Explorations in Cyber International Relations

Massachusetts Institute of Technology Harvard University

Introduction

David D. Clark

Computer Science and Artificial Intelligence Laboratory
Massachusetts Institute of Technology

October 1, 2011

This material is based on work supported by the U.S. Office of Naval Research, Grant No. N00014-09-1-0597. Any opinions, findings, conclusions or recommendations therein are those of the author(s) and do not necessarily reflect the views of the Office of Naval Research.

Citation: Clark, D. D. (2011). Introduction. *Daedalus, 140*(4), 5–16.

Unique Resource Identifier: https://doi.org/10.1162/DAED_a_00111

Introduction

David D. Clark

DAVID D. CLARK, a Fellow of the American Academy since 2002, is a Senior Research Scientist at the MIT Computer Science and Artificial Intelligence Laboratory. He served as Chief Protocol Architect in the development of the Internet. His current research looks at redefining the architectural underpinnings of the Internet and the relation of technology and architecture to economic, societal, and policy considerations. He is helping the U.S. National Science Foundation organize its Future Internet Design program.

This issue is concerned with the experience of using the Internet: how its character shapes the user experience and how our collective online participation raises larger societal and political questions. For most of us, the Internet has become indispensable. Whether we are sending email messages or searching the Web, it is a part of our daily lives. It seems to bring powerful benefits, and thus we use it, but it also appears to bring risks, limitations, and frustrations, causing some to react with mixed emotions. People fear loss of privacy and misuse of personal information; they fear the corruption of their computers by malicious software (malware); they fear the possible loss of precious information now stored online; and they resent the complexity of using all this technology. Some people refuse to use computers and the Internet for exactly these reasons, leading us to ask, why is the Internet what it is? What, or who, shapes its character? Are there technical factors that define what can and cannot be done on the Internet? How do the motivations of designers influence the character of the Internet? Are we "locked in" to a constrained set of capabilities, or is the future of the Internet open to many possibilities? Through a variety of essays, this issue explores that set of questions.

The original design goal of the Internet was modest: to facilitate the remote sharing of expensive computing equipment, at a time when computing was expensive. But even before the Internet became operational in 1983, the notion of its power as a tool for people to interact among themselves had taken hold. The first method to emerge was email, followed by an explosion of options: websites for sharing content, blogs, instant messaging, and "chatting"; shared participation in virtual worlds; social networking sites such as Facebook and Twitter; and so on – a seemingly endless array of tools to interact, collaborate, communicate, and learn.

Questions that center on the user experience are sometimes lost in other debates that arise around the Internet and its future. From a corporate perspective, the Internet is to a large extent driven by commerce, that is, the business of selling. Users are commonly the buyers, and only sometimes (as with auction sites such as eBay or job-search sites) the sellers. Consumers consume: they buy physical objects, which are then delivered to their doors, and they buy virtual products that exist only in digital form, such as music, video content, and movies. Even when users are not actively buying, but are simply "cruising the Web" or using social media, much of what they see is financed by embedded advertising. In this limited view of the Internet, what is needed is a stable and predictable platform that appeals to a set of users affluent enough to have a credit card. But this outlook does not address other aspects of the diverse Internet experience, such as participation in civic discourse or politics, or the simple social process of interacting with friends.

In Washington today, the Internet is increasingly viewed through the lens of security. There is talk about cyber-war, cyber-espionage, and attacks on critical infrastructure. This perspective is not concerned directly with what good might be done online, but with preventing bad outcomes that might cripple the utility of the Internet, for the needs of both the nation and the individual – and of course, for the business of selling.

This issue focuses on the user experience and the Internet as a platform for the wide-ranging endeavors of society because these subjects are sometimes drowned out by the loud voice of selling and the shrill call for security. For many of us, our real hope for the Internet is this broad aspiration, even if it must be financed by commerce. This issue explores the aspects of the Internet that will make it a hospitable platform for socially oriented activities and asks what we can learn from observing how the Internet is used today. A broad view of the Internet takes us beyond the commercial to the space of culture, politics, and – dare we hope – toward a still fragmentary and fragile global civil society.

