists and engineers as well as the members of their communities. Participants from technical fields were asked to reflect on instances in which they had encountered problems which they understood to be socio-technical in nature, and prior collaborative experiences to address such issues.

Individual presentations and group discussion suggested that workshop findings would be more easily presented as a research landscape, characterized by several key dimensions, each defined by an open ended set of questions, rather than a research agenda, defined by a single list of questions. The dimensions that seemed best to capture relevant past research and exciting and valuable future research were: Past Work, Issues, Themes, Methods and Approaches, and Integrated Case Studies.

It bears repeating that the workshop's horizon was not on all interesting and worthwhile research on network technology through the lens of humanities and social sciences but on research in both fields that offered exciting potential for mutual influence.

DIMENSIONS OF THE LANDSCAPE

PAST WORK

Issues, questions, themes, methods, and cases raised at the workshop build upon a significant body of past and ongoing research across the disciplines. Workshop participants acknowledged, particularly, a large body of work on key issues, including: *Network architectures and politics; Identity, Identifiability and Anonymity; Access to Networks and Networking; Security; Privacy. [We should probably have a footnote or appendix with some examples.]*

[2] See Appendix II for the workshop agenda.

ISSUES

The workshop identified numerous issues. We acknowledge that those listed below are diverse in generality, size (of existing body of work), and scope, including some overlap. It is also important to note that disciplines vary in the ways they apply the identical label. Below are a sample of the issues that generated greatest interest and sometimes disagreement among workshop participants.

Security: We can study security from multiple perspectives. What does security mean to network researchers in computer science and engineering? What does it mean to political scientists, philosophers, sociologists? How can we translate these definitions across fields? What happens, for instance, when different fields 'securitize' an issue? What are we aiming for when we strive for security? How can we achieve it? What are the tradeoffs between security and other values (e.g., free speech)? Among all the actors (or agents) on the network, including individuals, institutions, and governments, is everyone's security of equal value? And how does the value pf security differ among users, institutions, and national governments themselves? [Clark, Nissenbaum, Ohm]

Identifiers and identities: How should these be defined? What's at stake in different (technical) choices? Why is identity often posed as a panacea? Can we test this proposition? How should identity online mesh with identity and identities in relation to other spheres of interaction, particularly in the relation of the individual to governments, financial institutions, and other corporate entities, such as merchants and service providers? Can we embed application layer solutions (e.g., eBay reputational systems, or social network identities) in a more general network architecture or design? Are there alternative approaches up and down the layers, and can these approaches successfully migrate? What sorts of collaboration would be necessary to achieve this? [Clark]

Openness: This is a term that has been used with various meanings in relation to networks (specifically, the Internet). One of the most important meanings, with technical and societal implications, is the capacity for everyone to *join* the network. In the case of the Internet, this means, at least in theory that any machine able to communicate in TCP/IP is able to join. (Of course, there are other, mostly economic barriers.) This contrasts with a scenario in which the protocol itself is not open and in this case one would need the permission of a controlling authority or gatekeeper. Important questions follow from this observation. What values are at stake? What does openness mean? What are the trade-offs in open and closed networks or protocols? Open networks may promote organic growth, but also suffer, in the case of malicious actors, from a lack of vetting or barring mechanisms. Open systems therefore often bring up issues of trust and individual accountability and thus identity online (as discussed ealier). [Li] Whereas closed networks may be deployed or administered by a central authority, a network in which users have no choice but to participate (and participate according to enforceable rules of use), open systems can only invite users to participate. What are the mechanisms and conditions that encourage or compel participation in open systems? Must the design of open networks take into account incentive structures? Can open systems force participation at a technical level, as in Skype? And can such participation be largely opaque to users themselves, again as in Skype?

Trust: A system that limits the opportunity for users to do harm to one another is not the same as one that achieves the same result based on trust. This is the difference between trustworthy technical artifacts (so-called trusted systems, as in secure banking)

and technology which enables people to trust one another. How can we incorporate social values in the design of networks that actually promotes and sustains sociality? Do social networks, based on voluntary associations among users, point toward a model for trust networks more generally? Is the trust that pervades such networks durable over time and across platforms and between layers? [Mulligan, Chun]

Mechanisms of regulation, control, or enforcement ("the handoff" question): Workshop participants agreed that behavioral constraints and affordances could be embedded in a network environment at various different junctures and layers. For example, they can be built into the technology, expressed in law and policy, through social norms, through incentives structures. The picture is even more complicated than this because even within these different junctures (or modes) there are various possibilities, and network design choices may produce unintended points of control. For example, technical constraints can be imposed at different layers (e.g., physical versus application) or following different strategies (e.g., through post-hoc auditing or front-end vetting). Choosing mechanisms and points of control is a technical matter but ethical and political implications should be carefully considered. This issue covers a potentially huge terrain and offers great possibilities for collaboration among different approaches. It can be tackled thematically and also through detailed case-studies. [Benkler, Mulligan]

Local and global: The Internet is touted as a global network but its value and meaning is often local (culturally, geographically). This requires study of networks embedded in a variety of contexts. Research might therefore address who appropriates a network for what purposes, and how different economic, social, political, and cultural contexts make such action possible. This research may draw from human-computer interaction, but may also adopt a more anthropological or sociological lens in examining the everyday and local uses of a network. [Grinter] Such insights may inform network design, particularly those attempts to develop a network that is sensitive to local variations in use and deployment. Can network design also build upon the general geographic distribution that tends to characterize social network membership? Can we develop networks that are optimized according to the spatial and social distribution of our likely network associates? Might we adjust our approach to network search, for example, given network information about geographic hotspots for certain query strings? How can or should scientists and engineers take local political contexts into consideration when designing the features of networks and network services? [Kleinberg]

Privacy: This covers a universe of questions and issues, including, for example, privacy online, what it means, how much do we want or need; what other values it protects; what other values is clashes with; how design may mitigate these conflicts. What opportunities for monitoring and measurement do network design decisions create? Are privacy concerns inherent in the architecture of the network, or do secondary technologies (such a mass storage for data retention) play a more important role? Whom do network architectures empower to monitor user behavior and information? The current architecture of the Internet, for instances, puts Internet Service Providers in a uniquely powerful position to monitor all the activity of its subscribers. What are the minimal features of a network that commercial service providers require (e.g., location-based IP)? How can we break apart information that we would find desirable for the network to reveal from that which it must necessarily produce? [Ohm]

Conditions of participation (related to several other issues above): This rather obscure title refers to a set of questions about expectations network users may reasonably have

about the powers they have when they join a network and, conversely, what users may experience as part of normal participation in a network. To what kinds of activities are users legitimately expected to submit as a condition of participation? Specifically, researchers may wish to study whether traditional notions of real property have analogs in network space. Do users have a right to object to unsolicited email as long as they have signed up for email, or to having their systems used as "zombies," or having search "bots" visit their websites. Could we imagine a network in which only consensual associates could exchange packets? To what degree do the conditions of participation of social networks already follow this model?

Motivations for participation: Why do people join and participate in a network? Noshir Contractor's work on the creation, maintenance, dissolution, and reconstitution of networks focused precisely on this role of motivations. Can we design networks that take into account the various motivations of their users? What defines a successful network from this perspective? Should networks adjust to users' motivations, and if so, how might networks determine or allow users to specify their respective motivations? What other criteria figure in the success of a network? Judith Olson's work on remote scientific collaborations, for instance, delineates the myriad factors that may obtain in top-down successful network-based pursuits. On the other hand, what are the motivations for voluntary, collaborative online activities? What, for instance, are the social motivations of commons-based production on the Internet? [Benkler] If the degree to which certain network structures enable production of this sort has become clearer, there still remains much to explore about the micro-foundations of cooperation and collaborative production in general. For instance, can we develop networks that promote cooperation through solidarity rather than by reward or punishment? Can network design decisions help cultivate voluntary participation and behavior that conforms to the norms of the community without recourse to punitive mechanisms or technical restrictions?

THEMES

Certain ideas seemed to crop up across discussion of several of the issues and case studies. They seemed more appropriately to be understood as themes, rather than as issues.

Visibility and transparency: The concepts of visibility and transparency are salient in two, respects. The first we might describe as individual exposure and self-presentation on a network—that is the degree to which a user can or must reveal information about him- or herself at different layers of a network (MAC address, IP address, application account, etc.). [Chun] Trust, for instance, often requires some degree of exposure, as in reputational systems or social networking sites. Visibility in this sense may also refer to the ability to communicate or reveal one's motivation for participation or collaboration (as discussed above). But these concepts have another meaning in a related context: the ability to examine the inner workings of a network design, protocol, or application. Take, for example, the design decision to allow Web users to view page sources. Or technologies or software that are not black boxed (or open, in the terminology of the above discussion), and thus leave users free to tinker. Transparency of this sort has emerged as political value among certain coders and stakeholders. Might it not only encourage non-engineers to peer into the inner workings of systems and software (with the understanding that code is law), but also compel outside stakeholders to contribute to the dynamic development of the code that may become law.