The positive and negative aspects of email illustrate some of the issues that we must consider in attempting to understand how to make the Internet a hospitable place. Email was the first application that allowed users to interact. In the early days of the Internet, the user community was small and rather homogeneous, and email was an effective mechanism for communication. As the user community expanded, the phenomenon of bulk unsolicited email, otherwise called spam, emerged. The original designers of email were perhaps a bit naive in thinking that all users would be virtuous, polite, and trustworthy, but there were also two conscious design choices that, in retrospect, led to the proliferation of spam. The first was our preference that email be an "open" system whereby any-

one could send to anyone, like the phone system. One could print one's email address on a business card or have it listed in an organization's directory so that others could find and use it. The second was our resistance to requiring verification of one's identity in order to send email messages; we did not like the implications of mandatory identity cards or "Internet driver's licenses." But this openness allowed users to forge their identities. Once spam emerged, we realized that we had no good means to control or discipline spammers because they often operated outside the legal jurisdiction of targeted recipients, and they devised a variety of tricks to avoid detection and deterrence.

A number of important points can be drawn from this example. The first is that *the Internet* and *email* are two different things. One way to explain this difference is by analogy to other systems, perhaps the most accessible being the "information highway" that emerged in the 1990s. The highway analogy is apt in one respect: the Internet itself is a transport infrastructure over which all sorts of applications run, just as a highway is an infrastructure over which all sorts of vehicles run. Thus, it is the Internet that permits email to exist, but it is the particular design of the email application that defines and constrains the user experience. When we talk about *the Internet*, we need to clarify whether we mean only the transport infrastructure or the total experience – infrastructure and applications – that users perceive.

In the same way that email is distinct from the Internet infrastructure on which it sits, the Internet's many different applications (such as email, the Web, games, or Internet telephony) are distinct from each other. Each contains specific design features that create a different context to shape the user experience; different applications can provide very different experiences.

A number of design features influence the collective social experience that the Internet provides. Here, I highlight three that are illustrated in the case of email and that factor into a number of the papers in this volume: *identity*, *trust*, and *controlling bad behavior*. How each of these considerations is approached will to a large extent define the character of the various experiences that we have when we interact with others across the Internet.

In the real world, we have many ways to manage and track identity. We get to know people face to face, and recognition of physical cues is enough to evoke our knowledge of who someone is. In more structured situations, we use identity credentials (such as driver's licenses or passports) or third-party introductions. When we communicate across the Internet, none of the face-to-face cues are available, and we usually lack the more structured credentials. On the Internet, it sometimes seems as if users run around with bags on their heads.

This situation begs the question of why the Internet does not have some sort of built-in identity mechanism that would allow users to be sure of whom they are talking to. The answer is that different applications create different kinds of shared experiences, which have different requirements for identity. If we were to add an identity mechanism to the infrastructure – to the transport layer that defines the Internet – then it would work the same way for all applications, just as the width of a road constrains all vehicles. It would be the same for a consumer completing a banking application and a provider of sensitive information about medical conditions. It would be the same for a dissident in a repressive country, an investigative reporter, or a stockbroker

dealing with a client. That uniformity does not seem to match the needs of Western-style society. In some cases, the communicating parties need strong evidence of identity (for example, when a customer deals with a bank); in others, strong anonymity is crucial (as when we try to protect political speech).

The desire to vary mechanisms by circumstance suggests that tools to manage identity should not be built into the infrastructure layer of the Internet but into the applications themselves. Indeed, some applications contain strong identity tools. For example, banks go to considerable lengths to ensure that they are talking to known clients, and credit card companies act as trusted third parties to identify buyers and sellers in an online purchase. But to impose a uniform strong set of identity mechanisms on the Internet itself would have many negative social consequences.

Consider the above example of email. Given that the prevalence of spam makes email problematic today, should those responsible for email now redesign it so that every sender identifies himself with a valid identity credential issued by a trusted third party? I would argue that this measure would be overkill; it would not match the actual requirements of email as a social interaction tool. When we are introduced to people, we normally do not ask to see their driver's licenses. We use a social process that has been well honed over the ages, called "getting to know them." Over time, we build up a model of who people are and of the extent to which they are trustworthy in the role for which we know them. This process can be used with email as well. Imagine that all email users on the Internet had the option of constructing a credential that certifies who they are. Technically, such a credential is easy to create using encryption, and many people have done so already. From a face-to-face perspective, this approach sounds a little odd: what good does it do for me to have a certificate in which I assert that I am me? In the context of the Internet, it prevents others from impersonating me. If the credential is properly constructed, there is no way for someone else to forge it unless that person breaks into my computer (another issue to be considered). Using such a certificate does not tell you much the first time I send an email message to you. (Of course, if it mattered, you could call me on the phone or otherwise ask me if indeed I am the person who sent the email message.) The second time I send you a message, or the tenth or the hundredth, you can be assured that you are having a conversation with the same person. You are, in the phrasing used above, "getting to know me." This is one possible approach to the problem of identity, and different applications, with different social contexts, will call for different approaches.

Identity, while important, is not an end in itself. Identity is a mechanism that allows us to deal with the other two issues listed above. This brings us to the second issue: *trust*.

As we interact with others in various contexts, we make complex and subtle judgments about trust. We assess whether the parties are trustworthy, whether there are constraints that will limit bad action, and whether we should be confident, cautious, or fearful. We assess strangers on a bus, we decide whether we are at risk of being cheated by a checkout clerk, and we judge whether our friends are more or less trustworthy in different roles. Sometimes we judge wrongly; we may be deliberately misled by a "confidence artist," or we may simply not know someone well enough. But the ability to make and rely on useful assessments

about trust is fundamental to a working society.

On the Internet, the cues we rely on to gauge trust face to face are weak, and our judgments are prone to failure. The interplay between identity and trust is clear. If we cannot know for sure with whom we are talking, if everyone has a bag on his or her head, it makes little sense to assign different levels of trust to different people. A malicious person can (and, in fact, will) pretend to be a good friend. The design of an identity system must take into consideration what sort of identity is needed for appropriate assessment of trust in a particular context. The application designer must know which identity cues would be useful for different applications.

Just as there are many ways to construct and track identity, there are many ways to assess trust. We need the Internet equivalent of being able to tell when we are "going to a bad neighborhood." Will this website infest my machine with malicious software? Will it attempt to steal information about me? Will this merchant defraud me? Many application-specific mechanisms have been put in place to deal with these questions. For example, eBay's reputation system permits buyers and sellers to evaluate each other. Credit card companies not only keep track of buyers and sellers, they cover losses from fraudulent charges. In effect, they act as insurance companies, which relieves the buyers and sellers from having to make as strong a trust assessment as they might have to otherwise. Certainly, it is in the financial interest of credit card companies to provide this service; the sharing of risk allows markets to function, and it is also possible to make a profit through the business of providing insurance.

One of the interesting trends on the Internet is rating sites, where users give ratings of everything from hotels and restaurants to clothing and movies. Current schemes may have flaws, but they signal an important transition from an isolated, individual Internet experience to one embedded in a shared social context. In the real world, most of what we do is rooted in a shared context, but the original image of the "personal computer," for use at home, alone, seemed to decouple our respective experiences. Designers of applications have had to reconstruct that ability to share experiences and generate understanding of the world through interactive processes prevalent offline.

Some "social interaction" schemes can be abused, just as spammers abuse email. Most rating sites do not demand that users give a strong verification of their identity. They may require their users to give some bits of information about themselves, but instead of being identified by name, users choose a "handle" or pseudonym. What is to prevent a user from creating multiple pseudonyms and posting scores of bogus reviews, positive or negative, to change the rating of something?

Perhaps we would be more comfortable with reviews and ratings provided by our friends, people we know and (to an adequate degree) trust. Network-based constructs, such as social networking sites like Facebook, allow us to relate to our friends online. They capture a robust aspect of identity because users (usually) link their online identity to friends that they know in the real world. Given that these sites function on a basis of strong identity, not just pseudonyms, they might serve as the foundation for a rating scheme that allows the user to place a higher degree of trust in the ratings.