Incentives: Understanding the structure of incentives can shed light on relationships between architecture or design, on the one hand, and behavior or outcomes, on the other. An integrated study of existing incentives through empirical, ethnographic, historical, etc. methods is an important way of understanding what is already in place. One may also wish to disrupt, shape, or take advantage of naturally occurring incentive structures in order to achieve certain ends, for example, security or privacy, in the context of networks or network transactions. How might we determine the generalizability of an incentive structure of a specific network or application? Are incentive structures from one network or application appropriate, legitimate, or effective in another? [Contracter,Olson]

Networks as Experimental Environments: The Internet and Web have emerged as hugely important environments for studying individual and social behavior. There is plenty of scope for thinking about the needs and requirements of research online. Network engineers are also engaged in experimentation in such activities as PlanetLab and potentially GENI inspired systems. What is the relationship between those who intentionally and inadvertently use these systems and the designers and developers of these systems? Must networks users consent to participation? Is there something importantly different in the responsibilities designers and engineers have to users when the systems they put out for use are "experimental?" To what ethical code should academic network researchers hold themselves, and how might such a code compare to the one, if any, that obtains in commercial research? How can network engineers best communicate the value of their research to those who are likely to be involved in the experiment or later affected? Or. alternatively, if large-scale experimentation is simply not possible with consensual parties, should we set a grand challenge for network engineers and designers which asks that they determine how to do research on networks that itself solves the problem of network monitoring? [Peterson]

METHODS OR THEORETICAL FRAMEWORK

[We thought it might be good to provide one or two examples of articles or books that report on work following each of the respective methods or approaches. Help needed.]

Analytic:
Philosophy, political theory, social theory, legal theory, economics, comparative historical studies, critical theory (including a range of epistemological positions, from more traditional analytic philosophy to Actor Network Theory)

Empirical:

Qualitative: Interviews, surveys, fieldwork studies (in which researchers examine, for instance, networks in specific contexts of use or stakeholder incentives), content analysis

Quantitative: Statistical analysis, laboratory experiments, mathematical and agent-based modeling, pilot and deployment studies (following the building-deploying-learning-revising-rebuilding sequence)

In addition to these traditional methods and approaches to conducting research on the social, political, and ethical character of technologies, which may productively be applied to the case of networks, there are integrated approaches recently formulated specifically for the task of analyzing design for values and approaches to guiding design practice

taking values into consideration. These include “Values in Design” (VID) which generally refers to the study fine-grain design characteristics for values embodied in them or promoted or afforded by them. Values-at-Play, Value-Sensitive-Design, and Reflective Design include heuristics for taking values into consideration during the design process, that is, for taking values into consideration in the design practice.

An Exploratorium, a simulation environment in which researchers could explore mock-ups of different multi-layer models of networks noting the cross-layer impact of decisions, including user behavior and institutional design. For example, the Exploratorium might be used to simulate and explore different P2P schemes and their effect on both user and institutional behavior, or further explore empirical findings regarding incentive structures based on solidarity rather than rewards or punishments by adjusting institutional or technical parameters. [Olson, Contractor]

INTEGRATED CASE STUDIES

There is an important place for integrated case studies. In general, these would be rich multidisciplinary studies of events, mechanisms, applications, architectures, etc. relating to networks.

Web search: One example discussed at the workshop was search, search in networks (social search, web search), including, for example, algorithm design and privacy. Why do we take for granted the current model? Are they the best we can manage? Must algorithms tuned through machine learning be opaque to policy analysis? Is this a problem for values in design? [Kleinberg, Nissenbaum, Ohm]

Technology adoption by government agencies: Do agencies view technology as a procurement or policy question? What determines the perspective different agencies take, and what are the effect of this decision on the primacy of values in the adoption process? Which procedures open up the most productive spaces for discussion of values? [Mulligan]

Standards setting: Standard setting is an important site for determining socially relevant design features. There is often little reward for outsiders to participate in standard settings meetings. Why is this so and what about these meetings dissuades outside participation? How can outside stakeholder enter into or contribute to the debate? What are the social, bureaucratic, and epistemological conditions of participation? [Ohm, Mulligan, Elliot]

*Municipal wireless mesh network*s: A study of the deployment of a wireless mesh network in a municipality that calls upon engineers, social scientists, and policymakers. Researchers would consider the significance of local context and specific cultural, political, and motivational triggers.[Benkler]

Engineers’ response to assertions about values: How do engineers articulate the values at play in their selection of and approach to a technical problem? How do they respond to the assertion that values figure in their work? Do they resist this idea? Under what conditions do engineers reflect on values in design, and how might these reflections lead to different design choices? [Clark]

APPENDIX I

Workshop Participant Bios

APPENDIX II

Agenda