Not all users are nice or trustworthy. Internet applications must be constructed to detect and deal with "bad apples."

What are the options? If the law has been broken, perhaps the law enforcement tools that are open to the government can be used. But what if people online are simply rude or disruptive? How can a community protect itself? Spammers disrupt blogs by putting their spam messages into the comment sections of blogs. Disruptive players (called "griefers") interfere with multiplayer games through behavior or tactics that are irritating to other players.

A clear response to such behavior is shunning or expulsion from the community. The question for an application designer is whether the ability to shun or expel a user should be part of the system. Again, the issue of identity is key. If the application requires that users provide a strong indication of identity that is hard to forge or replicate, then a user can be ejected. Games that require users to sign up with a credit card can reject the card, which means that expelled players can return only as many times as they have different cards. A credit card company can refuse to serve a merchant, or refuse to authorize a payment to an overextended purchaser. On social networking sites such as Facebook, if the users have invested a great deal of effort constructing an identity that is linked to the identity of friends, ejection would be a painful punishment. But if a system does not require a user to present a strong form of identity, as many do not, then a user ejected under one pseudonym can obtain another and return.

The construction of online identity is an important aspect of forming a stable community. On the one hand, demands for strong confirmation of "real" identity may chill certain sorts of valid behavior, from political speech to searching for information on sensitive health issues. On the other hand, weak identity may make it hard to detect misbehavior (such as "ballot stuffing" on rating sites) or to eject misbehaving users.

As we begin to explore the experience of using the Internet, we might start by asking: who uses the Internet, and for what purposes? Who does not use the Internet, and why? In his essay, John B. Horrigan draws on survey data he gathered for the Pew Internet & American Life Project, and more recently, for the Federal Communications Commission's National Broadband Plan. Based on recent data, about two-thirds of American homes have broadband access, and people use the Internet for a wide and growing range of purposes, including sending email messages, using the Web, making or researching purchases, gathering news and weather information, watching a video or listening to music and radio, banking, playing games, and connecting with friends using social media tools.

On the other hand, about 22 percent of surveyed homes report that they do not use the Internet at all, citing reasons such as cost, inadequate digital literacy, lack of relevance, or deficient service in their area. The data reveal widespread concern about loss or misuse of personal information; 45 percent of non-users cite fears of bad things that might happen online.

More detailed data from Pew (reported elsewhere in this volume) make clear that the pool of non-users is not homogeneous across society. Non-users tend to be older and of lower socioeconomic status: the poor, the less educated, and the elderly are less likely to partake in the Internet experience. Horrigan observes that as more and more aspects of society move online, the costs of nonparticipation increase, to both non-users and society at large. Nonparticipation online can limit job opportunities – with 80 percent of Fortune 500 companies accepting only

online job applications – as well as access to online government services or health information. Horrigan concludes that society must address barriers to using the Internet, which are not just lack of hardware, but lack of mastery of the increasingly complex skills needed to participate: what he calls *digital literacy*. As the demands for skill level rise, the costs of exclusion may become increasingly significant.

A common fear is the loss of privacy and the misuse of personal information. Today, the rules about privacy are spelled out in often long, confusing "privacy policies" or "consent forms" offered by various providers of network services. In her essay, Helen Nissenbaum rejects this approach. Central to her argument is the observation that cyberspace is not a distinct space with its own distinct norms. Much of what we do on the Web (or on the Net generally) is a reflection of something we do in the real world. Norms from that context, including privacy standards, should be expected to hold in the equivalent online context. But currently there is no recognition of context and implied norms; thus, the privacy consent form must carry the total burden of defining the expectations of the parties who participate. To the extent that a policy tries to capture nuances, it becomes overlong and incomprehensible; to the extent that it aims for brevity and readability, it describes only the general nature of the policy and omits the details that matter in practice.

There are well-understood contexts in which all parties understand the norms that apply. Health care is governed both by laws and by commonly understood norms of behavior. Banking, whether online or offline, is similarly governed by both law and custom. Nissenbaum suggests that many other online behaviors could be understood in terms of prior offline analogues. For example, using a search engine might be analogous to using a library card catalog, which has a strong tradition of freedom from observation and tracking. Even if the online experience is somewhat novel, we can often find real-world analogues.

The crux of Nissenbaum's argument – that the online experience does not take place in a homogeneous and unique context but in a range of contexts that will develop different customary norms and governing laws – can be extended to attributes other than privacy. As I noted above, individual contexts will call for distinct approaches to identity. Coye Cheshire looks at the concept of trust online: how users decide if a service is trustworthy, whether to trust individuals they encounter online, and whether they can rely on the network and the services provided over it. He explores the meanings of *trust* and *trustworthy* in different contexts, observing that in instances where the risk is low, users will be willing to proceed in the face of considerable uncertainty about whether a website, a service, or an individual is trustworthy. A restaurant review may be malicious or hyped, but its accuracy is only minimally consequential for a prospective diner. In cases where the potential risk is high, tools are put in place to minimize uncertainty. Online banking bears a potential high risk, but banks have gone to considerable lengths to remove uncertainty from transactions and give users a high level of confidence that their banking services are trustworthy. Cheshire notes that mechanisms to enforce constraints on behavior (so that users can proceed without developing trust in the other parties) erode trust and the mechanisms by which it arises. Trust can emerge only in a context of ongoing interaction among parties where betrayal is possible. Cheshire's analysis of trust (and the distinction

with *trustworthy*) draws on the one hand from a range of writings on the subject, and on the other hand from experiments involving users in online contexts. He argues that in the future, the Internet will depend on social forms and institutional arrangements as much as technologies and systems. The Internet *is* the real world.

Fear of bad experiences online is an issue for users and a barrier for nonusers. Three essays in this volume deal with the problem of system and network security: protecting users, their computers, and the network from attack by malicious parties. The term *security* covers a range of concerns, including attacks by criminals on servers storing sensitive information, attempts to break into and subvert personal computers, and espionage carried out by states and powerful private-sector actors. Again, the essays collected in this volume focus on topics that are relevant to the experience of the individual user; they are less concerned with the potential of cyber-war and more so with the events that give users pause in their daily activities.

Vinton G. Cerf catalogs the many forms that these perils can take, providing insight into the roots of system insecurity as well as institutional approaches to improvement. Hazards include theft of personal information, spam, and the capture of one's computer by a remote operator, who then uses the computer to launch spam attacks on other users, or to flood a target site on the network with traffic to overload and disable it (a so-called denial of service attack). Given that perils can arise from both malicious acts and accidents, Cerf introduces the term *cyber-safety* to widen the scope of our objective beyond the more narrowly defined *cybersecurity*. Using a number of metaphors, including biology (viruses and infection), real-life analogues to understand the online experience (books versus e-books), and comparisons with offline mechanisms of protection (police and fire departments), he sketches the landscape of risk and response.

Deirdre K. Mulligan and Fred B. Schneider propose a new rationale by which society can improve its overall security posture. They first review past approaches to improving online security and examine why these approaches have failed. The attempt to provide security by building entirely invulnerable systems is simply impractical: today's systems are too complex, and the required level of effort would be too costly. Efforts to characterize the security problem as one of risk assessment and management (as we do in the offline world, using tools such as insurance) fail because we lack the tools and methods to quantify online risk. Finally, attempts to improve the landscape of security by using tools of deterrence to discourage misbehavior fail because we lack effective means of attribution and coherent means to pursue attackers across the jurisdictions of different states.

With this analysis as background, Mulligan and Schneider suggest a different way to think about improving security: through a new doctrine they call *public cybersecurity*. Their doctrine views the framework of public health and public health institutions as a model for cybersecurity. Just as good health is a benefit to all of society that must arise from the health of individuals, overall online security will improve by means of the steps individual users take to keep their own user contexts secure. But the benefit to any one user may not seem significant enough to justify his investment of effort and money into improving his own security. Security, like health, is a public issue, not an individual one. Thus, Mulligan and Schneider explore how the analogy of public health can be used to better

understand a large number of online issues, including system development, online surveillance, keeping systems up to date (installing "patches"), and isolation and quarantine of systems. Using the analogy to public health, this new public cybersecurity doctrine envisions a rational balance between the public interest in improved security and the rights of the individual.

L. Jean Camp considers the explicit question of whether and how we can motivate individual users to contribute to improved overall security. She uses two theoretical framings to explore this question. The first is *peer production*, in which users self-organize to create information (or other desired outcomes). She argues that users can be motivated to self-organize in ways that produce better system security, if the security challenge can be framed as a set of discrete tasks for which users can self-select based on skills and proclivity. She offers several examples, involving both technically skilled and ordinary users. The second regards the Internet as a common good, or a virtual *commons*. Using criteria developed by political economist Elinor Ostrom, she explores how the security problem can be framed in a way that allows users to self-regulate the commons. These two theories help model and define the circumstances under which user-centric efforts can be effective.

Several essays explore specific classes of behavior on the Internet. R. Kelly Garrett and Paul Resnick examine the experience of getting news and opinion online, questioning the hypothesis that personalization of news, made possible by the Internet, leads to increased political fragmentation. They reject the necessity of this outcome: personalization can take many forms, they observe, with different implications for social outcomes. Research suggests that people crave opinion reinforcement but do not go out of their way to avoid diverse viewpoints. If news is personalized along ideological lines, mirroring the ideologically segmented world of cable news today, it could indeed lead to increased fragmentation. By contrast, if personalization is used to expose willing readers to a range of viewpoints, selected perhaps for quality and thoughtfulness rather than bias, the result could expose readers to a more balanced selection of material. Research suggests that readers would be open to this sort of personalization. Narrow partisan channels *force* people to choose, but it is not clear that this is what people would prefer if given the choice.

The authors observe that "the technology and how people use it are still malleable; subtle architectural changes could have far-reaching implications for future news consumption patterns. [This] will require effort and creativity. . . . [T]echnologies that produce desirably diverse news streams may not emerge naturally." In understanding the Internet and the experiences that it provides, this observation is critical. As Garrett and Resnick and other essays remind us, the Internet is a built artifact. It is the way it is because people designed and built it to be that way. Thus, the future will be defined by those who choose to step up and design it. Originally, the designers of the Internet and its early applications (such as email) were researchers funded by the federal government. But with the success of the Internet, most of this effort has migrated to the commercial sector. And the motivations of the commercial sector may not perfectly align with one or another vision of preferred social outcomes.

In their essay, Kay Lehman Schlozman, Sidney Verba, and Henry E. Brady explore another specific class of online behavior: participation in the political process. A long-term concern with politi-

cal equality compelled the authors to understand whether the Internet might lower barriers to various sorts of political participation. Using survey data from the Pew Internet & American Life Project, they asked whether the Internet has changed political involvement in fundamental ways.

The survey focused on political participation as a function of socioeconomic status (SES) and age. Observing the striking power and durability of SES-based political inequality, the authors conclude that the Internet is not the "great leveler" that some optimists might hope. Online political behavior shows trends that are similar to traditional offline behavior. Significantly, these trends are not mirrored in other behaviors, such as participation in social networks, where SES is much less a predictor of participation.

As the Pew study confirms, given that those with lower SES are less likely to be on the Internet at all, that group suffers a double barrier: lack of access and the traditional bias against participation. Age is another factor: older populations are less likely to be online; even if older users online seem to be politically active, the overall level of online participation in older populations is low.

Schlozman, Verba, and Brady also look at new forms of political activity that arise online, such as political blogs and social networks. While the survey reveals that younger respondents (under age twenty-five) are heavy users of this technology, the authors do not find strong reasons to conclude that these new technologies may lead to a change in the nature of political participation. But they note that the Internet and its applications are young, and technical design decisions as well as changing user behavior are unpredictable.

Lee Sproull looks at a different sort of online activity: *prosocial* behavior; that is, activities intended to help people other than oneself, such as volunteering and supporting charities. (Volunteer activities include service projects, health support groups, peer production of information, and citizen science.) She catalogs various forms of observed online prosocial behavior and provides estimates of its prevalence. She notes that while all major types of online prosocial projects share a small number of attributes that derive from the underlying network technology and communications applications, each context for online behavior is a symbolically differentiated place on the Net, and different people seek out different places. Sproull then discusses the features of the online context that can facilitate or encourage such participation: the modularity and granularity of the task (how the work of one user is scoped, specified, and then aggregated), a site's social structure, and techniques to motivate participants. She also discusses the nuanced role of identity and trust in shaping and motivating the participant experience.

Finally, the essay by Yochai Benkler provides an analysis of two current events that illustrate how various actors, including the government and powerful private-sector players, engage to shape what happens on the Internet. The two cases are WikiLeaks and the ongoing struggle by the holders of copyright to repress the sharing of unauthorized copies of their material. The two cases have much in common: first, they both revolve around attempts to prevent access to material in the context of an open Internet that makes unregulated access the norm; and second, they demonstrate a complex and tangled interplay between the public and private sectors. His concern is that in both these cases, the approaches put in place would allow the blocking of access without the nor-

mal protections of U.S. law. More generally, his essay reminds us that not everyone has the same aspirations for the Internet, and that the future will be shaped by a tussle among those who care enough to advocate for their objectives.

A number of common themes run through the essays in this issue. First, the term *cyberspace* is potentially misleading on two grounds. The term suggests that the Internet is a distinct "space" or "place" to which we go online. Rather, our experience using the Internet is not separate and disconnected from our offline experiences. Much of what we do on the Internet has a close relationship to our offline behaviors. Additionally, the experience of using the Internet is not homogeneous and subject to consistent norms. Various aspects of using the Internet will differ in important ways, including the norms and expectations about behavior; the degree of uncertainty, risk, and significance; and the nature of the resulting interactions.

The Internet, as a low-level platform that supports a range of applications, is not the technology that creates or defines the user experience. The Internet itself, as distinct from the applications that run on it, was originally seen as a neutral platform intended to support as many patterns of interaction as possible. This generality implies that it supports both "good" and "bad" patterns of interaction, and "good" and "bad" experiences. It is the applications that have been built on top of this platform – email, the Web, Facebook, Internet telephony, search engines – that define the user experience. Each application is its own context with its own affordances and constraints.

The Internet is a built artifact. It is designed and engineered. One must not think of the Internet as fixed and exogenous; it can change and evolve, sometimes rapidly. As we consider the limits and benefits of the current Internet, we should not think only like observers or analysts, but also like designers and engineers. By most measures, the Internet is very young, and its designers have much to learn, including how to facilitate predictable, safe, and rich human interaction.

An important corollary to this last point is that the future character of the Internet will depend in large part on which parties choose to invest in shaping it. To the extent that the Internet is being designed and built by actors with commercial, profit-seeking motives, we may not see the emergence of applications that shape the social experience in ways leading to better civic engagement, prosocial activities, or news sources that offer personalization without polarization. The designers of the early Internet, mostly supported by research grants from the government, may have had different motivations from the designers of today.

Not only is the Internet a built artifact, it is a constraining artifact. In the offline world, people interact in a complex and open-ended environment that offers many different cues and signals for social interaction. In the online world of an Internet application, the context of interaction is bounded and limited by what the application designer has provided. It is a closed system, except to the extent that the application designer has intentionally created tools whereby users can evolve the social context. As a result, we should not view the resulting patterns of human behavior as only socially emergent. Online behavior is shaped by the capabilities and constraints of applications as well as by the socially centered, human factors that influence how people use those applications.

Past experience has taught us that while social interaction mediated by the mechanics of software is a constrained

and limited experience compared to its offline counterpart, the online version has powerful advantages. Thus, we can expect the Internet to be a compelling platform for interaction and engagement. Good design can help mitigate and compensate for limitations. As the essays contained herein remind us, the Internet is not a fixed artifact, but evolving and flexible. There are many possible futures for the Internet, depending on which actors choose to define that future. One of the goals of this issue is to stimulate debate about what that future should be.